

Supervisor Engine 6T Software Configuration Guide, Release 15.5SY

Cisco IOS Release 15.5SY

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

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

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

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

CONTENTS

Preface i

Audience i

Related Documentation i

Conventions i

Obtaining Documentation, Obtaining Support, and Security Guidelines ii

Product Overview 1-1

Supervisor Engine 6T Ports 1-1

USB Type A Port 1-1

Console ports 1-2

USB Type B Port 1-2

System Reset Button 1-2

Management Port 1-2

Management Port Operation Modes 1-3

LED Indicators 1-3

Status LED 1-3

System ID LED 1-4

Management Port LED 1-4

Status LED 1-4

System ID LED 1-4

Determining System Hardware Capacity 1-5

Command-Line Interfaces 2-1

Accessing the CLI 2-1

Accessing the CLI through the EIA/TIA-232 Console Interface 2-2

Accessing the CLI through Supervisor Engine 6T Software Configuration Guide, Release 15.4SYrough Telnet 2-2

Accessing the CLI through USB 2-3

Performing Command Line Processing 2-3

Performing History Substitution 2-4

Cisco IOS Command Modes 2-4

Displaying a List of Cisco IOS Commands and Syntax 2-6

Securing the CLI 2-7

ROM-Monitor Command-Line Interface 2-7

Smart Port Macros 3-1

Prerequisites for Smart Port Macros 3-1

Restrictions for Smart Port Macros 3-2

Information About Smart Port Macros	3-3
Information about Cisco-Provided Smart Port Macros	3-3
Information about User-Created Smart Port Macros	3-4
Default Settings for Smart Port Macros	3-4
How to Configure Smart Port Macros	3-4
Using the Cisco-Provided Smart Port Macros	3-4
Creating Smart Port Macros	3-13
Verifying the Smart Port Macro Configuration	3-15
Virtual Switching Systems	4-1
Prerequisites for VSS	4-1
Restrictions for VSS	4-2
General VSS Restrictions	4-2
VSL Restrictions	4-2
Multichassis EtherChannel (MEC) Restrictions	4-2
Dual-Active Detection Restrictions	4-3
Information About Virtual Switching Systems	4-4
VSS Overview	4-4
VSS Redundancy	4-12
Multichassis EtherChannels	4-15
Packet Handling	4-18
System Monitoring	4-22
Convert Port Mode	4-23
Dual-Active Detection	4-27
VSS Initialization	4-29
Default Settings for VSS	4-31
How to Configure a VSS	4-31
Configuring Easy VSS	4-31
Converting to a VSS	4-32
Displaying VSS Information	4-38
Converting a VSS to Standalone Chassis	4-39
Configuring VSS Parameters	4-40
Configuring Multichassis EtherChannels	4-49
Configuring Port Load Share Deferral on the Peer Switch	4-49
Configuring Dual-Active Detection	4-50
How to Upgrade a VSS	4-54
Performing a Fast Software Upgrade of a VSS	4-54
Performing an Enhanced Fast Software Upgrade of a VSS	4-55

Enhanced Fast Software Upgrade	5-1
Prerequisites for eFSU	5-2
Restrictions for eFSU	5-2
Information About eFSU	5-3
eFSU Operation	5-3
Outage Time and Support Considerations	5-4
Reserving Module Memory	5-4
Error Handling for eFSU Preload	5-5
Default Settings for eFSU	5-5
How to Perform an eFSU	5-5
eFSU Summarized Procedure	5-6
Preparing for the Upgrade	5-6
Copying the New Software Image	5-8
Loading the New Software onto the Standby Supervisor Engine	5-8
Displaying the Maximum Outage Time for Installed Modules (Optional)	5-10
Forcing a Switchover from Active to Standby	5-10
Accepting the New Software Version and Stopping the Rollback Process (Optional)	5-12
Committing the New Software to the Standby	5-12
Verifying the Software Installation	5-12
Aborting the Upgrade Process	5-13
How to Upgrade a Non-eFSU Image to an eFSU Image	5-14
Fast Software Upgrade	6-1
Stateful Switchover (SSO)	7-1
Prerequisites for SSO	7-1
Restrictions for SSO	7-2
General Restrictions	7-2
Configuration Mode Restrictions	7-2
Switchover Process Restrictions	7-2
Information About SSO	7-3
SSO Overview	7-3
SSO Operation	7-5
Route Processor Synchronization	7-6
SSO Operation	7-8
SSO-Aware Features	7-10
Default Settings for SSO	7-10
How to Configure SSO	7-10
Troubleshooting SSO	7-11

Possible SSO Problem Situations	7-11
SSO Troubleshooting	7-12
Verifying the SSO Configuration	7-12
Verifying that SSO Is Configured	7-12
Verifying that SSO Is Operating on the Device	7-13
Verifying SSO Features	7-13
Configuration Examples for SSO	7-16
Nonstop Forwarding (NSF)	8-1
Prerequisites for NSF	8-1
Restrictions for NSF	8-2
General Restrictions	8-2
Restrictions for BGP NSF	8-2
Restrictions for EIGRP NSF	8-2
Restrictions for OSPF NSF	8-2
Restrictions for IS-IS NSF	8-2
Restrictions for IPv6 NSF	8-3
Information About NSF	8-3
NSF Overview	8-3
Feature Interaction with NSF	8-4
Default Settings for NSF	8-9
How to Configure NSF	8-9
Configuring and Verifying BGP for NSF	8-9
Configuring and Verifying EIGRP NSF	8-10
Configuring and Verifying OSPF NSF	8-12
Configuring and Verifying IS-IS NSF	8-13
Troubleshooting Cisco Nonstop Forwarding	8-14
Configuration Examples for NSF	8-15
Example: Configuring BGP NSF	8-15
Example: Configuring BGP NSF Neighbor Device	8-15
Example: Verifying BGP NSF	8-16
Example: Configuring EIGRP NSF Converge Timer	8-16
Example: EIGRP Graceful-Restart Purge-Time Timer Configuration	8-16
Example: Configuring EIGRP NSF Route-Hold Timer	8-17
Example: Configuring EIGRP NSF Signal Timer	8-17
Example: Verifying EIGRP NSF	8-17
Example: Disabling EIGRP NSF Support	8-18
Example: Configuring OSPF NSF	8-18
Example: Verifying OSPF NSF	8-18

Example: Configuring IS-IS NSF	8-19
Example: Verifying IS-IS NSF	8-19
Route Processor Redundancy (RPR)	9-1
Prerequisites for RPR	9-1
Restrictions for RPR	9-1
General RPR Restrictions	9-2
Hardware Restrictions for RPR	9-2
Information About RPR	9-2
Supervisor Engine Redundancy Overview	9-3
RPR Operation	9-3
Supervisor Engine Configuration Synchronization	9-4
Default Settings for RPR	9-4
How to Configure RPR	9-4
Synchronizing the Supervisor Engine Configurations	9-4
Displaying the Redundancy States	9-4
Copying Files to the RP	9-5
Interface Configuration	10-1
Information About Interface Configuration	10-2
How to Configure a Range of Interfaces	10-2
How to Define and Use Interface-Range Macros	10-2
How to Configure Optional Interface Features	10-3
Configuring Ethernet Interface Speed and Duplex Mode	10-3
Configuring Jumbo Frame Support	10-6
Configuring IEEE 802.3x Flow Control	10-9
Configuring the Port Debounce Timer	10-10
Information About Online Insertion and Removal	10-11
How to Monitor and Maintain Interfaces	10-12
Monitoring Interface Status	10-12
Clearing Counters on an Interface	10-12
Resetting an Interface	10-13
Shutting Down and Restarting an Interface	10-13
How to Check Cable Status with the TDR	10-14
How to Convert Line Card to Operate in 40G Mode	10-14
Overview of Cisco CVR-4SFP10G-QSFP Adapter	10-15
Compatible Line Cards	10-16

UniDirectional Link Detection (UDLD) 11-1

Prerequisites for UDLD 11-1

Restrictions for UDLD 11-1

Information About UDLD 11-2

UDLD Overview 11-2

UDLD Aggressive Mode 11-3

Fast UDLD 11-4

Default Settings for UDLD 11-4

How to Configure UDLD 11-4

Enabling UDLD Globally 11-5

Enabling UDLD on LAN Interfaces 11-5

Disabling UDLD on Nonfiber-Optic LAN Interfaces 11-5

Disabling UDLD on Fiber-Optic LAN Interfaces 11-6

Configuring the UDLD Probe Message Interval 11-6

Configuring Fast UDLD 11-6

Resetting Disabled LAN Interfaces 11-7

Instant Access (IA) 12-1

Prerequisites for Instant Access 12-1

Restrictions for Instant Access 12-2

Information About Instant Access 12-6

Default Settings for Instant Access 12-6

How to Configure Instant Access 12-6

Configure Instant Access Staggered Initialization Mode 12-7

Enable FEX Auto-Config 12-7

Enable IA Client Preprovisioning 12-7

Configure Instant Access Port-Channel Interfaces 12-8

Configure Instant Access Channel Groups 12-9

Identify Connected IA Client Stack Modules 12-9

Renumbering FEX Switch-ID 12-11

Configure IA Clients 12-13

Display or Clear SDP and SRP Traffic 12-13

Upgrade FEX using In-Service Software Upgrade (ISSU) Process 12-13

Configure Optional Parameters for an IA Client 12-14

Configuring EnergyWise 13-1

Power Management 14-1

Power Management Overview 14-1

How to Enable or Disable Power Redundancy 14-2

How to Power Modules Off and On	14-3
How to Display System Power Status	14-3
How to Power Cycle Modules	14-5
Environmental Monitoring	15-1
Environmental Monitoring Overview	15-1
How to Determine Sensor Temperature Thresholds	15-2
How to Monitor the System Environmental Status	15-5
Information About LED Environmental Indications	15-7
Online Diagnostics	16-1
Prerequisites for Online Diagnostics	16-1
Restrictions for Online Diagnostics	16-1
Information About Online Diagnostics	16-2
Default Settings for Online Diagnostics	16-2
How to Configure Online Diagnostics	16-2
Setting Bootup Online Diagnostics Level	16-2
Configuring On-Demand Online Diagnostics	16-4
Scheduling Online Diagnostics	16-6
Configuring Health-Monitoring Diagnostics	16-6
How to Run Online Diagnostic Tests	16-7
Overview of Diagnostic Test Operation	16-8
Starting and Stopping Online Diagnostic Tests	16-8
Running All Online Diagnostic Tests	16-9
Displaying Online Diagnostic Tests and Test Results	16-9
How to Perform Memory Tests	16-29
How to Perform a Diagnostic Sanity Check	16-29
Onboard Failure Logging (OBFL)	17-1
Prerequisites for OBFL	17-1
Restrictions for OBFL	17-2
Information About OBFL	17-2
Overview of OBFL	17-2
Information about Data Collected by OBFL	17-2
Default Settings for OBFL	17-9
Enabling OBFL	17-10
Configuration Examples for OBFL	17-11
Enabling OBFL Message Logging: Example	17-11
OBFL Message Log: Example	17-11

OBFL Component Uptime Report: Example 17-11

OBFL Report for a Specific Time: Example 17-12

Switch Fabric Functionality 18-1

Prerequisites for Switch Fabric Functionality 18-1

Restrictions for Switch Fabric Functionality 18-1

Information About the Switch Fabric Functionality 18-2

Default Settings for Switch Fabric Functionality 18-2

Monitoring the Switch Fabric Functionality 18-2

Cisco IP Phone Support 19-1

Prerequisites for Cisco IP Phone Support 19-1

Restrictions for Cisco IP Phone Support 19-1

Information About Cisco IP Phone Support 19-2

Cisco IP Phone Connections 19-2

Cisco IP Phone Voice Traffic 19-3

Cisco IP Phone Data Traffic 19-4

Other Cisco IP Phone Features 19-4

Default Setting for Cisco IP Phone Support 19-4

How to Configure Cisco IP Phone Support 19-5

Configuring Voice Traffic Support 19-5

Configuring Data Traffic Support 19-6

Power over Ethernet (PoE) Support 20-1

Prerequisites for PoE 20-1

Restrictions for PoE 20-1

Information About PoE 20-2

Device Roles 20-2

PoE Overview 20-2

CPD-Based PoE Management 20-3

Inline Power IEEE Power Classification Override 20-4

LLDP Inline Power Negotiation for PoE+ (IEEE 802.3at) 20-4

How to Configure PoE Support 20-4

Displaying PoE Status 20-5

Configuring Per-Port PoE Support 20-5

Configuring PoE Power Priority 20-6

Configuring PoE Monitoring and Policing 20-8

Disabling LLDP Power Negotiation (IEEE 802.3at) 20-8

LAN Ports for Layer 2 Switching	21-1
Prerequisites for Layer 2 LAN Interfaces	21-1
Restrictions for Layer 2 LAN Interfaces	21-2
Information About Layer 2 Switching	21-2
Information about Layer 2 Ethernet Switching	21-2
Information about VLAN Trunks	21-4
Layer 2 LAN Port Modes	21-4
Default Settings for Layer 2 LAN Interfaces	21-5
How to Configure LAN Interfaces for Layer 2 Switching	21-5
Configuring a LAN Port for Layer 2 Switching	21-6
Enabling Out-of-Band MAC Address Table Synchronization	21-6
Configuring MAC Address Table Notification	21-7
Configuring a Layer 2 Switching Port as a Trunk	21-8
Configuring a LAN Interface as a Layer 2 Access Port	21-14
Configuring a Custom IEEE 802.1Q EtherType Field Value	21-15
Flex Links	22-1
Prerequisites for Flex Links	22-1
Restrictions for Flex Links	22-2
Information About Flex Links	22-2
Default Settings for Flex Links	22-4
How to Configure Flex Links	22-4
Monitoring Flex Links	22-6
EtherChannels	23-1
Prerequisites for EtherChannels	23-1
Restrictions for EtherChannels	23-2
Information About EtherChannels	23-3
EtherChannel Feature Overview	23-3
Information about EtherChannel Configuration	23-3
Information about LACP 1:1 Redundancy	23-6
Information about Port Channel Interfaces	23-7
Information about Load Balancing	23-7
Default Settings for EtherChannels	23-7
How to Configure EtherChannels	23-7
Configuring Port Channel Logical Interfaces	23-8
Configuring Channel Groups	23-9
Configuring the LACP System Priority and System ID	23-11
Configuring EtherChannel Load Balancing	23-11

Configuring the EtherChannel Hash-Distribution Algorithm	23-12
Configuring the EtherChannel Min-Links Feature	23-13
Configuring LACP 1:1 Redundancy	23-14
Configuring Auto Interleaved Port Priority For LACP Port Channels	23-15
Configuring LACP Port-Channel Standalone Disable	23-16
IEEE 802.1ak MVRP and MRP	24-1
Prerequisites for IEEE 802.1ak MVRP and MRP	24-1
Restrictions for IEEE 802.1ak MVRP and MRP	24-2
Information About IEEE 802.1ak MVRP and MRP	24-2
Overview	24-2
Dynamic VLAN Creation	24-4
MVRP Interoperability with VTP	24-4
MVRP Interoperation with Non-Cisco Devices	24-6
MVRP Interoperability with Other Software Features and Protocols	24-6
Default Settings for IEEE 802.1ak MVRP and MRP	24-8
How to Configure IEEE 802.1ak MVRP and MRP	24-8
Enabling MVRP	24-8
Enabling Automatic Detection of MAC Addresses	24-9
Enabling MVRP Dynamic VLAN Creation	24-9
Changing the MVRP Registrar State	24-10
Troubleshooting the MVRP Configuration	24-10
Configuration Examples for IEEE 802.1ak MVRP and MRP	24-11
Enabling MVRP	24-11
Enabling MVRP Automatic Detection of MAC Addresses	24-11
Enabling Dynamic VLAN Creation	24-12
Changing the MVRP Registrar State	24-12
VLAN Trunking Protocol (VTP)	25-1
Prerequisites for VTP	25-1
Restrictions for VTP	25-1
Information About VTP	25-2
VTP Overview	25-3
VTP Domains	25-3
VTP Modes	25-4
VTP Advertisements	25-4
VTP Authentication	25-5
VTP Version 2	25-5
VTP Version 3	25-6

VTP Pruning	25-7
VLAN Interaction	25-9
Default Settings for VTP	25-9
How to Configure VTP	25-10
Configuring VTP Global Parameters	25-10
Configuring the VTP Mode	25-15
Configuring VTP Mode on a Per-Port Basis	25-16
Displaying VTP Statistics	25-17
Virtual Local Area Networks (VLANs)	26-1
Prerequisites for VLANs	26-1
Restrictions for VLANs	26-2
Information About VLANs	26-2
VLAN Overview	26-2
VLAN Ranges	26-2
Default Settings for VLANs	26-3
How to Configure VLANs	26-4
Configurable VLAN Parameters	26-4
VLAN Locking	26-4
Creating or Modifying an Ethernet VLAN	26-5
Assigning a Layer 2 LAN Interface to a VLAN	26-6
Configuring the Internal VLAN Allocation Policy	26-6
Configuring VLAN Translation	26-7
Saving VLAN Information	26-9
Private VLANs	27-1
Prerequisites for Private VLANs	27-1
Restrictions for Private VLANs	27-2
Secondary and Primary VLANs	27-2
Private VLAN Ports	27-4
Limitations with Other Features	27-4
Information About Private VLANs	27-5
Private VLAN Domains	27-5
Private VLAN Ports	27-7
Primary, Isolated, and Community VLANs	27-7
Private VLAN Port Isolation	27-8
IP Addressing Scheme with Private VLANs	27-8
Private VLANs Across Multiple Switches	27-9
Private VLAN Interaction with Other Features	27-9

Default Settings for Private VLANs	27-10
How to Configure Private VLANs	27-10
Configuring a VLAN as a Private VLAN	27-11
Associating Secondary VLANs with a Primary VLAN	27-12
Mapping Secondary VLANs to the Layer 3 VLAN Interface of a Primary VLAN	27-13
Configuring a Layer 2 Interface as a Private VLAN Host Port	27-14
Configuring a Layer 2 Interface as a Private VLAN Promiscuous Port	27-15
Monitoring Private VLANs	27-16

Private Hosts 28-1

Prerequisites for Private Hosts	28-1
Restrictions for Private Hosts	28-1
General Private Host Restrictions	28-2
Private Host ACL Restrictions	28-2
Private Host VLAN on Trunk Port Restrictions	28-3
Private Host Interaction with Other Features	28-3
Private Host Spoofing Protection	28-3
Private Host Multicast Operation	28-4
Information About Private Hosts	28-4
Private Hosts Overview	28-4
Isolating Hosts in a VLAN	28-4
Restricting Traffic Flow (Using Private Hosts Port Mode and PACLs)	28-5
Port ACLs	28-7
Default Settings for Private Hosts	28-8
How to Configure Private Hosts	28-8
Configuration Summary	28-8
Detailed Configuration Steps	28-9
Configuration Examples	28-10

IEEE 802.1Q Tunneling 29-1

Prerequisites for 802.1Q Tunneling	29-1
Restrictions for 802.1Q Tunneling	29-1
Information About 802.1Q Tunneling	29-4
Default Settings for 802.1Q Tunneling	29-5
How to Configure 802.1Q Tunneling	29-6
Configuring 802.1Q Tunnel Ports	29-6
Configuring the Switch to Tag Native VLAN Traffic	29-6

Layer 2 Protocol Tunneling	30-1
Prerequisites for Layer 2 Protocol Tunneling	30-1
Restrictions for Layer 2 Protocol Tunneling	30-1
Information About Layer 2 Protocol Tunneling	30-2
Default Settings for Layer 2 Protocol Tunneling	30-2
How to Configure Layer 2 Protocol Tunneling	30-3
Spanning Tree Protocols	31-1
Prerequisites for Spanning Tree Protocols	31-1
Restrictions for Spanning Tree Protocols	31-2
Information About Spanning Tree Protocols	31-2
Information about STP	31-2
Information about IEEE 802.1w RSTP	31-13
Information about MST	31-18
Detecting Unidirectional Link Failure	31-25
Default Settings for Spanning Tree Protocols	31-25
Default STP Configuration	31-25
Default MST Configuration	31-26
How to Configure Spanning Tree Protocols	31-26
Configuring STP	31-26
Configuring MST	31-37
Optional STP Features	32-1
PortFast	32-2
Information about PortFast	32-2
Enabling PortFast	32-2
Bridge Assurance	32-4
Information about Bridge Assurance	32-4
Enabling Bridge Assurance	32-7
BPDU Guard	32-7
Information about BPDU Guard	32-7
Enabling BPDU Guard	32-7
PortFast Edge BPDU Filtering	32-9
Information about PortFast Edge BPDU Filtering	32-9
Enabling PortFast Edge BPDU Filtering	32-10
UplinkFast	32-11
Information about UplinkFast	32-11
Enabling UplinkFast	32-12
BackboneFast	32-13

Information about BackboneFast	32-13
Enabling BackboneFast	32-15
EtherChannel Guard	32-16
Information about EtherChannel Guard	32-16
Enabling EtherChannel Guard	32-16
Root Guard	32-17
Information about Root Guard	32-17
Enabling Root Guard	32-17
Loop Guard	32-17
Information about Loop Guard	32-17
Enabling Loop Guard	32-19
PVST Simulation	32-20
Information about PVST Simulation	32-20
Configuring PVST Simulation	32-21
Verifying the Optional STP Features	32-21
Using the show spanning-tree Commands	32-22
Examples of the show spanning-tree Commands	32-22
IP Unicast Layer 3 Switching	33-1
Prerequisites for Hardware Layer 3 Switching	33-1
Restrictions for Hardware Layer 3 Switching	33-2
Information About Layer 3 Switching	33-2
Hardware Layer 3 Switching	33-2
Layer 3-Switched Packet Rewrite	33-2
Default Settings for Hardware Layer 3 Switching	33-4
How to Configure Hardware Layer 3 Switching	33-4
Displaying Hardware Layer 3 Switching Statistics	33-5
Policy-Based Routing (PBR)	34-1
Prerequisites for PBR	34-1
Restrictions for PBR	34-2
Information About PBR	34-2
PBR Overview	34-2
PBR Recursive Next Hop for IPv4 Traffic	34-3
Default Settings for PBR	34-3
How to Configure PBR	34-3
Configuring PBR	34-4
Configuring Local PBR	34-5
Configuring PBR Recursive Next Hop	34-5

Configuration Examples for PBR	34-7
Equal Access Example	34-7
Differing Next Hops Example	34-8
Recursive Next-Hop IP Address: Example	34-8
Layer 3 Interfaces	35-1
Restrictions for Layer 3 Interfaces	35-1
How to Configure Subinterfaces on Layer 3 Interfaces	35-3
Unidirectional Ethernet (UDE) and Unidirectional Link Routing (UDLR)	36-1
Prerequisites for UDE and UDLR	36-1
Restrictions for UDE and UDLR	36-2
UDE Restrictions	36-2
UDLR Back-Channel Tunnel Restrictions	36-3
Information About UDE and UDLR	36-3
UDE and UDLR Overview	36-3
Information about UDE	36-3
Information about UDLR	36-4
Default Settings for UDE and UDLR	36-4
How to Configure UDE and UDLR	36-5
Configuring UDE	36-5
Configuring UDLR	36-6
Multiprotocol Label Switching (MPLS)	37-1
Prerequisites for MPLS	37-1
Restrictions for MPLS	37-1
Information About MPLS	37-2
MPLS Overview	37-2
IP to MPLS	37-4
MPLS to MPLS	37-4
MPLS to IP	37-4
MPLS VPN Forwarding	37-5
Recirculation	37-5
Hardware Supported Features	37-5
Supported MPLS Features	37-6
Default Settings for MPLS	37-7
How to Configure MPLS Features	37-7
Configuring MPLS	37-7
Configuring MUX-UNI Support on LAN Cards	37-7

Configuration Examples for MPLS 37-9

MPLS VPN Support 38-1

Prerequisites for MPLS VPN 38-1

Restrictions for MPLS VPN 38-2

Information About MPLS VPN Support 38-2

How to Configure MPLS VPNs 38-3

Configuration Example for MPLS VPNs 38-4

Ethernet over MPLS (EoMPLS) 39-1

Prerequisites for EoMPLS 39-1

Restrictions for EoMPLS 39-2

Information About EoMPLS 39-3

AToM Overview 39-3

EoMPLS Overview 39-3

Default Settings for EoMPLS 39-3

How to Configure EoMPLS 39-4

Configuring VLAN-Based EoMPLS 39-4

Configuring Port-Based EoMPLS 39-6

Configuring A-VPLS 40-1

Prerequisites for A-VPLS 40-1

Restrictions for A-VPLS 40-2

Information About A-VPLS 40-2

How to Configure A-VPLS 40-3

Enabling Load-Balancing with ECMP and FAT Pseudowires 40-3

Enabling Port-Channel Load-Balancing 40-3

Explicitly Specifying the PE Routers As Part of Virtual Ethernet Interface Configuration 40-4

Configuring an MPLS Traffic Engineering Tunnel 40-5

Configuring a GRE Tunnel 40-5

Routed Pseudo-Wire (RPW) and Routed VPLS 40-8

Ethernet Virtual Connections (EVCs) 41-1

Prerequisites for EVCs 41-1

Restrictions for EVCs 41-2

Information About EVCs 41-3

EVC Overview 41-3

Ethernet Flow Points 41-4

Service Instances and EFPs 41-4

Encapsulation (Flexible Service Mapping)	41-5
EFPs and MSTP	41-7
Bridge Domains	41-7
Rewrite Operations	41-9
Layer 3 and Layer 4 ACL Support	41-9
Advanced Frame Manipulation	41-9
Egress Frame Filtering	41-9
Default Settings for EVCs	41-9
How to Configure EVCs	41-10
Monitoring EVCs	41-14
Layer 2 over Multipoint GRE (L2omGRE)	42-1
Prerequisites for L2omGRE	42-1
Restrictions for L2omGRE	42-2
Information About L2omGRE	42-2
Default Settings for L2omGRE	42-3
How to Configure L2omGRE	42-3
Configuring a Loopback Interface	42-3
Configuring an mGRE Tunnel Interface	42-3
Configuring a VLAN Interface	42-4
L2omGRE Configuration Examples	42-5
Verifying the L2omGRE Configuration	42-5
Campus Fabric	43-1
Understanding Fabric Domain Elements	43-1
Campus Fabric Configuration Guidelines	43-2
How to Configure Fabric Overlay	43-2
Configuring Fabric Edge Devices	43-2
Configuring Control-Plane Devices	43-3
Configuring Border Devices	43-4
Example: Configuring Fabric Edge, Border and Control-Plane Devices	43-6
IPv4 Multicast Layer 3 Features	44-1
Prerequisites for IPv4 Multicast Layer 3 Features	44-1
Restrictions for IPv4 Multicast Layer 3 Features	44-1
Information About IPv4 Multicast Layer 3 Features	44-2
IPv4 Multicast Layer 3 Features Overview	44-2
Distributed MRIB and MFIB Infrastructure	44-3
Multicast Layer 3 Hardware Features Entries	44-4

Layer 3-Switched Multicast Statistics	44-4
Layer 3-Switched Multicast Packet Rewrite	44-5
Replication Modes	44-5
Local Egress Replication Mode	44-6
PIM-SM hardware register support	44-6
PIM-SM hardware SPT-switchover support	44-6
Control Plane Policing (CoPP)	44-7
Non-RPF Traffic Processing	44-7
Multicast Boundary	44-7
IPv4 Bidirectional PIM	44-8
Supported Multicast Features	44-8
Default Settings for IPv4 Multicast Layer 3 Features	44-15
How to Configure IPv4 Multicast Layer 3 Features	44-15
Enabling IPv4 Multicast Routing Globally	44-16
Enabling IPv4 PIM on Layer 3 Interfaces	44-16
Enabling IP Multicast Layer 3 Switching on Layer 3 Interfaces	44-16
Enabling IP MFIB forwarding on Layer 3 Interfaces	44-17
Configuring the Replication Mode	44-18
Configuring Multicast Boundary	44-18
Verifying Local Egress Replication	44-19
Displaying IPv4 Multicast PIM-SM register tunnel information	44-20
Displaying the IPv4 Multicast Routing Table	44-20
Displaying IPv4 MRIB Information	44-21
Displaying IPv4 MFIB Information	44-22
Viewing Directly Connected Entries	44-23
Displaying IPv4 Hardware Switching Information	44-24
Displaying IPv4 CoPP Information	44-25
Source-Specific Multicast with IGMPv3, IGMP v3lite, and URD	44-26
Configuring IPv4 Bidirectional PIM	44-26
Enabling IPv4 Bidirectional PIM Globally	44-27
Configuring the Rendezvous Point for IPv4 Bidirectional PIM Groups	44-27
Displaying IPv4 Bidirectional PIM Information	44-27
Using IPv4 Debug Commands	44-31
Redundancy for Multicast Traffic	44-31
IGMP Snooping for IPv4 Multicast Traffic	45-1
Prerequisites for IGMP Snooping	45-1
Restrictions for IGMP Snooping	45-1
General IGMP Snooping Restrictions	45-2
IGMP Snooping Querier Restrictions	45-2

Information About IGMP Snooping	45-2
IGMP Snooping Overview	45-3
Joining a Multicast Group	45-3
Leaving a Multicast Group	45-5
Information about the IGMP Snooping Querier	45-6
Information about IGMP Version 3 Support	45-6
Default Settings for IGMP Snooping	45-8
How to Configure IGMP Snooping	45-8
Enabling the IGMP Snooping Querier	45-9
Enabling IGMP Snooping	45-9
Configuring the IGMP Snooping Lookup Method	45-11
Configuring a Static Connection to a Multicast Receiver	45-11
Configuring a Multicast Router Port Statically	45-12
Configuring the IGMP Snooping Query Interval	45-12
Enabling IGMP Snooping Immediate-Leave Processing	45-13
Configuring IGMPv3 Snooping Explicit Host Tracking	45-13
Displaying IGMP Snooping Information	45-14
PIM Snooping	46-1
Prerequisites for PIM Snooping	46-1
Restrictions for PIM Snooping	46-2
Information About PIM Snooping	46-2
Default Settings for PIM Snooping	46-4
How to Configure PIM Snooping	46-5
Enabling PIM Snooping Globally	46-5
Enabling PIM Snooping in a VLAN	46-5
Disabling PIM Snooping Designated-Router Flooding	46-6
Multicast VLAN Registration (MVR)	47-1
Restrictions for MVR	47-1
Restrictions for MVR	47-1
Information About MVR	47-2
MVR Overview	47-2
Using MVR in a Multicast Television Application	47-3
Default MVR Configuration	47-5
How to Configure MVR	47-5
Configuring MVR Global Parameters	47-5
Configuring MVR Interfaces	47-6
Clearing MVR Counters	47-8

Displaying MVR Information 47-8

IPv4 IGMP Filtering 48-1

Prerequisites for IGMP Filtering 48-1

Restrictions for IGMP Filtering 48-1

Information About IGMP Filtering 48-2

IGMP Filtering Overview 48-3

IGMP Filter Precedence 48-4

Default Settings for IGMP Filtering 48-4

How to Configure IGMP Filters 48-4

Configuring IGMP Group and Channel Access Control 48-4

Configuring IGMP Group and Channel Limits 48-5

Configuring IGMP Version Filtering 48-5

Clearing IGMP Filtering Statistics 48-6

Verifying the IGMP Filtering Configuration 48-6

Displaying IGMP Filtering Configuration 48-6

Displaying IGMP Filtering Statistics 48-7

Configuration Examples for IGMP Filtering 48-8

IPv4 Router Guard 49-1

Prerequisites for Router Guard 49-1

Restrictions for Router Guard 49-1

Information About Router Guard 49-2

Default Settings for Router Guard 49-2

How to Configure Router Guard 49-2

Enabling Router Guard Globally 49-3

Disabling Router Guard on Ports 49-3

Clearing Router Guard Statistics 49-3

Verifying the Router Guard Configuration 49-4

Displaying Router Guard Configuration 49-4

Displaying Router Guard Interfaces 49-5

IPv4 Multicast VPN Support 50-1

Prerequisites for mVPNs 50-1

Restrictions for mVPNs 50-1

General Restrictions 50-2

mVPN with L3VPN over mGRE Restrictions 50-3

Information About mVPN 50-3

mVPN Overview 50-3

Multicast Routing and Forwarding and Multicast Domains	50-4
Multicast Distribution Trees	50-4
Multicast Tunnel Interfaces	50-7
PE Router Routing Table Support for mVPN	50-8
Multicast Distributed Switching Support	50-8
Hardware-Assisted IPv4 Multicast	50-8
Information About mVPN with L3VPN over mGRE	50-9
Default Settings for mVPNs	50-10
How to Configure mVPNs	50-11
Configuring a Multicast VPN Routing and Forwarding Instance	50-11
Configuring Multicast VRF Routing	50-16
Configuring Interfaces for Multicast Routing to Support mVPN	50-20
Configuring mVPN with L3VPN over mGRE	50-22
Configuration Examples for mVPNs	50-27
mVPN Configuration with Default MDTs Only	50-27
mVPN Configuration with Default and Data MDTs	50-29
Verifying the mVPN with L3VPN over mGRE Configuration	50-32
Configuration Sequence for mVPN with L3VPN over mGRE	50-33
IPv6 Multicast Support	51-1
Prerequisites for IPv6 Multicast	51-1
Restrictions for IPv6 Multicast	51-1
Information About IPv6 Multicast Support	51-2
Hardware-Supported IPv6 Layer 3 Multicast Features	51-2
Partially Hardware-Supported IPv6 Layer 3 Multicast Features	51-3
Software-Supported IPv6 Layer 3 Multicast Features	51-3
Unsupported IPv6 Layer 3 Multicast Features	51-3
How to Configure IPv6 Multicast Support	51-4
Verifying the IPv6 Multicast Layer 3 Configuration	51-4
Verifying MFIB Clients	51-5
Displaying the Switching Capability	51-5
Verifying the (S,G) Forwarding Capability	51-5
Verifying the (*,G) Forwarding Capability	51-5
Verifying the Subnet Entry Support Status	51-5
Verifying the Current Replication Mode	51-5
Displaying the Replication Mode Auto-Detection Status	51-5
Displaying the Replication Mode Capabilities	51-6
Displaying Subnet Entries	51-6
Displaying the IPv6 Multicast Summary	51-6

Displaying the NetFlow Hardware Forwarding Count	51-7
Displaying the FIB Hardware Bridging and Drop Counts	51-7
Displaying the Shared and Well-Known Hardware Adjacency Counters	51-8
IPv6 MLD Snooping	52-1
Prerequisites for MLD Snooping	52-1
Restrictions for MLD Snooping	52-2
General MLD Snooping Restrictions	52-2
MLD Snooping Querier Restrictions	52-2
Information About MLD Snooping	52-3
MLD Snooping Overview	52-3
MLD Messages	52-4
Source-Based Filtering	52-4
Explicit Host Tracking	52-4
MLD Snooping Proxy Reporting	52-5
Joining an IPv6 Multicast Group	52-5
Leaving a Multicast Group	52-7
Information about the MLD Snooping Querier	52-8
Default MLD Snooping Configuration	52-9
How to Configure MLD Snooping	52-9
Enabling the MLD Snooping Querier	52-10
Configuring the MLD Snooping Query Interval	52-10
Enabling MLD Snooping	52-11
Configuring a Static Connection to a Multicast Receiver	52-12
Configuring a Multicast Router Port Statically	52-12
Enabling Fast-Leave Processing	52-12
Enabling SSM Safe Reporting	52-13
Configuring Explicit Host Tracking	52-13
Configuring Report Suppression	52-14
Verifying the MLD Snooping Configuration	52-14
NetFlow Hardware Support	53-1
Prerequisites for NetFlow Hardware Support	53-1
Restrictions for NetFlow Hardware Support	53-1
Information About NetFlow Hardware Support	53-2
Default Settings for NetFlow Hardware Support	53-2
How to Configure NetFlow Hardware Support	53-2
Configuring Inactive Flow Aging	53-3
Configuring Fast Aging	53-3

Configuring Active Flow Aging	53-4
Verifying the NetFlow Table Aging Configuration	53-4
System Event Archive (SEA)	54-1
Information About the System Event Archive	54-1
How to Display the SEA Logging System	54-2
How to Copy the SEA To Another Device	54-3
Backplane Traffic Monitoring	55-1
Prerequisites for Backplane Traffic Monitoring	55-1
Restrictions for Backplane Traffic Monitoring	55-2
Information About Traffic Monitoring	55-2
Default Settings for Backplane Traffic Monitoring	55-2
How to Configure Backplane Traffic Monitoring	55-3
Local SPAN, RSPAN, and ERSPAN	56-1
Prerequisites for Local SPAN, RSPAN, and ERSPAN	56-1
Restrictions for Local SPAN, RSPAN, and ERSPAN	56-2
Feature Incompatibilities	56-2
Local SPAN, RSPAN, and ERSPAN Session Limits	56-3
Local SPAN, RSPAN, and ERSPAN Interface Limits	56-3
General Restrictions for Local SPAN, RSPAN, and ERSPAN	56-3
Restrictions for VSPAN	56-5
Restrictions for RSPAN	56-5
Restrictions for ERSPAN	56-6
Restrictions for Distributed Egress SPAN Mode	56-7
Information About Local SPAN, RSPAN, and ERSPAN	56-7
Local SPAN, RSPAN, and ERSPAN Overview	56-7
Local SPAN, RSPAN, and ERSPAN Sources	56-11
Local SPAN, RSPAN, and ERSPAN Destinations	56-12
Default Settings for Local SPAN, RSPAN, and ERSPAN	56-12
How to Configure Local SPAN, RSPAN, and ERSPAN	56-12
Configuring a Destination as an Unconditional Trunk (Optional)	56-13
Configuring Destination Trunk VLAN Filtering (Optional)	56-13
Configuring Destination Port Permit Lists (Optional)	56-15
Configuring the Egress SPAN Mode (Optional)	56-15
Configuring Local SPAN	56-16
Configuring RSPAN	56-20
Configuring ERSPAN	56-26

Configuring Source VLAN Filtering in Global Configuration Mode	56-30
Verifying the SPAN Configuration	56-31
Configuration Examples for SPAN	56-31
SNMP IfIndex Persistence	57-1
Prerequisites for SNMP IfIndex Persistence	57-1
Restrictions for SNMP IfIndex Persistence	57-1
Information About SNMP IfIndex Persistence	57-2
Default Settings for SNMP IfIndex Persistence	57-2
How to Configure SNMP IfIndex Persistence	57-2
Enabling SNMP IfIndex Persistence Globally	57-2
Disabling SNMP IfIndex Persistence Globally	57-3
Enabling and Disabling SNMP IfIndex Persistence on Specific Interfaces	57-3
Clearing SNMP IfIndex Persistence Configuration from a Specific Interface	57-4
Top-N Reports	58-1
Prerequisites for Top-N Reports	58-1
Restrictions for Top-N Reports	58-1
Information About Top-N Reports	58-2
Top-N Reports Overview	58-2
Top-N Reports Operation	58-2
Default Settings for Top-N Reports	58-3
How to Use Top-N Reports	58-3
Enabling Top-N Reports Creation	58-3
Displaying Top-N Reports	58-4
Clearing Top-N Reports	58-5
Layer 2 Traceroute Utility	59-1
Prerequisites for the Layer 2 Traceroute Utility	59-1
Restrictions for the Layer 2 Traceroute Utility	59-1
Information About the Layer 2 Traceroute Utility	59-2
How to Use the Layer 2 Traceroute Utility	59-3
Mini Protocol Analyzer	60-1
Prerequisites for the Mini Protocol Analyzer	60-1
Restrictions for the Mini Protocol Analyzer	60-1
Information About the Mini Protocol Analyzer	60-2
How to Configure the Mini Protocol Analyzer	60-2
Configuring a Capture Session	60-2

Filtering the Packets to be Captured	60-4
Starting and Stopping a Capture	60-5
Displaying and Exporting the Capture Buffer	60-7
Configuration Examples for the Mini Protocol Analyzer	60-7
General Configuration Examples	60-8
Filtering Configuration Examples	60-9
Operation Examples	60-10
Display Examples	60-10
Restrictions for PFC QoS	61-1
General Guidelines	61-2
PFC and DFC Guidelines	61-4
Class Map Command Restrictions	61-5
Policy Map Class Command Restrictions	61-5
Supported Granularity for CIR and PIR Rate Values	61-5
Supported Granularity for CIR and PIR Token Bucket Sizes	61-6
IP Precedence and DSCP Values	61-7
PFC QoS Overview	62-1
Classification, Marking, and Policing	63-1
Information About Classification, Marking, and Policing Policies	63-1
Classification, Marking, and Policing Policy Overview	63-1
Traffic Classification	63-2
Traffic Marking	63-3
Information about Policing	63-4
How to Configure Classification, Marking, and Policing Policies	63-7
Enabling Distributed Aggregate Policing	63-7
Configuring a Class Map	63-8
Configuring a Policy Map	63-9
Attaching a Policy Map to an Interface	63-17
Configuring Dynamic Per-Session Attachment of a Policy Map	63-19
Policy-Based Queueing	64-1
Prerequisites for Policy-Based Queueing	64-1
Restrictions for Policy-Based Queueing	64-2
Information About Policy-Based Queueing	64-4
Port-Based Queue Types	64-4
Queueing Policies	64-8
How to Configure Policy-Based Queueing	64-10

Configuring a Queueing Policy Class Map	64-11
Verifying a Queueing Policy Class Map	64-11
Configuring Queueing Policy Maps	64-11
Verifying a Queueing Policy Map	64-17
Attaching a Queueing Policy Map to an Interface	64-17
Configuration Examples for Policy-Based Queueing	64-18
Queueing Policy Sample Configuration	64-18
Queueing Policy Commands Supported by Each Queue Type	64-19
Queueing Policy Commands Sample Configurations for Each Queue Type	64-33
Queueing on FEX	64-45
QoS Global and Interface Options	65-1
How to Configure the Ingress LAN Port CoS Value	65-2
How to Configure Egress DSCP Mutation	65-3
Configuring Named DSCP Mutation Maps	65-3
Attaching an Egress DSCP Mutation Map to an Interface	65-4
How to Configure Ingress CoS Mutation on IEEE 802.1Q Tunnel Ports	65-4
Ingress CoS Mutation Configuration Guidelines and Restrictions	65-4
Configuring Ingress CoS Mutation Maps	65-6
Applying Ingress CoS Mutation Maps to IEEE 802.1Q Tunnel Ports	65-6
How to Configure DSCP Value Maps	65-7
Mapping Received CoS Values to Internal DSCP Values	65-7
Mapping Received IP Precedence Values to Internal DSCP Values	65-7
Configuring DSCP Markdown Values	65-8
Mapping Internal DSCP Values to Egress CoS Values	65-9
How to Configure Trusted Boundary with Cisco Device Verification	65-10
Legacy Configuration Procedures for Queueing-Only Mode	65-11
Legacy Configuration Procedures for VLAN-Based PFC QoS on Layer 2 LAN Ports	65-12
Legacy Configuration Procedures for Port Trust State	65-13
Legacy Configuration Procedures for DSCP-Based Queue Mapping	65-14
Enabling DSCP-Based Queue Mapping	65-14
Configuring Ingress DSCP-Based Queue Mapping	65-15
Mapping DSCP Values to Standard Receive-Queue Thresholds	65-15
Mapping DSCP Values to Standard Transmit-Queue Thresholds	65-16
Mapping DSCP Values to the Transmit Strict-Priority Queue	65-18
AutoQoS	66-1
Prerequisites for AutoQoS	66-1
Restrictions for AutoQoS	66-2

Information About AutoQoS	66-2
AutoQoS Support for a Cisco IP Phone	66-3
AutoQoS Support for Cisco IP Communicator	66-3
AutoQoS Support for Marked Traffic	66-4
Default Settings for AutoQoS	66-4
How to Configure AutoQoS	66-4
Configuring AutoQoS Support for a Cisco IP Phone	66-5
Configuring AutoQoS Support for Cisco IP Communicator	66-6
Configuring AutoQoS Support for Marked Traffic	66-7
MPLS QoS	67-1
Terminology	67-2
MPLS QoS Features	67-2
MPLS Experimental Field	67-3
Trust	67-3
Classification	67-3
Policing and Marking	67-3
Preserving IP ToS	67-4
EXP Mutation	67-4
MPLS DiffServ Tunneling Modes	67-4
MPLS QoS Overview	67-4
Specifying the QoS in the IP Precedence Field	67-4
MPLS QoS	67-5
MPLS Topology Overview	67-5
LERs at the Input Edge of an MPLS Network	67-6
LSRs in the Core of an MPLS Network	67-6
LERs at the Output Edge of an MPLS Network	67-7
LERs at the EoMPLS Edge	67-7
LERs at the IP Edge (MPLS, MPLS VPN)	67-8
LSRs at the MPLS Core	67-11
MPLS QoS Default Configuration	67-13
MPLS QoS Commands	67-15
MPLS QoS Restrictions	67-15
How to Configure MPLS QoS	67-16
Enabling Queueing-Only Mode	67-17
Configuring a Class Map to Classify MPLS Packets	67-17
Configuring a Policy Map	67-20
Displaying a Policy Map	67-24
Configuring MPLS QoS Egress EXP Mutation	67-24

Configuring EXP Value Maps	67-25
MPLS DiffServ Tunneling Modes	67-26
Short Pipe Mode	67-27
Uniform Mode	67-28
MPLS DiffServ Tunneling Restrictions and Usage Guidelines	67-30
How to Configure Short Pipe Mode	67-30
Ingress PE Router—Customer Facing Interface	67-30
Configuring Ingress PE Router—P Facing Interface	67-32
Configuring the P Router—Output Interface	67-33
Configuring the Egress PE Router—Customer Facing Interface	67-34
How to Configure Uniform Mode	67-34
Configuring the Ingress PE Router—Customer Facing Interface	67-35
Configuring the Ingress PE Router—P Facing Interface	67-36
Configuring the Egress PE Router—Customer Facing Interface	67-37
PFC QoS Statistics Data Export	68-1
Prerequisites for PFC QoS Statistics Data Export	68-1
Restrictions for PFC QoS Statistics Data Export	68-1
Information About PFC QoS Statistics Data Export	68-2
Default Settings for PFC QoS Statistics Data Export	68-2
How to Configure PFC QoS Statistics Data Export	68-2
Cisco IOS ACL Support	69-1
Restrictions for Cisco IOS ACLs	69-1
Restrictions for Layer 4 Operators in ACLs	69-2
Determining Layer 4 Operation Usage	69-2
Determining Logical Operation Unit Usage	69-3
Information About ACL Support	69-4
Policy-Based ACLs (PBACLs)	69-5
Restrictions for PBACLs	69-6
Information about PBACLs	69-6
How to Configure PBACLs	69-6
MAC ACLs	69-9
How to Configure Protocol-Independent MAC ACL Filtering	69-9
How to Enable VLAN-Based MAC QoS Filtering	69-10
ARP ACLs	69-12
Optimized ACL Logging	69-13
Restrictions for OAL	69-13
Information about OAL	69-13

How to Configure OAL	69-13
Dry Run Support for ACLs	69-15
Restrictions for Dry Run Support	69-15
Information About Dry Run Support	69-16
How to Configure Dry Run Support for ACLs	69-16
Hardware ACL Statistics	69-17
Restrictions for Hardware ACL Statistics	69-17
Information About Hardware ACL Statistics	69-17
How to Configure Hardware ACL Statistics	69-18
Cisco TrustSec (CTS)	70-1
Hardware Supported	70-4
AutoSecure	71-1
Prerequisites for AutoSecure	71-1
Restrictions for AutoSecure	71-2
Information About AutoSecure	71-2
AutoSecure Overview	71-2
Management Plane Security Enabled by AutoSecure	71-3
Forwarding Plane Security Enabled by AutoSecure	71-6
How to Configure AutoSecure	71-7
Configuring AutoSecure Parameters	71-7
Configuring Additional Security	71-8
Verifying AutoSecure	71-8
Examples for AutoSecure Configuration	71-9
MAC Address-Based Traffic Blocking	72-1
Port ACLs (PACLs)	73-1
Prerequisites for PACLs	73-1
Restrictions for PACLs	73-2
Information About PACLs	73-2
PACL Overview	73-2
EtherChannel and PACL Interactions	73-3
Dynamic ACLs (Applies to Merge Mode Only)	73-4
Trunk Ports	73-4
Layer 2 to Layer 3 Port Conversion	73-4
Port-VLAN Association Changes	73-4
PACL and VACL Interactions	73-4
How to Configure PACLs	73-7

Configuring IP and MAC ACLs on a Layer 2 Interface	73-7
Configuring Access-group Mode on Layer 2 Interface	73-8
Applying ACLs to a Layer 2 Interface	73-8
Applying ACLs to a Port Channel	73-9
Displaying an ACL Configuration on a Layer 2 Interface	73-9
VLAN ACLs (VACLs)	74-1
Prerequisites for VACLs	74-1
Restrictions for VACLs	74-2
Information About VACLs	74-2
How to Configure VACLs	74-3
Defining a VLAN Access Map	74-3
Configuring a Match Clause in a VLAN Access Map Sequence	74-3
Configuring an Action Clause in a VLAN Access Map Sequence	74-4
Applying a VLAN Access Map	74-4
Verifying VLAN Access Map Configuration	74-5
VLAN Access Map Configuration and Verification Examples	74-5
Configuring a Capture Port	74-6
Configuring VACL Logging	74-7
Policy-Based Forwarding (PBF)	75-1
Prerequisites for PBF	75-1
Restrictions for PBF	75-2
Information About PBF	75-2
Default Settings for PBF	75-2
How to Configure PBF	75-2
Monitoring PBF	75-3
Configuration Examples for PBF	75-3
Denial of Service (DoS) Protection	76-1
Security ACLs and VACLs	76-2
QoS Rate Limiting	76-2
Global Protocol Packet Policing	76-3
Prerequisites for Global Protocol Packet Policing	76-3
Restrictions for Global Protocol Packet Policing	76-3
Information About Global Protocol Packet Policing	76-5
How to Configure Single-Command Global Protocol Packet Policing	76-5
How to Configure Policy-Based Global Protocol Packet Policing	76-6
Unicast Reverse Path Forwarding (uRPF) Check	76-6

Prerequisites for uRPF Check	76-7
Restrictions for uRPF Check	76-7
Information about uRPF Check	76-7
How to Configure Unicast RPF Check	76-8
Configuring Sticky ARP	76-9
Monitoring Packet Drop Statistics	76-10
Prerequisites for Packet Drop Statistics	76-10
Restrictions for Packet Drop Statistics	76-10
Information About Packet Drop Statistics	76-10
How to Monitor Dropped Packets	76-10
Control Plane Policing (CoPP)	77-1
Prerequisites for CoPP	77-1
Restrictions for CoPP	77-2
Information About CoPP	77-3
Default Settings for CoPP	77-3
How to Configure CoPP	77-5
Configuring CoPP	77-5
Defining CoPP Traffic Classification	77-6
Monitoring CoPP	77-9
Dynamic Host Configuration Protocol (DHCP) Snooping	78-1
Prerequisites for DHCP Snooping	78-1
Restrictions for DHCP Snooping	78-2
DHCP Snooping Configuration Restrictions	78-2
DHCP Snooping Configuration Guidelines	78-2
Minimum DHCP Snooping Configuration	78-3
Information About DHCP Snooping	78-3
Overview of DHCP Snooping	78-4
Trusted and Untrusted Sources	78-4
DHCP Snooping Binding Database	78-5
Packet Validation	78-5
DHCP Snooping Option-82 Data Insertion	78-5
Overview of the DHCP Snooping Database Agent	78-7
Default Configuration for DHCP Snooping	78-8
How to Configure DHCP Snooping	78-9
Enabling DHCP Snooping Globally	78-9
Enabling DHCP Option-82 Data Insertion	78-10
Enabling the DHCP Option-82 on Untrusted Port Feature	78-10

Enabling DHCP Snooping MAC Address Verification	78-11
Enabling DHCP Snooping on VLANs	78-11
Configuring the DHCP Trust State on Layer 2 LAN Interfaces	78-12
Configuring Spurious DHCP Server Detection	78-13
Configuring DHCP Snooping Rate Limiting on Layer 2 LAN Interfaces	78-14
The DHCP Snooping Database Agent	78-14
Displaying the DHCP Snooping Binding Table	78-18
IP Source Guard	79-1
Prerequisites for IP Source Guard	79-1
Restrictions for IP Source Guard	79-2
Information About IP Source Guard	79-2
Overview of IP Source Guard	79-2
IP Source Guard Interaction with VLAN-Based Features	79-2
Channel Ports	79-3
Layer 2 and Layer 3 Port Conversion	79-3
IP Source Guard and Voice VLAN	79-3
IP Source Guard and Web-Based Authentication	79-3
Default Settings for IP Source Guard	79-3
How to Configure IP Source Guard	79-3
Displaying IP Source Guard PACL Information	79-5
Displaying IP Source Binding Information	79-6
Dynamic ARP Inspection (DAI)	80-1
Prerequisites for DAI	80-1
Restrictions for DAI	80-2
Information About DAI	80-3
Information about ARP	80-3
ARP Spoofing Attacks	80-3
DAI and ARP Spoofing Attacks	80-4
Interface Trust States and Network Security	80-4
Rate Limiting of ARP Packets	80-5
Relative Priority of ARP ACLs and DHCP Snooping Entries	80-6
Logging of Dropped Packets	80-6
Default Settings for DAI	80-6
How to Configure DAI	80-7
Enabling DAI on VLANs	80-7
Configuring DAI Hardware Acceleration	80-8
Configuring the DAI Interface Trust State	80-8

Applying ARP ACLs for DAI Filtering	80-9
Configuring ARP Packet Rate Limiting	80-10
Enabling DAI Error-Disabled Recovery	80-11
Enabling Additional Validation	80-11
Configuring DAI Logging	80-13
Displaying DAI Information	80-15
Configuration Examples for DAI	80-16
Two Switches Support DAI	80-16
One Switch Supports DAI	80-21
Traffic Storm Control	81-1
Prerequisites for Traffic Storm Control	81-1
Restrictions for Traffic Storm Control	81-2
Information About Traffic Storm Control	81-2
Default Setting for Traffic Storm Control	81-4
How to Enable Traffic Storm Control	81-4
Displaying Traffic Storm Control Settings	81-5
Unknown Unicast and Multicast Flood Control	82-1
Prerequisites for Unknown Traffic Flood Control	82-1
Restrictions for Unknown Traffic Flood Control	82-2
Information About Unknown Traffic Flood Control	82-2
Default Settings for Unknown Traffic Flood Control	82-2
How to Configure Unknown Traffic Flood Control	82-2
How to Configure UUFB or UMFB	82-2
How to Configure UUFRL	82-3
Configuration Examples for Unknown Traffic Flood Control	82-3
83-1	
IEEE 802.1X Port-Based Authentication	83-1
Prerequisites for 802.1X Authentication	83-1
Restrictions for 802.1X Authentication	83-2
802.1X Authentication	83-2
802.1X Host Mode	83-3
VLAN Assignment, Guest VLAN, Restricted VLAN, and Inaccessible Authentication Bypass	83-3
MAC Authentication Bypass	83-5
Web-Based Authentication	83-5
Network Edge Access Topology (NEAT) and Client Information Signalling Protocol (CISP)	83-5
Information About 802.1X Port-Based Authentication	83-6

802.1X Overview	83-6
802.1X Device Roles	83-7
Port-based Authentication Process	83-8
Authentication Initiation and Message Exchange	83-10
Ports in Authorized and Unauthorized States	83-12
802.1X Host Modes	83-13
802.1X Authentication with DHCP Snooping	83-15
802.1X Accounting	83-16
802.1X Authentication with VLAN Assignment	83-17
Multiple VLANs and VLAN User Distribution with VLAN Assignment	83-18
802.1X Authentication with Guest VLAN	83-19
802.1X Authentication with Restricted VLAN	83-20
802.1X Authentication with Inaccessible Authentication Bypass	83-21
802.1X Authentication with Voice VLAN Ports	83-22
802.1X Authentication with Port Security	83-23
802.1X Authentication with ACL Assignments and Redirect URLs	83-23
802.1X Authentication with Port Descriptors	83-26
802.1X Authentication with MAC Authentication Bypass	83-26
Network Admission Control Layer 2 IEEE 802.1X Validation	83-27
802.1X Authentication with Wake-on-LAN	83-28
MAC Move	83-28
MAC Replace	83-29
802.1x Supplicant and Authenticator Switches with Network Edge Access Topology (NEAT)	83-30
Default Settings for 802.1X Port-Based Authentication	83-31
How to Configure 802.1X Port-Based Authentication	83-32
Enabling 802.1X Authentication	83-33
Configuring Switch-to-RADIUS-Server Communication	83-34
Configuring 802.1X Authenticator Host Mode	83-35
Enabling Fallback Authentication	83-36
Enabling Periodic Reauthentication	83-38
Manually Reauthenticating the Client Connected to a Port	83-39
Initializing Authentication for the Client Connected to a Port	83-39
Removing 802.1X Client Information Globally	83-40
Removing 802.1X Client Information from an Interface	83-40
Clearing Authentication Sessions	83-40
Changing 802.1X Timeouts	83-40
Setting the Switch-to-Client Frame Retransmission Number	83-42
Setting the Reauthentication Number	83-43
Configuring IEEE 802.1X Accounting	83-43
Configuring VLAN User Distribution	83-44

Configuring a Guest VLAN	83-45
Configuring a Restricted VLAN	83-46
Configuring the Inaccessible Authentication Bypass Feature	83-47
Configuring MAC Authentication Bypass	83-49
Configuring NAC Layer 2 IEEE 802.1X Validation	83-50
Configuring NAC Agentless Audit Support	83-51
Configuring the Switch for DACLs or Redirect URLs	83-51
Configuring 802.1X Authentication with WoL	83-52
Enabling MAC Move	83-53
Enabling MAC Replace	83-53
Configuring NEAT Authenticator and Supplicant Switches	83-54
Disabling 802.1X Authentication on the Port	83-56
Resetting the 802.1X Configuration to the Default Values	83-57
Displaying Authentication Status and Information	83-57
Displaying 802.1X Status	83-57
Displaying Authentication Methods and Status	83-58
Displaying MAC Authentication Bypass Status	83-61
Web-Based Authentication	84-1
Prerequisites for Web-based Authentication	84-1
Restrictions for Web-based Authentication	84-1
Information About Web-Based Authentication	84-2
Web-Based Authentication Overview	84-2
Device Roles	84-3
Host Detection	84-3
Session Creation	84-4
Authentication Process	84-4
AAA Fail Policy	84-5
Customization of the Authentication Proxy Web Pages	84-5
Web-based Authentication Interactions with Other Features	84-5
Default Web-Based Authentication Configuration	84-7
How to Configure Web-Based Authentication	84-7
Web-based Authentication Configuration Task List	84-8
Configuring the Authentication Rule and Interfaces	84-8
Configuring AAA Authentication	84-9
Configuring Switch-to-RADIUS-Server Communication	84-9
Configuring the HTTP Server	84-11
Configuring an AAA Fail Policy	84-13
Configuring the Web-based Authentication Parameters	84-14

Removing Web-based Authentication Cache Entries 84-14

Displaying Web-Based Authentication Status 84-15

Port Security 85-1

Prerequisites for Port Security 85-1

Restrictions for Port Security 85-1

Information About Port Security 85-2

Port Security with Dynamically Learned and Static MAC Addresses 85-3

Port Security with Sticky MAC Addresses 85-3

Port Security with IP Phones 85-4

Default Port Security Configuration 85-4

How to Configure Port Security 85-4

Enabling Port Security 85-4

Configuring the Port Security Violation Mode on a Port 85-6

Configuring the Maximum Number of Secure MAC Addresses on a Port 85-7

Enabling Port Security with Sticky MAC Addresses on a Port 85-8

Configuring a Static Secure MAC Address on a Port 85-8

Configuring Secure MAC Address Aging on a Port 85-9

Verifying the Port Security Configuration 85-10

Lawful Intercept 86-1

Prerequisites for Lawful Intercept 86-1

Restrictions for Lawful Intercept 86-1

General Configuration Restrictions 86-2

MIB Guidelines 86-3

Information About Lawful Intercept 86-3

Lawful Intercept Overview 86-4

Benefits of Lawful Intercept 86-4

CALEA for Voice 86-5

Network Components Used for Lawful Intercept 86-5

Lawful Intercept Processing 86-7

Lawful Intercept MIBs 86-7

How to Configure Lawful Intercept Support 86-9

Security Considerations 86-9

Accessing the Lawful Intercept MIBs 86-9

Restricting Access to the Lawful Intercept MIBs 86-10

Configuring SNMPv3 86-10

Online Diagnostic Tests 1-1

Global Health-Monitoring Tests 1-2

TestAsicSync	1-2
TestEARLInternalTables	1-3
TestErrorCounterMonitor	1-3
TestFexModeLoopback	1-4
TestIntPortLoopback	1-4
TestL3TcamMonitoring	1-4
TestLtlFpoeMemoryConsistency	1-5
TestMacNotification	1-6
TestPortTxMonitoring	1-6
TestScratchRegister	1-7
TestSnrMonitoring	1-7
TestUnusedPortLoopback	1-8
Per-Port Tests	1-8
TestActiveToStandbyLoopback	1-9
TestCCPLoopback	1-9
TestDataPortLoopback	1-10
TestDCPLoopback	1-10
TestL2CTSLoopback	1-11
TestL3CTSLoopback	1-11
TestLoopback	1-12
TestMediaLoopback	1-12
TestMgmtPortsLoopback	1-13
TestNetflowInlineRewrite	1-13
TestNonDisruptiveLoopback	1-14
TestNPLoopback	1-14
TestTransceiverIntegrity	1-15
PFC Layer 2 Tests	1-15
TestBadBpduTrap	1-15
TestDontConditionalLearn	1-16
TestMatchCapture	1-16
TestNewIndexLearn	1-17
DFC Layer 2 Tests	1-17
TestBadBpdu	1-18
TestCapture	1-18
TestConditionalLearn	1-19
TestDontLearn	1-19
TestIndexLearn	1-20
TestNewLearn	1-20
TestPortSecurity	1-21

TestProtocolMatchChannel	1-21
TestStaticEntry	1-22
TestTrap	1-22
PFC Layer 3 Tests	1-22
TestAclDeny	1-23
TestAclPermit	1-23
TestAclRedirect	1-24
TestDQUP	1-24
TestInbandEdit	1-25
TestIPv4FibShortcut	1-25
TestIPv6FibShortcut	1-26
TestL3Capture2	1-26
TestMPLSFibShortcut	1-27
TestNATFibShortcut	1-27
TestNetflowShortcut	1-28
TestRBAcl	1-28
DFC Layer 3 Tests	1-28
TestAclDeny	1-29
TestAclPermit	1-29
TestAclRedirect	1-30
TestInbandEdit	1-30
TestIPv4FibShortcut	1-31
TestIPv6FibShortcut	1-31
TestL3Capture2	1-32
TestMPLSFibShortcut	1-32
TestNATFibShortcut	1-33
TestNetflowShortcut	1-33
TestRBAcl	1-34
Replication Engine Tests	1-34
TestEgressSpan	1-34
TestIngressSpan	1-35
TestL3VlanMet	1-35
Fabric Tests	1-35
TestFabricCh0Health	1-36
TestFabricCh1Health	1-36
TestFabricExternalSnake	1-37
TestFabricFlowControlStatus	1-37
TestFabricInternalSnake	1-38
TestFabricVlanLoopback	1-38

TestFexFabricLinkStatus	1-39
TestSynchedFabChannel	1-39
Exhaustive Memory Tests	1-39
TestAclQosTcam	1-40
TestAsicMemory	1-40
TestEarlMemOnBootup	1-41
Service Module Tests	1-41
TestPcLoopback	1-41
TestPortASICLoopback	1-42
Stress Tests	1-42
TestEobcStressPing	1-42
TestMicroburst	1-43
TestNVRAMBatteryMonitor	1-43
TestTrafficStress	1-44
General Tests	1-44
ScheduleSwitchover	1-44
TestCFRW	1-45
TestFirmwareDiagStatus	1-45
TestOBFL	1-45
TestRwEngineOverSubscription	1-46
TestVDB	1-46
Critical Recovery Tests	1-46
TestTxPathMonitoring	1-47
ViSN Tests	1-47
TestRslHm	1-47
TestVSActiveToStandbyLoopback	1-48
TestVslBridgeLink	1-48
TestVslLocalLoopback	1-49
TestVslStatus	1-49

INDEX

Preface

This preface describes who should read the *Supervisor Engine 2T Software Configuration Guide*, Release 15.1SY, and its document conventions.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the switches supported in Cisco IOS Release 15.4SY.

Related Documentation

See the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

For information about MIBs, go to this URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

Conventions

This document uses the following conventions:

Convention	Description
boldface font	Commands, command options, and keywords are in boldface .
<i>italic</i> font	Arguments for which you supply values are in <i>italics</i> .
[]	Elements in square brackets are optional.
{ x y z }	Alternative keywords are grouped in braces and separated by vertical bars.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
screen font	Terminal sessions and information the system displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen</i> font	Arguments for which you supply values are in <i>italic screen</i> font.
→	This pointer highlights an important line of text in an example.

Convention	Description
^	The symbol ^ represents the key labeled Control—for example, the key combination ^D in a screen display means hold down the Control key while you press the D key.
< >	Nonprinting characters, such as passwords are in angle brackets.

Notes use the following conventions:



Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the publication.

Cautions use the following conventions:



Caution

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

Obtaining Documentation, Obtaining Support, and Security Guidelines

For information on obtaining documentation, obtaining support, providing documentation feedback, security guidelines, and also recommended aliases and general Cisco documents, see the monthly *What's New* in Cisco Product Documentation, which also lists all new and revised Cisco technical documentation, at:

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

1

Product Overview

- [Supervisor Engine 6T Ports, page 1-1](#)
- [LED Indicators, page 1-3](#)
- [Determining System Hardware Capacity, page 1-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- For complete information about the supported chassis, modules, and software features, see the *Release Notes for Cisco IOS Release 15.2SY*:
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.1SY/release_notes.html
and see the Catalyst 6500 Series *Switch Supervisor Engine Guide*:

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Supervisor Engine 6T Ports

USB Type A Port

The USB Type A port is the only external storage interface for this switch. The port is connected to the route processor, which allows the Cisco IOS software to access the port. The port supports Cisco USB flash drives with capacities from 128 MB to 8 GB (USB devices with port densities of 128 MB, 256 MB, 1 GB, 4 GB, and 8 GB are supported). Cisco IOS software provides standard file system access to the flash device: read, write, erase, and copy. The software also provides the ability to format the flash device with a FAT file system (FAT32 and FAT16).

Addressing it : disk0:

Console ports

- EIA/TIA-232 (RS-232) port with RJ-45 connector
- USB port

RJ45: The console port is an RJ-45 port that provides universal asynchronous receiver/transmitter (UART) support to access the route processor with a serial console running at 9600 baud rate with 8 bits for data, no parity bit, and 1 stop bit.

USB Type B Port

The USB 2.0 port Type B serves as a second console connection to the route processor. The USB console port connection uses a USB Type A to Type B cable. The USB console interface speeds are the same as the RJ-45 console interface speeds. Windows PCs need a driver for the USB port.

The USB-prefer mode is the default, but it can be overridden using the command-line interface (CLI). When this port is in USB-prefer mode, the RJ-45 console port will be disabled if both ports are connected. For more information on using the CLI to configure the USB console interface, see the Catalyst 6500 software guide.

System Reset Button

This reset access button is used to reset the system. Pressing the button brings down the route processor and all port card modular slots.

Front Panel mode supports

- 10G*8 and 2*40G (Default mode)
- 10G*8 - 10Gigabit * 8 ports (Performance Mode)
- 40G*2 - 40Gigabit * 2 ports (Performance Mode)
- 10G*16 - 10Gigabit * 16 ports (Port Numbers ranging from: 1 to 8 and 11 to 18)

Management Port

The management port is a 10/100/1000 copper Ethernet port directly connected to the route processor. The switch also has a fiber port that can be used as the Ethernet Management port. It supports TFTP/SCP image downloading, network management, SNMP, Telnet, and SSH connections. The management port is isolated from other ports in the system in a dedicated management VRF; it is not part of the EARL forwarding logic. The management port provides direct access to the CPU, even when the system is heavily loaded.

The management port is a Layer 3 port in host mode, and only accepts traffic that terminates on the router. This port does not route packets between itself and other ports. The port processes only the following packet types and properly enqueues them:

- Address Resolution Protocol (ARP)
- IPv4 unicast
- Cisco Discovery Protocol (CDP)

- Link Layer Discovery Protocol (LLDP)

Management Port Operation Modes

The Ethernet management port (mgmt0) can be operated in the following modes:

- RJ45 (Default mode)
- SFP

To change from the default mode to SFP, run the following command:

```
Router(config)#int mgmt0
Router(config-if)#media-type ?
  rj45 Use RJ45 connector
  sfp Use SFP connector

Router(config-if)#media-type sf
Router(config-if)#media-type sfp
Router(config-if)#
```

If media-type sfp is configured, the rommon variable SFP_MGMT=1 is configured. When the system is in rommon, the SFP operation mode functions.

If no media-type sfp or the default media-type configuration is performed, the rommon variable SFP_MGMT=1 is unconfigured. When the system is in rommon, the RJ45 operation mode functions.



Note

By default, mgmt0 port exists in the “no shutdown” mode.

LED Indicators

You can use the switch LEDs to monitor switch activity and performance. You can also monitor the status of the fan tray assembly, and the power supplies.

Status LED

The status LED indicates the status of the system.

Table 1-1 Status LED Indicator

Color/State	Description
Off	System is not operational.
Green	System is operating normally without alarms.
Amber	System has triggered a minor environmental alarm.
Red	System has triggered a major environmental alarm.

System ID LED

The System ID (blue beacon) LED can be provisioned by the operator to indicate that the switch needs attention.

Color/State	Description
Blinking Blue	The system needs attention.

Management Port LED

This table describes the management port LEDs.

Color/State	Description
Off	Port is not provisioned.
Amber	Port is provisioned, but administratively not operational.
Green	Port is linked up.
Alternating Green and Amber	A port fault is detected, or the port beacon has been provisioned by the operator.

Status LED

The status LED indicates the status of the system.

Color/State	Description
Off	System is not operational.
Green	System is operating normally without alarms.
Amber	System has triggered a minor environmental alarm.
Red	System has triggered a major environmental alarm.

System ID LED

The System ID (blue beacon) LED can be provisioned by the operator to indicate that the switch needs attention.

Color/State	Description
Blinking Blue	The system needs attention.

Determining System Hardware Capacity

You can determine the system hardware capacity by entering the **show platform hardware capacity** command. This command displays the current system utilization of the hardware resources and displays a list of the currently available hardware capacities, including the following:

- Hardware forwarding table utilization
- Switch fabric utilization
- CPU(s) utilization
- Memory device (flash, DRAM, NVRAM) utilization

This example shows how to display EOBC-related statistics for the route processor, the switch processor, and the DFCs:

Router#**show platform hardware capacity eobc**

EOBC Resources

Module		Packets/sec	Total packets	Dropped packets
1	Rx:	37	927814	1
	Tx:	38	937931	0
2	Rx:	21	481277	0
	Tx:	21	491303	0
3 RP	Rx:	148	3367923	0
	Tx:	146	3331458	0
4 RP	Rx:	18	499657	0
	Tx:	18	497068	0
5	Rx:	2	6581834	0
	Tx:	2	104755	3
7	Rx:	54	1275866	0
	Tx:	54	1285205	0

This example shows how to display the current and peak switching utilization:

Router#**show platform hardware capacity cpu**

CPU Resources

CPU utilization: Module	5 seconds	1 minute	5 minutes
1	9% / 0%	9%	10%
2	6% / 0%	4%	4%
3 RP	3% / 0%	4%	4%
4 RP	1% / 0%	1%	1%
5	1% / 1%	2%	2%

7	1% / 0%	2%	2%	
Processor memory: Module	Bytes:	Total	Used	%Used
1		1583178372	194864020	12%
2		1583178372	179324784	11%
3 RP		2369406172	416547644	18%
4 RP		2329657520	376045324	16%
5		937408864	164029640	17%
7		1583178372	147337052	9%
I/O memory: Module	Bytes:	Total	Used	%Used
1		268435456	34353840	13%
2		268435456	34353840	13%
3 RP		469762048	280646296	60%
4 RP		469762048	280646296	60%
7		268435456	34353796	13%

This example shows how to display information about the total capacity, the bytes used, and the percentage that is used for the flash and NVRAM resources present in the system:

Router#show platform hardware capacity flash

Flash/NVRAM Resources

Usage: Module Device	Bytes:	Total	Used	%Used
1 dfc#1-bootdisk:		1011761152	14139392	1%
2 dfc#2-bootdisk:		1011761152	30588928	3%
3 RP nvram:		2702560	36576	1%
3 RP const_nvram:		1048556	60576	6%
3 RP bootdisk:		3848056832	100683776	3%
4 slavenvram:		2702560	36576	1%
4 slaveconst_nvram:		1048556	60576	6%
4 slavebootdisk:		3848056832	92282880	2%
5 dfc#5-bootflash:		15990785	15990784	100%
7 dfc#7-bootdisk:		1024557056	53723136	5%

This example shows how to display the capacity and utilization of the PFC and DFCs present in the system:

Router#show platform hardware capacity forwarding

L2 Forwarding Resources

MAC Table usage:	Module	Collisions	Total	Used	%Used
1 (E1)		0	131072	1000	1%
1 (E2)		0	131072	1000	1%
2		0	131072	1000	1%
3		0	131072	1000	1%
4		0	131072	1000	1%
5		0	131072	1000	1%
7		0	131072	1000	1%

L3 Forwarding Resources

FIB TCAM usage:	Total	Used	%Used
72 bits (IPv4, MPLS, EoM)	262144	18	1%
144 bits (IP mcast, IPv6)	131072	7	1%
288 bits (IPv6 mcast)	65536	1	1%

detail:	Protocol	Used	%Used
	IPv4	16	1%
	MPLS	1	1%
	EoM	1	1%
	IPv6	1	1%
	IPv4 mcast	6	1%
	IPv6 mcast	1	1%

```

Adjacency usage:
Total      Used      %Used
1048576    32009      3%

Forwarding engine load:
Module      pps    peak-pps      peak-time
1 (E1)      0        7 02:49:50 UTC Sat Mar 26 2016
1 (E2)      0        4 03:31:21 UTC Sat Mar 26 2016
2           0        7 02:48:25 UTC Sat Mar 26 2016
3           7       10 02:48:51 UTC Sat Mar 26 2016
4           0        7 02:48:31 UTC Sat Mar 26 2016
5           0        3 02:49:02 UTC Sat Mar 26 2016
7           0       10 06:05:41 UTC Sat Mar 26 2016

```

This example shows how to display the interface resources:

Router#show platform hardware capacity interface

Interface Resources

Interface drops:

```

Module      Total drops:  Tx      Rx      Highest drop port:  Tx  Rx
no modules reported drops

```

Interface buffer sizes:

```

Mod/Port      Bytes:  Tx buffer      Rx buffer
4/1           252695040      960000
4/2           252695040      960000
4/3           252695040      960000
4/4           252695040      960000
4/5           252695040      960000
4/6           252695040      960000
4/7           252695040      960000
4/8           252695040      960000
4/9           1007952384     3989248
4/10          1007952384     3989248

```


2

Command-Line Interfaces

- [Accessing the CLI, page 2-1](#)
- [Performing Command Line Processing, page 2-3](#)
- [Performing History Substitution, page 2-4](#)
- [Cisco IOS Command Modes, page 2-4](#)
- [Displaying a List of Cisco IOS Commands and Syntax, page 2-6](#)
- [Securing the CLI, page 2-7](#)
- [ROM-Monitor Command-Line Interface, page 2-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Accessing the CLI

- [Accessing the CLI through the EIA/TIA-232 Console Interface, page 2-1](#)
- [Accessing the CLI through Telnet, page 2-2](#)
- [Accessing the CLI through USB, page 2-3](#)

Accessing the CLI through the EIA/TIA-232 Console Interface

**Note**

EIA/TIA-232 was known as recommended standard 232 (RS-232) before its acceptance as a standard by the Electronic Industries Alliance (EIA) and Telecommunications Industry Association (TIA).

Perform initial configuration over a connection to the EIA/TIA-232 console interface. See the *Catalyst 6500 Series Switch Module Installation Guide* for console interface cable connection procedures.

To make a console connection, perform this task:

	Command	Purpose
Step 1	Press Return.	Brings up the prompt.
Step 2	Router> enable	Initiates enable mode enable.
Step 3	Password: <i>password</i> Router#	Completes enable mode enable.
Step 4	Router# quit	Exits the session when finished.

After making a console connection, you see this display:

Press Return for Console prompt

```
Router> enable
Password:
Router#
```

Accessing the CLI through Telnet



Note

Before you can make a telnet connection to the switch, you must configure an IP address.

The switch supports up to eight simultaneous telnet sessions. Telnet sessions disconnect automatically after remaining idle for the period specified with the **exec-timeout** command.

To make a telnet connection to the switch, perform this task:

	Command	Purpose
Step 1	telnet {hostname ip_addr}	Makes a telnet connection from the remote host to the switch you want to access.
Step 2	Password: <i>password</i> Router#	Initiates authentication. Note If no password has been configured, press Return.
Step 3	Router> enable	Initiates enable mode enable.
Step 4	Password: <i>password</i> Router#	Completes enable mode enable.
Step 5	Router# quit	Exits the session when finished.

This example shows how to open a Telnet session to the switch:

```
unix_host% telnet Router_1
Trying 172.20.52.40...
Connected to 172.20.52.40.
Escape character is '^]'.
```

```
User Access Verification
```

```
Password:
Router_1> enable
Password:
Router_1#
```


Accessing the CLI through USB

The USB console port connection uses a USB Type A to 5-pin mini-Type B cable. The USB console interface speeds are the same as the RJ-45 console interface speeds.

If you use the USB mini-Type B console port, the Cisco Windows USB device driver must be installed on any PC connected to the console port (for operation with Microsoft Windows). Mac OS X or Linux do not require special drivers.

The 4-pin mini-Type B connector resembles the 5-pin mini-Type B connectors. They are not compatible. Use only the 5-pin mini-Type B.

Figure 2-1 5-pin USB Mini-Type B Port.



With the Cisco Windows USB device driver, you can connect and disconnect the USB cable from the console port without affecting Windows HyperTerminal operations.

The console output always goes to both the RJ-45 and the USB console connectors, but the console input is active on only one of the console connectors at any one time. The USB console takes precedence over the RJ-45 console. When a cable is connected into the USB console port, the RJ-45 console port becomes inactive. Conversely, when the USB cable is disconnected from the USB console port, the RJ-45 port becomes active.

You can use the command-line interface (CLI) to configure an inactivity timeout which reactivates the RJ-45 console if the USB console has been activated and no input activity has occurred on the USB console for a specified time.

After the USB console deactivates due to inactivity, you cannot use the CLI to reactivate it. Disconnect and reconnect the USB cable to reactivate the USB console.

Performing Command Line Processing

Commands are not case sensitive. You can abbreviate commands and parameters if the abbreviations contain enough letters to be different from any other currently available commands or parameters. You can scroll through the last 20 commands stored in the history buffer, and enter or edit the command at the prompt. [Table 2-1](#) lists the keyboard shortcuts for entering and editing commands.

Table 2-1 Keyboard Shortcuts

Keystrokes	Purpose
Press Ctrl-B or press the left arrow key	Moves the cursor back one character. Note The arrow keys function only on ANSI-compatible terminals such as VT100s.
Press Ctrl-F or press the right arrow key	Moves the cursor forward one character. Note The arrow keys function only on ANSI-compatible terminals such as VT100s.
Press Ctrl-A	Moves the cursor to the beginning of the command line.
Press Ctrl-E	Moves the cursor to the end of the command line.
Press Esc B	Moves the cursor back one word.
Press Esc F	Moves the cursor forward one word.

Performing History Substitution

The history buffer stores the last 20 commands you entered. History substitution allows you to access these commands without retyping them, by using special abbreviated commands. [Table 2-2](#) lists the history substitution commands.

Table 2-2 History Substitution Commands

Command	Purpose
Ctrl-P or the up arrow key.	Recalls commands in the history buffer, beginning with the most recent command. Repeat the key sequence to recall successively older commands. Note The arrow keys function only on ANSI-compatible terminals such as VT100s.
Ctrl-N or the down arrow key.	Returns 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. Note The arrow keys function only on ANSI-compatible terminals such as VT100s.
Router# show history	While in EXEC mode, lists the last several commands you have just entered.

Cisco IOS Command Modes


Note

For complete information about Cisco IOS command modes, see the *Cisco IOS Configuration Fundamentals Configuration Guide* at this URL:

http://www.cisco.com/en/US/docs/ios-xml/ios/fundamentals/configuration/15_sy/fundamentals-15-sy-book.html

The Cisco IOS user interface is divided into many different modes. The commands available to you depend on which mode you are currently in. To get a list of the commands in a given mode, type a question mark (?) at the system prompt. See the “[Displaying a List of Cisco IOS Commands and Syntax](#)” section on page 2-6.

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

The configuration modes allow you to make changes to the running configuration. If you later save the configuration, these commands are stored across reboots. You must start at global configuration mode. From global configuration mode, you can enter interface configuration mode, subinterface configuration mode, and a variety of protocol-specific modes.



Note

You can enter EXEC mode commands by entering the **do** keyword before the EXEC mode command.

ROM-monitor mode is a separate mode used when the switch cannot boot properly. For example, the switch might enter ROM-monitor mode if it does not find a valid system image when it is booting, or if its configuration file is corrupted at startup. See the “[ROM-Monitor Command-Line Interface](#)” section on page 2-7.

Table 2-3 lists and describes frequently used Cisco IOS modes.

Table 2-3 *Frequently Used Cisco IOS Command Modes*

Mode	Description of Use	How to Access	Prompt
User EXEC	Connect to remote devices, change terminal settings on a temporary basis, perform basic tests, and display system information.	Log in.	Router>
Privileged EXEC (enable)	Set operating parameters. The privileged command set includes the commands in user EXEC mode, as well as the configure command. Use this command to access the other command modes.	From the user EXEC mode, enter the enable command and the enable password.	Router#
Global configuration	Configure features that affect the system as a whole.	From the privileged EXEC mode, enter the configure terminal command.	Router (config)#

Table 2-3 *Frequently Used Cisco IOS Command Modes (continued)*

Mode	Description of Use	How to Access	Prompt
Interface configuration	Many features are enabled for a particular interface. Interface commands enable or modify the operation of an interface.	From global configuration mode, enter the interface type slot/port command.	Router (config-if) #
Console configuration	From the directly connected console or the virtual terminal used with Telnet, use this configuration mode to configure the console interface.	From global configuration mode, enter the line console 0 command.	Router (config-line) #

The Cisco IOS command interpreter, called the EXEC, interprets and executes the commands you enter. You can abbreviate commands and keywords by entering just enough characters to make the command unique from other commands. For example, you can abbreviate the **show** command to **sh** and the **configure terminal** command to **conf t**.

When you type **exit**, the switch backs out one level. To exit configuration mode completely and return to privileged EXEC mode, press **Ctrl-Z**.

Displaying a List of Cisco IOS Commands and Syntax

In any command mode, you can display a list of available commands by entering a question mark (?).

```
Router> ?
```

To display a list of commands that begin with a particular character sequence, type in those characters followed by the question mark (?). Do not include a space. This form of help is called word help because it completes a word for you.

```
Router# co?
collect  configure  connect  copy
```

To display keywords or arguments, enter a question mark in place of a keyword or argument. Include a space before the question mark. This form of help is called command syntax help because it reminds you which keywords or arguments are applicable based on the command, keywords, and arguments you have already entered.

For example:

```
Router# configure ?
memory          Configure from NV memory
network          Configure from a TFTP network host
overwrite-network Overwrite NV memory from TFTP network host
terminal         Configure from the terminal
<cr>
```

To redisplay a command you previously entered, press the up arrow key or **Ctrl-P**. You can continue to press the up arrow key to see the last 20 commands you entered.



Tip

If you are having trouble entering a command, check the system prompt, and enter the question mark (?) for a list of available commands. You might be in the wrong command mode or using incorrect syntax.

Enter **exit** to return to the previous mode. Press **Ctrl-Z** or enter the **end** command in any mode to immediately return to privileged EXEC mode.

Securing the CLI

Securing access to the CLI prevents unauthorized users from viewing configuration settings or making configuration changes that can disrupt the stability of your network or compromise your network security. You can create a strong and flexible security scheme for your switch by configuring one or more of these security features:

- Protecting access to privileged EXEC commands—At a minimum, you should configure separate passwords for the user EXEC and privileged EXEC (enable) IOS command modes. You can further increase the level of security by configuring username and password pairs to limit access to CLI sessions to specific users. For more information, see this publication:

http://www.cisco.com/en/US/docs/ios/security/configuration/guide/sec_cfg_sec_4cli.html

- Controlling switch access with RADIUS, TACACS+, or Kerberos—For a centralized and scalable security scheme, you can require users to be authenticated and authorized by an external security server running either Remote Authentication Dial-In User Service (RADIUS), Terminal Access Controller Access-Control System Plus (TACACS+), or Kerberos. For more information, see this publication:

http://www.cisco.com/en/US/docs/ios-xml/ios/security/config_library/15-sy/secdata-15-sy-library.html

- Configuring a secure connection with SSH or HTTPS—To prevent eavesdropping of your configuration session, you can use a Secure Shell (SSH) client or a browser that supports HTTP over Secure Socket Layer (HTTPS) to make an encrypted connection to the switch. For more information, see this publication:

http://www.cisco.com/en/US/docs/ios-xml/ios/security/config_library/15-sy/secdata-15-sy-library.html

For more information about HTTPS, see “HTTPS - HTTP Server and Client with SSL 3.0” at this URL:

http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_cfg_sec_4cli.html

- Copying configuration files securely with SCP—To prevent eavesdropping when copying configuration files or image files to or from the switch, you can use the Secure Copy Protocol (SCP) to perform an encrypted file transfer. For more information about SCP, see “Secure Copy” at this URL:

http://www.cisco.com/en/US/docs/ios-xml/ios/sec_usr_ssh/configuration/15-sy/sec-usr-ssh-sec-copy.html

ROM-Monitor Command-Line Interface

The ROM-monitor is a ROM-based program that executes upon platform power-up, reset, or when a fatal exception occurs. The switch enters ROM-monitor mode if it does not find a valid software image, if the NVRAM configuration is corrupted, or if the configuration register is set to enter ROM-monitor mode. From the ROM-monitor mode, you can load a software image manually from flash memory, from a network server file, or from bootflash.

You can also enter ROM-monitor mode by restarting and pressing the **Break** key during the first 60 seconds of startup.

**Note**

The **Break** key is always enabled for 60 seconds after rebooting, regardless of whether the **Break** key is configured to be off by configuration register settings.

To access the ROM-monitor mode through a terminal server, you can escape to the Telnet prompt and enter the **send break** command for your terminal emulation program to break into ROM-monitor mode.

Once you are in ROM-monitor mode, the prompt changes to rommon 1>. Enter a question mark (?) to see the available ROM-monitor commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

3

Smart Port Macros

- [Prerequisites for Smart Port Macros, page 3-1](#)
- [Restrictions for Smart Port Macros, page 3-1](#)
- [Information About Smart Port Macros, page 3-3](#)
- [Default Settings for Smart Port Macros, page 3-4](#)
- [How to Configure Smart Port Macros, page 3-4](#)
- [Verifying the Smart Port Macro Configuration, page 3-15](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Smart Port Macros

None.

Restrictions for Smart Port Macros

- You can display all of the macros on the switch by using the **show parser macro** user EXEC command. Display the contents of a specific macro by using the **show parser macro name macro-name** user EXEC command.
- You cannot edit a macro. If the name following the **macro name** command is an existing macro's name, that macro is replaced by the new macro.

- If a description already exists for a macro, the **macro description** command appends any description that you enter to the existing description; it does not replace it. The entered descriptions are separated by the pipe (“|”) character.
- The maximum macro description length is 256 characters. When the description string becomes longer than 256 characters, the oldest descriptions are deleted to make room for new ones.
- User-created recursive macros are not supported. You cannot define a macro that calls another macro.
- Each user-created macro can have up to three keyword-value pairs.
- A macro definition can contain up to 3,000 characters. Line endings count as two characters.
- When creating a macro, do not use the **exit** or **end** commands or change the command mode by using **interface interface-id**. This could cause commands that follow **exit**, **end**, or **interface interface-id** to execute in a different command mode. When creating a macro, all CLI commands should be in the same configuration mode.
- When creating a macro that requires the assignment of unique values, use the **parameter value** keywords to designate values specific to the interface. Keyword matching is case sensitive. All matching occurrences of the keyword are replaced with the corresponding value. Any full match of a keyword, even if it is part of a larger string, is considered a match and is replaced by the corresponding value.
- Macro names are case sensitive. For example, the commands **macro name Sample-Macro** and **macro name sample-macro** will result in two separate macros.
- Some macros might contain keywords that require a parameter value. You can use the **macro global apply macro-name ?** global configuration command or the **macro apply macro-name ?** interface configuration command to display a list of any required values in the macro. If you apply a macro without entering the keyword values, the commands are invalid and are not applied.
- When a macro is applied globally to a switch or to a switch interface, the existing configuration on the interface is retained. This is helpful when applying an incremental configuration.
- If you modify a macro definition by adding or deleting commands, the changes are not reflected on the interface where the original macro was applied. You need to reapply the updated macro on the interface to apply the new or changed commands.
- You can use the **macro global trace macro-name** global configuration command or the **macro trace macro-name** interface configuration command to apply and debug a macro to find any syntax or configuration errors. If a command fails because of a syntax error or a configuration error, the macro continues to apply the remaining commands.
- Some CLI commands are specific to certain interface types. If a macro is applied to an interface that does not accept the configuration, the macro will fail the syntax check or the configuration check, and the switch will return an error message.
- Applying a macro to an interface range is the same as applying a macro to a single interface. When you use an interface range, the macro is applied sequentially to each interface within the range. If a macro command fails on one interface, it is still applied to the remaining interfaces.
- When you apply a macro to a switch or a switch interface, the macro name is automatically added to the switch or interface. You can display the applied commands and macro names by using the **show running-config** user EXEC command.

Information About Smart Port Macros

- [Information about Cisco-Provided Smart Port Macros, page 3-3](#)
- [Information about User-Created Smart Port Macros, page 3-3](#)

Information about Cisco-Provided Smart Port Macros

Use the **show parser macro** user EXEC command to display the Cisco-provided smart port macros and the commands they contain.

Table 3-1 *Cisco-Provided Smart Port Macros*

Macro Name	Description
cisco-global	Use this global configuration macro to enable load balancing across VLANs, provide rapid convergence of spanning-tree instances and to enable port error recovery.
cisco-desktop	Use this interface configuration macro for increased network security and reliability when connecting a desktop device, such as a PC, to a switch port.
cisco-phone	Use this interface configuration macro when connecting a desktop device such as a PC with a Cisco IP phone to a switch port. This macro is an extension of the cisco-desktop macro and provides the same security and resiliency features, but with the addition of dedicated voice VLANs to ensure proper treatment of delay-sensitive voice traffic.
cisco-switch	Use this interface configuration macro for Layer 2 connections between devices like switches and routers.
cisco-router	Use this interface configuration macro for Layer 3 connections between devices like switches and routers.

Cisco also provides a collection of pretested, Cisco-recommended baseline configuration templates for Catalyst switches. The online reference guide templates provide the CLI commands that you can use to create smart port macros based on the usage of the port. You can use the configuration templates to create smart port macros to build and deploy Cisco-recommended network designs and configurations.

Information about User-Created Smart Port Macros

Smart port macros provide a convenient way to save and share common configurations. You can use smart port macros to enable features and settings based on the location of a switch in the network and for mass configuration deployments across the network.

Each smart port macro is a user-defined set of Cisco IOS CLI commands. When you apply a smart port macro on an interface, the CLI commands within the macro are configured on the interface. When the macro is applied to an interface, the existing interface configurations are not lost. The new commands are added to the interface and are saved in the running configuration file.

Default Settings for Smart Port Macros

This example shows how to list the Cisco-provided smart port macros that are provided by default:

```
Router# show parser macro brief
      default global      : cisco-global
      default interface:  cisco-desktop
      default interface:  cisco-phone
      default interface:  cisco-switch
      default interface:  cisco-router
```

How to Configure Smart Port Macros

- [Using the Cisco-Provided Smart Port Macros, page 3-4](#)
- [Creating Smart Port Macros, page 3-13](#)

Using the Cisco-Provided Smart Port Macros

- [Using the cisco-global Smart Port Macro, page 3-4](#)
- [Using the cisco-desktop Smart Port Macro, page 3-5](#)
- [Using the cisco-phone Smart Port Macro, page 3-7](#)
- [Using the cisco-switch Smart Port Macro, page 3-9](#)
- [Using the cisco-router Smart Port Macro, page 3-11](#)

Using the cisco-global Smart Port Macro

- [Displaying the Contents of the cisco-global Smart Port Macro, page 3-4](#)
- [Applying the cisco-global Smart Port Macro, page 3-5](#)

Displaying the Contents of the cisco-global Smart Port Macro

```
Router# show parser macro name cisco-global
Macro name : cisco-global
Macro type : default global
# Enable dynamic port error recovery for link state
# failures
errdisable recovery cause link-flap
errdisable recovery interval 60

# VTP requires Transparent mode for future 802.1x Guest VLAN
# and current Best Practice
vtp domain [smartports]
vtp mode transparent

# Config Cos to DSCP mappings
platform qos map cos-dscp 0 8 16 26 32 46 48 56

# Enable aggressive mode UDLD on all fiber uplinks
udld aggressive

# Enable Rapid PVST+ and Loopguard
spanning-tree mode rapid-pvst
```

```
spanning-tree loopguard default
spanning-tree extend system-id
```

Applying the cisco-global Smart Port Macro

To apply the cisco-global smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# macro global apply cisco-global	Applies the cisco-global smart port macro.
Step 3	Router(config)# end	Returns to privileged EXEC mode.

This example shows how to apply the cisco-global smart port macro and display the name of the applied macro:

```
Router# configure terminal
Router(config)# macro global apply cisco-global
Changing VTP domain name from previous_domain_name to [smartports]
Setting device to VTP TRANSPARENT mode.
Router(config)# end
Router# show parser macro description
Global Macro(s): cisco-global

Interface      Macro Description(s)
-----
Router#
```

Using the cisco-desktop Smart Port Macro

- [Displaying the Contents of the cisco-desktop Smart Port Macro, page 3-6](#)
- [Applying the cisco-desktop Smart Port Macro, page 3-6](#)

Displaying the Contents of the cisco-desktop Smart Port Macro

```
Router# show parser macro name cisco-desktop
Macro name : cisco-desktop
Macro type : default interface
# macro keywords $AVID
# Basic interface - Enable data VLAN only
# Recommended value for access vlan (AVID) should not be 1
switchport
switchport access vlan $AVID
switchport mode access

# Enable port security limiting port to a single
# MAC address -- that of desktop
switchport port-security
switchport port-security maximum 1

# Ensure port-security age is greater than one minute
# and use inactivity timer
switchport port-security violation restrict
switchport port-security aging time 2
# Configure port as an edge network port
spanning-tree portfast
spanning-tree bpduguard enable
```

Applying the cisco-desktop Smart Port Macro

To apply the cisco-desktop smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface type slot/port	Selects the interface to configure.
Step 3	Router(config-if)# macro apply cisco-desktop \$AVID access_vlan_ID	Applies the cisco-desktop smart port macro. The recommended range for <i>access_vlan_ID</i> is 2–4094.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to apply the cisco-desktop smart port macro to Gigabit Ethernet port 1/1 with VLAN 2 specified as the access VLAN and how to verify the result:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/1
Router(config-if)# macro apply cisco-desktop $AVID 2
%Warning: portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION

%Portfast has been configured on GigabitEthernet1/1 but will only
have effect when the interface is in a non-trunking mode.
Router(config)# end
Router# show parser macro description interface gigabitethernet 1/1
Global Macro(s): cisco-global

Interface      Macro Description(s)
-----
Gi1/1          cisco-desktop
-----
```

```
Router# show running-config interface gigabitethernet 1/1
Building configuration...

Current configuration : 307 bytes
!
interface GigabitEthernet1/1
 switchport
 switchport access vlan 2
 switchport mode access
 switchport port-security
 switchport port-security aging time 2
 switchport port-security violation restrict
 shutdown
 macro description cisco-desktop
 spanning-tree portfast
 spanning-tree bpduguard enable
end

Router#
```

Using the cisco-phone Smart Port Macro

- [Displaying the Contents of the cisco-phone Smart Port Macro, page 3-7](#)
- [Applying the cisco-phone Smart Port Macro, page 3-8](#)

Displaying the Contents of the cisco-phone Smart Port Macro

```
Router# show parser macro name cisco-phone
Macro name : cisco-phone
Macro type : default interface
# macro keywords $AVID $VVID
# VoIP enabled interface - Enable data VLAN
# and voice VLAN (VVID)
# Recommended value for access vlan (AVID) should not be 1
switchport
switchport access vlan $AVID
switchport mode access

# Update the Voice VLAN (VVID) value which should be
# different from data VLAN
# Recommended value for voice vlan (VVID) should not be 1
switchport voice vlan $VVID

# Enable port security limiting port to a 3 MAC
# addressess -- One for desktop and two for phone
switchport port-security
switchport port-security maximum 3

# Ensure port-security age is greater than one minute
# and use inactivity timer
switchport port-security violation restrict
switchport port-security aging time 2
# Enable auto-qos to extend trust to attached Cisco phone
auto qos voip cisco-phone

# Configure port as an edge network port
spanning-tree portfast
spanning-tree bpduguard enable
```

Applying the cisco-phone Smart Port Macro

To apply the cisco-phone smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 3	Router(config-if)# macro apply cisco-phone \$AVID <i>access_vlan_ID \$VVID voice_vlan_ID</i>	Applies the cisco-phone smart port macro. The recommended range for <i>access_vlan_ID</i> is 2–4094. The recommended range for <i>voice_vlan_ID</i> is 2–4094.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

When applying the cisco-phone smart port macro, note the following information:

- Some of the generated commands are in the category of PFC QoS commands that are applied to [all ports controlled by a port ASIC](#). When one of these generated commands is applied, PFC QoS displays the messages caused by application of the command to all the ports controlled by the port ASIC. Depending on the module, these commands are applied to as many as 48 ports. See the “Number of port groups” and “Port ranges per port group” listed for each module in the *Release Notes for Cisco IOS Release 15.2SY*:
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.1SY/release_notes.html
- You might see messages that instruct you to configure other ports to trust CoS. You must do so to enable the generated QoS commands.
- You might not be able to apply the cisco-phone smart port macro and other macros on ports that are controlled by the same port ASIC because of conflicting port trust state requirements.

This example shows how to apply the cisco-phone smart port macro to Gigabit Ethernet port 2/2 with VLAN 2 specified as the access VLAN and how to verify the result:

```
Router# configure terminal
Router(config)# interface gigabitethernet 2/2
Router(config-if)# macro apply cisco-phone $AVID 2 $VVID 3
Hardware QoS is enabled
Propagating cos-map to inband port
Propagating cos-map configuration to: [port list not shown]
```

[Output for other ports controlled by the same port ASIC omitted]

```
Warning: rcv cosmap will not be applied in hardware.
To modify rcv cosmap in hardware, all of the interfaces below
must be put into 'trust cos' state:
[port list not shown]
%Warning: portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION

%Portfast has been configured on GigabitEthernet1/2 but will only
have effect when the interface is in a non-trunking mode.
Router(config)# end
```

```
Router# show parser macro description interface gigabitethernet 2/2
Global Macro(s): cisco-global
```

```
Interface      Macro Description(s)
-----
```



```

Gi2/2          cisco-phone
-----

Router# show running-config interface gigabitethernet 2/2
Building configuration...

Building configuration...

Current configuration : 307 bytes
!
interface GigabitEthernet1/2
Building configuration...

Current configuration : 1336 bytes
!
interface GigabitEthernet2/2
 switchport
 switchport access vlan 2
 switchport mode access
 switchport voice vlan 3
 switchport port-security
 switchport port-security maximum 3
 switchport port-security aging time 2
 switchport port-security violation restrict
 shutdown

[QoS queuing commands omitted: these vary according to port type]

platform qos trust cos
auto qos voip cisco-phone
macro description cisco-phone
spanning-tree portfast
spanning-tree bpduguard enable
end

Router#

```

Using the cisco-switch Smart Port Macro

- [Displaying the Contents of the cisco-switch Smart Port Macro, page 3-9](#)
- [Applying the cisco-switch Smart Port Macro, page 3-10](#)

Displaying the Contents of the cisco-switch Smart Port Macro

```

Router# show parser macro name cisco-switch
Macro name : cisco-switch
Macro type : default interface
# macro keywords $NVID
# Do not apply to EtherChannel/Port Group
# Access Uplink to Distribution

# Define unique Native VLAN on trunk ports
# Recommended value for native vlan (NVID) should not be 1
switchport
switchport trunk native vlan $NVID

# Update the allowed VLAN range (VRANGE) such that it
# includes data, voice and native VLANs
# switchport trunk allowed vlan VRANGE

# Hardcode trunk and disable negotiation to

```

```
# speed up convergence
switchport trunk encapsulation dot1q
switchport mode trunk
switchport nonegotiate

# 802.1w defines the link as pt-pt for rapid convergence
spanning-tree link-type point-to-point

Router#
```

Applying the cisco-switch Smart Port Macro

To apply the cisco-switch smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 3	Router(config-if)# macro apply cisco-switch \$NVID native_vlan_ID	Applies the cisco-switch smart port macro. The recommended range for <i>native_vlan_ID</i> is 2–4094.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to apply the cisco-switch smart port macro to Gigabit Ethernet port 1/4 with VLAN 4 specified as the native VLAN and how to verify the result:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/4
Router(config-if)# macro apply cisco-switch $NVID 4
Router(config-if)# end
Router# show parser macro description interface gigabitethernet 1/4
Interface      Macro Description(s)
-----
Gil1/4         cisco-switch
-----

Router# show running-config interface gigabitethernet 1/4
Building configuration...

Current configuration : 247 bytes
!
interface GigabitEthernet1/4
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 4
 switchport mode trunk
 switchport nonegotiate
 shutdown
 macro description cisco-switch
 spanning-tree link-type point-to-point
end

Router#
```

Using the cisco-router Smart Port Macro

- [Displaying the Contents of the cisco-router Smart Port Macro, page 3-11](#)
- [Applying the cisco-router Smart Port Macro, page 3-11](#)

Displaying the Contents of the cisco-router Smart Port Macro

```
Router# show parser macro name cisco-router
Macro name : cisco-router
Macro type : default interface
# macro keywords $NVID
# Do not apply to EtherChannel/Port Group
# Access Uplink to Distribution
switchport

# Define unique Native VLAN on trunk ports
# Recommended value for native vlan (NVID) should not be 1
switchport trunk native vlan $NVID

# Update the allowed VLAN range (VRANGE) such that it
# includes data, voice and native VLANs
# switchport trunk allowed vlan VRANGE

# Hardcode trunk and disable negotiation to
# speed up convergence
switchport trunk encapsulation dot1q
switchport mode trunk
switchport nonegotiate

# Configure qos to trust this interface
auto qos voip trust
platform qos trust dscp

# Ensure fast access to the network when enabling the interface.
# Ensure that switch devices cannot become active on the interface.
spanning-tree portfast
spanning-tree bpduguard enable

Router#
```

Applying the cisco-router Smart Port Macro

To apply the cisco-router smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 3	Router(config-if)# macro apply cisco-router \$NVID native_vlan_ID	Applies the cisco-router smart port macro. The recommended range for <i>native_vlan_ID</i> is 2–4094.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

**Note**

The cisco-router smart port macro includes the **auto qos voip trust** command. When entered on a port configured with the **switchport** command, the **auto qos voip trust** command generates and applies the **platform qos trust cos** command to the port, but the cisco-router smart port macro changes the port trust state to trust DSCP with the **platform qos trust dscp** command. When you apply the cisco-router smart port macro, ignore messages that instruct you to enter the **platform qos trust cos** command on other ports controlled by the port ASIC.

This example shows how to apply the cisco-router smart port macro to Gigabit Ethernet port 1/5 and how to verify the result:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/5
Router(config-if)# macro apply cisco-router $NVID 5
Hardware QoS is enabled
Propagating cos-map to inband port
Propagating cos-map configuration to: [port list not shown]
```

[Output for other ports controlled by the same port ASIC omitted]

[Output from temporarily applied trust CoS command omitted]

```
%Warning: portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION
```

```
%Portfast has been configured on GigabitEthernet1/5 but will only
have effect when the interface is in a non-trunking mode.
Router(config-if)# end
Router# show parser macro description interface gigabitethernet 1/5
Interface      Macro Description(s)
```

```
-----
Gil/5          cisco-router
-----
```

```
Router# show running-config interface gigabitethernet 1/5
Building configuration...
```

```
Current configuration : 1228 bytes
!
interface GigabitEthernet1/5
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 5
 switchport mode trunk
 switchport nonegotiate
 shutdown
 wrp-queue bandwidth 20 100 200
```

[QoS queuing commands omitted: these vary according to port type]

```
platform qos trust dscp
auto qos voip trust
macro description cisco-router
spanning-tree portfast
spanning-tree bpduguard enable
end
```

```
Router#
```

Creating Smart Port Macros

- [Creating Smart Port Macros, page 3-13](#)
- [Applying User-Created Smart Port Macros, page 3-14](#)

Creating Smart Port Macros

To create a smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# macro name <i>macro-name</i>	Creates a macro. Macro names are case sensitive. For example, the commands macro name Sample-Macro and macro name sample-macro will result in two separate macros. A macro definition can contain up to 3,000 characters. Line endings count as two characters. There is no prompt displayed in macro creation mode. Enter the macro commands on separate lines. Use the # character at the beginning of a line to enter a comment within the macro. Use the @ character to end the macro. Do not use the exit or end commands or change the command mode with the interface interface-id in a macro. This could cause any commands following exit , end , or interface interface-id to execute in a different command mode. For best results, all commands in a macro should be in the same configuration mode. Each user-created macro can have up to three keyword-value pairs.
Step 3	# macro keywords <i>keyword1 keyword2 keyword3</i>	(Optional) You can create a help string to describe the keywords that you define in the macro. You can enter up to three help string comments in a macro.
Step 4	end	Returns to privileged EXEC mode.
Step 5	show parser macro name <i>macro-name</i>	Verifies that the macro was created.



Note

The **no** form of the **macro name** global configuration command only deletes the macro definition. It does not affect the configuration of those interfaces on which the macro is already applied.

This example shows how to create a macro that defines the Layer 2 access VLAN and the number of secure MAC addresses and also includes two help string keywords by using # **macro keywords**:

```
Router(config)# macro name test
#macro keywords $VLANID $MAX
switchport access vlan $VLANID
switchport port-security maximum $MAX
@
```

Applying User-Created Smart Port Macros

To apply a smart port macro, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# default interface <i>interface-id</i>	(Optional) Clears all configuration from the specified interface.
Step 3	Router(config)# interface <i>interface_id</i>	(Required for interface macros.) Specifies the interface on which to apply the macro and enters interface configuration mode.
Step 4	Router(config)# macro [global] { apply trace } <i>macro-name</i> [<i>keyword value</i>] [<i>keyword value</i>] [<i>keyword value</i>]	Applies or traces and applies each individual command defined in the macro. For global macros: - To find any syntax or configuration errors, enter the macro global trace <i>macro-name</i> command to apply and debug the macro. - To display a list of any keyword-value pairs defined in the macro, enter the macro global apply <i>macro-name</i> ? command. For interface macros: - To find any syntax or configuration errors, enter the macro trace <i>macro-name</i> command to apply and debug the macro. - To display a list of any keyword-value pairs defined in the macro, enter the macro apply <i>macro-name</i> ? command. To successfully apply the macro, you must enter any required keyword-value pairs. Keyword matching is case sensitive. In the commands that the macro applies, all matching occurrences of keywords are replaced with the corresponding values.
Step 5	Router(config)# end	Returns to privileged EXEC mode.

You can delete a global macro-applied configuration on a switch only by entering the **no** version of each command that is in the macro. You can delete all configurations on an interface by entering the **default interface** *interface_id* interface configuration command.

This example shows how to apply the user-created macro called `snmp`, to set the host name address to **test-server** and to set the IP precedence value to 7:

```
Router(config)# macro global apply snmp ADDRESS test-server VALUE 7
```

This example shows how to debug the user-created macro called `snmp` by using the **macro global trace** global configuration command to find any syntax or configuration errors in the macro when it is applied to the switch:

```
Router(config)# macro global trace snmp VALUE 7
Applying command...'snmp-server enable traps port-security'
Applying command...'snmp-server enable traps linkup'
Applying command...'snmp-server enable traps linkdown'
Applying command...'snmp-server host'
%Error Unknown error.
Applying command...'snmp-server ip precedence 7'
```

This example shows how to apply the user-created macro called `desktop-config` and to verify the configuration:

```
Router(config)# interface gigabitethernet1/2
Router(config-if)# macro apply desktop-config
Router(config-if)# end
Router# show parser macro description
Interface      Macro Description
-----
Gil/2         desktop-config
-----
```

This example shows how to apply the user-created macro called `desktop-config` and to replace all occurrences of `vlan` with VLAN ID 25:

```
Router(config-if)# macro apply desktop-config vlan 25
```

Verifying the Smart Port Macro Configuration

Table 3-2 Commands to Display Smartports Macros

Command	Purpose
<code>show parser macro</code>	Displays all configured macros.
<code>show parser macro name macro-name</code>	Displays a specific macro.
<code>show parser macro brief</code>	Displays the configured macro names.
<code>show parser macro description [interface interface-id]</code>	Displays the macro description for all interfaces or for a specified interface.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

4

Virtual Switching Systems

- [Prerequisites for VSS, page 4-1](#)
- [Restrictions for VSS, page 4-2](#)
- [Information About Virtual Switching Systems, page 4-4](#)
- [Default Settings for VSS, page 4-31](#)
- [How to Configure a VSS, page 4-31](#)
- [How to Upgrade a VSS, page 4-54](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for VSS

The VSS configurations in the startup-config file must match on both chassis.

Restrictions for VSS

- [General VSS Restrictions, page 4-2](#)
- [VSL Restrictions, page 4-2](#)
- [Multichassis EtherChannel \(MEC\) Restrictions, page 4-2](#)
- [Dual-Active Detection Restrictions, page 4-3](#)

General VSS Restrictions

- VSS mode does not support supervisor engine redundancy within a chassis.
- If you configure a new value for switch priority, the change takes effect only after you save the configuration file and perform a restart.
- Out-of-band MAC address table synchronization among DFC-equipped switching modules (the **mac address-table synchronize** command) is enabled automatically in VSS mode, which is the recommended configuration.
- Because the output of the **show running-config** command on ICS supervisor engines could be out of sync with the active supervisor engine, ICS supervisor engines do not support the **show running-config** command. The active and standby supervisor engines support the **show running-config** command.

VSL Restrictions

- For line redundancy, we recommend configuring at least two ports per switch for the VSL. For module redundancy, the two ports can be on different switching modules in each chassis.
- The **no platform qos channel-consistency** command is automatically applied when you configure the VSL. Do not remove this command.
- VSL ports cannot be Mini Protocol Analyzer sources (the **monitor ... capture** command). Monitor capture sessions cannot be started if a source is the VSL on the port channel of the standby switch. The following message is displayed when a remote VSL port channel on the standby switch is specified and you attempt to start the monitor capture:

```
% remote VSL port is not allowed as capture source
```

The following message is displayed when a scheduled monitor capture start fails because a source is a remote VSL port channel:

```
Packet capture session 1 failed to start. A source port is a remote VSL.
```

Multichassis EtherChannel (MEC) Restrictions

- All links in an MEC must terminate locally on the active or standby chassis of the same virtual domain.
- For an MEC using the LACP control protocol, the *minlinks* command argument defines the minimum number of physical links in each chassis for the MEC to be operational.
- For an MEC using the LACP control protocol, the *maxbundle* command argument defines the maximum number of links in the MEC across the whole VSS.

- MEC supports LACP 1:1 redundancy. For additional information about LACP 1:1 redundancy, refer to the [“Information about LACP 1:1 Redundancy” section on page 23-6](#).
- An MEC can be connected to another MEC in a different VSS domain.
- Ports on the supervisor engines are not stateful and will experience a reset across switchovers (see the [“Switchover Process Restrictions” section on page 7-2](#)).
- With an MEC that includes supervisor engine ports configured on a VSS that supports the VSS Quad-Sup SSO (VS4O) feature, be aware of [CSCuh51005](#).

Dual-Active Detection Restrictions

- If Flex Links are configured on the VSS, use PAgP dual-active detection.
- For dual-active detection link redundancy, configure at least two ports per switch for dual-active detection. For module redundancy, the two ports can be on different switching modules in each chassis, and should be on different modules than the VSL, if feasible.
- When you configure dual-active fast hello mode, all existing configurations are removed automatically from the interface except for these commands:
 - **description**
 - **logging event**
 - **load-interval**
 - **rcv-queue cos-map**
 - **rcv-queue queue-limit**
 - **rcv-queue random-detect**
 - **rcv-queue threshold**
 - **wrr-queue bandwidth**
 - **wrr-queue cos-map**
 - **wrr-queue queue-limit**
 - **wrr-queue random-detect**
 - **wrr-queue threshold**
 - **priority-queue cos-map**
- Only these configuration commands are available on dual-active detection fast hello ports:
 - **default**
 - **description**
 - **dual-active**
 - **exit**
 - **load-interval**
 - **logging**
 - **no**
 - **shutdown**

- ASIC-specific QoS commands are not configurable on dual-active detection fast hello ports directly, but are allowed to remain on the fast hello port if the commands were configured on another non-fast hello port in that same ASIC group. For a list of these commands, see [Chapter 23, “Restrictions for PFC QoS.”](#)

Information About Virtual Switching Systems

- [VSS Overview, page 4-4](#)
- [VSS Redundancy, page 4-13](#)
- [Multichassis EtherChannels, page 4-15](#)
- [Packet Handling, page 4-18](#)
- [System Monitoring, page 4-22](#)
- [Convert Port Mode, page 4-23](#)
- [Dual-Active Detection, page 4-27](#)
- [VSS Initialization, page 4-29](#)

VSS Overview


- [VSS Topology, page 4-4](#)
- [Key Concepts, page 4-5](#)
- [VSS Functionality, page 4-7](#)
- [Hardware Requirements, page 4-10](#)
- [Information about VSL Topology, page 4-12](#)

VSS Topology

Network operators increase network reliability by configuring switches in redundant pairs and by provisioning links to both switches in the redundant pair. [Figure 4-1](#) shows a typical network configuration. Redundant network elements and redundant links can add complexity to network design and operation. Virtual switching simplifies the network by reducing the number of network elements and hiding the complexity of managing redundant switches and links.

VSS mode combines a pair of switches into a single network element. VSS mode manages the redundant links, which externally act as a single port channel.

VSS mode simplifies network configuration and operation by reducing the number of Layer 3 routing neighbors and by providing a loop-free Layer 2 topology.

Figure 4-1 Typical Network Design


Key Concepts

- [Virtual Switching System, page 4-5](#)
- [Active and Standby Chassis, page 4-6](#)
- [Virtual Switch Link, page 4-6](#)
- [Multichassis EtherChannel \(MEC\), page 4-7](#)

Virtual Switching System

A VSS combines a pair of switches into a single network element. For example, a VSS in the distribution layer of the network interacts with the access and core networks as if it were a single switch. See [Figure 4-2](#).

An access switch connects to both chassis of the VSS using one logical port channel. VSS mode manages redundancy and load balancing on the port channel. This capability enables a loop-free Layer 2 network topology. VSS mode also simplifies the Layer 3 network topology because VSS mode reduces the number of routing peers in the network.

Figure 4-2 VSS in the Distribution Network

Active and Standby Chassis


When you create or restart a VSS, the peer chassis negotiate their roles. One chassis becomes the active chassis, and the other chassis becomes the standby.

The active chassis controls the VSS. It runs the Layer 2 and Layer 3 control protocols for the switching modules on both chassis. The active chassis also provides management functions for the VSS, such as module online insertion and removal (OIR) and the console interface.

The active and standby chassis perform packet forwarding for ingress data traffic on their locally hosted interfaces. However, the standby chassis sends all control traffic to the active chassis for processing.

You can defer the traffic load on a multichassis EtherChannel (MEC) chassis to address traffic recovery performance during the standby chassis startup. For example, [Figure 4-3](#) represents network layout where a VSS (active and standby switches) is interacting with an upstream switch (switch 2) and a downstream switch (switch 1).


Figure 4-3 Switch Interconnected Through VSS



Virtual Switch Link

For the two chassis of the VSS to act as one network element, they need to share control information and data traffic.

The virtual switch link (VSL) is a special link that carries control and data traffic between the two chassis of a VSS, as shown in [Figure 4-4](#). The VSL is implemented as an EtherChannel with up to eight links. The VSL gives control traffic higher priority than data traffic so that control messages are never discarded. Data traffic is load balanced among the VSL links by the EtherChannel load-balancing algorithm.

Figure 4-4 Virtual Switch Link


181322

Multichassis EtherChannel (MEC)

An EtherChannel, which is configured on a port channel interface, is two or more physical links that combine to form one logical link. Layer 2 protocols operate on the EtherChannel as a single logical entity.

A MEC is a port channel with member ports on both chassis of the VSS. A connected non-VSS device views the MEC as a standard EtherChannel. See [Figure 4-5](#).

VSS mode supports a maximum of 512 EtherChannels. This limit applies to the combined total of regular EtherChannels and MECs. Because the VSL requires two EtherChannel numbers (one for each chassis), there are 510 user-configurable EtherChannels.

Figure 4-5 VSS with MEC

181323

**Note**

Ports on the supervisor engines are not stateful and will experience a reset across switchovers (see the [“Switchover Process Restrictions”](#) section on page 7-2).

VSS Functionality

- [Redundancy and High Availability, page 4-8](#)
- [Packet Handling, page 4-8](#)
- [System Management, page 4-8](#)
- [VSS Quad-Sup SSO \(VS4O\), page 4-9](#)
- [Interface Naming Convention, page 4-10](#)
- [Software Features, page 4-10](#)

Redundancy and High Availability

In VSS mode, supervisor engine redundancy operates between the active and standby chassis, using stateful switchover (SSO) and nonstop forwarding (NSF). The peer chassis exchange configuration and state information across the VSL and the standby supervisor engine runs in hot standby mode.

The standby chassis monitors the active chassis using the VSL. If it detects failure, the standby chassis initiates a switchover and takes on the active role. When the failed chassis recovers, it takes on the standby role.

**Note**

Ports on the supervisor engines are not stateful and will experience a reset across switchovers (see the [“Switchover Process Restrictions” section on page 7-2](#)).

If the VSL fails completely, the standby chassis assumes that the active chassis has failed, and initiates a switchover. After the switchover, if both chassis are active, the dual-active detection feature detects this condition and initiates recovery action. For additional information about dual-active detection, see the [“Dual-Active Detection” section on page 4-27](#).

Packet Handling

The active supervisor engine runs the Layer 2 and Layer 3 protocols and features for the VSS and manages the DFC modules for both chassis.

The VSS uses VSL to communicate protocol and system information between the peer chassis and to carry data traffic between the chassis when required.

Both chassis perform packet forwarding for ingress traffic on their interfaces. If possible, ingress traffic is forwarded to an outgoing interface on the same chassis to minimize data traffic that must traverse the VSL.

Because the standby chassis is actively forwarding traffic, the active supervisor engine distributes updates to the standby supervisor engine PFC and all standby chassis DFCs.

System Management

The active supervisor engine acts as a single point of control for the VSS. For example, the active supervisor engine handles OIR of switching modules on both chassis. The active supervisor engine uses VSL to send messages to and from local ports on the standby chassis.

The command console on the active supervisor engine is used to control both chassis. In virtual switch mode, the command console on the standby supervisor engine blocks attempts to enter configuration mode.

The standby chassis runs a subset of system management tasks. For example, the standby chassis handles its own power management.

VSS Quad-Sup SSO (VS40)


Release 15.1(1)SY1 and later releases support the VSS Quad-Sup SSO (VS40) feature.

Best Practice for VSL Links Between Supervisors

If there is a VSL link from ICA to ICA supervisor engine, ICS to ICS supervisor engine, and line card to line card, the control link will always be from the line card. The next preferred control link is from ICS to ICS VSL link, and the least preferred control link is from ICA to ICA VSL link.

In case of an ISSU image upgrade, if there is no cross VSL link, the control link will always be from the ICS supervisor engine. If there is a cross VSL link, the control link can be from any of the ICA supervisor engines.

Figure 4-6 Typical VSS Quad-Supervisor Configuration



These are the quad-supervisor VSS roles:

- In-chassis active (ICA) supervisor engines—The VSS active supervisor engine in one chassis and the VSS standby supervisor engine in the other chassis are ICA supervisor engines.

If the VSS active ICA supervisor engine crashes, a switchover to the standby ICA supervisor engine in other chassis occurs. Both VSS chassis remain active. All switching modules remain active.

In the chassis with the previously active ICA supervisor engine, an SSO switchover from the previously active ICA supervisor engine to the ICS standby supervisor engine occurs and the ICS takes over as the ICA. The failed ICA reloads and becomes the ICS.

You can verify the switchover mode of the supervisor engines by entering the **show module** command.

- In-chassis standby (ICS) supervisor engines—The other supervisor engines are ICS supervisor engines. The supervisor engine uplinks ports are available to forward traffic.



Note ICS supervisor engines do not support **show** commands. To avoid tracebacks, do not issue **show** commands on ICS supervisor engines.

If the supervisor engine PFC modes do not match then an ICS supervisor engine is reset to ROMMON. Configure both chassis to run in the same PFC mode.

ICS supervisor engines boot in SSO standby mode, which fully initializes and configures the ICS and maintains stateful feature and user session information.

To verify the switchover mode of the supervisor engines, enter the **show module** command.

When not in VSS quad supervisor engine mode, if you insert a supervisor engine to be an ICS, the supervisor engine resets to update the supervisor engine number and then reboots before going online.

Quad-supervisor SSO (VS4O) supports eFSU upgrades. You can upgrade or downgrade your VSS system using ISSU. See [“How to Upgrade a VSS” section on page 4-54](#) for more information about eFSU upgrades.

Interface Naming Convention

In VSS mode, interfaces are specified using switch number (in addition to slot and port), because the same slot numbers are used on both chassis. For example, the **interface 1/5/4** command specifies port 4 of the switching module in slot 5 of switch 1. The **interface 2/5/4** command specifies port 4 on the switching module in slot 5 of switch 2.

Software Features

With some exceptions, VSS mode has feature parity with non-VSS mode. Major exceptions include:

- VSS mode does not support supervisor engine redundancy within a chassis.
- Port-based QoS and ACLs can be applied to any physical port, except VSL ports. ACLs can be applied to no more than 2,046 ports.

Hardware Requirements

- [Chassis and Modules, page 4-10](#)
- [VSL Hardware Requirements, page 4-11](#)
- [PFC, DFC, and CFC Requirements, page 4-11](#)
- [Multichassis EtherChannel Requirements, page 4-11](#)
- [Service Module Support, page 4-12](#)

Chassis and Modules

Table 4-1 VSS Hardware Requirements

Hardware	Count	Requirements
Chassis	2	All chassis supported with Supervisor Engine 6T in Cisco IOS Release 15.4SY support VSS mode. Note The two chassis need not be identical.
Supervisor Engines	2	Either two C6800-SUP6T or two C6800-SUP6T-XLsupervisor engines. The two supervisor engines must match exactly.
Switching Modules	2+	VSS mode support as shown in the Release Notes. In VSS mode, unsupported switching modules remain powered off.

VSL Hardware Requirements

The VSL EtherChannel supports only 40-Gigabit and 10-Gigabit Ethernet ports. The ports can be located on the supervisor engine (recommended) or on one of the following switching modules:

- WS-X6904-40G-2T
- WS-X6908-10GE
- WS-X6816-10T-2T
- WS-X6816-10G-2T
- C6800-32P-10G and C6800-32P-10G-XL
- C6800-16P-10G and C6800-16P-10G-XL
- C6800-8P-10G and C6800-8P-10G-XL

Use any of the 10-Gigabit or 40-Gigabit Ethernet ports on the supervisor engines to create the VSL between the two chassis.

You can add more physical links to the VSL EtherChannel by using the 40-Gigabit or 10-Gigabit Ethernet ports on switching modules that support the VSL.



Note

- If you use 10-Gigabit Ethernet ports to form VSL, you can bundle a maximum of 8 links to the VSL ether channel.
- If you use 40-Gigabit Ethernet ports to form VSL, you can bundle a maximum of 4 links to the VSL ether channel.
- For switching modules **WS-X6816-10T-2T** and **WS-X6816-10G-2T**, when using the ports on these switching module as VSL links, you must operate the ports in performance mode, not in oversubscription mode. Enter the ***no hw-module switch x slot y oversubscription port-group num*** command when configuring the switching module. If you enter the ***no hw-module switch switch_number slot slot_number oversubscription*** command to configure non-oversubscription mode (performance mode), then only ports 1, 5, 9, and 13 are configurable; the other ports on the module are disabled.
- Port-groups are independent of each other and one or more port-groups can operate in non-oversubscribed mode for VSL with the unused ports administratively shutdown, while the others can still operate in oversubscribed mode.

PFC, DFC, and CFC Requirements

Switching modules with a DFC4, or DFC4XL support VSS mode.

With a PFC4, the VSS will automatically operate in PFC4 mode, even if some of the modules have a DFC4XL. With a PFC4XL, but some modules equipped with a DFC4, you need to configure the VSS to operate in PFC4 mode. The **platform hardware vsl pfc mode non-xl** configuration command sets the system to operate in PFC4 mode after the next restart. See the [“SSO Dependencies” section on page 4-29](#) for further details about this command.

Multichassis EtherChannel Requirements

Physical links from any module with a DFC4, or DFC4XL can be used to implement a Multichassis EtherChannel (MEC).

Service Module Support

**Note**

Service modules are not supported in Supervisor Engine 6T.

- Application Control Engine (ACE):
 - ACE20-MOD-K9
 - ACE30-MOD-K9
- ASA Services Module: WS-SVC-ASA-SM1-K9
- Firewall Services Module (FWSM): WS-SVC-FWM-1-K9
- Network Analysis Module (NAM):
 - WS-SVC-NAM-1
 - WS-SVC-NAM-2
 - WS-SVC-NAM3-6G-K9
- Wireless Services Module (WiSM):
 - WS-SVC-WISM-1-K9
 - WS-SVC-WISM2

**Note**

Before deploying a service module in VSS mode, upgrade the module to the minimum supported release in standalone mode. See the service module release notes for information about the minimum required service module software version.

Information about VSL Topology

A VSS is two chassis that communicate using the VSL, which is a special port group. Configure both of the 10-Gigabit Ethernet ports on the supervisor engines as VSL ports. Optionally, you can also configure the VSL port group to contain switching module 40- or 10-Gigabit Ethernet ports. This configuration provides additional VSL capacity. See [Figure 4-7](#) for an example topology.

Figure 4-7 VSL Topology Example

VSS Redundancy

- [Overview, page 4-13](#)
- [RPR and SSO Redundancy, page 4-13](#)
- [Failed Chassis Recovery, page 4-14](#)
- [VSL Failure, page 4-15](#)
- [User Actions, page 4-15](#)

Overview

A VSS operates stateful switchover (SSO) between the active and standby supervisor engines. Compared to standalone mode, VSS mode has the following important differences in its redundancy model:

- The active and standby supervisor engines are hosted in separate chassis and use the VSL to exchange information.
- The active supervisor engine controls both chassis of the VSS. The active supervisor engine runs the Layer 2 and Layer 3 control protocols and manages the switching modules on both chassis.
- The active and standby chassis both perform data traffic forwarding.

If the active supervisor engine fails, the standby supervisor engine initiates a switchover and assumes the active role.

RPR and SSO Redundancy

The VSS normally runs stateful switchover (SSO) between the active and standby supervisor engines (see [Figure 4-8](#)). The VSS determines the role of each supervisor engine during initialization.

Figure 4-8 Chassis Roles in VSS Mode

The VSS uses the VSL link to synchronize configuration data from the active to the standby supervisor engine. Also, protocols and features that support high availability synchronize their events and state information to the standby supervisor engine.

VSS mode operates with stateful switchover (SSO) redundancy if it meets the following requirements:

- Both supervisor engines are running the same software version.
- The VSL-related configuration in the two chassis matches.
- The PFC mode matches.
- SSO and nonstop forwarding (NSF) are configured on both chassis.

See the [“SSO Dependencies” section on page 4-29](#) for additional details about the requirements for SSO redundancy on a VSS. See [Chapter 8, “Nonstop Forwarding \(NSF\)”](#) for information about configuring SSO and NSF.

With SSO redundancy, the supervisor engine in the standby chassis runs in hot standby state and is always ready to assume control following a fault on the active supervisor engine. Configuration, forwarding, and state information are synchronized from the active supervisor engine to the redundant supervisor engine at startup and whenever changes to the active supervisor engine configuration occur. If a switchover occurs, traffic disruption is minimized.

If a VSS does not meet the requirements for SSO redundancy, the VSS uses route processor redundancy (RPR). In RPR mode, the active supervisor engine does not synchronize configuration changes or state information with the standby. The standby supervisor engine is only partially initialized and the switching modules on the standby supervisor are not powered up. If a switchover occurs, the standby supervisor engine completes its initialization and powers up the switching modules. Traffic is disrupted for approximately 2 minutes.

Failed Chassis Recovery

If the active chassis or supervisor engine fails, the VSS initiates a stateful switchover (SSO) and the former standby supervisor engine assumes the active role. The failed chassis performs recovery action by reloading the supervisor engine.

If the standby chassis or supervisor engine fails, no switchover is required. The failed chassis performs recovery action by reloading the supervisor engine.

The VSL links are unavailable while the failed chassis recovers. After the chassis reloads, it becomes the new standby chassis and the VSS reinitializes the VSL links between the two chassis.

The switching modules on the failed chassis are unavailable during recovery, so the VSS operates only with the MEC links that terminate on the active chassis. The bandwidth of the VSS is reduced until the failed chassis has completed its recovery and become operational again. Any devices that are connected only to the failed chassis experience an outage.



Note

The VSS may experience a brief data path disruption when the switching modules in the standby chassis become operational after the SSO.

After the SSO, much of the processing power of the active supervisor engine is consumed in bringing up a large number of ports simultaneously in the standby chassis. As a result, some links might be brought up before the supervisor engine has configured forwarding for the links, causing traffic to those links to be lost until the configuration is complete. This condition is especially disruptive if the link is an MEC link. Two methods are available to reduce data disruption following an SSO:

- You can configure the VSS to activate non-VSL ports in smaller groups over a period of time rather than all ports simultaneously. For information about deferring activation of the ports, see the [“Configuring Deferred Port Activation During Standby Recovery” section on page 4-48](#).

- You can defer the load sharing of the peer switch's MEC member ports during reestablishment of the port connections. See the [“Failed Chassis MEC Recovery” section on page 4-17](#) for details about load share deferral.

VSL Failure

To ensure fast recovery from VSL failures, fast link notification is enabled in virtual switch mode on all port channel members (including VSL ports) whose hardware supports fast link notification.



Note

Fast link notification is not compatible with link debounce mechanisms. In virtual switch mode, link debounce is disabled on all port channel members.

If a single VSL physical link goes down, the VSS adjusts the port group so that the failed link is not selected.

If the standby chassis detects complete VSL link failure, it initiates a stateful switchover (SSO). If the active chassis has failed (causing the VSL links to go down), the scenario is chassis failure, as described in the previous section.

If only the VSL has failed and the active chassis is still operational, this is a dual-active scenario. The VSS detects that both chassis are operating in active mode and performs recovery action. See the [“Dual-Active Detection” section on page 4-27](#) for additional details about the dual-active scenario.

User Actions

From the active chassis command console, you can initiate a VSS switchover or a reload.

If you enter the **reload** command from the command console, the entire VSS performs a reload.

To reload only the VSS standby supervisor engine:

- In a two supervisor VSS setup, use **redundancy reload shelf[1/2]** command.
- In a four supervisor VSS setup, use **redundancy reload peer** command,

To force a switchover from the active to the standby supervisor engine, use the **redundancy force-switchover** command.

Multichassis EtherChannels

- [Overview, page 4-15](#)
- [MEC Failure Scenarios, page 4-16](#)

Overview


A multichassis EtherChannel is an EtherChannel with ports that terminate on both chassis of the VSS (see [Figure 4-9](#)). A VSS MEC can connect to any network element that supports EtherChannel (such as a host, server, router, or switch).

At the VSS, an MEC is an EtherChannel with additional capability: the VSS balances the load across ports in each chassis independently. For example, if traffic enters the active chassis, the VSS will select an MEC link from the active chassis. This MEC capability ensures that data traffic does not unnecessarily traverse the VSL.

Each MEC can optionally be configured to support either PAgP or LACP. These protocols run only on the active chassis. PAgP or LACP control packets destined for an MEC link on the standby chassis are sent across VSL.

An MEC can support up to eight active physical links, which can be distributed in any proportion between the active and standby chassis.

Figure 4-9 MEC Topology



MEC Failure Scenarios

- [Single MEC Link Failure, page 4-16](#)
- [All MEC Links to the Active Chassis Fail, page 4-17](#)
- [All MEC Links to the Standby Chassis Fail, page 4-17](#)
- [All MEC Links Fail, page 4-17](#)
- [Standby Chassis Failure, page 4-17](#)
- [Active Chassis Failure, page 4-17](#)
- [Failed Chassis MEC Recovery, page 4-17](#)



Note

Configure the MEC with at least one link to each chassis. This configuration conserves VSL bandwidth (traffic egress link is on the same chassis as the ingress link), and increases network reliability (if one VSS supervisor engine fails, the MEC is still operational).

Single MEC Link Failure

If a link within the MEC fails (and other links in the MEC are still operational), the MEC redistributes the load among the operational links, as in a regular port.

All MEC Links to the Active Chassis Fail

If all links to the active chassis fail, the MEC becomes a regular EtherChannel with operational links to the standby chassis.

Data traffic terminating on the active chassis reaches the MEC by crossing the VSL to the standby chassis. Control protocols continue to run in the active chassis. Protocol messages reach the MEC by crossing the VSL.

All MEC Links to the Standby Chassis Fail

If all links fail to the standby chassis, the MEC becomes a regular EtherChannel with operational links to the active chassis.

Control protocols continue to run in the active chassis. All control and data traffic from the standby chassis reaches the MEC by crossing the VSL to the active chassis.

All MEC Links Fail

If all links in an MEC fail, the logical interface for the EtherChannel is set to unavailable. Layer 2 control protocols perform the same corrective action as for a link-down event on a regular EtherChannel.

On adjacent switches, routing protocols and Spanning Tree Protocol (STP) perform the same corrective action as for a regular EtherChannel.

Standby Chassis Failure

If the standby chassis fails, the MEC becomes a regular EtherChannel with operational links on the active chassis. Connected peer switches detect the link failures, and adjust their load-balancing algorithms to use only the links to the active chassis.

Active Chassis Failure

Active chassis failure results in a stateful switchover (SSO). See the [“VSS Redundancy” section on page 4-13](#) for details about SSO on a VSS. After the switchover, the MEC is operational on the new active chassis. Connected peer switches detect the link failures (to the failed chassis), and adjust their load-balancing algorithms to use only the links to the new active chassis.

Failed Chassis MEC Recovery

When a failed chassis returns to service as the new standby chassis, protocol messages reestablish the MEC links between the recovered chassis and connected peer switches.

Although the recovered chassis' MEC links are immediately ready to receive unicast traffic from the peer switch, received multicast traffic may be lost for a period of several seconds to several minutes. To reduce this loss, you can configure the port load share deferral feature on MEC port channels of the peer switch. When load share deferral is configured, the peer's deferred MEC port channels will establish with an initial load share of 0. During the configured deferral interval, the peer's deferred port channels are capable of receiving data and control traffic, and of sending control traffic, but are unable to forward data traffic to the VSS. See the [“Configuring Port Load Share Deferral on the Peer Switch” section on page 4-49](#) for details about configuring port load share deferral.

Packet Handling

- [Packet Handling Overview, page 4-18](#)
- [Traffic on the VSL, page 4-18](#)
- [Layer 2 Protocols, page 4-19](#)
- [Layer 3 Protocols, page 4-19](#)
- [SPAN Support with VSS, page 4-21](#)

Packet Handling Overview

In VSS mode, the active supervisor engine runs the Layer 2 and Layer 3 protocols and features for the VSS and manages the DFC modules for both chassis.

The VSS uses the VSL to communicate system and protocol information between the peer chassis and to carry data traffic between the two chassis.

Both chassis perform packet forwarding for ingress traffic on their local interfaces. VSS mode minimizes the amount of data traffic that must traverse the VSL.

Traffic on the VSL

The VSL carries data traffic and in-band control traffic between the two chassis. All frames forwarded over the VSL link are encapsulated with a special 32-byte header, which provides information for the VSS to forward the packet on the peer chassis.

The VSL transports control messages between the two chassis. Messages include protocol messages that are processed by the active supervisor engine, but received or transmitted by interfaces on the standby chassis. Control traffic also includes module programming between the active supervisor engine and switching modules on the standby chassis.

The VSS needs to transmit data traffic over the VSL under the following circumstances:

- Layer 2 traffic flooded over a VLAN (even for dual-homed links).
- Packets processed by software on the active supervisor engine where the ingress interface is on the standby chassis.
- The packet destination is on the peer chassis, such as the following examples:
 - Traffic within a VLAN where the known destination interface is on the peer chassis.
 - Traffic that is replicated for a multicast group and the multicast receivers are on the peer chassis.
 - The known unicast destination MAC address is on the peer chassis.
 - The packet is a MAC notification frame destined for a port on the peer chassis.

VSL also transports system data, such as NetFlow export data and SNMP data, from the standby chassis to the active supervisor engine.

To preserve the VSL bandwidth for critical functions, the VSS uses strategies to minimize user data traffic that must traverse the VSL. For example, if an access switch is dual-homed (attached with an MEC terminating on both VSS chassis), the VSS transmits packets to the access switch using a link on the same chassis as the ingress link.

Traffic on the VSL is load-balanced with the same global hashing algorithms available for EtherChannels (the default algorithm is source-destination IP).

Layer 2 Protocols

- [Layer 2 Protocol Overview, page 4-19](#)
- [Spanning Tree Protocol, page 4-19](#)
- [Virtual Trunk Protocol, page 4-19](#)
- [EtherChannel Control Protocols, page 4-19](#)
- [Multicast Protocols, page 4-19](#)

Layer 2 Protocol Overview

The active supervisor engine runs the Layer 2 protocols (such as STP and VTP) for the switching modules on both chassis. Protocol messages that are transmitted and received on the standby chassis switching modules must traverse the VSL to reach the active supervisor engine.

Spanning Tree Protocol

The active chassis runs Spanning Tree Protocol (STP). The standby chassis redirects STP BPDUs across the VSL to the active chassis.

The STP bridge ID is commonly derived from the chassis MAC address. To ensure that the bridge ID does not change after a switchover, the VSS continues to use the original chassis MAC address for the STP Bridge ID.

Virtual Trunk Protocol

Virtual Trunk Protocol (VTP) uses the IP address of the switch and local current time for version control in advertisements. After a switchover, VTP uses the IP address of the newly active chassis.

EtherChannel Control Protocols

Link Aggregation Control Protocol (LACP) and Port Aggregation Protocol (PAgP) packets contain a device identifier. The VSS defines a common device identifier for both chassis to use.

A new PAgP enhancement has been defined for assisting with dual-active scenario detection. For additional information, see the [“Dual-Active Detection” section on page 4-27](#).

Multicast Protocols

With Release 15.1(1)SY1 and later releases, fast-redirect optimization makes multicast traffic redirection between inter-chassis or intra-chassis line cards faster for Layer 2 trunk and Layer3 multichassis EtherChannel or distributed EtherChannel in case of member port link failure and recovery. This operation occurs mainly when a member port link goes down (port leaves the EtherChannel) and when the member port link goes up (port joins or rejoins the EtherChannel). Fast-redirect does not take effect when you add or remove a member port due to a configuration change or during system boot up.

Layer 3 Protocols

- [Layer 3 Protocol Overview, page 4-20](#)
- [IPv4, page 4-20](#)
- [IPv6, MPLS, and VPLS, page 4-20](#)

- [IPv4 Multicast, page 4-20](#)
- [Software Features, page 4-21](#)

Layer 3 Protocol Overview

The RP on the active supervisor engine runs the Layer 3 protocols and features for the VSS. Both chassis perform packet forwarding for ingress traffic on their interfaces. If possible, ingress traffic is forwarded to an outgoing interface on the same chassis, to minimize data traffic that must traverse the VSL.

Because the standby chassis is actively forwarding traffic, the active supervisor engine distributes updates to the standby supervisor engine PFC and all standby chassis DFCs.

IPv4

The supervisor engine on the active chassis runs the IPv4 routing protocols and performs any required software forwarding.

Routing updates received on the standby chassis are redirected to the active chassis across the VSL.

Hardware forwarding is distributed across all DFCs on the VSS. The supervisor engine on the active chassis sends FIB updates to all local DFCs, remote DFCs, and the standby supervisor engine PFC.

All hardware routing uses the router MAC address assigned by the active supervisor engine. After a switchover, the original MAC address is still used.

The supervisor engine on the active chassis performs all software forwarding (for protocols such as IPX) and feature processing (such as fragmentation and TTL exceed). If a switchover occurs, software forwarding is disrupted until the new active supervisor engine obtains the latest CEF and other forwarding information.

In virtual switch mode, the requirements to support non-stop forwarding (NSF) are the same as in standalone mode. See [Chapter 8, “Nonstop Forwarding \(NSF\).”](#)

From a routing peer perspective, EtherChannels remain operational during a switchover (only the links to the failed chassis are down).

The VSS implements path filtering by storing only local paths (paths that do not traverse the VSL) in the FIB entries. Therefore, IP forwarding performs load sharing among the local paths. If no local paths to a given destination are available, the VSS updates the FIB entry to include remote paths (reachable by traversing the VSL).

IPv6, MPLS, and VPLS

The VSS supports IPv6 unicast, MPLS, and VPLS.

IPv4 Multicast

The IPv4 multicast protocols run on the active supervisor engine. Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) protocol packets received on the standby supervisor engine are transmitted across VSL to the active chassis.

The active supervisor engine sends IGMP and PIM protocol packets to the standby supervisor engine in order to maintain Layer 2 information for stateful switchover (SSO).

The active supervisor engine distributes multicast FIB and adjacency table updates to the standby supervisor engine and switching module DFCs.

For Layer 3 multicast in the VSS, learned multicast routes are stored in hardware in the standby supervisor engine. After a switchover, multicast forwarding continues, using the existing hardware entries.

**Note**

To avoid multicast route changes as a result of the switchover, we recommend that all links carrying multicast traffic be configured as MEC rather than Equal Cost Multipath (ECMP).

In virtual switch mode, the active chassis does not program the multicast expansion table (MET) on the standby chassis. The standby supervisor engine programs the outgoing interface hardware entries for all local multicast receivers.

If all switching modules on the active chassis and standby chassis are egress capable, the multicast replication mode is set to egress mode; otherwise, the mode is set to ingress mode.

In egress replication mode, replication is distributed to DFCs that have ports in outgoing VLANs for a particular flow. In ingress mode, replication for all outgoing VLANs is done on the ingress DFC.

For packets traversing VSL, all Layer 3 multicast replication occurs on the ingress chassis. If there are multiple receivers on the egress chassis, replicated packets are forwarded over the VSL.

Software Features

Software features run only on the active supervisor engine. Incoming packets to the standby chassis that require software processing are sent across the VSL.

For features supported in hardware, the ACL configuration is sent to the TCAM manager on the active supervisor engine, the standby supervisor engine, and all DFCs.

SPAN Support with VSS

The VSS supports all SPAN features for non-VSL interfaces. The VSS supports SPAN features on VSL interfaces with the following limitations:

- VSL ports cannot be a SPAN destination.
- VSL ports cannot be an RSPAN, ERSPAN, or egress-only SPAN source.
- If a VSL port is configured as a local SPAN source, the SPAN destination interface must be on the same chassis as the source interface.
- SPAN copies are always made on the chassis where the ingress port is located.
- Two VSLs cannot share the same SPAN session.
- A pair of LTL indices are used to avoid duplicate SPAN copies across VSL interfaces.

The number of SPAN sessions available to a VSS is the same as for a single chassis running in standalone mode.

With a VSL port as a SPAN source, the following limitations apply:

- The SPAN destination must be on the same chassis.
- Port channel interfaces cannot be the SPAN destination.

System Monitoring

- [Power Management, page 4-22](#)
- [Environmental Monitoring, page 4-22](#)
- [File System Access, page 4-22](#)
- [Diagnostics, page 4-22](#)
-
- [Network Management, page 4-23](#)

Power Management

You can control power-related functions for the standby chassis from the active chassis. For example, to control power to the modules and slots on the standby chassis, use the following commands:

(no) power enable switch <1/2> module <1-7> for C6807-XL chassis.

(no) power enable switch <1/2> module <1-13> for C6500 E-series chassis.

Use the **show power switch<1/2>** command to see the current power settings and status.

Environmental Monitoring

Environmental monitoring runs on both supervisor engines. The standby chassis reports notifications to the active supervisor engine. The active chassis gathers log messages for both chassis. The active chassis synchronizes the calendar and system clock to the standby chassis.

File System Access

You can access file systems of both chassis from the active chassis. Prefix the device name with the switch number and slot number to access directories on the standby chassis. For example, the command **dir sw2-slot6-disk0:** lists the contents of disk0 on the standby chassis (assuming switch 2 is the standby chassis). You can access the standby chassis file system only when VSL is operational.

Diagnostics

You can use the **diagnostic schedule** and **diagnostic start** commands on a VSS. In virtual switch mode, these commands require an additional parameter, which specifies the chassis to apply the command.

When you configure a VSL port on a switching module or a supervisor engine module, the diagnostics suite incorporates additional tests for the VSL ports.

Use the **show diagnostic content** command to display the diagnostics test suite for a module.

VSL Diagnostics

The following VSL-specific diagnostics tests are disruptive:

- TestVSetActiveToStandbyLoopback
- TestVslBridgeLink
- TestVslLocalLoopback

The following VSL-specific diagnostics test is available for VSL ports on switching modules or the supervisor engine. This test is not disruptive:

- TestVslStatus

Network Management

- [Telnet over SSH Sessions and the Web Browser User Interface, page 4-23](#)
- [SNMP, page 4-23](#)
- [Console Connections, page 4-23](#)

Telnet over SSH Sessions and the Web Browser User Interface

VSS mode supports remote access using Telnet over SSH sessions and the Cisco web browser user interface.

All remote access is directed to the active supervisor engine, which manages the VSS.

A VSS switchover disconnects Telnet over SSH sessions and web browser sessions.

SNMP

The SNMP agent runs on the active supervisor engine. CISCO-VIRTUAL-SWITCH-MIB is the MIB for VSS mode and contains the following main components:

- cvsGlobalObjects — Domain #, Switch #, Switch Mode
- cvsCoreSwitchConfig — Switch Priority
- cvsChassisTable — Chassis Role and Uptime
- cvsVSLConnectionTable — VSL Port Count, Operational State
- cvsVSLStatsTable — Total Packets, Total Error Packets
- cvsVSLPortStatsTable — TX/RX Good, Bad, Bi-dir and Uni-dir Packets

Console Connections

Connect console cables to both supervisor engine console ports. The console on the standby chassis adds the characters “-stdby” to the command line prompt to indicate that the chassis is operating in standby mode. You cannot enter configuration mode on the standby chassis console.

The following example shows the prompt on the standby console:

```
Router-stdby> show switch virtual
Switch mode                : Virtual Switch
Virtual switch domain number : 100
Local switch number        : 1
Local switch operational role: Virtual Switch Standby
Peer switch number         : 2
Peer switch operational role : Virtual Switch Active
```

Convert Port Mode

Supervisor 6T supports the following port modes:

- Default mode (with 8x10G and 2x40G ports operating in over-subscription mode)

- 16x10G ports in Over-subscription mode
- 8x10G ports in Performance mode
- 2x40G ports in Performance mode

**Note**

- The entire systems reloads after any mode conversion.
- Direct mode conversion from non-default mode to non-default mode is not supported.
- To convert from a non-default mode to another non-default mode, convert a supervisor from non-default to default mode and then default to another non-default mode. For example, to move from 8x10G Performance mode to 2x40G Performance mode, convert 8x10G Performance mode to default mode and then from default mode to 2x40G Performance mode.

Convert to 10G Over-subscription Mode (16x10G Ports)

To configure a SUP from default mode to the 10G over-subscription mode (16x10G Ports), use the following commands.

For non-VSS system:

```
hw-module slot {sup slot no.} operation-mode port-group 2 TenGigabitEthernet
```

For VSS system:

```
hw-module switch {switch-id} slot {sup slot no.} operation-mode port-group 2  
TenGigabitEthernet
```

Example

```
Router(config)#hw-module switch 2 slot 3 operation-mode port-group 2 TenGigabitEthernet  
Operation mode change will reload the entire system. Do you want to continue? [confirm]
```

To configure the SUP back to default mode from the 10G over-subscription mode (16x10G Ports), use the following commands.

For non-VSS system:

```
no/default hw-module slot {sup slot no.} operation-mode port-group 2 TenGigabitEthernet
```

For VSS system:

```
no/default hw-module switch {switch-id} slot {sup slot no.} operation-mode port-group 2  
TenGigabitEthernet
```

Example

```
Router(config)#no hw-module switch 2 slot 3 operation-mode port-group 2 TenGigabitEthernet  
Operation mode change will reload the entire system. Do you want to continue? [confirm]
```

Convert to 10G Performance Mode

To configure a SUP from default mode to the 10G performance mode, use the following commands.

For non-VSS system:

```
hw-module slot {sup slot no.} performance TenGigabitEthernet
```

For VSS system:

```
hw-module switch {switch-id} slot {sup slot no.} performance TenGigabitEthernet
```


Example

```
Router(config)#hw-module switch 1 slot 3 performance TenGigabitEthernet
Performance mode change will reload the entire system. Do you want to continue? [confirm]
```

To configure the SUP back to default mode from the 10G performance mode, use the following commands.

For non-VSS system:

```
no/default hw-module slot {sup slot no.} performance
```

For VSS system:

```
no/default hw-module switch {switch-id} slot {sup slot no.} performance
```

Example

```
Router(config)#no hw-module switch 1 slot 3 performance
Performance mode change will reload the entire system. Do you want to continue? [confirm]
```

Convert to 40G Performance Mode

To configure a SUP from default mode to the 40G performance mode, use the following commands.

For non-VSS system:

```
hw-module slot {sup slot no.} performance FortyGigabitEthernet
```

For VSS system:

```
hw-module switch {switch-id} slot {sup slot no.} performance FortyGigabitEthernet
```

Example

```
Router(config)#hw-module switch 1 slot 4 performance FortyGigabitEthernet
Performance mode change will reload the entire system. Do you want to continue? [confirm]
```

To configure the SUP back to default mode from the 40G performance mode, use the following commands.

For non-VSS system:

```
no/default hw-module slot {sup slot no.} performance
```

For VSS system:

```
no/default hw-module switch {switch-id} slot {sup slot no.} performance
```

Example

```
Router(config)#no hw-module switch 1 slot 4 performance
Performance mode change will reload the entire system. Do you want to continue? [confirm]
```

Verifying Supervisor Mode after Conversion

Use the following commands to verify the mode of the SUP.

For non-VSS system:

```
show hw-module slot {sup slot no.} operation mode
show hw-module slot {sup slot no.} oversubscription
show hw-module slot {sup slot no.} port-group-mapping
show interfaces status module {sup slot no.}
```

For VSS system:

```
show hw-module switch {switch-id} slot {sup slot no.} operation mode
show hw-module switch {switch-id} slot {sup slot no.} oversubscription
show hw-module switch {switch-id} slot {sup slot no.} port-group-mapping
show interfaces status switch {switch-id} module {sup slot no.}
```

Example (16x10G Oversubscription Mode)

```
Router#show hw-module switch 2 slot 3 operation mode
Switch 2 module 3 port group 1 is running in TenGigabitEthernet mode
Switch 2 module 3 port group 2 is running in TenGigabitEthernet mode
```

```
Router#show hw-module switch 2 slot 3 oversubscription
port-group oversubscription-mode
1 enabled
2 enabled
```

```
Router#show hw-module switch 2 slot 3 port-group-mapping
port-group ports
1 1, 2, 3, 4, 5, 6, 7, 8,
2 11, 12, 13, 14, 15, 16, 17, 18,
```

```
Router#show interfaces status switch 2 module 3
```

```
Port Name Status Vlan Duplex Speed Type
Te2/3/1 disabled routed full 10G No Connector
Te2/3/2 disabled routed full 10G No Connector
Te2/3/3 disabled routed full 10G No Connector
Te2/3/4 disabled routed full 10G No Connector
Te2/3/5 disabled routed full 10G No Connector
Te2/3/6 disabled routed full 10G No Connector
Te2/3/7 disabled routed full 10G No Connector
Te2/3/8 disabled routed full 10G No Connector
Te2/3/11 disabled routed full 10G No Connector
Te2/3/12 disabled routed full 10G No Connector
Te2/3/13 disabled routed full 10G No Connector
Te2/3/14 disabled routed full 10G No Connector
Te2/3/15 disabled routed full 10G No Connector
Te2/3/16 disabled routed full 10G No Connector
Te2/3/17 disabled routed full 10G No Connector
Te2/3/18 disabled routed full 10G No Connector
```

Example (8x10G Performance Mode)

```
Router#show hw-module switch 1 slot 3 operation mode
Switch 1 module 3 port group 1 is running in TenGigabitEthernet mode
Switch 1 module 3 port group 2 is running in TenGigabitEthernet mode
```

```
Router#show hw-module switch 1 slot 3 oversubscription
Module is in performance mode for native port Tengigabitethernet
```

```
Router#show hw-module switch 1 slot 3 port-group-mapping
port-group ports
1 5, 6, 7, 8,
2 1, 2, 3, 4,
```

```
Router#show interfaces status switch 1 module 3
```

```
Port Name Status Vlan Duplex Speed Type
Te1/3/1 disabled routed full 10G No Connector
Te1/3/2 disabled routed full 10G No Connector
Te1/3/3 disabled routed full 10G No Connector
Te1/3/4 disabled routed full 10G No Connector
```

```

Tel/3/5 disabled routed full 10G No Connector
Tel/3/6 disabled routed full 10G No Connector
Tel/3/7 disabled routed full 10G No Connector
Tel/3/8 disabled routed full 10G No Connector
mgmt0 disabled routed a-full a-100 10/100/1000BaseT

```

Example (2x40G Performance Mode)

```

Router#show hw-module switch 1 slot 4 operation mode
Switch 1 module 4 port group 1 is running in FortyGigabitEthernet mode
Switch 1 module 4 port group 2 is running in FortyGigabitEthernet mode

```

```

Router#show hw-module switch 1 slot 4 oversubscription
Module is in performance mode for native port Fortygigabitethernet

```

```

Router#show hw-module switch 1 slot 4 port-group-mapping
port-group ports
1 10,
2 9,

```

```

Router#show interfaces status switch 1 module 4

```

```

Port Name Status Vlan Duplex Speed Type
Fo1/4/9 disabled routed full 40G No Connector
Fo1/4/10 disabled routed full 40G No Connector

```

Dual-Active Detection

- [Dual-Active Detection Overview, page 4-27](#)
- [Dual-Active Detection Using Enhanced PAgP, page 4-28](#)
- [Dual-Active Detection Using Dual-Active Fast Hello Packets, page 4-28](#)
- [Recovery Actions, page 4-28](#)

Dual-Active Detection Overview

If the VSL fails, the standby chassis cannot determine the state of the active chassis. To ensure that switchover occurs without delay, the standby chassis assumes the active chassis has failed and initiates switchover to take over the active role.

If the original active chassis is still operational, both chassis are now active. This situation is called a *dual-active scenario*. A dual-active scenario can have adverse affects on network stability, because both chassis use the same IP addresses, SSH keys, and STP bridge ID. The VSS must detect a dual-active scenario and take recovery action.

The VSS supports these following methods for detecting a dual-active scenario:

- Enhanced PAgP—Uses PAgP messaging over the MEC links to communicate between the two chassis through a neighbor switch.
- dual-active fast-hello—Uses special hello messages over a backup Ethernet connection.
- dual-homed FEX—Uses SDP packets over the dual-homed RSL links to communicate between the two chassis.

You can configure any of the detection methods to be active at the same time.

For line redundancy, we recommend dedicating at least two ports per switch for dual-active detection. For module redundancy, the two ports can be on different switching modules in each chassis, and should be on different modules than the VSL links, if feasible.

Dual-Active Detection Using Enhanced PAgP

If a VSS MEC terminates on a Cisco switch, you can run the port aggregation protocol (PAgP) on the MEC. If enhanced PAgP is running on an MEC between the VSS and another switch running Release 12.2(33)SXH1 or a later release, the VSS can use enhanced PAgP to detect a dual-active scenario.

The MEC must have at least one port on each chassis of the VSS. In VSS mode, PAgP messages include a new type length value (TLV) that contains the ID of the VSS active switch. Only switches in VSS mode send the new TLV.

When the VSS standby chassis detects VSL failure, it initiates SSO and becomes VSS active. Subsequent PAgP messages to the connected switch from the newly VSS active chassis contain the new VSS active ID. The connected switch sends PAgP messages with the new VSS active ID to both VSS chassis.

If the formerly active chassis is still operational, it detects the dual-active scenario because the active ID in the PAgP messages changes. This chassis initiates recovery actions as described in the [“Recovery Actions” section on page 4-28](#).

Dual-Active Detection Using Dual-Active Fast Hello Packets

To use the dual-active fast hello packet detection method, you must provision a direct Ethernet connection between the two VSS chassis. You can dedicate up to four non-VSL links for this purpose.

The two chassis periodically exchange special Layer 2 dual-active hello messages containing information about the switch state. If the VSL fails and a dual-active scenario occurs, each switch recognizes from the peer's messages that there is a dual-active scenario and initiates recovery actions as described in the [“Recovery Actions” section on page 4-28](#). If a switch does not receive an expected dual-active fast hello message from the peer before the timer expires, the switch assumes that the link is no longer capable of dual-active detection. For more information, see the [“Configuring Enhanced PAgP Dual-Active Detection” section on page 4-50](#).

Recovery Actions

An active chassis that detects a dual-active condition shuts down all of its non-VSL interfaces (except interfaces configured to be excluded from shutdown) to remove itself from the network, and waits in recovery mode until the VSL links have recovered. You might need to physically repair the VSL failure. When the shut down chassis detects that VSL is operational again, the chassis reloads and returns to service as the standby chassis.

Loopback interfaces are also shut down in recovery mode. Do not configure loopback interfaces while in recovery mode, because any new loopback interfaces configured in recovery mode will not be shut down.



Note

If the running configuration of the chassis in recovery mode has been changed without saving, the chassis will not automatically reload. In this situation, you must save the running configuration and then reload manually.

VSS Initialization

- [VSS Initialization Overview, page 4-29](#)
- [Virtual Switch Link Protocol, page 4-29](#)
- [SSO Dependencies, page 4-29](#)
- [Initialization Procedure, page 4-30](#)

VSS Initialization Overview

A VSS is formed when the two chassis and the VSL link between them become operational. The peer chassis communicate over the VSL to negotiate the chassis roles.

If only one chassis becomes operational, it assumes the active role. The VSS forms when the second chassis becomes operational and both chassis bring up their VSL interfaces.

Virtual Switch Link Protocol

The Virtual Switch Link Protocol (VSLP) consists of several protocols that contribute to virtual switch initialization. The VSLP includes the following protocols:

- **Role Resolution Protocol**—The peer chassis use Role Resolution Protocol (RRP) to negotiate the role (active or standby) for each chassis.
- **Link Management Protocol**—The Link Management Protocol (LMP) runs on all VSL links, and exchanges information required to establish communication between the two chassis. LMP identifies and rejects any unidirectional links. If LMP flags a unidirectional link, the chassis that detects the condition brings the link down and up to restart the VSLP negotiation. VSL moves the control traffic to another port if necessary.

SSO Dependencies

For the VSS to operate with SSO redundancy, the VSS must meet the following conditions:

- **Identical software versions**—Both supervisor engine modules on the VSS must be running the identical software version.
- **VSL configuration consistency**—During the startup sequence, the standby chassis sends virtual switch information from the startup-config file to the active chassis. The active chassis ensures that the following information matches correctly on both chassis:
 - Switch virtual domain
 - Switch virtual node
 - Switch priority
 - VSL port channel: switch virtual link identifier
 - VSL ports: channel-group number, shutdown, total number of VSL ports
 - Power redundancy-mode
 - Power enable on VSL modules

If the VSS detects a mismatch, it prints out an error message on the active chassis console and the standby chassis comes up in RPR mode.

After you correct the configuration file, save the file by entering the **copy running-config startup-config** command on the active chassis, and then restart the standby chassis.

- PFC mode check—If both supervisor engines are provisioned with PFC4, the VSS will automatically operate in PFC4 mode, even if some of the switching modules are equipped with DFC4XLs.

However, if the supervisor engines are provisioned with PFC4XL and there is a mixture of DFC4 and DFC4XL switching modules, the system PFC mode will depend on how the DFC4XL and DFC4XL switching modules are distributed between the two chassis.

Each chassis in the VSS determines its system PFC mode. If the supervisor engine of a given chassis is provisioned with PFC4XL and all the switching modules in the chassis are provisioned with DFC4XL, the PFC mode for the chassis is PFC4XL. However, if any of the switching modules is provisioned with DFC4, the chassis PFC mode will be set to PFC4. If there is a mismatch between the PFC modes of two chassis, the VSS will come up in RPR mode instead of SSO mode. You can prevent this situation by using the **platform hardware vsl pfc mode non-xl** command to force the VSS to operate in PFC4 mode after the next reload.

- SSO and NSF enabled—SSO and NSF must be configured and enabled on both chassis. For detailed information on configuring and verifying SSO and NSF, see [Chapter 8, “Nonstop Forwarding \(NSF\)”](#).

If these conditions are not met, the VSS operates in RPR redundancy mode. For a description of SSO and RPR, see the [“VSS Redundancy” section on page 4-13](#).

Initialization Procedure

- [VSL Initialization, page 4-30](#)
- [System Initialization, page 4-30](#)
- [VSL Down, page 4-31](#)

VSL Initialization

A VSS is formed when the two chassis and the VSL link between them become operational. Because both chassis need to be assigned their role (active or standby) before completing initialization, VSL is brought online before the rest of the system is initialized. The initialization sequence is as follows:

1. The VSS initializes all cards with VSL ports, and then initializes the VSL ports.
2. The two chassis communicate over VSL to negotiate their roles (active or standby).
3. The active chassis completes the boot sequence, including the consistency check described in the [“SSO Dependencies” section on page 4-29](#).
4. If the consistency check completed successfully, the standby chassis comes up in SSO standby mode. If the consistency check failed, the standby chassis comes up in RPR mode.
5. The active chassis synchronizes configuration and application data to the standby chassis.

System Initialization

If you boot both chassis simultaneously, the VSL ports become active, and the chassis will come up as active and standby. If priority is configured, the higher priority switch becomes active.

If you boot up only one chassis, the VSL ports remain inactive, and the chassis comes up as active. When you subsequently boot up the other chassis, the VSL links become active, and the new chassis comes up as standby.

VSL Down

If the VSL is down when both chassis try to boot up, the situation is similar to a dual-active scenario. One of the chassis becomes active and the other chassis initiates recovery from the dual-active scenario. For further information, see the [“Configuring Dual-Active Detection” section on page 4-50](#).

Default Settings for VSS

None.

How to Configure a VSS

- [Configuring Easy VSS, page 4-31](#)
- [Converting to a VSS, page 4-32](#)
- [Displaying VSS Information, page 4-38](#)
- [Converting a VSS to Standalone Chassis, page 4-39](#)
- [Configuring VSS Parameters, page 4-40](#)
- [Configuring Multichassis EtherChannels, page 4-49](#)
- [Configuring Port Load Share Deferral on the Peer Switch, page 4-49](#)
- [Configuring Dual-Active Detection, page 4-50](#)
-
- [How to Upgrade a VSS, page 4-54](#)

Configuring Easy VSS

Easy VSS mode allows the conversion of standalone switches into one virtual switching system (VSS) using a single command.

The following prerequisites must be fulfilled for successful configuration of Easy VSS:

- Both the switches must be in standalone mode
- Easy VSS mode must be enabled on both the switches
- CDP should be running between interfaces
- Any manual configurations done using **switch virtual domain** *domain id* command will be lost

To configure Easy VSS mode, perform the following tasks:

Converting to a VSS

- [VSS Conversion Overview, page 4-32](#)
- [Backing Up the Standalone Configuration, page 4-33](#)
- [Configuring SSO and NSF, page 4-33](#)
- [Assigning Virtual Switch Domain and Switch Numbers, page 4-34](#)
- [Configuring the VSL Port Channel, page 4-35](#)
- [Configuring the VSL Ports, page 4-35](#)
- [Verifying the PFC Operating Mode, page 4-36](#)
- [Converting the Chassis to Virtual Switch Mode, page 4-37](#)
- [Auto-Configuring the Standby VSL Information, page 4-37](#)
- [\(Optional\) Configuring Standby Chassis Modules, page 4-38](#)

VSS Conversion Overview

The standalone mode is the default operating mode (a single chassis switch). VSS mode combines two standalone switches into one virtual switching system (VSS), operating in VSS mode.



Note


When you convert two standalone switches into one VSS, all non-VSL configuration settings on the standby chassis revert to default settings.

To convert two standalone chassis into a VSS, perform the following major activities:

- Save the standalone configuration files.
- Configure SSO and NSF on each chassis.
- Configure each chassis as a VSS.
- Convert to a VSS.
- Configure the peer VSL information.

In the procedures that follow, the example commands assume the configuration shown in [Figure 4-10](#).

Figure 4-10 Example VSS



Two chassis, A and B, are converted into a VSS with virtual switch domain 100. 10-Gigabit Ethernet port 5/1 on Switch 1 is connected to 10-Gigabit Ethernet port 5/2 on Switch 2 to form the VSL.

Backing Up the Standalone Configuration

Save the configuration files for both chassis. These files are needed to revert to standalone mode from virtual switch mode.

Switch 1 Task

	Command	Purpose
Step 1	Switch-1# copy running-config startup-config	(Optional) Saves the running configuration to startup configuration.
Step 2	Switch-1# copy startup-config disk0:old-startup-config	Copies the startup configuration to a backup file.

Switch 2 Task

	Command	Purpose
Step 1	Switch-2# copy running-config startup-config	(Optional) Saves the running configuration to the startup configuration file.
Step 2	Switch-2# copy startup-config disk0:old-startup-config	Copies the startup configuration to a backup file.

Configuring SSO and NSF

SSO and NSF must be configured and enabled on both chassis.

Switch 1 Task

	Command	Purpose
Step 1	Switch-1(config)# redundancy	Enters redundancy configuration mode.
Step 2	Switch-1(config-red)# mode sso	Configures SSO. When this command is entered, the redundant supervisor engine is reloaded and begins to work in SSO mode.
Step 3	Switch-1(config-red)# exit	Exits redundancy configuration mode.
Step 4	Switch-1(config)# router ospf processID	Enables an OSPF routing process, which places the router in router configuration mode.
Step 5	Switch-1(config-router)# nsf	Enables NSF operations for OSPF.
Step 6	Switch-1(config-router)# end	Exits to privileged EXEC mode.
Step 7	Switch-1# show running-config	Verifies that SSO and NSF are configured and enabled.
Step 8	Switch-1# show redundancy states	Displays the operating redundancy mode.

Switch 2 Task

	Command	Purpose
Step 1	Switch-2(config)# redundancy	Enters redundancy configuration mode.
Step 2	Switch-2(config-red)# mode sso	Configures SSO. When this command is entered, the redundant supervisor engine is reloaded and begins to work in SSO mode.
Step 3	Switch-2(config-red)# exit	Exits redundancy configuration mode.
Step 4	Switch-2(config)# router ospf <i>processID</i>	Enables an OSPF routing process, which places the router in router configuration mode.
Step 5	Switch-2(config-router)# nsf	Enables NSF operations for OSPF.
Step 6	Switch-2(config-router)# end	Exits to privileged EXEC mode.
Step 7	Switch-2# show running-config	Verifies that SSO and NSF are configured and enabled.
Step 8	Switch-2# show redundancy states	Displays the operating redundancy mode.

For detailed information on configuring and verifying SSO and NSF, see [Chapter 8, “Nonstop Forwarding \(NSF\).”](#)

Assigning Virtual Switch Domain and Switch Numbers

Configure the same virtual switch domain number on both chassis. The virtual switch domain is a number between 1 and 255, and must be unique for each VSS in your network (the domain number is incorporated into various identifiers to ensure that these identifiers are unique across the network). Within the VSS, you must configure one chassis to be switch number 1 and the other chassis to be switch number 2.

Switch 1 Task

	Command	Purpose
Step 1	Switch-1(config)# switch virtual domain 100	Configures the virtual switch domain on Chassis A.
Step 2	Switch-1(config-vs-domain)# switch 1	Configures Chassis A as virtual switch number 1.
Step 3	Switch-1(config-vs-domain)# exit	Exits config-vs-domain.

Switch 2 Task

	Command	Purpose
Step 1	Switch-2(config)# switch virtual domain 100	Configures the virtual switch domain on Chassis B.
Step 2	Switch-2(config-vs-domain)# switch 2	Configures Chassis B as virtual switch number 2.
Step 3	Switch-2(config-vs-domain)# exit	Exits config-vs-domain.

**Note**

The switch number is not stored in the startup or running configuration, because both chassis use the same configuration file (but must not have the same switch number).

Configuring the VSL Port Channel

The VSL is configured with a unique port channel on each chassis. During the conversion, the VSS configures both port channels on the active chassis. If the standby chassis VSL port channel number has been configured for another use, the VSS comes up in RPR mode. To avoid this situation, check that both port channel numbers are available on both of the chassis.

Check the port channel number by using the **show running-config interface port-channel** command. The command displays an error message if the port channel is available for VSL. For example, the following command shows that port channel 20 is available on Switch 1:

```
Switch-1 # show running-config interface port-channel 20
% Invalid input detected at '^' marker.
```

Switch 1 Task

	Command	Purpose
Step 1	Switch-1(config)# interface port-channel 10	Configures port channel 10 on Switch 1.
Step 2	Switch-1(config-if)# switch virtual link 1	Associates Switch 1 as owner of port channel 10.
Step 3	Switch-1(config-if)# no shutdown	Activates the port channel.
Step 4	Switch-1(config-if)# exit	Exits interface configuration.

Switch 2 Task

	Command	Purpose
Step 1	Switch-2(config)# interface port-channel 20	Configures port channel 20 on Switch 2.
Step 2	Switch-2(config-if)# switch virtual link 2	Associates Switch 2 as owner of port channel 20.
Step 3	Switch-2(config-if)# no shutdown	Activates the port channel.
Step 4	Switch-2(config-if)# exit	Exits interface configuration mode.

Configuring the VSL Ports

You must add the VSL physical ports to the port channel. In the following example, 10-Gigabit Ethernet ports 3/1 and 3/2 on Switch 1 are connected to 10-Gigabit Ethernet ports 5/2 and 5/3 on Switch 2. For VSL line redundancy, configure the VSL with at least two ports per chassis. For module redundancy, the two ports can be on different switching modules in each chassis.

Switch 1 Task

	Command	Purpose
Step 1	Switch-1(config)# interface range tengigabitethernet 3/1-2	Enters configuration mode for interface range tengigabitethernet 3/1-2 on Switch 1.
Step 2	Switch-1(config-if)# channel-group 10 mode on	Adds this interface to channel group 10.
Step 3	Switch-1(config-if)# no shutdown	Activates the port.

Switch 2 Task

	Command	Purpose
Step 1	Switch-2(config)# interface range tengigabitethernet 5/2-3	Enters configuration mode for interface range tengigabitethernet 5/2-3 on Switch 2.
Step 2	Switch-2(config-if)# channel-group 20 mode on	Adds this interface to channel group 20.
Step 3	Switch-2(config-if)# no shutdown	Activates the port.

Verifying the PFC Operating Mode

Ensure that the PFC operating mode matches on both chassis. Enter the **show platform hardware pfc mode** command on each chassis to display the current PFC mode. If only one of the chassis is in PFC4XL mode, you can configure it to use PFC4 mode with the **platform hardware vsl pfc mode non-xl** command.

Switch 1 Task

	Command	Purpose
Step 1	Switch-1# show platform hardware pfc mode	Ensures that the PFC operating mode matches on both chassis, to ensure that the VSS comes up in SSO redundancy mode.
Step 2	Switch-1(config)# platform hardware vsl pfc mode non-xl	(Optional) Sets the PFC operating mode to PFC4 on Chassis A.

Switch 2 Task

	Command	Purpose
Step 3	Switch-2# show platform hardware pfc mode	Ensures that the PFC operating mode matches on both chassis, to ensure that the VSS comes up in SSO redundancy mode.
Step 4	Switch-2(config)# platform hardware vsl pfc mode non-xl	(Optional) Sets the PFC operating mode to PFC4 on Chassis B.

Converting the Chassis to Virtual Switch Mode

Conversion to VSS mode requires a restart for both chassis. After the reboot, commands that specify interfaces with *module_#/port_#* now include the switch number. For example, a port on a switching module is specified by *switch_#/module_#/port_#*.

Before restarting, the VSS converts the startup configuration to use the *switch_#/module_#/port_#* convention. A backup copy of the startup configuration file is saved on the RP. This file is assigned a default name, but you are also prompted to override the default name if you want to change it.

Switch 1 Task

Command	Purpose
Switch-1# switch convert mode virtual	Converts Switch 1 to virtual switch mode. After you enter the command, you are prompted to confirm the action. Enter yes. The system creates a converted configuration file, and saves the file to the RP bootflash.

Switch 2 Task

Command	Purpose
Switch-2# switch convert mode virtual	Converts Switch 2 to virtual switch mode. After you enter the command, you are prompted to confirm the action. Enter yes. The system creates a converted configuration file, and saves the file to the RP bootflash.

After you confirm the command (by entering **yes** at the prompt), the running configuration is automatically saved as the startup configuration and the chassis reboots. After the reboot, the chassis is in virtual switch mode, so you must specify interfaces with three identifiers (*switch_#/module_#/port_#*).

Auto-Configuring the Standby VSL Information

The two chassis now form a VSS, and the system will auto-configure the standby VSL. After the merge has completed successfully, enter all configuration commands for the VSS on the active chassis. The startup configuration file is automatically synchronized to the standby chassis after the standby chassis reaches the ready state. The VSSmode automatically merges the configuration information on the standby chassis.

All non-VSL interface configurations on the standby chassis revert to the default configuration and non-VSL related configurations are not merged. If you fail to perform any of the required configurations, you will have to repeat the configuration on the active chassis. Auto-configuration merges these commands for the standby chassis:

- **hw-module switch** *number slot number*
- **switch virtual domain** *number*
- **switch** *number priority priority*

- **power redundancy-mode combined switch** *number*
- **no power enable switch** *num module number*
- **interface port-channel** *num switch virtual link number*
- **interface** *type switch_#/slot_#/port_# channel-group number mode on*

(Optional) Configuring Standby Chassis Modules

After the reboot, each chassis contains the module provisioning for its own slots. In addition, the modules from the standby chassis are automatically provisioned on the active chassis with default configuration.

Configurations for the standby chassis modules revert to their default settings (for example, no IP addresses).

You can view the module provisioning information in the configuration file, by entering the **show startup-config** command (after you have saved the configuration).



Note

Do not delete or modify this section of the configuration file. In Cisco IOS Release 12.2(50)SY and later releases, you can no longer add module provisioning entries using the **module provision** CLI command. When a module is not present, the provisioning entry for that module can be cleared using the **no slot** command with the **module provision** CLI command. Note that the VSS setup does not support the **module clear-config** command.

The following example shows the module provisioning information from a configuration file:

```
module provision switch 1
  slot 1 slot-type 148 port-type 60 number 4   virtual-slot 17
  slot 2 slot-type 137 port-type 31 number 16  virtual-slot 18
  slot 3 slot-type 227 port-type 60 number 8   virtual-slot 19
  slot 4 slot-type 225 port-type 61 number 48  virtual-slot 20
  slot 5 slot-type 82  port-type 31 number 2   virtual-slot 21
module provision switch 2
  slot 1 slot-type 148 port-type 60 number 4   virtual-slot 33
  slot 2 slot-type 227 port-type 60 number 8   virtual-slot 34
  slot 3 slot-type 137 port-type 31 number 16  virtual-slot 35
  slot 4 slot-type 225 port-type 61 number 48  virtual-slot 36
  slot 5 slot-type 82  port-type 31 number 2   virtual-slot 37
```

Displaying VSS Information

These commands display basic information about the VSS:

Command	Purpose
show switch virtual	Displays the virtual switch domain number, and the switch number and role for each of the chassis.
show switch virtual role	Displays the role, switch number, and priority for each of the chassis in the VSS.
show switch virtual link	Displays the status of the VSL.

The following example shows the information output from these commands:

```
Router# show switch virtual
Switch mode : Virtual Switch
Virtual switch domain number : 100
Local switch number : 1
Local switch operational role: Virtual Switch Active
Peer switch number : 2
Peer switch operational role : Virtual Switch Standby

Router# show switch virtual role
Switch Switch Status Preempt Priority Role Session ID
Number Oper (Conf) Oper (Conf) Local Remote
-----
LOCAL 1 UP FALSE(N) 100(100) ACTIVE 0 0
REMOTE 2 UP FALSE(N) 100(100) STANDBY 8158 1991

In dual-active recovery mode: No

Router# show switch virtual link
VSL Status: UP
VSL Uptime: 4 hours, 26 minutes
VSL SCP Ping: Pass OK
VSL ICC (Ping): Pass
VSL Control Link: Te 1/5/1
```

Converting a VSS to Standalone Chassis

- [Copying the VSS Configuration to a Backup File, page 4-39](#)
- [Converting the Active Chassis to Standalone, page 4-39](#)
- [Converting the Peer Chassis to Standalone, page 4-40](#)

Copying the VSS Configuration to a Backup File

Save the configuration file from the active chassis. You may need this file if you convert to virtual switch mode again. You only need to save the file from the active chassis, because the configuration file on the standby chassis is identical to the file on the active chassis.

	Command	Purpose
Step 1	Switch-1# copy running-config startup-config	(Optional) Saves the running configuration to startup configuration. This step is only required if you there are unsaved changes in the running configuration that you want to preserve.
Step 2	Switch-1# copy startup-config disk0:vs-startup-config	Copies the startup configuration to a backup file.

Converting the Active Chassis to Standalone

When you convert the active chassis to standalone mode, the active chassis removes the provisioning and configuration information related to VSL links and the peer chassis modules, saves the configuration file, and performs a reload. The chassis comes up in standalone mode with only the provisioning and configuration data relevant to the standalone system.

The standby chassis of the VSS becomes active. VSL links on this chassis are down because the peer is no longer available.

To convert the active chassis to standalone mode, perform this task on the active chassis:

Command	Purpose
Switch-1# switch convert mode stand-alone	Converts Switch 1 to standalone mode. After you enter the command, you are prompted to confirm the action. Enter yes .

Converting the Peer Chassis to Standalone

When you convert the new active chassis to standalone mode, the chassis removes the provisioning and configuration information related to VSL links and the peer chassis modules, saves the configuration file and performs a reload. The chassis comes up in standalone mode with only its own provisioning and configuration data.

To convert the peer chassis to standalone, perform this task on the standby chassis:

Command	Purpose
Switch-2# switch convert mode stand-alone	Converts Switch 2 to standalone mode. After you enter the command, you are prompted to confirm the action. Enter yes .

Configuring VSS Parameters

- [Configuring VSL Switch Priority, page 4-41](#)
- [Configuring the PFC Mode, page 4-42](#)
- [Configuring a VSL, page 4-42](#)
- [Configuring VSL Encryption, page 4-43](#)
- [Displaying VSL Information, page 4-45](#)
- [Configuring VSL QoS, page 4-46](#)
- [Subcommands for VSL Port Channels, page 4-46](#)
- [Subcommands for VSL Ports, page 4-47](#)
- [Configuring the Router MAC Address Assignment, page 4-47](#)

Configuring VSL Switch Priority

To configure the switch priority, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain 100	Enters configuration mode for the virtual switch domain.
Step 2	Router(config-vs-domain)# switch [1 2] priority [priority_num]	Configures the priority for the chassis. The switch with the higher priority assumes the active role. The range is 1 (lowest priority) to 255 (highest priority); the default is 100. Note • The new priority value only takes effect after you save the configuration and perform a reload of the VSS. • If the higher priority switch is currently in standby state, you can make it the active switch by initiating a switchover. Enter the redundancy force-switchover command. • The show switch virtual role command displays the operating priority and the configured priority for each switch in the VSS. • The no form of the command resets the priority value to the default priority value of 100. The new value takes effect after you save the configuration and perform a reload.



Note

If you make configuration changes to the switch priority, the changes only take effect after you save the running configuration to the startup configuration file and perform a reload. The **show switch virtual role** command shows the operating and configured priority values. You can manually set the standby switch to active using the **redundancy force-switchover** command.

This example shows how to configure virtual switch priority:

```
Router(config)# switch virtual domain 100
Router(config-vs-domain)# switch 1 priority 200
Router(config-vs-domain)# exit
```

This example shows how to display priority information for the VSS:

```
Router# show switch virtual role
```

Switch	Switch Number	Status	Preempt Oper (Conf)	Priority Oper (Conf)	Role	Session ID	
						Local	Remote
LOCAL	1	UP	FALSE(N)	100(200)	ACTIVE	0	0
REMOTE	2	UP	FALSE(N)	100(100)	STANDBY	8158	1991

In dual-active recovery mode: No

Configuring the PFC Mode

If you have a mixture of DFC4 and DFC4XL switching modules in the VSS, set the PFC mode by performing this task:

Command	Purpose
Router(config)# platform hardware vsl pfc mode non-xl	Sets the PFC configuration mode for the VSS to PFC4. Note This command requires a system reload before it takes effect.

This example shows how to set the PFC configuration mode for the VSS to PFC4. You can wait until the next maintenance window to perform the **reload** command.

```
Router(config)# platform hardware vsl pfc mode non-xl
Router(config)# end
Router# reload
```

If all the supervisor engines and switching modules in the VSS are XL, the following warning is displayed if you set the PFC mode to PFC4:

```
Router(config)# platform hardware vsl pfc mode non-xl
PFC Preferred Mode: PFC4XL. The discrepancy between Operating Mode and
Preferred Mode could be due to PFC mode config. Your System has all PFC4XL modules.
Remove ' platform hardware vsl pfc mode non-xl ' from global config.
```

This example shows how to display the operating and configured PFC modes:

```
Router# show platform hardware pfc mode
PFC operating mode : PFC4
Configured PFC operating mode : PFC4
```

Configuring a VSL

To configure a port channel to be a VSL, perform this task:

	Command	Purpose
Step 1	Router(config) # interface port-channel <i>channel_num</i>	Enters configuration mode for the specified port channel.
Step 2	Router(config-if) # switch virtual link <i>switch_num</i>	Assigns the port channel to the virtual link for the specified switch.


Note

We recommend that you configure the VSL prior to converting the chassis into a VSS.

This example shows how to configure the VSL:

```
Switch-1(config)# interface port-channel 10
Switch-1(config-if)# switch virtual link 1
Switch-1(config-if)# no shutdown
Switch-1(config)# interface tenGigabitEthernet 5/1
Switch-1(config-if)# channel-group 10 mode on
Switch-1(config-if)# no shutdown
```

```
Switch-2(config)# interface port-channel 25
Switch-2(config-if)# switch virtual link 2
Switch-2(config-if)# no shutdown
Switch-2(config-if)# interface tenGigabitEthernet 5/2
Switch-2(config-if)# channel-group 25 mode on
Switch-2(config-if)# no shutdown
```

Configuring VSL Encryption

- [VSL Encryption Overview, page 4-43](#)
- [VSL Encryption Restrictions, page 4-43](#)
- [Configuring the VSL Encryption Key, page 4-44](#)
- [Enabling VSL Encryption, page 4-44](#)
- [Displaying the VSL Encryption State, page 4-45](#)

VSL Encryption Overview

Cisco IOS Release 15.4SY supports HW-based encryption on a VSL configured on a Supervisor Engine 6T or any VSL capable line cards. VSL encryption uses an encryption key that you manually configure. The encryption key is stored securely.

VSL Encryption Restrictions

- VSL encryption requires a MACSec license on each chassis.
- The chassis must be rebooted to configure an encryption key or enable VSL encryption.
- You enter the encryption key on the active chassis. You cannot enter the encryption key on the standby chassis.
- If it is acceptable to send the key as plain text over the VSL to the other chassis, then you can allow one chassis to send the key to the other chassis. For maximum security, configure the encryption key on each chassis.
- There are no **show** commands that display the encryption key.
- You cannot remove the encryption key while VSL encryption is enabled.
- The following commands take effect after a reboot:
 - To remove the encryption key, enter the **clear switch pmk** EXEC mode command.
 - To disable VSL encryption, enter the **no vsl-encryption** virtual switch domain configuration submode command.
- If the encryption key and VSL encryption state on the two chassis do not match, the VSL does not transition to the link-up state.

- In VSS mode, you cannot configure the FIPS encryption mode without VSL encryption. To avoid a system shutdown, enable VSL encryption before you enable FIPS encryption mode. ([CSCts96040](#), [CSCtx58304](#))

Configuring the VSL Encryption Key

To configure the VSL encryption key, perform this task:

Command	Purpose
Router# switch pmk encryption_key	Configures the VSL encryption key. <i>encryption_key</i> is a hexadecimal string up to 32 characters (256 bits). - You will be asked if you want to automatically synchronize the encryption key. If you do not automatically synchronize the encryption key, configure the same encryption key on the other chassis.

This example show how to configure a VSL encryption key:

```
Router# switch pmk encryption_key
Key effective only upon reboot and will override old VSL PMK.
Key needs to be provisioned on both VSS switches.
Warning - Sending the key to standby will cause the key to be sent over an unencrypted VSL
link.
Do you want to automatically synchronize the key [yes/no]?
```

Enabling VSL Encryption

To enable VSL encryption, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain domain_id	Enters VSS configuration mode.
Step 2	Router(config-vs-domain)# vsl-encryption	Enables VSL encryption.

This example shows how to enable VSL encryption:

```
Router(config)# switch virtual domain domain_id
Router(config-vs-domain)# vsl-encryption
```



Note

- If you manually configure the encryption key on each chassis:
 - Reboot the active chassis; the standby chassis becomes active.
 - Configure the encryption key on new active chassis.
 - Reboot the new active chassis.
- If you allow the encryption key to be sent to the standby chassis, reboot the active chassis.

Displaying the VSL Encryption State

This example shows how to display the VSL encryption state:

```
Router# show switch virtual link | include Encryption
VSL Encryption : Configured Mode - On, Operational Mode - On
```

Displaying VSL Information

To display information about the VSL, perform one of these tasks:

Command	Purpose
Router# show switch virtual link	Displays information about the VSL.
Router# show switch virtual link port-channel	Displays information about the VSL port channel.
Router# show switch virtual link port	Displays information about the VSL ports.

This example shows how to display VSL information:

```
Router# show switch virtual link
VSL Status : UP
VSL Uptime : 1 day, 3 hours, 39 minutes
VSL SCP Ping : Pass
VSL ICC Ping : Pass
VSL Control Link : Te 1/5/1

Router# show switch virtual link port-channel
VSL Port Channel Information

Flags:  D - down          P - bundled in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby  (LACP only)
        R - Layer3       S - Layer2
        U - in use       N - not in use, no aggregation
        f - failed to allocate aggregator

        M - not in use, no aggregation due to minimum links not met
        m - not in use, port not aggregated due to minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated

Group  Port-channel  Protocol  Ports
-----+-----+-----+-----
10     Po10(RU)        -         Te1/5/4(P) Te1/5/5(P)
20     Po20(RU)        -         Te2/5/4(P) Te2/5/5(P)

Router# show switch virtual link port
VSL Link Info          : Configured: 2 Operational: 1

Interface  State          Peer
           MAC           Switch  Interface
-----
Te1/5/4    operational  0013.5fcb.1480  2    Te2/5/4
Te1/5/5    link_down    -              -    -

Interface          Last operational      Current packet      Last Diag      Time since
Failure state      State                  Result              Last Diag
-----
```

```

Te1/5/4 No failure      Hello bidir      Never ran      7M:51S
Te1/5/5 No failure      No failure       Never ran      7M:51S

```

Interface	State	Hello Tx (T4) ms			Hello Rx (T5*) ms		
		Cfg	Cur	Rem	Cfg	Cur	Rem
Te1/5/4	operational	500	500	404	5000	5000	4916
Te1/5/5	link_down	500	-	-	500000	-	-
Te2/5/4	operational	500	500	404	500000	500000	499916
Te2/5/5	link_down	500	-	-	500000	-	-

*T5 = min_rx * multiplier

Configuring VSL QoS

The VSS automatically configures VSL ports for trust CoS, using default CoS mappings (you cannot change the mappings on VSL ports).

For switching modules that support per-ASIC configuration, the VSL configuration applies to all ports on the same ASIC (including any non-VSL ports).

The VSS disables the QoS commands on VSL ports (and any non-VSL ports on the same ASIC). For example, you cannot use QoS queuing or map commands on VSL ports.

To ensure that all eight QoS receive queues are enabled for the 10-Gigabit Ethernet ports on the supervisor engine, enter the **platform qos 10g-only** global configuration command.

In Cisco IOS Release 12.2(50)SY and later releases, when the **platform qos 10g-only** command is entered and only one of the two 10-Gigabit Ethernet ports on the supervisor engine is a VSL port, the non-VSL 10-Gigabit Ethernet port can be configured for QoS.

Subcommands for VSL Port Channels

On a VSL port channel, only a subset of interface subcommands are available in the command console. [Table 4-2](#) describes the available interface subcommands.

Table 4-2 Interface Subcommands for VSL Port Channels

Subcommand	Description
default	Sets a command to its defaults.
description	Enters a text description for the interface.
exit	Exits from interface configuration mode.
load-interval	Specifies interval for load calculation for an interface.
logging	Configures logging for interface.
platform	Specifies platform-specific command.
no	Disables a command, or sets the command defaults.
shutdown	Shuts down the selected interface.

Table 4-2 Interface Subcommands for VSL Port Channels (continued)

Subcommand	Description
switch virtual link	Specifies the switch associated with this port channel.
vslp	Specifies VSLP interface configuration commands.

Subcommands for VSL Ports

If a port is included in a VSL port channel, only a subset of interface subcommands are available in the command console. [Table 4-3](#) describes the available interface subcommands.

Table 4-3 Interface Subcommands for VSL Ports

Subcommand	Description
channel-group	Adds the interface to the specified channel group.
default	Sets a command to its defaults.
description	Adds a description to the interface.
exit	Exits from interface configuration mode.
load-interval	Specifies interval for load calculation for an interface.
logging	Configures logging for the interface.
no	Disables a command, or sets the command defaults.
shutdown	Shuts down the selected interface.

Configuring the Router MAC Address Assignment

When the VSS is started for the first time, the initial active supervisor engine assigns a router MAC address for the VSS. By default, the supervisor engine assigns a MAC address from its own chassis. After a switchover to the second chassis, the VSS continues to use the MAC address from the previously active chassis as the router MAC address.

In the rare case where both chassis later become inactive and then start up with the second supervisor engine becoming the initial active supervisor engine, the VSS will start up with a router MAC address from the second chassis. Other Layer 2 hosts that do not respond to GARP and are not directly connected to the VSS will retain the earlier router MAC address of the VSS, and will not be able to communicate with the VSS. To avoid this possibility, you can configure the VSS to assign a router MAC address from a reserved pool of addresses with the domain ID encoded in the last octet of the MAC address, or you can specify a MAC address.

**Note**

If you change the router MAC address, you must reload the virtual switch for the new router MAC address to take effect.

To specify that the router MAC address is assigned from a reserved pool of domain-based addresses, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain <i>domain_id</i>	Enters VSS configuration mode.
Step 2	Router(config-vs-domain)# mac address use-virtual	The router MAC address is assigned from a reserved pool of domain-based addresses. Note The no form of this command reverts to the default setting, using a MAC address from the backplane of the initial active chassis.

To specify a router MAC address, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain <i>domain_id</i>	Enters VSS configuration mode.
Step 2	Router(config-vs-domain)# mac address <i>mac_address</i>	The router MAC address is specified in three 2-byte hexadecimal numbers.

This example shows how to configure router MAC address assignment from a reserved pool of domain-based addresses:

```
Router(config)# switch virtual domain 255
Router(config-vs-domain)# mac address use-virtual
```

The following example shows how to specify the router MAC address in hexadecimal format:

```
Router(config)# switch virtual domain 255
Router(config-vs-domain)# mac address 0123.4567.89ab
```

Configuring Deferred Port Activation During Standby Recovery

Instead of allowing all ports to be activated simultaneously when a failed chassis is restarted as the standby chassis, you can configure the system to defer activation of non-VSL ports and then activate the ports in groups over a period of time.

To specify deferred port activation, perform this task:

Command	Purpose
Router(config)# switch virtual domain 1	Enters VSS configuration mode.

Command	Purpose
Router(config-vs-domain)# standby port delay <i>delay-time</i>	Specifies that the port activation will be initially deferred and then performed in cycles. For <i>delay-time</i> , specify the period in seconds before port activation will begin. The range is 30 to 3600.
Router(config-vs-domain)# standby port bringup <i>number cycle-time</i>	Specifies the number of ports to be activated per cycle and the waiting time between cycles. For <i>number</i> , specify the number of ports to be activated per cycle. The range is 1 to 100. The default value is 1 port. For <i>cycle-time</i> , specify the period in seconds between cycles. The range is 1 to 10. The default value is 1 second.

This example shows how to configure port activation to be deferred by 120 seconds, then activated in groups of 20 ports every 5 seconds:

```
Router(config)# switch virtual domain 1
Router(config-vs-domain)# standby port delay 120
Router(config-vs-domain)# standby port bringup 20 5
```

Configuring Multichassis EtherChannels

Configure multichassis EtherChannels (MECs) as you would for a regular EtherChannel. The VSS will recognize that the EtherChannel is an MEC when ports from both chassis are added to the EtherChannel. You can verify the MEC configuration by entering the **show etherchannel** command.

One VSS supports a maximum of 512 port channels.



Note

Releases earlier than Cisco IOS Release 12.2(50)SY support a maximum of 128 port channels.

Configuring Port Load Share Deferral on the Peer Switch

To configure the load share deferral feature for a port channel, perform this task on the switch that is an MEC peer to the VSS:

	Command	Purpose
Step 1	Router(config)# port-channel load-defer <i>time</i>	(Optional) Configures the port load share deferral interval for all port channels. <i>time</i>—The time interval during which load sharing is initially 0 for deferred port channels. The range is 1 to 1800 seconds; the default is 120 seconds.
Step 2	Router(config)# interface port-channel <i>channel-num</i>	Enters interface configuration mode for the port channel.
Step 3	Router(config-if)# port-channel port load-defer	Enables port load share deferral on the port channel.

This example shows how to configure the load share deferral feature on port channel 10 of the switch that is an MEC peer to the VSS:

```
Router(config)# port-channel load-defer 60
Router(config)# interface port-channel 10
Router(config-if)# port-channel port load-defer
This will enable the load share deferral feature on this port-channel.
```

**Note**

To provide the best support for multicast traffic, configure the load share deferral feature on all EtherChannels that have member ports on more than one module.

Configuring Dual-Active Detection

- [Configuring Enhanced PAgP Dual-Active Detection, page 4-50](#)
- [Configuring Fast Hello Dual-Active Detection, page 4-51](#)
- [Configuring the Exclusion List, page 4-52](#)
- [Displaying Dual-Active Detection, page 4-53](#)

Configuring Enhanced PAgP Dual-Active Detection

If enhanced PAgP is running on the MECs between the VSS and its access switches, the VSS can use enhanced PAgP messaging to detect a dual-active scenario.

By default, PAgP dual-active detection is enabled. However, the enhanced messages are only sent on port channels with trust mode enabled (see the trust mode description below).

**Note**

Before changing PAgP dual-active detection configuration, ensure that all port channels with trust mode enabled are in administrative down state. Use the **shutdown** command in interface configuration mode for the port channel. Remember to use the **no shutdown** command to reactivate the port channel when you are finished configuring dual-active detection.

To enable or disable PAgP dual-active detection, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain <i>domain_id</i>	Enters virtual switch submode.
Step 2	Router(config-vs-domain)# dual-active detection pagp	Enables sending of the enhanced PAgP messages.

You must configure trust mode on the port channels that will detect PAgP dual-active detection. By default, trust mode is disabled.

**Note**

If PAgP dual-active detection is enabled, you must place the port channel in administrative down state before changing the trust mode. Use the **shutdown** command in interface configuration mode for the port channel. Remember to use the **no shutdown** command to reactivate the port channels when you are finished configuring trust mode on the port channel.

To configure trust mode on a port channel, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain <i>domain_id</i>	Enters virtual switch submenu.
Step 2	Router(config-vs-domain)# dual-active detection pagp trust channel-group <i>group_number</i>	Enables trust mode for the specified port channel.

This example shows how to enable PAGP dual-active detection:

```
Router(config)# interface port-channel 20
Router(config-if)# shutdown
Router(config-if)# exit
Router(config)# switch virtual domain 100
Router(config-vs-domain)# dual-active detection pagp
Router(config-vs-domain)# dual-active detection pagp trust channel-group 20
Router(config-vs-domain)# exit
Router(config)# interface port-channel 20
Router(config-if)# no shutdown
Router(config-if)# exit
```

This example shows the error message if you try to enable PAGP dual-active detection when a trusted port channel is not shut down first:

```
Router(config)# switch virtual domain 100
Router(config-vs-domain)# dual-active detection pagp
Trusted port-channel 20 is not administratively down.
To change the pagp dual-active configuration, "shutdown" these port-channels first.
Remember to "no shutdown" these port-channels afterwards.
```

This example shows the error message if you try to configure trust mode for a port channel that is not shut down first:

```
Router(config)# switch virtual domain 100
Router(config-vs-domain)# dual-active detection pagp trust channel-group 20
Trusted port-channel 20 is not administratively down. To change the pagp dual-active trust
configuration, "shutdown" the port-channel first. Remember to "no shutdown" the
port-channel afterwards.
```

Configuring Fast Hello Dual-Active Detection

Fast hello dual-active detection is enabled by default; however, you must configure dual-active interface pairs to act as fast hello dual-active messaging links.

To configure fast hello dual-active detection, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain <i>domain_id</i>	Enters the virtual switch submenu.
Step 2	Router(config-vs-domain)# dual-active detection fast-hello	Enables the fast hello dual-active detection method. Fast hello dual-active detection is enabled by default.
Step 3	Router(config-vs-domain)# exit	Exits virtual switch submenu.
Step 4	Router(config)# interface <i>type switch/slot/port</i>	Selects the interface to configure. This interface must be directly connected to the other chassis and must not be a VSL link.

	Command	Purpose
Step 5	Router(config-if)# dual-active fast-hello	Enables fast hello dual-active detection on the interface, automatically removes all other configuration from the interface, and restricts the interface to dual-active configuration commands.
Step 6	Router(config-if)# no shutdown	Activates the interface.

When you configure fast hello dual-active interface pairs, note the following information:

- You can configure a maximum of four interfaces on each chassis to connect with the other chassis in dual-active interface pairs.
- Each interface must be directly connected to the other chassis and must not be a VSL link. We recommend using links from a switching module not used by the VSL.
- Each interface must be a physical port. Logical ports such as an SVI are not supported.
- Configuring fast hello dual-active mode will automatically remove all existing configuration from the interface and will restrict the interface to fast hello dual-active configuration commands.
- Unidirectional link detection (UDLD) will be disabled on fast hello dual-active interface pairs.

This example shows how to configure an interface for fast hello dual-active detection:

```
Router(config)# switch virtual domain 255
Router(config-vs-domain)# dual-active detection fast-hello
Router(config-vs-domain)# exit
Router(config)# interface fastethernet 1/2/40
Router(config-if)# dual-active fast-hello
WARNING: Interface FastEthernet1/2/40 placed in restricted config mode. All extraneous
configs removed!

Router(config-if)# no shutdown
Router(config-if)# exit
Router(config)# exit
Router# show run interface fastethernet 1/2/40
interface FastEthernet1/2/40
  no switchport
  no ip address
  dual-active fast-hello
end
```

Configuring the Exclusion List

When a dual-active scenario is detected, part of the recovery action is for the chassis to shut down all of its non-VSL interfaces. You can specify one or more interfaces to be excluded from this action (for example, to exclude the interface you use for remote access to the chassis).

To specify interfaces that are not to be shut down by dual-active recovery, perform this task:

	Command	Purpose
Step 1	Router(config)# switch virtual domain <i>domain_id</i>	Enters virtual switch submode.
Step 2	Router(config-vs-domain)# dual-active exclude interface <i>type switch/slot/port</i>	Specifies an interface to exclude from shutting down in dual-active recovery.

When you configure the exclusion list, note the following information:

- The interface must be a physical port configured with an IP address.
- The interface must not be a VSL port.
- The interface must not be in use for fast hello dual-active detection.

This example shows how to configure an interface as an exclusion:

```
Router(config)# switch virtual domain 100
Router(config-vs-domain)# dual-active exclude interface gigabitethernet 1/5/5
```

Displaying Dual-Active Detection

To display information about dual-active detection, perform this task:

Command	Purpose
Router# show switch virtual dual-active [pagp fast-hello summary]	Displays information about dual-active detection configuration and status.

This example shows how to display the summary status for dual-active detection:

```
Router# show switch virtual dual-active summary
Pagp dual-active detection enabled: Yes
Fast-hello dual-active detection enabled: Yes
```

No interfaces excluded from shutdown in recovery mode

In dual-active recovery mode: No

This example shows how to display information for fast-hello dual-active detection:

```
Router# show switch virtual dual-active fast-hello
Fast-hello dual-active detection enabled: Yes
```

Fast-hello dual-active interfaces:

```
Port      State (local only)
-----
Gi1/4/47  Link dn
Gi2/4/47  -
```

This example shows how to display PAgP status and the channel groups with trust mode enabled:

```
Router# show pagp dual-active
PAgP dual-active detection enabled: Yes
PAgP dual-active version: 1.1
```

Channel group 3 dual-active detect capability w/nbrs Dual-Active trusted group: No

	Dual-Active	Partner	Partner	Partner
Port	Detect Capable	Name	Port	Version
Fa1/2/33	No	None	None	N/A

Channel group 4

Dual-Active trusted group: Yes

No interfaces configured in the channel group

Channel group 5

Dual-Active trusted group: Yes

Channel group 5 is not participating in PAgP

```
Channel group 10 dual-active detect capability w/nbrs Dual-Active trusted group: Yes
      Dual-Active   Partner
Port   Detect Capable Name          Port   Version
Gi1/6/1 Yes          partner-1   Gi1/5/1   1.1
Gi2/5/1 Yes          partner-1   Gi1/5/2   1.1
```

```
Channel group 11 dual-active detect capability w/nbrs Dual-Active trusted group: No
      Dual-Active   Partner
Port   Detect Capable Name          Port   Version
Gi1/6/2 Yes          partner-1   Gi1/3/1   1.1
Gi2/5/2 Yes          partner-1   Gi1/3/2   1.1
```

```
Channel group 12 dual-active detect capability w/nbrs Dual-Active trusted group: Yes
      Dual-Active   Partner
Port   Detect Capable Name          Port   Version
Fa1/2/13 Yes          partner-1   Fa1/2/13   1.1
Fa1/2/14 Yes          partner-1   Fa1/2/14   1.1
Gi2/1/15 Yes          partner-1   Fa1/2/15   1.1
Gi2/1/16 Yes          partner-1   Fa1/2/16   1.1
```

**Note**

The **show switch virtual dual-active pagp** command displays the same output as the **show pagp dual-active** command.

How to Upgrade a VSS

- [Performing a Fast Software Upgrade of a VSS, page 4-54](#)
- [Performing an Enhanced Fast Software Upgrade of a VSS, page 4-55](#)

Performing a Fast Software Upgrade of a VSS

The FSU of a VSS is similar to the RPR-based standalone chassis FSU described in [Chapter 6, “Fast Software Upgrade.”](#) While the standalone chassis upgrade is initiated by reloading the standby supervisor engine, the VSS upgrade is initiated by reloading the standby chassis. During the FSU procedure, a software version mismatch between the active and the standby chassis causes the system to boot in RPR redundancy mode, which is stateless and causes a hard reset of the all modules. As a result, the FSU procedure requires system downtime corresponding to the RPR switchover time.

**Note**

VSS mode supports only one supervisor engine in each chassis.

To perform an FSU of a VSS, perform this task:

	Command	Purpose
Step 1	Router# copy tftp disk_name	Uses TFTP to copy the new software image to flash memory on the active and standby chassis (disk0: and slavedisk0:). Answer the prompts to identify the name and location of the new software image.
Step 2	Router# config terminal	Enters global configuration mode.
Step 3	Router(config)# no boot system	Removes any previously assigned boot variables.

	Command	Purpose
Step 4	Router(config)# config-register 0x2102	Sets the configuration register.
Step 5	Router(config)# boot system flash <i>device:file_name</i>	Configures the chassis to boot the new image.
Step 6	Router(config)# end	Returns to privileged EXEC mode.
Step 7	Router# copy running-config startup-config	Saves the configuration.
Step 8	Router# redundancy reload peer	Reloads the standby chassis and brings it back online running the new version of the Cisco IOS software. Due to the software version mismatch between the two chassis, the standby chassis will be in RPR redundancy mode. Note Before reloading the standby chassis, make sure you wait long enough to ensure that all configuration synchronization changes have completed.
Step 9	Router# redundancy force-switchover	Forces the standby chassis to assume the role of the active chassis running the new Cisco IOS image. The modules are reloaded and the module software is downloaded from the new active chassis. The old active chassis reboots with the new image and becomes the standby chassis.

This example shows how to perform an FSU:

```
Router# config terminal
Router(config)# no boot system
Router(config)# config-register 0x2102
Router(config)# boot system flash disk0:image_name
Router(config)# end
Router# copy running-config startup-config
Router# redundancy reload peer
Router# redundancy force-switchover
```

Performing an Enhanced Fast Software Upgrade of a VSS

An eFSU uses the same commands and software infrastructure as an in-service software upgrade (ISSU). The eFSU differs from an ISSU in that it resets the modules, which results in a brief traffic interruption. The eFSU sequence for a VSS follows the same logical steps as the single-chassis eFSU described in the [Chapter 5, “Enhanced Fast Software Upgrade,”](#) but the procedure applies to the VSS active and VSS standby supervisor engine in each chassis, instead of two supervisor engines in one chassis. During an eFSU, the VSS standby chassis, including the supervisor engine and modules, is upgraded and brought up in a stateful switchover (SSO) mode. The eFSU process then forces a switchover and performs the same upgrade on the other chassis, which becomes the new VSS standby.



Note

VSS mode supports only one supervisor engine in each chassis. If another supervisor resides in the chassis it will act as the DFC.

This section contains the following topics:

- [eFSU Restrictions and Guidelines, page 4-56](#)

- [eFSU Stages for a VSS Upgrade, page 4-56](#)
- [Configuring and Performing an eFSU Upgrade, page 4-58](#)
- [eFSU Upgrade Example, page 4-60](#)

eFSU Restrictions and Guidelines

When performing an eFSU, note the following guidelines and restrictions:

- There must be at least one VSL link between supervisor engine in each chassis.
- Images from different feature sets, regardless of release, fail the eFSU compatibility check.
- The new image file must reside in the file system of the supervisor engine in each chassis before the eFSU can be initiated. The **issu** commands will accept only global file system names (for example, disk0: or bootdisk:). The **issu** commands will not accept switch number-specific file system names (for example, sw1-slot5-disk0:).
- When preparing for the eFSU, do not change the boot variable. Although a boot variable change is required in the FSU (RPR) procedure, changing the boot variable in the eFSU procedure will cause the CurrentVersion variable to be inconsistent, preventing execution of the eFSU.
- The **issu** commands for a VSS eFSU upgrade are similar to those for a single-chassis (standalone) eFSU, as described in the [Chapter 5, “Enhanced Fast Software Upgrade,”](#) with the following differences:
 - Where the standalone **issu** commands accept an argument of slot number, the VSS **issu** commands accept a switch and slot number, in the format of *switch/slot* (for example, 1/5 refers to switch 1, slot 5).
 - For a normal VSS eFSU, it is not necessary to specify a switch or slot number when entering the VSS **issu** commands.
- You cannot change the rollback timer period during the eFSU process.
- During the eFSU process, do not perform any manual switchover other than those caused by the **issu** commands.
- During the eFSU process, do not perform an online insertion or removal (OIR) of any module.
- During an eFSU downgrade, if the process is aborted (either due to an MCL error or by entering the **abortversion** command) just after executing the **loadversion** command, the SSO VSS standby is reloaded with the original image but the SSO VSS standby's ICS is not because the bootvar of the SSO VSS standby's ICS is not modified during an abort executed after the **loadversion** command.

eFSU Stages for a VSS Upgrade

The eFSU sequence consists of several stages, each explicitly initiated by entering a specific **issu** command in the CLI. At each stage, you can verify the system status or roll back the upgrade before moving to the next stage.

The following sections describe the eFSU stages for a VSS upgrade:

- [Preparation, page 4-57](#)
- [Loadversion Stage, page 4-57](#)
- [Runversion Stage, page 4-57](#)
- [Acceptversion Stage \(Optional\), page 4-57](#)
- [Commitversion Stage, page 4-58](#)
- [Abortversion \(Optional\), page 4-58](#)

Preparation

Before you can initiate the eFSU process, the upgrade image must reside in the file system of the supervisor engine in each chassis; otherwise, the initial command will be rejected. The VSS must be in a stable operating state with one chassis in the VSS active state and the other chassis in the hot VSS standby state.

Loadversion Stage

The eFSU process begins when you enter the **issu loadversion** command specifying the location in memory of the new upgrade images on the VSS active and VSS standby chassis. Although the **issu loadversion** command allows you to specify the switch and slot number of the VSS active and VSS standby chassis, it is not necessary to do so. When you enter the **issu loadversion** command, the entire VSS standby chassis, including the supervisor engine and modules, is reloaded with the new upgrade image. Because the VSS standby chassis modules are unavailable while reloading, the throughput of the VSS is temporarily reduced by 50 percent during this stage. After reloading, the VSS standby chassis boots with the new image and initializes in SSO mode, restoring traffic throughput. In this state, the VSS standby chassis runs a different software version than the VSS active chassis, which requires the VSS active chassis to communicate with modules running different image versions between the two chassis.

Runversion Stage

When the VSS standby chassis is successfully running the new image in SSO mode, you can enter the **issu runversion** command. This command forces a switchover in which the upgraded VSS standby chassis takes over as the new VSS active chassis. The formerly VSS active chassis reloads and initializes as the new VSS standby chassis in SSO mode, running the old image. As in the loadversion stage, the throughput of the VSS is temporarily reduced during the VSS standby chassis reload, and the VSS standby chassis runs a different software version than the VSS active chassis.

Acceptversion Stage (Optional)

When you enter the **issu runversion** command, a switchover to the chassis running the new image occurs, which starts an automatic rollback timer as a safeguard to ensure that the upgrade process does not cause the VSS to be nonoperational. Before the rollback timer expires, you must either accept or commit the new software image. If the timer expires, the upgraded chassis reloads and reverts to the previous software version. To stop the rollback timer, enter the **issu acceptversion** command. Prior to starting the eFSU process, you can disable the rollback timer or configure it to a value up to two hours (the default is 45 minutes).

Operating with an upgraded VSS active chassis, this stage allows you to examine the functionality of the new software image. When you are satisfied that the new image is acceptable, enter the **issu commitversion** command to complete the upgrade process.

Commitversion Stage

To apply the upgrade image to the second chassis, completing the eFSU, enter the **issu commitversion** command. The VSS standby chassis is reloaded and booted with the new upgrade image, initializing again as the VSS standby chassis. As in the loadversion stage, the throughput of the VSS is temporarily reduced while the modules are reloaded and initialized. After the successful reload and reboot of the VSS standby chassis, the VSS upgrade process is complete.

Abortversion (Optional)

At any time before you enter the **issu commitversion** command, you can roll back the upgrade by entering the **issu abortversion** command. The upgrade process also aborts automatically if the software detects a failure. The rollback process depends on the current state. If the eFSU is aborted before you enter the **issu runversion** command, the VSS standby chassis is reloaded with the old image. If the eFSU is aborted after the **issu runversion** command, a switchover is executed. The VSS standby chassis, running the old image, becomes the VSS active chassis. The formerly VSS active chassis is then reloaded with the old image, completing the rollback.

Configuring and Performing an eFSU Upgrade

The following sections describe how to configure and perform an eFSU upgrade:

- [Changing the eFSU Rollback Timer, page 4-58](#)
- [Performing an eFSU Upgrade, page 4-59](#)
- [Aborting an eFSU Upgrade, page 4-59](#)

Changing the eFSU Rollback Timer

To view or change the eFSU rollback timer, perform the following task before beginning an upgrade:

	Command	Purpose
Step 1	Router# config terminal	Enters configuration mode.
Step 2	Router(config)# issu set rollback-timer {seconds hh:mm:ss}	(Optional) Sets the rollback timer to ensure that the upgrade process does not leave the VSS nonoperational. If the timer expires, the software image reverts to the previous software image. To stop the timer, you must either accept or commit the new software image. The timer duration can be set with one number (<i>seconds</i>), indicating the number of seconds, or as hours, minutes, and seconds with a colon as the delimiter (<i>hh:mm:ss</i>). The range is 0 to 7200 seconds (2 hours); the default is 2700 seconds (45 minutes). A setting of 0 disables the rollback timer.
Step 3	Router(config)# exit	Returns to privileged EXEC mode.
Step 4	Router# show issu rollback timer	Displays the current rollback timer value.

This example shows how to set the eFSU rollback timer to one hour using both command formats:

```
Router# config terminal
Router(config)# issu set rollback-timer 3600
% Rollback timer value set to [ 3600 ] seconds
Router(config)# issu set rollback-timer 01:00:00
```

```
% Rollback timer value set to [ 3600 ] seconds
Router(config)#
```

Performing an eFSU Upgrade

To perform an eFSU upgrade (or downgrade) of a VSS, perform this task:

	Command	Purpose
Step 1	Router# copy tftp <i>disk_name</i>	Uses TFTP to copy the new software image to flash memory on the VSS active and VSS standby chassis (disk0: and slavedisk0:) and to the ICS's, if they exist. Answer the prompts to identify the name and location of the new software image.
Step 2	Router# show issu state [<i>switch/slot</i>] [detail]	(Optional) Verifies that the VSS is ready to run the eFSU. Note You can use the show issu state command at any stage of the upgrade to verify the progress and status of the upgrade.
Step 3	Router# issu loadversion [<i>active_switch/slot</i>] <i>active-image</i> [<i>standby_switch/slot</i>] <i>standby-image</i>	Starts the upgrade process by loading the new software image onto the VSS standby chassis. The image name includes the path of the target image to be loaded, in the format <i>devicename:filename</i> . It may take several seconds for the new image to load and for the VSS standby chassis to transition to SSO mode.
Step 4	Router# issu runversion	Forces a switchover, causing the VSS standby chassis to become VSS active and begin running the new software. The previously VSS active chassis becomes VSS standby and boots with the old image.
Step 5	Router# issu acceptversion	(Optional) Halts the rollback timer to ensure that the new software image is not automatically aborted during the upgrade process.
Step 6	Router# issu commitversion	Loads the new software image onto the VSS standby chassis.
Step 7	Router# show issu state [<i>switch/slot</i>] [detail]	Verifies the status of the upgrade process. If the upgrade was successful, both the VSS active and VSS standby chassis are running the new software version.

For an example of the eFSU upgrade sequence, see the “[eFSU Upgrade Example](#)” section on page 4-60.

Aborting an eFSU Upgrade

To manually abort the eFSU and roll back the upgrade, perform this task:

Command	Purpose
Router# issu abortversion	Stops the upgrade process and forces a rollback to the previous software image.

This example shows how to abort an eFSU upgrade for a VSS:

```
Router# issu abortversion
```

eFSU Upgrade Example

This example shows how to perform and verify an eFSU upgrade for a VSS.

Verify the System Readiness

After copying the new image files into the file systems of the active and VSS standby chassis, enter the **show issu state detail** command and the **show redundancy status** command to check that the VSS is ready to perform the eFSU. One chassis must be in the active state and the other chassis in the hot VSS standby state. Both chassis should be in the ISSU Init state and in SSO redundancy state. In the example, both chassis are running an “oldversion” image.

```
Router# show issu state detail
      Slot = 1/2
      RP State = Active
      ISSU State = Init
      Boot Variable = disk0:s72033-oldversion.v1,12;
      Operating Mode = sso
      Primary Version = N/A
      Secondary Version = N/A
      Current Version = disk0:s72033-oldversion.v1
      Variable Store = PrstVbl

      Slot = 2/7
      RP State = Standby
      ISSU State = Init
      Boot Variable = disk0:s72033-oldversion.v1,12;
      Operating Mode = sso
      Primary Version = N/A
      Secondary Version = N/A
      Current Version = disk0:s72033-oldversion.v1

Router# show redundancy status
  my state = 13 -ACTIVE
  peer state = 8 -STANDBY HOT
    Mode = Duplex
    Unit = Secondary
    Unit ID = 18

Redundancy Mode (Operational) = sso
Redundancy Mode (Configured) = sso
Redundancy State = sso
  Maintenance Mode = Disabled
  Communications = Up

  client count = 132
  client_notification_TMR = 30000 milliseconds
    keep_alive TMR = 9000 milliseconds
    keep_alive count = 0
    keep_alive threshold = 18
    RF debug mask = 0x0
```

Load the New Image to the VSS Standby Chassis

Enter the **issu loadversion** command to start the upgrade process. In this step, the VSS standby chassis reboots, reloads with the new image, and initializes as the VSS standby chassis in SSO redundancy mode, running the new image. This step is complete when the chassis configuration is synchronized, as indicated by the “Bulk sync succeeded” message.

```
Router# issu loadversion disk0:s72033-newversion.v2
```

```
000133: Aug  6 16:17:44.486 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet1/2/4, changed state to down
000134: Aug  6 16:17:43.507 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet2/7/4, changed state to down
000135: Aug  6 16:17:43.563 PST: %LINK-3-UPDOWN: Interface TenGigabitEthernet2/7/4,
changed state to down
000136: Aug  6 16:17:44.919 PST: %LINK-3-UPDOWN: Interface TenGigabitEthernet1/2/4,
changed state to down
```

(Deleted many interface and protocol down messages)

```
%issu loadversion executed successfully, Standby is being reloaded
```

(Deleted many interface and protocol down messages, then interface and protocol up messages)

```
0000148: Aug  6 16:27:54.154 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet1/2/5, changed state to up
000149: Aug  6 16:27:54.174 PST: %LINK-3-UPDOWN: Interface TenGigabitEthernet2/7/5,
changed state to up
000150: Aug  6 16:27:54.186 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet2/7/5, changed state to up
000151: Aug  6 16:32:58.030 PST: %HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync
succeeded
```

Verify the New Image on the VSS Standby Chassis

You can now enter the **show issu state detail** command and the **show redundancy** command to check that both chassis are in the ISSU Load Version state and SSO redundancy state. In this example, the VSS standby chassis is now running the “newversion” image.

```
Router# show issu state detail
      Slot = 1/2
      RP State = Active
      ISSU State = Load Version
      Boot Variable = disk0:s72033-oldversion.v1,12
      Operating Mode = sso
      Primary Version = disk0:s72033-oldversion.v1
      Secondary Version = disk0:s72033-newversion.v2
      Current Version = disk0:s72033-oldversion.v1
      Variable Store = PrstVbl

      Slot = 2/7
      RP State = Standby
      ISSU State = Load Version
      Boot Variable =
disk0:s72033-newversion.v2,12;disk0:s72033-oldversion.v1,12
      Operating Mode = sso
      Primary Version = disk0:s72033-oldversion.v1
      Secondary Version = disk0:s72033-newversion.v2
      Current Version = disk0:s72033-newversion.v2

Router# show redundancy status
  my state = 13 -ACTIVE
  peer state = 8 -STANDBY HOT
    Mode = Duplex
    Unit = Secondary
    Unit ID = 18

Redundancy Mode (Operational) = sso
Redundancy Mode (Configured) = sso
Redundancy State = sso
  Maintenance Mode = Disabled
  Communications = Up

  client count = 132
  client_notification_TMR = 30000 milliseconds
    keep_alive TMR = 9000 milliseconds
    keep_alive count = 1
  keep_alive threshold = 18
  RF debug mask = 0x0
```

Execute a Switchover to the New Image

When the VSS standby chassis is successfully running the new image in SSO redundancy state, enter the **issu runversion** command to force a switchover. The upgraded VSS standby chassis takes over as the new active chassis, running the new image. The formerly active chassis reloads and initializes as the new VSS standby chassis in SSO mode, running the old image (in case the software upgrade needs to be aborted and the old image restored). This step is complete when the chassis configuration is synchronized, as indicated by the “Bulk sync succeeded” message.

```
Router# issu runversion
This command will reload the Active unit.  Proceed ? [confirm]
(Deleted many lines)

Download Start
```

```

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
(Deleted many lines)
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Download Completed! Booting the image.
Self decompressing the image :
#####
(Deleted many lines)
##### [OK]
running startup....

(Deleted many lines)

000147: Aug  6 16:53:43.199 PST: %HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync
succeeded

```

Verify the Switchover

You can now enter the **show issu state detail** command and the **show redundancy** command to check that both chassis are in the ISSU Run Version state and SSO redundancy state. In this example, the active chassis is now running the “newversion” image.

```

Router# show issu state detail
          Slot = 2/7
          RP State = Active
          ISSU State = Run Version
          Boot Variable =
disk0:s72033-newversion.v2,12;disk0:s72033-oldversion.v1,12
          Operating Mode = sso
          Primary Version = disk0:s72033-newversion.v2
          Secondary Version = disk0:s72033-oldversion.v1
          Current Version = disk0:s72033-newversion.v2
          Variable Store = PrstVbl

          Slot = 1/2
          RP State = Standby
          ISSU State = Run Version
          Boot Variable = disk0:s72033-oldversion.v1,12
          Operating Mode = sso
          Primary Version = disk0:s72033-newversion.v2
          Secondary Version = disk0:s72033-oldversion.v1
          Current Version = disk0:s72033-oldversion.v1

Router# show redundancy status
my state = 13 -ACTIVE
peer state = 8 -STANDBY HOT
Mode = Duplex
Unit = Primary
Unit ID = 39

Redundancy Mode (Operational) = sso
Redundancy Mode (Configured) = sso
Redundancy State = sso
Maintenance Mode = Disabled
Communications = Up

client count = 134
client_notification_TMR = 30000 milliseconds
keep_alive TMR = 9000 milliseconds
keep_alive count = 1
keep_alive threshold = 18
RF debug mask = 0x0

```

Commit the New Image to the VSS Standby Chassis

When the active chassis is successfully running the new image in the SSO redundancy state, you can enter either the **issu acceptversion** command to stop the rollback timer and hold this state indefinitely, or the **issu commitversion** command to continue with the eFSU. To continue, enter the **issu commitversion** command to upgrade the VSS standby chassis and complete the eFSU sequence. The VSS standby chassis reboots, reloads with the new image, and initializes as the VSS standby chassis in the SSO redundancy state, running the new image. This step is complete when the chassis configuration is synchronized, as indicated by the “Bulk sync succeeded” message.

```
Router# issu commitversion
Building configuration...
[OK]
000148: Aug  6 17:17:28.267 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet2/7/4, changed state to down
000149: Aug  6 17:17:28.287 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet1/2/4, changed state to down
```

(Deleted many interface and protocol down messages)

%issu commitversion executed successfully

(Deleted many interface and protocol down messages, then interface and protocol up messages)

```
000181: Aug  6 17:41:51.086 PST: %LINEPROTO-5-UPDOWN: Line protocol on Interface
TenGigabitEthernet1/2/5, changed state to up
000182: Aug  6 17:42:52.290 PST: %HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync
succeeded
```

Verify That the Upgrade is Complete

You can now enter the **show issu state detail** command and the **show redundancy** command to check the results of the eFSU. In this example, both chassis are now running the “newversion” image, indicating that the eFSU was successful. Because the eFSU has completed, the two chassis will be once again in the ISSU Init Version state, as they were before the eFSU was initiated.

```
Router# show issu state detail
Slot = 2/7
RP State = Active
ISSU State = Init
Boot Variable =
disk0:s72033-newversion.v2,12;disk0:s72033-oldversion.v1,12
Operating Mode = sso
Primary Version = N/A
Secondary Version = N/A
Current Version = disk0:s72033-newversion.v2
Variable Store = PrstVbl

Slot = 1/2
RP State = Standby
ISSU State = Init
Boot Variable =
disk0:s72033-newversion.v2,12;disk0:s72033-oldversion.v1,12
Operating Mode = sso
Primary Version = N/A
Secondary Version = N/A
Current Version = disk0:s72033-newversion.v2

Router# show redundancy status
```



```
my state = 13 -ACTIVE
peer state = 8  -STANDBY HOT
  Mode = Duplex
  Unit = Primary
  Unit ID = 39

Redundancy Mode (Operational) = sso
Redundancy Mode (Configured)  = sso
Redundancy State               = sso
  Maintenance Mode = Disabled
Communications = Up

client count = 134
client_notification_TMR = 30000 milliseconds
  keep_alive TMR = 9000 milliseconds
  keep_alive count = 1
  keep_alive threshold = 18
  RF debug mask = 0x0
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

Participate in the [Technical Documentation Ideas forum](#)

5

Enhanced Fast Software Upgrade

- [Prerequisites for eFSU, page 5-1](#)
- [Restrictions for eFSU, page 5-2](#)
- [Information About eFSU, page 5-3](#)
- [Default Settings for eFSU, page 5-5](#)
- [How to Perform an eFSU, page 5-5](#)
- [How to Upgrade a Non-eFSU Image to an eFSU Image, page 5-14](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - The
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

For detailed information on 6800IA ISSU upgrade procedure, see the following documents:

<http://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/117564-technote-is-su-00.html>

http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-6800ia-switch/white_paper_c11-728265.html

Prerequisites for eFSU

None.

Restrictions for eFSU

- [SX SY EFSU Compatibility Matrix](#)

- The Release 15.0(1)SY no payload encryption (NPE) images do not support the hitless ACL update feature or the **[no] platform hardware acl update-mode hitless** command.

The Release 15.0(1)SY1 and later no payload encryption (NPE) images support hitless ACL update and the **platform hardware acl update-mode hitless** command is configured by default (because it is the default, the command does not appear in the configuration file).

In other releases and images that support the hitless ACL update feature, the **platform hardware acl update-mode hitless** command is configured by default.

With NPE images, to avoid problems when doing a downgrade from Release 15.0(1)SY1 or later to Release 15.0(1)SY, do not disable the hitless ACL update feature (**no platform hardware acl update-mode hitless**), because the CLI does not exist in the Release 15.0(1)SY NPE images, and configuring the nondefault condition causes the CLI to appear in the Release 15.0(1)SY1 configuration file, which then, as an unsupported command, causes problems with Release 15.0(1)SY.

The hitless ACL update feature consumes TCAM resources. If TCAM utilization is high, enabling the hitless ACL update feature to support a downgrade might cause TCAM conflicts with other configured features.

- eFSU requires two supervisor engines, one active and one standby.
- Both the active and standby supervisor engines must have enough flash memory to store both the old and new software images prior to the upgrade process.
- Images from different features sets, regardless of release, fail the eFSU compatibility check.
- When downgrading with eFSU to an earlier version of Cisco IOS Software, the configuration files fail to synchronize and the standby supervisor engine reloads unless you disable any features or functions that are not supported in the earlier version before you start the process. Remove any configuration commands that are not available in the earlier version.
- During an eFSU upgrade, the modules are restarted.
- The switch examines the old and new software images and automatically performs the appropriate process (eFSU) to upgrade the software image:
 - For a patch upgrade, if the module software is the same in both the old and the new software images, because no module software upgrade is required, the eFSU upgrades only the supervisor engine software. The system downtime is from 0 to 3 seconds.
 - If the module software in the images is different, the modules are restarted or reset during the upgrade process. System downtime depends on whether the modules support eFSU (see the [“Outage Time and Support Considerations”](#) section on page 5-4 for more information).
- The eFSU upgrade feature works with NSF/SSO. Software features that do not support NSF/SSO stop operating until they come back online after the switchover that occurs during the software upgrade.
- All modules that support eFSU preload must have at least 512 MB of memory, with enough memory free to hold the new software image. If there is insufficient free memory, eFSU does not attempt the preload, but instead resets the modules during the switchover.
- Online insertion and replacement (OIR) is not supported during an eFSU. If you attempt to insert a new module in the switch while the upgrade is active, the switch does not provide power for the module. When the upgrade ends, the switch resets the newly inserted module.

- Do not perform a manual switchover between supervisor engines during the upgrade.
- Make sure that the configuration register is set to allow autoboot (the lowest byte of the register should be set to 2).
- Before you enter the **issu abortversion** command (to abort a software upgrade), make sure that the standby supervisor engine is Up (STANDBY HOT [in SSO] or COLD [in RPR]).
- The Fast Software Upgrade (FSU) process supports upgrade from earlier releases. During this process, the module software image is also upgraded on those modules that support eFSU.
- Release 15.5(1)SY3 does not support Smart install feature. Before performing In-Service Software Upgrade (ISSU), remove vstack and related configurations from the running configuration.

**Note**

eFSU will not proceed in IA systems if **issu runversion fex** command is not executed.

Information About eFSU

- [eFSU Operation, page 5-3](#)
- [Outage Time and Support Considerations, page 5-4](#)
- [Reserving Module Memory, page 5-4](#)
- [Error Handling for eFSU Preload, page 5-5](#)

**Note**

The switch supports eFSU in VSS mode. See the “[Restrictions for VSS](#)” section on [page 4-2](#) for more information.

eFSU Operation

eFSU is an enhanced software upgrade procedure. Non-eFSU (FSU) software upgrades require system downtime, because a software version mismatch between the active and the standby supervisor engines forces the system to boot in RPR redundancy mode, which is stateless and causes a hard reset of the all modules.

eFSU enables an increase in network availability by reducing the downtime caused by software upgrades. eFSU does this by:

- Bringing up the standby supervisor engine in SSO mode even when the active and the standby supervisor engines have different software versions, or with VSS configured, when the supervisor engines in the two chassis have different software versions.

During an eFSU, new software is loaded onto the standby supervisor engine while the active supervisor engine continues to operate using the previous software. As part of the upgrade, the standby processor reaches the SSO Standby Hot stage, a switchover occurs, and the standby becomes active, running the new software. In previous releases Supervisor Engines running different software versions ran in the Route Processor Redundancy Mode.

You can continue with the upgrade to load the new software onto the other processor, or you can abort the upgrade and resume operation with the old software.

- Preloading new module software into memory on supported modules to avoid a hard reset.

If the new software release contains new module software, eFSU preloads the new module software onto any modules in the switch that support eFSU preload. When the switchover occurs between the active and standby supervisor engines, the modules are restarted with the new software image.

The WS-X67xx modules support eFSU preload. All other modules undergo a hard reset at switchover, and the software image loads after the module restarts.

During a software upgrade, the switch performs the following steps automatically on modules that support eFSU preload:

- Reserves the necessary memory for the new Cisco IOS software image on each module.
- Preloads a new software image onto the modules as part of the **issu loadversion** command.
- Restarts the modules with the new software image when a switchover occurs (**issu runversion**).
- During the restart, the software features and routing protocols are not available.
- If a rollback or abort occurs, to minimize disruption, the switch preloads the original software version onto the module. Once the rollback or abort is completed, the module is restarted with the original software version.



Note

All modules that support eFSU preload must have at least 512 MB of memory, with enough memory free to hold the new software image. If there is insufficient free memory, eFSU does not attempt the preload, but instead resets the modules during the switchover.

Outage Time and Support Considerations

During an eFSU upgrade, modules are restarted or reset after the switchover that occurs between the supervisor engines. Because the modules are restarted or reset, any links attached to the modules go up and down and traffic processing is disrupted until protocols and software features are brought back online. The length of time that module processing is disrupted (outage time) depends on whether the eFSU process was able to preload a new software image onto the module.

- For modules that support eFSU preload, the outage time for an eFSU module warm reload is faster than an RPR mode module reload.
- For modules that do not support eFSU preload, the outage time for module reload is similar to an RPR mode module reload.

Once the new software is loaded (**issu loadversion**), you can use the **show issu outage slot all** command to display the maximum outage time for installed modules. See the [“Displaying the Maximum Outage Time for Installed Modules \(Optional\)”](#) section on page 5-10 for a command example.

Reserving Module Memory

On modules that support eFSU, the supervisor engine automatically reserves memory on the module to store the new software image (decompressed format). The amount of memory needed varies according to the module type.

Although we do not recommend it, you can enter the following command to keep the switch from reserving memory for the software preload (where *slot-num* specifies which slot the module is installed in):

```
no mdr download reserve memory image slot slot-num
```

**Note**

All modules that support eFSU preload must have at least 512 MB of memory, with enough memory free to hold the new software image. If there is insufficient free memory, eFSU does not attempt the preload, but instead resets the modules during the switchover.

To display whether or not the memory reservation was successful on a module, use the **show issu outage slot all** command. See the [“Displaying the Maximum Outage Time for Installed Modules \(Optional\)”](#) section on page 5-10 for a command example.

Error Handling for eFSU Preload

If problems occur during eFSU preload, the switch takes the following actions:

- Module crash during loadversion—The module is reset when a switchover occurs.
- Module not active when eFSU started—No power is provided to the module during the software upgrade, and the module is reset when the process ends. The same action is applied to a module that is inserted into the switch after the software upgrade process has begun.
- Module crash during run version or during rollback—The module boots with the software image version that corresponds to the software image that is present on the active supervisor engine.

Default Settings for eFSU

None.

How to Perform an eFSU

- [eFSU Summarized Procedure, page 5-6](#)
- [Preparing for the Upgrade, page 5-6](#)
- [Copying the New Software Image, page 5-8](#)
- [Loading the New Software onto the Standby Supervisor Engine, page 5-8](#)
- [Displaying the Maximum Outage Time for Installed Modules \(Optional\), page 5-10](#)
- [Forcing a Switchover from Active to Standby, page 5-10](#)
- [Accepting the New Software Version and Stopping the Rollback Process \(Optional\), page 5-12](#)
- [Committing the New Software to the Standby, page 5-12](#)
- [Verifying the Software Installation, page 5-12](#)
- [Aborting the Upgrade Process, page 5-14](#)

eFSU Summarized Procedure

This task is a summarized procedure for an eFSU:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Router# copy tftp disk_name	Uses TFTP to copy the new software image to flash memory on the active and standby supervisor engines (disk0: and slavedisk0:). Answer the prompts to identify the name and location of the new software image.
Step 3	Router# show version in image Router# show bootvar Router# show redundancy Router# show issu state [detail]	These show commands verify that the switch is ready to run eFSU. The show version and show bootvar commands verify the boot image settings. The show redundancy and show issu state commands verify that redundancy mode is enabled and that SSO and NSF are configured. Note Use the show redundancy and show issu state commands throughout the upgrade to verify the status of the upgrade.
Step 4	Router# issu loadversion active-slot active-image standby-slot standby-image	Starts the upgrade process and loads the new software image onto the standby supervisor engine. It may take several seconds for the new image to load and for the standby supervisor engine to transition to SSO mode.
Step 5	Router# show issu outage slot all	(Optional) Displays the maximum outage time for installed modules. Enter the command on the switch processor of the supervisor engine.
Step 6	Router# issu runversion	Forces a switchover, which causes the standby supervisor engine to become active and begin running the new software. The previously active processor becomes standby and boots with the old image.
Step 7	Router# issu acceptversion	(Optional) Halts the rollback timer to ensure that the new software image is not automatically aborted during the upgrade process.
Step 8	Router# issu commitversion	Loads the new software image onto the standby supervisor engine in the specified slot.
Step 9	Router# show redundancy Router# show issu state [detail]	Verifies the status of the upgrade process. If the upgrade was successful, both the active and standby supervisor engines are running the new software version.

Preparing for the Upgrade

- [Verifying the Boot Image Version and Boot Variable, page 5-7](#)
- [Verifying Redundancy Mode, page 5-7](#)
- [Verifying eFSU State, page 5-8](#)

**Note**

Before attempting to perform a software upgrade, be sure to review the [“Restrictions for eFSU”](#) section on page 5-2.

Verifying the Boot Image Version and Boot Variable

Before starting, enter the **show version** and **show bootvar** commands to verify the boot image version and BOOT environment variable, as shown in the following examples:

```
Router# show version | in image
BOOT variable = disk0:image_name;
CONFIG_FILE variable =
BOOTLDR variable =
Configuration register is 0x2002

Standby is up
Standby has 1048576K/65536K bytes of memory.

Standby BOOT variable = disk0:image_name;
Standby CONFIG_FILE variable =
Standby BOOTLDR variable =
```

Verifying Redundancy Mode

Verify that redundancy mode is enabled and that NSF and SSO are configured. The following command example shows how to verify redundancy:

```
Router# show redundancy
Redundant System Information :
-----
    Available system uptime = 45 minutes
Switchovers system experienced = 0
    Standby failures = 0
    Last switchover reason = none

    Hardware Mode = Duplex
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso
    Maintenance Mode = Disabled
    Communications = Up

Current Processor Information :
-----
    Active Location = slot 6
    Current Software state = ACTIVE
    Uptime in current state = 44 minutes
    Image Version = Cisco IOS Software, image_details
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled Wed 18-Feb-09 12:48 by kchristi
    BOOT = disk0:image_name;
    CONFIG_FILE =
    BOOTLDR =
    Configuration register = 0x2002

Peer Processor Information :
-----
    Standby Location = slot 5
    Current Software state = STANDBY HOT
    Uptime in current state = 28 minutes
```

```

Image Version = Cisco IOS Software, image_details
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled image_details

BOOT = disk0:image_name ;
CONFIG_FILE =
BOOTLDR =
Configuration register = 0x2002

```

Verifying eFSU State

Verify that the the ISSU state is **Init**, rather than an intermediate eFSU upgrade state. Enter this command:

```

Router# show issu state detail

Slot = 6
RP State = Active
ISSU State = Load Version
Boot Variable = disk0:image_name
Operating Mode = sso
Primary Version = disk0:sierra.0217
Secondary Version = disk0:sierra.0217
Current Version = disk0:sierra.0217
Variable Store = PrstVbl
ROMMON CV = [disk0:image_name]

Slot = 5
RP State = Standby
ISSU State = Load Version
Boot Variable = disk0:image_name
Operating Mode = sso
Primary Version = disk0:image_name
Secondary Version = disk0:image_name
Current Version = disk0:image_name

```

Copying the New Software Image

Before starting the eFSU process, copy the new software image to flash memory (disk0: and slavedisk0:) on the active and standby supervisor engines.

Loading the New Software onto the Standby Supervisor Engine

Enter the **issu loadversion** command to start the upgrade process. This command reboots the standby supervisor engine and loads the new software image onto the standby supervisor engine. When the download is complete, you are prompted to enter the **runversion** command.



Note

Do not automatically disable the features that are not common to both images. During the standby initialization, after you enter the **issu loadversion** command, if there are any enabled features that are not supported on the standby supervisor engine, a message is displayed that states that the standby supervisor engine cannot initialize while this feature is enabled, and the standby supervisor engine is forced to RPR (in the load-version state).

```

Router# issu loadversion device:filename
%issu loadversion executed successfully, Standby is being reloaded

```

When execution of the **issu loadversion** command completes, the standby supervisor engine is loaded with the new software image and the supervisor engine is in SSO mode. The **issu loadversion** command might take several seconds to complete. If you enter the **show** commands too soon, you might not see the information that you need.

These examples show how to check the status of the upgrade using the **show redundancy** and **show issu state detail** commands:

```
Router# show redundancy
Redundant System Information :
-----
    Available system uptime = 1 hour, 0 minutes
    Switchovers system experienced = 0
        Standby failures = 1
    Last switchover reason = none

    Hardware Mode = Duplex
    Configured Redundancy Mode = sso
    Operating Redundancy Mode = sso
    Maintenance Mode = Disabled
    Communications = Up

Current Processor Information :
-----
    Active Location = slot 6
    Current Software state = ACTIVE
    Uptime in current state = 59 minutes
    Image Version = Cisco IOS Software, image_details
    Technical Support: http://www.cisco.com/techsupport
    Copyright (c) 1986-2009 by Cisco Systems, Inc.
    Compiled ...

    BOOT = disk0:image_name
    CONFIG_FILE =
    BOOTLDR =
    Configuration register = 0x2002

Peer Processor Information :
-----
    Standby Location = slot 5
    Current Software state = STANDBY HOT
    Uptime in current state = 3 minutes
    Image Version = Cisco IOS Software, image_name
    Technical Support: http://www.cisco.com/techsupport
    Copyright (c) 1986-2009 by Cisco Systems, Inc.
    Compiled ...

    BOOT = disk0:image_name
    CONFIG_FILE =
    BOOTLDR =
    Configuration register = 0x2002

Router# show issu state detail
    Slot = 6
    RP State = Active
    ISSU State = Load Version
    Boot Variable = disk0:image_name
    Operating Mode = sso
    Primary Version = disk0:image_name
    Secondary Version = disk0:image_name
    Current Version = disk0:image_name
    Variable Store = PrstVbl
    ROMMON CV = [disk0:image_name]

    Slot = 5
```

```

RP State = Standby
ISSU State = Load Version
Boot Variable = disk0:image_name
Operating Mode = sso
Primary Version = disk0:image_name
Secondary Version = disk0:image_name
Current Version = disk0:image_name

```

Displaying the Maximum Outage Time for Installed Modules (Optional)

Once the new software is downloaded, you can enter the **show issu outage slot all** command on the switch processor to display the maximum outage time for the installed modules:

```

Router# show issu outage slot all
Slot # Card Type                                MDR Mode      Max Outage Time
-----
1 CEF720 8 port 10GE with DFC                  WARM_RELOAD    300 secs
2 96-port 10/100 Mbps RJ45                     RELOAD         360 secs
4 CEF720 48 port 1000mb SFP                    RELOAD         360 secs

Slot # Reason                                Error Number
-----
1 PLATFORM_INIT                             3
2 PLATFORM_INIT                             3
4 PREDOWNLOAD_LC_MIMIMUM_MEMORY_FAILURE      5
Router#

```

Forcing a Switchover from Active to Standby

Enter the **issu runversion** command to force a switchover between the active and standby supervisor engines. The standby supervisor engine, which has the new software image loaded, becomes active. The previously active supervisor engine becomes the standby and boots with the old software image (in case the software upgrade needs to be aborted and the old image restored).

```
Router# issu runversion
```

This command will reload the Active unit. Proceed ? [confirm] **y**

A switchover between the supervisor engines occurs now. The previous standby supervisor engine becomes active and is running the new software version. The previous active supervisor engine, now the standby supervisor engine, boots with the old software.



Note

At this point, the new active supervisor engine is running the new software image and the standby is running the old software image. You should verify the state of the active and standby supervisor engines as shown in the following examples (**show redundancy** and **show issu state detail**).

```

Router# show redundancy
-----
Available system uptime = 1 hour, 9 minutes
Switchovers system experienced = 1
Standby failures = 0
Last switchover reason = user forced

Hardware Mode = Duplex
Configured Redundancy Mode = sso
Operating Redundancy Mode = sso

```

```

Maintenance Mode = Disabled
Communications = Up

Current Processor Information :
-----
      Active Location = slot 5
      Current Software state = ACTIVE
      Uptime in current state = 7 minutes
      Image Version = Cisco IOS Software, image_details
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled ...
      BOOT = disk0:image_name
      CONFIG_FILE =
      BOOTLDR =
      Configuration register = 0x2002

Peer Processor Information :
-----
      Standby Location = slot 6
      Current Software state = STANDBY HOT
      Uptime in current state = 0 minutes
      Image Version = Cisco IOS Software, image_details
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled Wed 18-Feb-09 12:48 by kchristi
      BOOT = disk0:image_name
      CONFIG_FILE =
      BOOTLDR =
      Configuration register = 0x2002

Router# show issu state detail
      Slot = 5
      RP State = Active
      ISSU State = Run Version
      Boot Variable = disk0:image_name
      Operating Mode = sso
      Primary Version = disk0:image_name
      Secondary Version = disk0:image_name
      Current Version = disk0:image_name
      Variable Store = PrstVbl
      ROMMON CV = [disk0:image_name]

      Slot = 6
      RP State = Standby
      ISSU State = Run Version
      Boot Variable = disk0:image_name
      Operating Mode = sso
      Primary Version = disk0:image_name
      Secondary Version = disk0:image_name
      Current Version = disk0:image_name

```

**Note**

To complete the upgrade process, enter the **issu acceptversion** (optional) and **issu commitversion** commands (as described in the following sections).

Accepting the New Software Version and Stopping the Rollback Process (Optional)

You must either accept or commit the new software image, or the rollback timer will expire and stop the upgrade process. If that occurs, the software image reverts to the previous software version. The rollback timer is a safeguard to ensure that the upgrade process does not leave the switch nonoperational.



Note

New features that are not supported by the previous image are allowed to be enabled only after you enter the **issu commitversion** command.

The following command sequence shows how the **issu acceptversion** command stops the rollback timer to enable you to examine the functionality of the new software image. When you are satisfied that the new image is acceptable, enter the **issu commitversion** command to end the upgrade process.

```
Router# show issu rollback-timer
      Rollback Process State = In progress
      Configured Rollback Time = 00:45:00
      Automatic Rollback Time = 00:37:28

Router# issu acceptversion
% Rollback timer stopped. Please issue the commitversion command.
```

View the rollback timer to see that the rollback process has been stopped:

```
Router# show issu rollback-timer
      Rollback Process State = Not in progress
      Configured Rollback Time = 00:45:00
```

Committing the New Software to the Standby

Enter the **issu commitversion** command to load the new software image onto the standby supervisor engine and complete the software upgrade process. In the following example, the new image is loaded onto the standby supervisor engine in slot 5:

```
Router# issu commitversion
Building configuration...
[OK]
%issu commitversion executed successfully
```



Note

The software upgrade process is now complete. Both the active and standby supervisor engines are running the new software version.

Verifying the Software Installation

You should verify the status of the software upgrade. If the upgrade was successful, both the active and standby supervisor engines are running the new software version.

```
Router# show redundancy
Redundant System Information :
-----
      Available system uptime = 1 hour, 17 minutes
Switchovers system experienced = 1
      Standby failures = 1
```

```

Last switchover reason = user forced

      Hardware Mode = Duplex
Configured Redundancy Mode = sso
      Operating Redundancy Mode = sso
      Maintenance Mode = Disabled
      Communications = Up

Current Processor Information :
-----
      Active Location = slot 5
      Current Software state = ACTIVE
      Uptime in current state = 15 minutes
      Image Version = Cisco IOS Software, image_name
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled ...

      BOOT = disk0:image_name
      CONFIG_FILE =
      BOOTLDR =
      Configuration register = 0x2002

Peer Processor Information :
-----
      Standby Location = slot 6
      Current Software state = STANDBY HOT
      Uptime in current state = 0 minutes
      Image Version = Cisco IOS Software, image_details
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled ...

      BOOT = disk0:image_name
      CONFIG_FILE =
      BOOTLDR =
      Configuration register = 0x2002

Router# show issu state detail

      Slot = 5
      RP State = Active
      ISSU State = Init
      Boot Variable = disk0:image_name
      Operating Mode = sso
      Primary Version = N/A
      Secondary Version = N/A
      Current Version = disk0:image_name
      Variable Store = PrstVbl
      ROMMON CV = [disk0:image_name ]

      Slot = 6
      RP State = Standby
      ISSU State = Init
      Boot Variable = disk0:image_name
      Operating Mode = sso
      Primary Version = N/A
      Secondary Version = N/A
      Current Version = disk0:image_name

```

Aborting the Upgrade Process

You can manually abort the software upgrade at any stage by entering the **issu abortversion** command. The upgrade process also aborts on its own if the software detects a failure.

If you abort the process after you enter the **issu loadversion** command, the standby supervisor engine is reset and reloaded with the original software.

The following is an example of the **issu abortversion slot image** command that shows how to abort the software upgrade process:

```
Router# issu abortversion 6 c7600s72033
```

**Note**

Before you enter the **issu abortversion** command, make sure that the standby supervisor engine is Up (STANDBY HOT [in SSO] or COLD [in RPR]).

How to Upgrade a Non-eFSU Image to an eFSU Image

If the new Cisco IOS software image does not support eFSU, you must manually upgrade the software image. To do so, you must upgrade the software image on the standby supervisor engine and then perform a manual switchover so that the standby takes over processing with the new image. You can then upgrade the software image on the previously active, and now standby, supervisor engine. For more information, see the “[eFSU Summarized Procedure](#)” section on [page 5-6](#).

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

6

Fast Software Upgrade



Note

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- Supported only with redundant supervisor engines. Cisco IOS software is upgraded on the standby RP, and a manual switchover is performed. The new Cisco IOS image can then be upgraded on the other RP.
- During the upgrade process, different images will be loaded on the RPs for a very short period of time. If a switchover occurs during this time, the device will recover in RPR mode.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

To upgrade or downgrade the Cisco IOS image, perform this task:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# copy { ftp: http:// https:// rcp: scp: tftp: } <i>device:filename</i>	Copies a Cisco IOS image onto the flash device of the active RP.
Step 3	Router# copy { ftp: http:// https:// rcp: scp: tftp: } slavedevice: <i>filename</i>	Copies a Cisco IOS image onto the flash device of the standby RP.
Step 4	Router# configure terminal	Enters global configuration mode.
Step 5	Router(config)# no boot system flash [<i>flash-fs:</i>] [<i>partition-number:</i>] [<i>filename</i>]	(Optional) Clears any existing system flash boot image specification.
Step 6	Router(config)# boot system flash [<i>flash-fs:</i>] [<i>partition-number:</i>] [<i>filename</i>]	Specifies the filename of stored image in flash memory.
Step 7	Router(config)# config-register 0x2102	Sets the configuration register setting to the default value.

	Command	Purpose
Step 8	Router(config)# exit	Exits global configuration mode and returns the router to privileged EXEC mode.
Step 9	Router# copy running-config startup-config	Saves the configuration changes to the startup configuration file.
Step 10	hw-module {module <i>standby_slot</i> } reset	Resets and reloads the standby processor with the specified Cisco IOS image, and executes the image.
Step 11	redundancy force-switchover	Forces a switchover to the standby RP.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

7

Stateful Switchover (SSO)

- [Prerequisites for SSO, page 7-1](#)
- [Restrictions for SSO, page 7-1](#)
- [Information About SSO, page 7-3](#)
- [Default Settings for SSO, page 7-10](#)
- [How to Configure SSO, page 7-10](#)
- [Troubleshooting SSO, page 7-11](#)
- [Verifying the SSO Configuration, page 7-12](#)
- [Configuration Examples for SSO, page 7-16](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- SSO and NSF do not support IPv6 multicast traffic.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for SSO

None.

Restrictions for SSO

- [General Restrictions, page 7-2](#)
- [Configuration Mode Restrictions, page 7-2](#)

- [Switchover Process Restrictions, page 7-2](#)

General Restrictions

- Two RPs must be installed in the chassis, each running the same version of the Cisco IOS software.
- Both RPs must run the same Cisco IOS image. If the RPs are operating different Cisco IOS images, the system reverts to RPR mode even if SSO is configured.
- Configuration changes made through SNMP may not be automatically configured on the standby RP after a switchover occurs.
- Load sharing between dual processors is not supported.
- Enhanced Object Tracking (EOT) is not stateful switchover-aware and cannot be used with HSRP, Virtual Router Redundancy Protocol (VRRP), or Gateway Load Balancing Protocol (GLBP) in SSO mode.
- Multicast is not SSO-aware and restarts after switchover; therefore, multicast tables and data structures are cleared upon switchover.

Configuration Mode Restrictions

- The configuration registers on both RPs must be set the same for the networking device to behave the same when either RP is rebooted.
- During the startup (bulk) synchronization, configuration changes are not allowed. Before making any configuration changes, wait for a message similar to the following:

```
%HA-5-MODE:Operating mode is sso, configured mode is sso.
```

Switchover Process Restrictions

- If any changes to the fabric configuration happen simultaneously with an RP switchover, the chassis is reset and all line cards are reset.
- If the switch is configured for SSO mode, and the active RP fails before the standby is ready to switch over, the switch will recover through a full system reset.
- During SSO synchronization between the active and standby RPs, the configured mode will be RPR. After the synchronization is complete, the operating mode will be SSO. If a switchover occurs before the synchronization is complete, the switchover will be in RPR mode.
- If a switchover occurs before the bulk synchronization step is complete, the new active RP may be in inconsistent states. The switch will be reloaded in this case.
- Switchovers in SSO mode will not cause the reset of any line cards.
- Interfaces on the RP itself are not stateful and will experience a reset across switchovers. In particular, the GE interfaces on the RPs are reset across switchovers and do not support SSO.
- Any line cards that are not online at the time of a switchover (line cards not in Cisco IOS running state) are reset and reloaded on a switchover.

Information About SSO

- [SSO Overview, page 7-3](#)
- [SSO Operation, page 7-5](#)
- [Route Processor Synchronization, page 7-6](#)
- [SSO Operation, page 7-8](#)
- [SSO-Aware Features, page 7-10](#)

SSO Overview

The switch supports fault resistance by allowing a redundant supervisor engine to take over if the primary supervisor engine fails. Cisco SSO (frequently used with NSF) minimizes the time a network is unavailable to its users following a switchover while continuing to forward IP packets. The switch supports route processor redundancy (RPR). For more information, see [Chapter 9, “Route Processor Redundancy \(RPR\).”](#)

SSO is particularly useful at the network edge. Traditionally, core routers protect against network faults using router redundancy and mesh connections that allow traffic to bypass failed network elements. SSO provides protection for network edge devices with dual Route Processors (RPs) that represent a single point of failure in the network design, and where an outage might result in loss of service for customers.


SSO has many benefits. Because the SSO feature maintains stateful feature information, user session information is maintained during a switchover, and line cards continue to forward network traffic with no loss of sessions, providing improved network availability. SSO provides a faster switchover than RPR by fully initializing and fully configuring the standby RP, and by synchronizing state information, which can reduce the time required for routing protocols to converge. Network stability may be improved with the reduction in the number of route flaps had been created when routers in the network failed and lost their routing tables.

SSO is required by the Cisco Nonstop Forwarding (NSF) feature (see [Chapter 8, “Nonstop Forwarding \(NSF\)”](#)).

Figure 7-1 illustrates how SSO is typically deployed in service provider networks. In this example, Cisco NSF with SSO is primarily at the access layer (edge) of the service provider network. A fault at this point could result in loss of service for enterprise customers requiring access to the service provider network.

For Cisco NSF protocols that require neighboring devices to participate in Cisco NSF, Cisco NSF-aware software images must be installed on those neighboring distribution layer devices. Additional network availability benefits might be achieved by applying Cisco NSF and SSO features at the core layer of your network; however, consult your network design engineers to evaluate your specific site requirements.


Figure 7-1 Cisco NSF with SSO Network Deployment: Service Provider Networks



72134

Additional levels of availability may be gained by deploying Cisco NSF with SSO at other points in the network where a single point of failure exists. [Figure 7-2](#) illustrates an optional deployment strategy that applies Cisco NSF with SSO at the enterprise network access layer. In this example, each access point in the enterprise network represents another single point of failure in the network design. In the event of a switchover or a planned software upgrade, enterprise customer sessions would continue uninterrupted through the network.

Figure 7-2 Cisco NSF with SSO Network Deployment: Enterprise Networks



SSO Operation

SSO establishes one of the RPs as the active processor while the other RP is designated as the standby processor. SSO fully initializes the standby RP, and then synchronizes critical state information between the active and standby RP.

During an SSO switchover, the line cards are not reset, which provides faster switchover between the processors. The following events cause a switchover:

- A hardware failure on the active supervisor engine
- Clock synchronization failure between supervisor engines
- A manual switchover or shutdown

An SSO switchover does not interrupt Layer 2 traffic. An SSO switchover preserves FIB and adjacency entries and can forward Layer 3 traffic after a switchover. SSO switchover duration is between 0 and 3 seconds.

Route Processor Synchronization

- [Synchronization Overview, page 7-6](#)
- [Bulk Synchronization During Initialization, page 7-6](#)
- [Synchronization of Startup Configuration, page 7-6](#)
- [Incremental Synchronization, page 7-7](#)

Synchronization Overview

In networking devices running SSO, both RPs must be running the same configuration so that the standby RP is always ready to assume control if the active RP fails. SSO synchronizes the configuration information from the active RP to the standby RP at startup and whenever changes to the active RP configuration occur. This synchronization occurs in two separate phases:

- While the standby RP is booting, the configuration information is synchronized in bulk from the active RP to the standby RP.
- When configuration or state changes occur, an incremental synchronization is conducted from the active RP to the standby RP.

Bulk Synchronization During Initialization

When a system with SSO is initialized, the active RP performs a chassis discovery (discovery of the number and type of line cards and fabric cards, if available, in the system) and parses the startup configuration file.

The active RP then synchronizes this data to the standby RP and instructs the standby RP to complete its initialization. This method ensures that both RPs contain the same configuration information.

Even though the standby RP is fully initialized, it interacts only with the active RP to receive incremental changes to the configuration files as they occur. Executing CLI commands on the standby RP is not supported.

Synchronization of Startup Configuration

During system startup, the startup configuration file is copied from the active RP to the standby RP. Any existing startup configuration file on the standby RP is overwritten.

The startup configuration is a text file stored in the NVRAM of the RP. It is synchronized whenever you perform the following operations:

- CLI command **copy system:running-config nvram:startup-config** is used.
- CLI command **copy running-config startup-config** is used.
- CLI command **write memory** is used.
- CLI command **copy filename nvram:startup-config** is used.
- SNMP SET of MIB variable ccCopyEntry in CISCO_CONFIG_COPY MIB is used.
- System configuration is saved using the **reload** command.
- System configuration is saved following entry of a forced switchover CLI command.

Incremental Synchronization

- [Incremental Synchronization Overview, page 7-7](#)
- [CLI Commands, page 7-7](#)
- [SNMP SET Commands, page 7-7](#)
- [Routing and Forwarding Information, page 7-7](#)
- [Chassis State, page 7-7](#)
- [Line Card State, page 7-7](#)
- [Counters and Statistics, page 7-8](#)

Incremental Synchronization Overview

After both RPs are fully initialized, any further changes to the running configuration or active RP states are synchronized to the standby RP as they occur. Active RP states are updated as a result of processing feature information, external events (such as the interface becoming up or down), or user configuration commands (using CLI commands or Simple Network Management Protocol [SNMP]) or other internal events.

CLI Commands

CLI changes to the running configuration are synchronized from the active RP to the standby RP. In effect, the CLI command is run on both the active and the standby RP.

SNMP SET Commands

Configuration changes caused by an SNMP set operation are synchronized on a case-by-case basis. Currently only two SNMP configuration set operations are supported:

- **shut** and **no-shut** (of an interface)
- **link up/down trap enable/disable**

Routing and Forwarding Information

Routing and forwarding information is synchronized to the standby RP:

- State changes for SSO-aware features (for example, SNMP) are synchronized to the standby RP.
- Cisco Express Forwarding updates to the Forwarding Information Base (FIB) are synchronized to the standby RP.

Chassis State

Changes to the chassis state due to line card insertion or removal are synchronized to the standby RP.

Line Card State

Changes to the line card states are synchronized to the standby RP. Line card state information is initially obtained during bulk synchronization of the standby RP. Following bulk synchronization, line card events, such as whether the interface is up or down, received at the active processor are synchronized to the standby RP.

Counters and Statistics

The various counters and statistics maintained in the active RP are not synchronized because they may change often and because the degree of synchronization they require is substantial. The volume of information associated with statistics makes synchronizing them impractical.

**Note**

Not synchronizing counters and statistics between RPs may create problems for external network management systems that monitor this information.

SSO Operation

- [SSO Conditions, page 7-8](#)
- [Switchover Time, page 7-8](#)
- [Online Removal of the Active RP, page 7-9](#)
- [Fast Software Upgrade, page 7-9](#)
- [Core Dump Operation, page 7-9](#)

SSO Conditions

An automatic or manual switchover may occur under the following conditions:

- A fault condition that causes the active RP to crash or reboot—automatic switchover
- The active RP is declared dead (not responding)—automatic switchover
- The CLI is invoked—manual switchover

The user can force the switchover from the active RP to the standby RP by using a CLI command. This manual procedure allows for a “graceful” or controlled shutdown of the active RP and switchover to the standby RP. This graceful shutdown allows critical cleanup to occur.

**Note**

This procedure should not be confused with the graceful shutdown procedure for routing protocols in core routers—they are separate mechanisms.

**Caution**

The SSO feature introduces a number of new command and command changes, including commands to manually cause a switchover. The **reload** command does not cause a switchover. The **reload** command causes a full reload of the box, removing all table entries, resetting all line cards, and interrupting nonstop forwarding.

Switchover Time

The time required by the device to switch over from the active RP to the standby RP is between zero and three seconds.

Although the newly active processor takes over almost immediately following a switchover, the time required for the device to begin operating again in full redundancy (SSO) mode can be several minutes, depending on the platform. The length of time can be due to a number of factors including the time needed for the previously active processor to obtain crash information, load code and microcode, and synchronize configurations between processors.

On DFC-equipped switching modules, forwarding information is distributed, and packets forwarded from the same line card should have little to no forwarding delay; however, forwarding packets between line cards requires interaction with the RP, meaning that packet forwarding might have to wait for the switchover time.

Online Removal of the Active RP

Online removal of the active RP automatically forces a stateful switchover to the standby RP.

Fast Software Upgrade

You can use Fast Software Upgrade (FSU) to reduce planned downtime. With FSU, you can configure the system to switch over to a standby RP that is preloaded with an upgraded Cisco IOS software image. FSU reduces outage time during a software upgrade by transferring functions to the standby RP that has the upgraded Cisco IOS software preinstalled. You can also use FSU to downgrade a system to an older version of Cisco OS or have a backup system loaded for downgrading to a previous image immediately after an upgrade.

SSO must be configured on the networking device before performing FSU.

**Note**

During the upgrade process, different images will be loaded on the RPs for a short period of time. During this time, the device will operate in RPR mode.

Core Dump Operation

In networking devices that support SSO, the newly active primary processor runs the core dump operation after the switchover has taken place. Not having to wait for dump operations effectively decreases the switchover time between processors.

Following the switchover, the newly active RP will wait for a period of time for the core dump to complete before attempting to reload the formerly active RP. The time period is configurable. For example, on some platforms an hour or more may be required for the formerly active RP to perform a coredump, and it might not be site policy to wait that much time before resetting and reloading the formerly active RP. In the event that the core dump does not complete within the time period provided, the standby is reset and reloaded regardless of whether it is still performing a core dump.

The core dump process adds the slot number to the core dump file to identify which processor generated the file content.

**Note**

Core dumps are generally useful only to your technical support representative. The core dump file, which is a very large binary file, must be transferred using the TFTP, FTP, or remote copy protocol (rcp) server and subsequently interpreted by a Cisco Technical Assistance Center (TAC) representative that has access to source code and detailed memory maps.

SSO-Aware Features

A feature is SSO-aware if it maintains, either partially or completely, undisturbed operation through an RP switchover. State information for SSO-aware features is synchronized from active to standby to achieve stateful switchover for those features.

The dynamically created state of SSO-unaware features is lost on switchover and must be reinitialized and restarted on switchover.

The output of the **show redundancy clients** command displays the SSO-aware features (see the [“Verifying SSO Features” section on page 7-13](#)).

Default Settings for SSO

None.

How to Configure SSO



Note

See [Chapter 6, “Fast Software Upgrade,”](#) for information about how to copy images onto the switch. During the upgrade process, different images will be loaded on the RPs for a very short period of time. If a switchover occurs during this time, the device will recover in RPR mode.

Either the SSO or RPR redundancy mode is always configured. The SSO redundancy mode is configured by default. To revert to the default SSO redundancy mode from the RPR redundancy mode, perform this task:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# redundancy	Enters redundancy configuration mode.
Step 4	Router(config)# mode sso	Sets the redundancy configuration mode to SSO on both the active and standby RP. Note After configuring SSO mode, the standby RP will automatically reset.
Step 5	Router(config-red)# end	Exits redundancy configuration mode and returns the switch to privileged EXEC mode.
Step 6	Router# copy running-config startup-config	Saves the configuration changes to the startup configuration file.

This example shows how to configure the SSO redundancy mode:

```
Router> enable
Router# configure terminal
Router(config)# redundancy
Router(config)# mode sso
```

```
Router(config-red)# end
Router# copy running-config startup-config
Router#
```

Troubleshooting SSO

- [Possible SSO Problem Situations, page 7-11](#)
- [SSO Troubleshooting, page 7-12](#)

Possible SSO Problem Situations

- The standby RP was reset, but there are no messages describing what happened—To display a log of SSO events and clues as to why a switchover or other event occurred, enter the **show redundancy history** command on the newly active RP:

```
Router# show redundancy history
```

- The **show redundancy states** command shows an operating mode that is different than what is configured on the networking device—On certain platforms the output of the **show redundancy states** command displays the actual operating redundancy mode running on the device, and not the configured mode as set by the platform. The operating mode of the system can change depending on system events. For example, SSO requires that both RPs on the networking device be running the same software image; if the images are different, the device will not operate in SSO mode, regardless of its configuration.

For example, during the upgrade process different images will be loaded on the RPs for a short period of time. If a switchover occurs during this time, the device will recover in RPR mode.

- Reloading the device disrupts SSO operation—The SSO feature introduces a number of commands, including commands to manually cause a switchover. The reload command is not an SSO command. This command causes a full reload of the box, removing all table entries, resetting all line cards, and thereby interrupting network traffic forwarding. To avoid reloading the box unintentionally, use the **redundancy force-switchover** command.
- During a software upgrade, the networking device appears to be in a mode other than SSO—During the software upgrade process, the show redundancy command indicates that the device is running in a mode other than SSO.

This is normal behavior. Until the FSU procedure is complete, each RP will be running a different software version. While the RPs are running different software versions, the mode will change to either RPR. The device will change to SSO mode once the upgrade has completed.

- The previously active processor is being reset and reloaded before the core dump completes—Use the **crashdump-timeout** command to set the maximum time that the newly active processor waits before resetting and reloading the previously active processor.
- Issuing a “send break” does not cause a system switchover—This is normal operation. Using “send break” to break or pause the system is not recommended and may cause unpredictable results. To initiate a manual switchover, use the **redundancy force-switchover** command.

In Cisco IOS software, you can enter ROM monitor mode by restarting the switch and then pressing the Break key or issuing a “send break” command from a telnet session during the first 60 seconds of startup. The send break function can be useful for experienced users or for users under the direction of a Cisco Technical Assistance Center (TAC) representative to recover from certain system problems or to evaluate the cause of system problems.

SSO Troubleshooting

The following commands may be used as needed to troubleshoot the SSO feature. These commands do not have to be entered in any particular order.

Command	Purpose
Router(config-red)# crashdump-timeout [<i>mm</i> <i>hh:mm</i>]	Sets the longest time that the newly active RP will wait before reloading the formerly active RP.
Router# debug redundancy { <i>all</i> <i>ui</i> <i>clk</i> <i>hub</i> }	Debugs redundancy on the networking device.
Router# show diag [<i>slot-number</i> <i>chassis</i> <i>subslot slot/subslot</i>] [<i>details</i> <i>summary</i>]	Displays hardware information.
Router# show redundancy [<i>clients</i> <i>counters</i> <i>debug-log</i> <i>handover</i> <i>history</i> <i>switchover history</i> <i>states</i> <i>inter-device</i>]	Displays the redundancy configuration mode of the RP. Also displays information about the number of switchovers, system uptime, processor uptime, and redundancy state, and reasons for any switchovers.
Router# show version	Displays image information for each RP.

Verifying the SSO Configuration

- [Verifying that SSO Is Configured](#)
- [Verifying that SSO Is Operating on the Device](#)
- [Verifying SSO Features](#)

Verifying that SSO Is Configured

In the following example, the **show redundancy** command is used to verify that SSO is configured on the device.

```
Router> enable
Router# show redundancy
Redundant System Information :
-----
    Available system uptime = 3 days, 4 hours, 35 minutes
    Switchovers system experienced = 0
        Standby failures = 1
    Last switchover reason = none

    Hardware Mode = Duplex
    Configured Redundancy Mode = sso
    Operating Redundancy Mode = sso
    Maintenance Mode = Disabled
    Communications = Up

Current Processor Information :
-----
    Active Location = slot 5
    Current Software state = ACTIVE
    Uptime in current state = 3 days, 4 hours, 35 minutes
    Image Version = Cisco IOS Software, s2t54 Software ...

Synced to ...
Copyright (c) 1986-2011 by Cisco Systems, Inc.
```

```

Compiled ...
                                BOOT = disk0:0726_c4,12
                                CONFIG_FILE =
                                BOOTLDR =
                                Configuration register = 0x2102

Peer Processor Information :
-----
                                Standby Location = slot 6
                                Current Software state = STANDBY HOT
                                Uptime in current state = 3 hours, 55 minutes
                                Image Version = Cisco IOS Software, s2t54 Software ...

Synced to ...
Copyright (c) 1986-2011 by Cisco Systems, Inc.
Compiled ...
                                BOOT = disk0:0726_c4,12
                                CONFIG_FILE =
                                BOOTLDR =
                                Configuration register = 0x2102

Router#

```

Verifying that SSO Is Operating on the Device

In the following example, the **show redundancy** command with the **states** keyword is used to verify that SSO is configured on the device.

```

Router# show redundancy states
my state = 13 -ACTIVE
  peer state = 8 -STANDBY HOT
    Mode = Duplex
    Unit = Primary
    Unit ID = 5

Redundancy Mode (Operational) = sso
Redundancy Mode (Configured) = sso
Redundancy State = sso
  Maintenance Mode = Disabled
  Manual Swact = enabled
  Communications = Up

  client count = 135
  client_notification_TMR = 30000 milliseconds
  keep_alive TMR = 9000 milliseconds
  keep_alive count = 1
  keep_alive threshold = 18
  RF debug mask = 0x0

Router#

```

Verifying SSO Features

Enter the **show redundancy clients** command to display the list of features that have registered as SSO features.

```

Router# show redundancy clients
clientID = 0      clientSeq = 0      RF_INTERNAL_MSG
clientID = 1319   clientSeq = 1      Cat6k Platform First
clientID = 29     clientSeq = 60     Redundancy Mode RF
clientID = 139    clientSeq = 61     IfIndex

```

clientID = 3300	clientSeq = 62	Persistent Variable
clientID = 25	clientSeq = 68	CHKPT RF
clientID = 1515	clientSeq = 69	HAL RF
clientID = 3100	clientSeq = 73	MCM
clientID = 77	clientSeq = 80	Event Manager
clientID = 1328	clientSeq = 81	Cat6k Asic API RF Cl
clientID = 1334	clientSeq = 82	Cat6k AUTOSHUT RF Cl
clientID = 1333	clientSeq = 83	Cat6k OVERSUB RF Cli
clientID = 1302	clientSeq = 84	Cat6k Fabric Manager
clientID = 1331	clientSeq = 86	Cat6k Inline Power
clientID = 1303	clientSeq = 88	Cat6k OIR
clientID = 518	clientSeq = 89	PM Port Data
clientID = 1306	clientSeq = 93	Cat6k QoS Manager
clientID = 1501	clientSeq = 98	Cat6k CWAN HA
clientID = 1503	clientSeq = 99	CWAN VLAN RF Client
clientID = 1310	clientSeq = 100	Cat6k Feature Manage
clientID = 1700	clientSeq = 101	Cat6k L3 Lif
clientID = 78	clientSeq = 102	TSPTUN HA
clientID = 305	clientSeq = 103	Multicast ISSU Conso
clientID = 304	clientSeq = 104	IP multicast RF Clie
clientID = 22	clientSeq = 105	Network RF Client
clientID = 88	clientSeq = 106	HSRP
clientID = 114	clientSeq = 107	GLBP
clientID = 225	clientSeq = 108	VRRP
clientID = 1505	clientSeq = 111	Cat6k SPA TSM
clientID = 1509	clientSeq = 114	Cat6k Online Diag HA
clientID = 1337	clientSeq = 116	Cat6k MPLS RF Client
clientID = 75	clientSeq = 120	Tableid HA
clientID = 1338	clientSeq = 124	Cat6k CTS Manager
clientID = 512	clientSeq = 126	LAN-Switch BD Manage
clientID = 501	clientSeq = 127	LAN-Switch VTP VLAN
clientID = 513	clientSeq = 128	LAN-Switch IDBHAL
clientID = 71	clientSeq = 129	XDR RRP RF Client
clientID = 24	clientSeq = 130	CEF RRP RF Client
clientID = 146	clientSeq = 132	BFD RF Client
clientID = 301	clientSeq = 135	MRIB RP RF Client
clientID = 306	clientSeq = 139	MFIB RRP RF Client
clientID = 1504	clientSeq = 146	Cat6k CWAN Interface
clientID = 1507	clientSeq = 147	CWAN LTL Mgr HA RF C
clientID = 520	clientSeq = 151	RFS RF
clientID = 210	clientSeq = 152	Auth Mgr
clientID = 5	clientSeq = 153	Config Sync RF clien
clientID = 138	clientSeq = 155	MDR SM
clientID = 1308	clientSeq = 156	Cat6k Local Target L
clientID = 1351	clientSeq = 157	RF VS Client
clientID = 1358	clientSeq = 158	Cat6k VSlot
clientID = 502	clientSeq = 162	LAN-Switch Port Mana
clientID = 514	clientSeq = 163	SWITCH_VLAN_HA
clientID = 1313	clientSeq = 165	Cat6k Platform
clientID = 1318	clientSeq = 166	Cat6k Power
clientID = 23	clientSeq = 171	Frame Relay
clientID = 49	clientSeq = 172	HDLIC
clientID = 72	clientSeq = 173	LSD HA Proc
clientID = 113	clientSeq = 174	MFI STATIC HA Proc
clientID = 1335	clientSeq = 180	C6K EFP RF client
clientID = 200	clientSeq = 181	ETHERNET OAM RF
clientID = 207	clientSeq = 183	ECFM RF
clientID = 202	clientSeq = 184	ETHERNET LMI RF
clientID = 208	clientSeq = 186	LLDP
clientID = 20	clientSeq = 193	IPROUTING NSF RF cli
clientID = 21	clientSeq = 197	PPP RF
clientID = 1352	clientSeq = 201	C6K_provision_rf_cli
clientID = 1307	clientSeq = 202	Cat6k IDPROM
clientID = 74	clientSeq = 206	MPLS VPN HA Client

clientID = 34	clientSeq = 208	SNMP RF Client
clientID = 1502	clientSeq = 209	CWAN APS HA RF Clie
clientID = 52	clientSeq = 210	ATM
clientID = 35	clientSeq = 219	History RF Client
clientID = 90	clientSeq = 231	RSVP HA Services
clientID = 250	clientSeq = 243	EEM Server RF CLIENT
clientID = 252	clientSeq = 245	EEM POLICY-DIR RF CL
clientID = 54	clientSeq = 247	SNMP HA RF Client
clientID = 73	clientSeq = 248	LDP HA
clientID = 76	clientSeq = 249	IPRM
clientID = 57	clientSeq = 250	ARP
clientID = 50	clientSeq = 257	FH_RF_Event_Detector
clientID = 1508	clientSeq = 263	CWAN LTL SP RF Clie
clientID = 1304	clientSeq = 267	Cat6k Ehc
clientID = 1305	clientSeq = 271	Cat6k PAGP/LACP
clientID = 503	clientSeq = 272	Spanning-Tree Protoc
clientID = 1309	clientSeq = 273	CMRP RF Client
clientID = 1311	clientSeq = 275	Cat6k L3 Manager
clientID = 1317	clientSeq = 276	Cat6k CAPI
clientID = 1506	clientSeq = 277	CWAN SRP RF Client
clientID = 83	clientSeq = 284	AC RF Client
clientID = 145	clientSeq = 285	VFI Mgr
clientID = 84	clientSeq = 286	AToM manager
clientID = 85	clientSeq = 287	SSM
clientID = 87	clientSeq = 291	SLB RF Client
clientID = 504	clientSeq = 294	Switch SPAN client
clientID = 507	clientSeq = 295	Switch Backup Interf
clientID = 105	clientSeq = 298	DHCP Snooping
clientID = 1510	clientSeq = 304	Call-Home RF
clientID = 203	clientSeq = 307	MVRP RF
clientID = 151	clientSeq = 310	IP Tunnel RF
clientID = 94	clientSeq = 311	Config Verify RF cli
clientID = 516	clientSeq = 314	EnergyWise rf client
clientID = 508	clientSeq = 316	Port Security Client
clientID = 509	clientSeq = 317	LAN-Switch IP Host T
clientID = 515	clientSeq = 318	SISF table
clientID = 135	clientSeq = 322	IKE RF Client
clientID = 136	clientSeq = 323	IPSEC RF Client
clientID = 130	clientSeq = 324	CRYPTO RSA
clientID = 400	clientSeq = 326	IP Admission RF Clie
clientID = 3099	clientSeq = 335	ISSU process
clientID = 4005	clientSeq = 338	ISSU Test Client
clientID = 93	clientSeq = 342	Network RF 2 Client
clientID = 1320	clientSeq = 343	Cat6k PF_ML_RP
clientID = 510	clientSeq = 345	LAN-Switch PAGP/LACP
clientID = 511	clientSeq = 346	LAN-Switch Private V
clientID = 1321	clientSeq = 347	PM SP client
clientID = 1322	clientSeq = 348	VLAN Mapping
clientID = 1315	clientSeq = 350	Cat6k Clear Counter
clientID = 141	clientSeq = 352	DATA DESCRIPTOR RF C
clientID = 1000	clientSeq = 361	CTS HA
clientID = 1001	clientSeq = 362	Keystore
clientID = 3150	clientSeq = 363	SIA SD RF CLIENT
clientID = 3151	clientSeq = 364	SIA SB RF CLIENT
clientID = 3152	clientSeq = 365	SIA SCL RF CLIENT
clientID = 3153	clientSeq = 366	SIA SVE RF CLIENT
clientID = 3154	clientSeq = 367	SIA TCP RF CLIENT
clientID = 1332	clientSeq = 373	PCLC
clientID = 1367	clientSeq = 375	Cat6k ITASCA_RP
clientID = 4032	clientSeq = 379	ACL handle RF Client
clientID = 4020	clientSeq = 381	IOS Config ARCHIVE
clientID = 4021	clientSeq = 382	IOS Config ROLLBACK
clientID = 1339	clientSeq = 404	Cat6k blue beacon RF
clientID = 1362	clientSeq = 405	VS HA

```

clientID = 517      clientSeq = 406      LAN-Switch IDBHAL2
clientID = 1336     clientSeq = 415      Cat6k NTI SUP SI swi
clientID = 65000    clientSeq = 416      RF_LAST_CLIENT

```

Configuration Examples for SSO

This example configures the SSO redundancy mode :

```

Router# configure terminal
Router(config)# redundancy
Router(config-red)# mode sso
Router(config-red)# exit
Router# copy running-config startup-config

```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

8

Nonstop Forwarding (NSF)

- [Prerequisites for NSF, page 8-1](#)
- [Restrictions for NSF, page 8-1](#)
- [Information About NSF, page 8-2](#)
- [Default Settings for NSF, page 8-8](#)
- [How to Configure NSF, page 8-8](#)
- [Configuration Examples for NSF, page 8-14](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- Stateful switchover (SSO) and nonstop forwarding (NSF) do not support IPv6 multicast traffic.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for NSF

None.

Restrictions for NSF

- [General Restrictions, page 8-2](#)
- [Restrictions for BGP NSF, page 8-2](#)
- [Restrictions for EIGRP NSF, page 8-2](#)
- [Restrictions for OSPF NSF, page 8-2](#)

- [Restrictions for IS-IS NSF, page 8-2](#)
- [Restrictions for IPv6 NSF, page 8-2](#)

General Restrictions

- NSF requires SSO (see [Chapter 7, “Stateful Switchover \(SSO\)”](#)).
- The Hot Standby Routing Protocol (HSRP) is not supported with Cisco Nonstop Forwarding with Stateful Switchover. Do not use HSRP with Cisco Nonstop Forwarding with Stateful Switchover.

Restrictions for BGP NSF

- All neighboring devices participating in BGP NSF must be NSF-capable, having been configured for BGP graceful restart as described in the [“Configuring and Verifying BGP for NSF” section on page 8-8](#).

Restrictions for EIGRP NSF

- All neighboring devices participating in EIGRP NSF operation must be NSF-capable or NSF-aware.
- An NSF-aware router cannot support two NSF-capable peers performing an NSF restart operation at the same time. However, both neighbors will reestablish peering sessions after the NSF restart operation is complete.

Restrictions for OSPF NSF

- OSPF NSF for virtual links is not supported.
- All OSPF networking devices on the same network segment must be NSF-aware (that is, running an NSF software image).
- OSPF NSF for sham links is not supported.

Restrictions for IS-IS NSF

- For IETF IS-IS, all neighboring devices must be running an NSF-aware software image.

Restrictions for IPv6 NSF

- IPv6 must be enabled on your router for IPv6 NSF to be supported.

Information About NSF

- [NSF Overview, page 8-3](#)
- [Feature Interaction with NSF, page 8-3](#)

NSF Overview

NSF works with SSO to minimize the amount of time a network is unavailable to its users following a switchover. The main objective of Cisco NSF is to continue forwarding IP packets following a route processor (RP) switchover.

Usually, when a networking device restarts, all routing peers of that device detect that the device went down and then came back up. This transition results in what is called a routing flap, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. Cisco NSF helps to suppress routing flaps in SSO-enabled devices, thus reducing network instability.

Cisco NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. With Cisco NSF, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards while the standby RP assumes control from the failed active RP during a switchover. The ability of line cards to remain up through a switchover and to be kept current with the Forwarding Information Base (FIB) on the active RP is key to Cisco NSF operation.

The Cisco NSF feature has several benefits, including the following:

- Improved network availability—NSF continues forwarding network traffic and application state information so that user session information is maintained after a switchover.
- Overall network stability—Network stability may be improved with the reduction in the number of route flaps that had been created when routers in the network failed and lost their routing tables.
- Neighboring routers do not detect link flapping—Because the interfaces remain up across a switchover, neighboring routers do not detect a link flap (that is, the link does not go down and come back up).
- Prevents routing flaps—Because SSO continues forwarding network traffic in the event of a switchover, routing flaps are avoided.
- No loss of user sessions—User sessions established prior to the switchover are maintained.

A networking device is NSF-aware if it is running NSF-compatible software. A device is NSF-capable if it has been configured to support NSF and would rebuild routing information from NSF-aware or NSF-capable neighbors.

CEF is always enabled on the switch and cannot be disabled. The routing protocols depend on CEF to continue forwarding packets during switchover while the routing protocols rebuild the Routing Information Base (RIB) tables. Once the routing protocols have converged, CEF updates the FIB table and removes stale route entries and CEF updates the line cards with the new FIB information.

Feature Interaction with NSF

- [Cisco Express Forwarding, page 8-4](#)
- [Routing Protocol Operation, page 8-4](#)
- [BGP Operation, page 8-4](#)
- [EIGRP Operation, page 8-5](#)
- [IS-IS Operation, page 8-6](#)
- [OSPF Operation, page 8-7](#)
- [IPv6 Routing Protocol Operation, page 8-7](#)

Cisco Express Forwarding

A key element of NSF is packet forwarding. In a Cisco networking device, packet forwarding is provided by CEF. CEF is always enabled on the switch and cannot be disabled. CEF maintains the FIB, and uses the FIB information that was current at the time of the switchover to continue forwarding packets during a switchover. This feature reduces traffic interruption during the switchover.

During normal NSF operation, CEF on the active RP synchronizes its current FIB and adjacency databases with the FIB and adjacency databases on the standby RP. Upon switchover of the active RP, the standby RP initially has FIB and adjacency databases that are mirror images of those that were current on the active RP. For platforms with intelligent line cards, the line cards will maintain the current forwarding information over a switchover; for platforms with forwarding engines, CEF will keep the forwarding engine on the standby RP current with changes that are sent to it by CEF on the active RP. In this way, the line cards or forwarding engines will be able to continue forwarding after a switchover as soon as the interfaces and a data path are available.

As the routing protocols start to repopulate the RIB on a prefix-by-prefix basis, the updates in turn cause prefix-by-prefix updates to CEF, which it uses to update the FIB and adjacency databases. Existing and new entries will receive the new version (“epoch”) number, indicating that they have been refreshed. The forwarding information is updated on the line cards or forwarding engine during convergence. The RP signals when the RIB has converged. The software removes all FIB and adjacency entries that have an epoch older than the current switchover epoch. The FIB now represents the newest routing protocol forwarding information.

Routing Protocol Operation

The routing protocols run only on the active RP, and they receive routing updates from their neighbor routers. Routing protocols do not run on the standby RP. Following a switchover, the routing protocols request that the NSF-aware neighbor devices send state information to help rebuild the routing tables. Alternately, the IS-IS protocol can be configured to synchronize state information from the active to the standby RP to help rebuild the routing table on the NSF-capable device in environments where neighbor devices are not NSF-aware.

For NSF operation, the routing protocols depend on CEF to continue forwarding packets while the routing protocols rebuild the routing information.

BGP Operation

When a NSF-capable router begins a BGP session with a BGP peer, it sends an OPEN message to the peer. Included in the message is a declaration that the NSF-capable device has “graceful restart capability.” Graceful restart is the mechanism by which BGP routing peers avoid a routing flap following a switchover. If the BGP peer has received this capability, it is aware that the device sending the message is NSF-capable. Both the NSF-capable router and its BGP peer(s) need to exchange the graceful restart capability in their OPEN messages, at the time of session establishment. If both the peers do not exchange the graceful restart capability, the session will not be graceful restart capable.

If the BGP session is lost during the RP switchover, the NSF-aware BGP peer marks all the routes associated with the NSF-capable router as stale; however, it continues to use these routes to make forwarding decisions for a set period of time. This functionality means that no packets are lost while the newly active RP is waiting for convergence of the routing information with the BGP peers.

After an RP switchover occurs, the NSF-capable router reestablishes the session with the BGP peer. In establishing the new session, it sends a new graceful restart message that identifies the NSF-capable router as having restarted.

At this point, the routing information is exchanged between the two BGP peers. Once this exchange is complete, the NSF-capable device uses the routing information to update the RIB and the FIB with the new forwarding information. The NSF-aware device uses the network information to remove stale routes from its BGP table. Following that, the BGP protocol is fully converged.

If a BGP peer does not support the graceful restart capability, it will ignore the graceful-restart capability in an OPEN message but will establish a BGP session with the NSF-capable device. This function will allow interoperability with non-NSF-aware BGP peers (and without NSF functionality), but the BGP session with non-NSF-aware BGP peers will not be graceful restart capable.

**Note**

BGP support in NSF requires that neighbor networking devices be NSF-aware; that is, the devices must have the graceful restart capability and advertise that capability in their OPEN message during session establishment. If an NSF-capable router discovers that a particular BGP neighbor does not have graceful restart capability, it will not establish an NSF-capable session with that neighbor. All other neighbors that have graceful restart capability will continue to have NSF-capable sessions with this NSF-capable networking device.

EIGRP Operation

EIGRP NSF capabilities are exchanged by EIGRP peers in hello packets. The NSF-capable router notifies its neighbors that an NSF restart operation has started by setting the restart (RS) bit in a hello packet. When an NSF-aware router receives notification from an NSF-capable neighbor that an NSF-restart operation is in progress, the NSF-capable and NSF-aware routers immediately exchange their topology tables. The NSF-aware router sends an end-of-table (EOT) update packet when the transmission of its topology table is complete. The NSF-aware router then performs the following actions to assist the NSF-capable router:

- The EIGRP hello hold timer is expired to reduce the time interval set for hello packet generation and transmission. This allows the NSF-aware router to reply to the NSF-capable router more quickly reducing the amount of time required for the NSF-capable router to rediscover neighbors and rebuild the topology table.
- The route-hold timer is started. This timer is used to set the period of time that the NSF-aware router will hold known routes for the NSF-capable neighbor. This timer is configured with the **timers nsf route-hold** command. The default time period is 240 seconds.
- The NSF-aware router notes in the peer list that the NSF-capable neighbor is restarting, maintains adjacency, and holds known routes for the NSF-capable neighbor until the neighbor signals that it is ready for the NSF-aware router to send its topology table or the route-hold timer expires. If the route-hold timer expires on the NSF-aware router, the NSF-aware router will discard held routes and treat the NSF-capable router as a new router joining the network and reestablishing adjacency accordingly.
- The NSF-aware router will continue to send queries to the NSF-capable router which is still in the process of converging after switchover, effectively extending the time before a stuck-in-active (SIA) condition can occur.

When the switchover operation is complete, the NSF-capable router notifies its neighbors that it has reconverged and has received all of their topology tables by sending an EOT update packet to the assisting routers. The NSF-capable then returns to normal operation. The NSF-aware router will look for alternate paths (go active) for any routes that are not refreshed by the NSF-capable (restarting router). The NSF-aware router will then return to normal operation. If all paths are refreshed by the NSF-capable router, the NSF-aware router will immediately return to normal operation.

**Note**

NSF-aware routers are completely compatible with non-NSF aware or capable neighbors in an EIGRP network. A non-NSF aware neighbor will ignore NSF capabilities and reset adjacencies and otherwise maintain the peering sessions normally.

IS-IS Operation

The IS-IS protocol can be configured to use state information that has been synchronized between the active and the standby RP to recover route information following a switchover instead of information received from peer devices.

When an IS-IS NSF-capable router performs an RP switchover, it must perform two tasks in order to resynchronize its Link State Database with its IS-IS neighbors. First, it must relearn the available IS-IS neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the Link State Database for the network.

The IS-IS NSF feature offers two options when configuring NSF:

- Internet Engineering Task Force (IETF) IS-IS
- Cisco IS-IS

If neighbor routers on a network segment are NSF-aware, meaning that neighbor routers are running a software version that supports the IETF Internet draft for router restartability, they will assist an IETF NSF router which is restarting. With IETF, neighbor routers provide adjacency and link-state information to help rebuild the routing information following a switchover. A benefit of IETF IS-IS configuration is operation between peer devices based on a proposed standard.

**Note**

If you configure IETF on the networking device, but neighbor routers are not IETF-compatible, NSF will abort following a switchover.

If the neighbor routers on a network segment are not NSF-aware, you must use the Cisco configuration option. The Cisco IS-IS configuration transfers both protocol adjacency and link-state information from the active to the standby RP. A benefit of Cisco configuration is that it does not rely on NSF-aware neighbors.

IETF IS-IS Configuration

Using the IETF IS-IS configuration, as quickly as possible after an RP switchover, the NSF-capable router sends IS-IS NSF restart requests to neighboring NSF-aware devices. Neighbor networking devices recognize this restart request as a cue that the neighbor relationship with this router should not be reset, but that they should initiate database resynchronization with the restarting router. As the restarting router receives restart request responses from routers on the network, it can begin to rebuild its neighbor list.

Once this exchange is complete, the NSF-capable device uses the link-state information to remove stale routes, update the RIB, and update the FIB with the new forwarding information. IS-IS is then fully converged.

The switchover from one RP to the other happens within seconds. IS-IS reestablishes its routing table and resynchronizes with the network within a few additional seconds. At this point, IS-IS waits for a specified interval before it will attempt a second NSF restart. During this time, the new standby RP will boot up and synchronize its configuration with the active RP. The IS-IS NSF operation waits for a specified interval to ensure that connections are stable before attempting another restart of IS-IS NSF. This functionality prevents IS-IS from attempting back-to-back NSF restarts with stale information.

Cisco IS-IS Configuration

Using the Cisco configuration option, full adjacency and LSP information is saved, or “checkpointed,” to the standby RP. Following a switchover, the newly active RP maintains its adjacencies using the checkpointed data, and can quickly rebuild its routing tables.

**Note**

Following a switchover, Cisco IS-IS NSF has complete neighbor adjacency and LSP information; however, it must wait for all interfaces that had adjacencies prior to the switchover to come up. If an interface does not come up within the allocated interface wait time, the routes learned from these neighbor devices are not considered in routing table recalculation. IS-IS NSF provides a command to extend the wait time for interfaces that, for whatever reason, do not come up in a timely fashion.

The switchover from one RP to the other happens within seconds. IS-IS reestablishes its routing table and resynchronizes with the network within a few additional seconds. At this point, IS-IS waits for a specified interval before it will attempt a second NSF restart. During this time, the new standby RP will boot up and synchronize its configuration with the active RP. Once this synchronization is completed, IS-IS adjacency and LSP data is checkpointed to the standby RP; however, a new NSF restart will not be attempted by IS-IS until the interval time expires. This functionality prevents IS-IS from attempting back-to-back NSF restarts.

OSPF Operation

When an OSPF NSF-capable router performs an RP switchover, it must perform two tasks in order to resynchronize its Link State Database with its OSPF neighbors. First, it must relearn the available OSPF neighbors on the network without causing a reset of the neighbor relationship. Second, it must re-acquire the contents of the Link State Database for the network.

As quickly as possible after an RP switchover, the NSF-capable router sends an OSPF NSF signal to neighboring NSF-aware devices. Neighbor networking devices recognize this signal as a cue that the neighbor relationship with this router should not be reset. As the NSF-capable router receives signals from other routers on the network, it can begin to rebuild its neighbor list.

Once neighbor relationships are reestablished, the NSF-capable router begins to resynchronize its database with all of its NSF-aware neighbors. At this point, the routing information is exchanged between the OSPF neighbors. Once this exchange is complete, the NSF-capable device uses the routing information to remove stale routes, update the RIB, and update the FIB with the new forwarding information. The OSPF protocols are then fully converged.

**Note**

OSPF NSF requires that all neighbor networking devices be NSF-aware. If an NSF-capable router discovers that it has non-NSF-aware neighbors on a particular network segment, it will disable NSF capabilities for that segment. Other network segments composed entirely of NSF-capable or NSF-aware routers will continue to provide NSF capabilities.

The OSPF RFC 3623 Graceful Restart feature allows you to configure IETF NSF in multivendor networks. For more information, see the [OSPF RFC 3623 Graceful Restart](#) document.

IPv6 Routing Protocol Operation

IPv6 support for NSF includes the following features:

- [Nonstop Forwarding and Graceful Restart for MP-BGP IPv6 Address Family](#), page 8-8
- [Nonstop Forwarding for IPv6 RIP](#), page 8-8

- [Nonstop Forwarding for IPv6 Static Routes, page 8-8](#)

Nonstop Forwarding and Graceful Restart for MP-BGP IPv6 Address Family

The switch supports the graceful restart capability for IPv6 BGP unicast and VPNv6 address families, enabling Cisco NSF functionality for BGP IPv6. The BGP graceful restart capability allows the BGP routing table to be recovered from peers without keeping the TCP state.

NSF continues forwarding packets while routing protocols converge, therefore avoiding a route flap on switchover. Forwarding is maintained by synchronizing the FIB between the active and standby RP. On switchover, forwarding is maintained using the FIB. The RIB is not kept synchronized; therefore, the RIB is empty on switchover. The RIB is repopulated by the routing protocols and subsequently informs FIB about RIB convergence by using the NSF_RIB_CONVERGED registry call. The FIB tables are updated from the RIB, removing any stale entries. The RIB starts a failsafe timer during RP switchover, in case the routing protocols fail to notify the RIB of convergence.

The Cisco BGP address family identifier (AFI) model is modular and scalable, and supports multiple AFIs and subsequent address family identifier (SAFI) configurations.

For information about how to configure the IPv6 BGP graceful restart capability, see the “[Implementing Multiprotocol BGP for IPv6](#)” document.

Nonstop Forwarding for IPv6 RIP

RIP registers as an IPv6 NSF client. Doing so has the benefit of using RIP routes installed in the Cisco Express Forwarding table until RIP has converged on the standby.

Nonstop Forwarding for IPv6 Static Routes

Cisco NSF supports IPv6 static routes.

Default Settings for NSF

None.

How to Configure NSF

- [Configuring and Verifying BGP for NSF, page 8-8](#) (optional)
- [Configuring and Verifying EIGRP NSF, page 8-10](#) (optional)
- [Configuring and Verifying OSPF NSF, page 8-11](#) (optional)
- [Configuring and Verifying IS-IS NSF, page 8-12](#) (optional)
- [Troubleshooting Cisco Nonstop Forwarding, page 8-14](#) (optional)

Configuring and Verifying BGP for NSF

- [Configuring BGP for NSF, page 8-9](#)
- [Verifying NSF for BGP, page 8-9](#)

Configuring BGP for NSF

Perform this task to configure BGP for NSF. Repeat this task on each BGP NSF peer device:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# router bgp <i>autonomous-system-number</i>	Enables a BGP routing process, and enters router configuration mode.
Step 4	Router(config-router)# bgp graceful-restart [restart-time <i>seconds</i> stalepath-time <i>seconds</i>]	Enables the BGP graceful restart capability, which starts NSF for BGP.

This example shows how to configure BGP for NSF:

```
Router> enable
Router# configure terminal
Router(config)# router bgp 120
Router(config-router)# bgp graceful-restart
```

Verifying NSF for BGP

Perform this task to verify that the graceful restart function is configured on the SSO-enabled networking device and on the neighbor devices:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# show running-config	Displays the contents of the current running configuration file. Verify that the phrase “bgp graceful-restart” appears in the BGP configuration of the SSO-enabled router. Repeat this step on each of the BGP neighbors.
Step 3	Router# show ip bgp neighbors [<i>ip-address</i> advertised-routes dampened-routes flap-statistics paths [<i>reg-exp</i>] received prefix-filter received-routes routes policy [<i>detail</i>]]]	Displays information about BGP and TCP connections to neighbors. On the SSO device and the neighbor device, this command verifies that the graceful restart function is shown as both advertised and received, and confirms the address families that have the graceful restart capability. If no address families are listed, then BGP NSF also will not occur.

This example shows how to NSF for BGP:

```
Router> enable
Router# configure terminal
Router# show running-config
Router# show ip bgp neighbors
```

Configuring and Verifying EIGRP NSF

- [Configuring EIGRP for NSF, page 8-10](#)
- [Verifying EIGRP for NSF, page 8-11](#)

Configuring EIGRP for NSF



Note

- An NSF-aware router must be completely converged with the network before it can assist an NSF-capable router in an NSF restart operation.
- Distributed platforms that run a supporting version of Cisco IOS software can support full NSF capabilities. These routers can perform a restart operation and can support other NSF capable peers.
- Single processor platforms that run a supporting version of Cisco IOS software support only NSF awareness. These routers maintain adjacency and hold known routes for the NSF-capable neighbor until it signals that it is ready for the NSF-aware router to send its topology table or the route-hold timer expires.

Perform this task to configure EIGRP for NSF. Repeat this procedure on each EIGRP NSF peer device:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# router eigrp <i>as-number</i>	Enables an EIGRP routing process, and enters router configuration mode.
Step 4	Router(config-router)# nsf [{ cisco ietf } interface wait <i>seconds</i> interval <i>minutes</i> t3 [adjacency manual <i>seconds</i>]	(Optional) Enables EIGRP NSF support on an NSF capable router. Enter this command on only NSF-capable routers. NSF awareness is enabled by default when a supporting version of Cisco IOS software is installed on a router that supports NSF capability or NSF awareness.
Step 5	Router(config-router)# timers nsf converge <i>seconds</i>	Adjusts the maximum time that restarting router will wait for the EOT notification from an NSF-capable or NSF-aware peer.
Step 6	Router(config-router)# timers nsf route-hold <i>seconds</i>	Sets the route-hold timer to determine how long an NSF-aware router that is running EIGRP will hold routes for an inactive peer.
Step 7	Router(config-router)# timers nsf signal <i>seconds</i>	Adjusts the maximum time for the initial restart period.

This example shows how to configure EIGRP for NSF:

```
Router> enable
Router# configure terminal
Router(config)# router eigrp 109
Router(config-router)# nsf
Router(config-router)# timers nsf converge 60
```

```
Router(config-router)# timers nsf route-hold 120
Router(config-router)# timers nsf signal seconds
```

Verifying EIGRP for NSF

Perform this task to verify that NSF awareness or capability or both are enabled on the SSO-enabled networking device and on the neighbor devices.

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# show ip protocols	Displays the parameters and current state of the active routing protocol process. Repeat this step on each of the EIGRP neighbors.

This example shows how to verify EIGRP for NSF:

```
Router> enable
Router# show ip protocols
```

Configuring and Verifying OSPF NSF

- [Configuring OSPF for NSF, page 8-11](#)
- [Verifying OSPF for NSF, page 8-12](#)

Configuring OSPF for NSF



Note

All peer devices participating in OSPF NSF must be made OSPF NSF aware; NSF awareness is enabled by default when a supporting version of Cisco IOS software is installed on a router that supports NSF capability or NSF awareness.

Perform this task to configure OSPF for NSF:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# router ospf process-id [vrf vpn-name]	Enables an OSPF routing process, and places the router in router configuration mode.
Step 4	Router(config-router)# nsf [{cisco ietf} interface wait seconds interval minutes t3 [adjacency manual seconds]	Enables EIGRP NSF support on an NSF capable router. • Enter this command on NSF-capable routers only.

This example shows how to configure OSPF for NSF:

```
Router> enable
Router# configure terminal
Router(config)# router ospf 12
Router(config-router)# nsf
```

Verifying OSPF for NSF

Perform this task to verify OSPF for NSF:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# show ip ospf [<i>process-id</i>]	Displays general information about OSPF routing processes.

This example shows how to verify OSPF for NSF:

```
Router> enable
Router# show ip ospf
```

Configuring and Verifying IS-IS NSF

- [Configuring NSF for IS-IS, page 8-12](#)
- [Verifying NSF for IS-IS, page 8-13](#)

Configuring NSF for IS-IS

Perform this task to configure NSF for IS-IS:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# router isis <i>area-tag</i>	Enables the IS-IS routing protocol to specify an IS-IS process, and places the router in router configuration mode.
Step 4	Router(config-router)# nsf [{ cisco ietf } interface <i>wait seconds</i> interval <i>minutes</i> t3 [adjacency manual seconds]	Enables NSF operation for IS-IS. • ietf—Enables IS-IS in homogeneous network where adjacencies with networking devices supporting IETF draft-based restartability is guaranteed. • cisco—Runs IS-IS in heterogeneous networks that might not have adjacencies with NSF-aware networking devices.

	Command	Purpose
Step 5	Router(config-router)# nsf interval <i>minutes</i>	Configures the minimum time between Cisco NSF restart attempts.
Step 6	Router(config-router)# nsf t3 { <i>manual seconds</i> adjacency }	Specifies the methodology used to determine how long IETF Cisco NSF will wait for the link-state packet (LSP) database to synchronize before generating overloaded link-state information for itself and flooding that information out to its neighbors.
Step 7	Router(config-router)# nsf interface wait <i>seconds</i>	Specifies how long a Cisco NSF restart will wait for all interfaces with IS-IS adjacencies to come up before completing the restart.

This example shows how to configure NSF for IS-IS:

```
Router> enable
Router# configure terminal
Router(config)# router isis cisco1
Router(config-router)# nsf ietf
Router(config-router)# nsf interval 2
Router(config-router)# nsf t3 manual 40
Router(config-router)# nsf interface wait 15
```

Verifying NSF for IS-IS

Perform this task to verify NSF for IS-IS:

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# show running-config	Displays the contents of the current running configuration file.
Step 3	Router# show isis nsf	Displays current state information regarding IS-IS NSF.

This example shows how to verify NSF for IS-IS:

```
Router> enable
Router# show running-config
Router# show isis nsf
```

Troubleshooting Cisco Nonstop Forwarding

To troubleshoot Cisco Nonstop Forwarding, use the following commands as needed:

Command	Purpose
Router# debug eigrp nsf	Displays notifications and information about NSF events for an EIGRP routing process.
Router# debug ip eigrp notifications	Displays information and notifications for an EIGRP routing process. This output includes NSF notifications and events.
Router# debug isis nsf [detail]	Displays information about the IS-IS state during a Cisco NSF restart.
Router# debug ospf nsf [detail]	Displays debugging messages related to OSPF Cisco NSF commands.
Router# show cef nsf	Displays the current NSF state of CEF on both the active and standby RPs.
Router# show cef state	Displays the CEF state on a networking device.
Router# show clns neighbors	Display both end-system and intermediate system neighbors.
Router# show ip bgp	Displays entries in the BGP routing table.
Router# show ip bgp neighbor	Displays information about the TCP and BGP connections to neighbor devices.\
Router# show ip cef	Displays entries in the FIB that are unresolved, or displays a FIB summary.
Router# show ip eigrp neighbors [interface-type as-number static detail]	To display detailed information about neighbors discovered by EIGRP.
Router# show ip ospf	Displays general information about OSPF routing processes.
Router# show ip ospf neighbor [detail]	Displays OSPF-neighbor information on a per-interface basis.
Router# show ip protocols	Displays the parameters and current state of the active routing protocol process. The status of EIGRP NSF configuration and support is displayed in the output.
Router# show isis database [detail]	Displays the IS-IS link-state database.
Router# show isis nsf	Displays the current state information regarding IS-IS Cisco NSF.

Configuration Examples for NSF

- [Example: Configuring BGP NSF, page 8-15](#)
- [Example: Configuring BGP NSF Neighbor Device, page 8-15](#)
- [Example: Verifying BGP NSF, page 8-15](#)
- [Example: Configuring EIGRP NSF Converge Timer, page 8-16](#)
- [Example: EIGRP Graceful-Restart Purge-Time Timer Configuration, page 8-16](#)
- [Example: Configuring EIGRP NSF Route-Hold Timer, page 8-16](#)
- [Example: Configuring EIGRP NSF Signal Timer, page 8-16](#)

- [Example: Disabling EIGRP NSF Support, page 8-17](#)
- [Example: Verifying EIGRP NSF, page 8-16](#)
- [Example: Configuring OSPF NSF, page 8-17](#)
- [Example: Verifying OSPF NSF, page 8-17](#)
- [Example: Configuring IS-IS NSF, page 8-18](#)
- [Example: Verifying IS-IS NSF, page 8-18](#)

Example: Configuring BGP NSF

The following example shows how to configure BGP NSF on a networking device.

```
Router# configure terminal
Router(config)# router bgp 590
Router(config-router)# bgp graceful-restart
```

Example: Configuring BGP NSF Neighbor Device

The following example shows how to configure BGP NSF on a neighbor router. All devices supporting BGP NSF must be NSF-aware, meaning that these devices recognize and advertise graceful restart capability.

```
Router# configure terminal
Router(config)# router bgp 770
Router(config-router)# bgp graceful-restart
```

Example: Verifying BGP NSF

Verify that “bgp graceful-restart” appears in the BGP configuration of the SSO-enabled router by entering the **show running-config** command.

```
Router# show running-config

router bgp 120
bgp graceful-restart
 neighbor 10.2.2.2 remote-as 300
```

On the SSO device and the neighbor device, verify that the graceful restart function is shown as both advertised and received, and confirm the address families that have the graceful restart capability. If no address families are listed, then BGP NSF also will not occur.

```
Router# show ip bgp neighbors x.x.x.x

BGP neighbor is 192.168.2.2, remote AS YY, external link
  BGP version 4, remote router ID 192.168.2.2
  BGP state = Established, up for 00:01:18
  Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh:advertised and received(new)
  Address family IPv4 Unicast:advertised and received
  Address family IPv4 Multicast:advertised and received
  Graceful Restart Capabilty:advertised and received
    Remote Restart timer is 120 seconds
  Address families preserved by peer:
    IPv4 Unicast, IPv4 Multicast
```

```
Received 1539 messages, 0 notifications, 0 in queue
Sent 1544 messages, 0 notifications, 0 in queue
Default minimum time between advertisement runs is 30 seconds
```

Example: Configuring EIGRP NSF Converge Timer

The **timers nsf converge** command is used to adjust the maximum time that a restarting router will wait for the EOT notification from an NSF-capable or NSF-aware peer. The following example shows how to set the converge timer to one minute.

```
Router# configure terminal
Router(config)# router eigrp 101
Router(config-router)# timers nsf converge 60
```

Example: EIGRP Graceful-Restart Purge-Time Timer Configuration

The **timers graceful-restart purge-time** command is used to set the route-hold timer that determines how long an NSF-aware router that is running EIGRP will hold routes for an inactive peer. The following example shows how to set the route-hold timer to two minutes:

```
Router(config-router)# timers graceful-restart purge-time 120
```

Example: Configuring EIGRP NSF Route-Hold Timer

The **timers nsf route-hold** command is used to set the maximum period of time that an NSF-aware router will hold known routes for an NSF-capable neighbor during a switchover operation. The following example shows how to set the route-hold timer to two minutes.

```
Router# configure terminal
Router(config)# router eigrp 101
Router(config-router)# timers nsf route-hold 120
```

Example: Configuring EIGRP NSF Signal Timer

The **timers nsf signal** command is used to adjust the maximum time for the initial restart period. The following example shows how to set the signal timer to 10 seconds.

```
Router# configure terminal
Router(config)# router eigrp 101
Router(config-router)# timers nsf signal 10
```

Example: Verifying EIGRP NSF

Verify that EIGRP NSF support is present in the installed Cisco IOS software image by entering the **show ip protocols** command. “EIGRP NSF-aware route hold timer is...” is displayed in the output when either NSF awareness or capability is supported. This line displays the default or user-defined value for the route-hold timer. “EIGRP NSF...” is displayed in the output only when the NSF capability is supported. This line will also print “disabled” or “enabled” depending on the status of the EIGRP NSF feature.

```
Router# show ip protocols
```

```

Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
  EIGRP NSF enabled
    NSF signal timer is 20s
    NSF converge timer is 120s
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.4.9.0/24
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170

```

Example: Disabling EIGRP NSF Support

EIGRP NSF capability is enabled by default on distributed platforms that run a supporting version of Cisco IOS software. The **nsf** command used to enable or disable the EIGRP NSF capability. The following example shows how to disable NSF capability:

```

Router# configure terminal
Router(config)# router eigrp 101
Router(config-router)# no nsf

```

Example: Configuring OSPF NSF

The following example shows how to configure OSPF NSF on a networking device:

```

Router# configure terminal
Router(config)# router ospf 400
Router(config-router)# nsf

```

Example: Verifying OSPF NSF

To verify NSF for OSPF, you must check that the NSF function is configured on the SSO-enabled networking device. Verify that “nsf” appears in the OSPF configuration of the SSO-enabled device by entering the **show running-config** command:

```

Router# show running-config

router ospf 120
log-adjacency-changes
nsf
network 192.168.20.0 0.0.0.255 area 0
network 192.168.30.0 0.0.0.255 area 1
network 192.168.40.0 0.0.0.255 area 2

```

Next, use the **show ip ospf** command to verify that NSF is enabled on the device.

```

Router> show ip ospf

```

```

Routing Process "ospf 1" with ID 192.168.2.1 and Domain ID 0.0.0.1
Supports only single TOS(TOS0) routes
Supports opaque LSA
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 0. Checksum Sum 0x0
Number of opaque AS LSA 0. Checksum Sum 0x0
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
External flood list length 0
Non-Stop Forwarding enabled, last NSF restart 00:02:06 ago (took 44 secs)
Area BACKBONE(0)
Number of interfaces in this area is 1 (0 loopback)
Area has no authentication
SPF algorithm executed 3 times

```

Example: Configuring IS-IS NSF

The following example shows how to configure Cisco proprietary IS-IS NSF operation on a networking device:

```

Router# configure terminal
Router(config)# router isis
Router(config-router)# nsf cisco

```

The following example shows how to configure IS-IS NSF for IETF operation on a networking device:

```

Router# configure terminal
Router(config)# router isis
Router(config-router)# nsf ietf

```

Example: Verifying IS-IS NSF

Verify that NSF appears in the IS-IS configuration of the SSO-enabled device by entering the **show running-config** command. The display will show either Cisco IS-IS or IETF IS-IS configuration. The following example indicates that the device uses the Cisco implementation of IS-IS NSF:

```

Router# show running-config

router isis
nsf cisco

```

If the NSF configuration is set to **cisco**, use the **show isis nsf** command to verify that NSF is enabled on the device. Using the Cisco configuration, the display output will be different on the active and standby RPs. The following example shows output for the Cisco configuration on the active RP. In this example, note the presence of the phrase “NSF restart enabled”:

```

Router# show isis nsf

NSF is ENABLED, mode 'cisco'

RP is ACTIVE, standby ready, bulk sync complete
NSF interval timer expired (NSF restart enabled)
Checkpointing enabled, no errors
Local state:ACTIVE, Peer state:STANDBY HOT, Mode:SSO

```

The following example shows sample output for the Cisco configuration on the standby RP. In this example, note the presence of the phrase “NSF restart enabled”:

```
Router# show isis nsf

NSF enabled, mode 'cisco'
RP is STANDBY, chkpt msg receive count:ADJ 2, LSP 7
NSF interval timer notification received (NSF restart enabled)
Checkpointing enabled, no errors
Local state:STANDBY HOT, Peer state:ACTIVE, Mode:SSO
```

The following example shows sample output for the IETF IS-IS configuration on the networking device:

```
Router# show isis nsf

NSF is ENABLED, mode IETF
NSF pdb state:Inactive
NSF L1 active interfaces:0
NSF L1 active LSPs:0
NSF interfaces awaiting L1 CSNP:0
Awaiting L1 LSPs:
NSF L2 active interfaces:0
NSF L2 active LSPs:0
NSF interfaces awaiting L2 CSNP:0
Awaiting L2 LSPs:
Interface:Serial3/0/2
    NSF L1 Restart state:Running
    NSF p2p Restart retransmissions:0
    Maximum L1 NSF Restart retransmissions:3
    L1 NSF ACK requested:FALSE
    L1 NSF CSNP requested:FALSE
    NSF L2 Restart state:Running
    NSF p2p Restart retransmissions:0
    Maximum L2 NSF Restart retransmissions:3
    L2 NSF ACK requested:FALSE
Interface:GigabitEthernet2/0/0
    NSF L1 Restart state:Running
    NSF L1 Restart retransmissions:0
    Maximum L1 NSF Restart retransmissions:3
    L1 NSF ACK requested:FALSE
    L1 NSF CSNP requested:FALSE
    NSF L2 Restart state:Running
    NSF L2 Restart retransmissions:0
    Maximum L2 NSF Restart retransmissions:3
    L2 NSF ACK requested:FALSE
    L2 NSF CSNP requested:FALSE
Interface:Loopback1
    NSF L1 Restart state:Running
    NSF L1 Restart retransmissions:0
    Maximum L1 NSF Restart retransmissions:3
    L1 NSF ACK requested:FALSE
    L1 NSF CSNP requested:FALSE
    NSF L2 Restart state:Running
    NSF L2 Restart retransmissions:0
    Maximum L2 NSF Restart retransmissions:3
    L2 NSF ACK requested:FALSE
    L2 NSF CSNP requested:FALSE
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

9

Route Processor Redundancy (RPR)

- [Prerequisites for RPR, page 9-1](#)
- [Restrictions for RPR, page 9-1](#)
- [Information About RPR, page 9-2](#)
- [Default Settings for RPR, page 9-4](#)
- [How to Configure RPR, page 9-4](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - In route processor redundancy (RPR) redundancy mode, the ports on a supervisor engine in standby mode are disabled.
 - RPR supports IPv6 multicast traffic.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for RPR

None.

Restrictions for RPR

- [General RPR Restrictions, page 9-2](#)
- [Hardware Restrictions for RPR, page 9-2](#)

General RPR Restrictions

- When a redundant supervisor engine is in standby mode, the two Gigabit Ethernet interfaces on the standby supervisor engine are always active.
- Supervisor engine redundancy does not provide supervisor engine mirroring or supervisor engine load balancing. Only one supervisor engine is active.
- Configuration changes made through SNMP are not synchronized to the standby supervisor engine. After you configure the switch through SNMP, copy the running-config file to the startup-config file on the active supervisor engine to trigger synchronization of the startup-config file on the standby supervisor engine.
- Supervisor engine switchover takes place after the failed supervisor engine completes a core dump. A core dump can take up to 15 minutes. To get faster switchover time, disable core dump on the supervisor engines.
- You cannot perform configuration changes during the startup (bulk) synchronization. If you attempt to make configuration changes during this process, the following message is generated:

```
Config mode locked out till standby initializes
```
- If configuration changes occur at the same time as a supervisor engine switchover, these configuration changes are lost.

Hardware Restrictions for RPR

- Cisco IOS supports redundant configurations where the supervisor engines are identical. If they are not identical, one will boot first and become active and hold the other supervisor engine in a reset condition.
- Each supervisor engine must have the resources to run the switch on its own, which means all supervisor engine resources are duplicated, including all flash devices.
- Make separate console connections to each supervisor engine. Do not connect a Y cable to the console ports.
- Except during an FSU, both supervisor engines must have the same system image (see the [“Copying Files to the RP”](#) section on page 9-5).
- The configuration register must be set to 0x2102 (`config-register 0x2102`).

**Note**

There is no support for booting from the network.

Information About RPR

- [Supervisor Engine Redundancy Overview, page 9-3](#)
- [RPR Operation, page 9-3](#)
- [Supervisor Engine Configuration Synchronization, page 9-4](#)

Supervisor Engine Redundancy Overview

The switch supports fault resistance by allowing a standby supervisor engine to take over if the primary supervisor engine fails. RPR supports a switchover time of 2 or more minutes.

The following events cause a switchover:

- A hardware failure on the active supervisor engine
- Clock synchronization failure between supervisor engines
- A manual switchover

RPR Operation

RPR supports the following features:

- Auto-startup and bootvar synchronization between active and standby supervisor engines
- Hardware signals that detect and decide the active or standby status of supervisor engines
- Clock synchronization every 60 seconds from the active to the standby supervisor engine
- A standby supervisor engine that is booted but not all subsystems are up: if the active supervisor engine fails, the standby supervisor engine become fully operational
- An operational supervisor engine present in place of the failed unit becomes the standby supervisor engine
- Support for fast software upgrade (FSU) (see [Chapter 6, “Fast Software Upgrade”](#).)

When the switch is powered on, RPR runs between the two supervisor engines. The supervisor engine that boots first becomes the RPR active supervisor engine. The route processor (RP) and Policy Feature Card (PFC) become fully operational. The RP and PFC on the standby supervisor engine come out of reset but are not operational.

In a switchover, the standby supervisor engine become fully operational and the following occurs:

- All switching modules power up again
- Remaining subsystems on the RP (including Layer 2 and Layer 3 protocols) are brought up
- Access control lists (ACLs) are reprogrammed into supervisor engine hardware

**Note**

In a switchover, there is a disruption of traffic because some address states are lost and then restored after they are dynamically redetermined.

Supervisor Engine Configuration Synchronization



Note

Configuration changes made through SNMP are not synchronized to the standby supervisor engine. After you configure the switch through SNMP, copy the running-config file to the startup-config file on the active supervisor engine to trigger synchronization of the startup-config file on the standby supervisor engine.

During RPR mode operation, the startup-config files and the config-register configurations are synchronized by default between the two supervisor engines. In a switchover, the new active supervisor engine uses the current configuration.

Default Settings for RPR

None.

How to Configure RPR

- [Synchronizing the Supervisor Engine Configurations, page 9-4](#)
- [Displaying the Redundancy States, page 9-4](#)
- [Copying Files to the RP, page 9-5](#)

Synchronizing the Supervisor Engine Configurations

During normal operation, the startup-config and config-registers configuration are synchronized by default between the two supervisor engines. In a switchover, the new active supervisor engine uses the current configuration.



Note

Do not change the default auto-sync configuration.

Displaying the Redundancy States

To display the redundancy states, perform this task:

Command	Purpose
Router# show redundancy states	Displays the redundancy states.

This example shows how to display the redundancy states:

```
Router# show redundancy states
my state = 13 -ACTIVE
    peer state = 8 -STANDBY HOT
        Mode = Duplex
        Unit = Primary
```

```

Unit ID = 1

Redundancy Mode (Operational) = Route Processor Redundancy
Redundancy Mode (Configured)  = Route Processor Redundancy
Split Mode = Disabled
Manual Swact = Enabled
Communications = Up

client count = 11
client_notification_TMR = 30000 milliseconds
keep_alive TMR = 9000 milliseconds
keep_alive count = 0
keep_alive threshold = 18
RF debug mask = 0x0

```

In this example, the system cannot enter the redundancy state because the second supervisor engine is disabled or missing:

```

Router# show redundancy states
my state = 13 -ACTIVE
peer state = 1 -DISABLED
Mode = Simplex
Unit = Primary
Unit ID = 1

Redundancy Mode (Operational) = rpr
Redundancy Mode (Configured)  = rpr
Redundancy State               = Non Redundant
Maintenance Mode = Disabled
Communications = Down          Reason: Simplex mode

client count = 11
client_notification_TMR = 30000 milliseconds
keep_alive TMR = 4000 milliseconds
keep_alive count = 0
keep_alive threshold = 7
RF debug mask = 0x0

```

Copying Files to the RP

Use the following command to copy a file to the **bootdisk:** device on an active RP:

```
Router# copy source_device:source_filename bootdisk:target_filename
```

Use the following command to copy a file to the **bootdisk:** device on a standby RP:

```
Router# copy source_device:source_filename slavebootdisk:target_filename
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

10

Interface Configuration

- [Information About Interface Configuration](#), page 10-2
- [How to Configure a Range of Interfaces](#), page 10-2
- [How to Define and Use Interface-Range Macros](#), page 10-2
- [How to Configure Optional Interface Features](#), page 10-3
- [Information About Online Insertion and Removal](#), page 10-11
- [How to Monitor and Maintain Interfaces](#), page 10-11
- [How to Check Cable Status with the TDR](#), page 10-14
- [How to Convert Line Card to Operate in 40G Mode](#), page 10-14

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Information About Interface Configuration

Many features in the software are enabled on a per-interface basis. When you enter the **interface** command, you must specify the following information:

- Interface type:
 - Fast Ethernet (use the **fastethernet** keyword)
 - Gigabit Ethernet (use the **gigabitethernet** keyword)
 - 10-Gigabit Ethernet (use the **tengigabitethernet** keyword)
 - 40-Gigabit Ethernet (use the **fortygigabitethernet** keyword)
- Slot number—The slot in which the module is installed. On switches supported by Cisco IOS Release 15.4SY, slots are numbered starting with 1 from top to bottom.
- Port number—The physical port number on the module. On switches supported by Cisco IOS Release 15.4SY, the port numbers always begin with 1. When facing the rear of the switch, ports are numbered from the left to the right.

You can identify ports from the physical location. You also can use **show** commands to display information about a specific port, or all the ports.

See this document for information about the **interface** command:

<http://www.cisco.com/en/US/docs/ios-xml/ios/interface/command/ir-i1.html#GUID-0D6BDFCD-3FBB-4D26-A274-C1221F8592DF>

How to Configure a Range of Interfaces

The interface-range configuration mode allows you to configure multiple interfaces with the same configuration parameters. After you enter the interface-range configuration mode, all command parameters you enter are attributed to all interfaces within that range until you exit out of the interface-range configuration mode. See this document for information about the **interface range** command:

<http://www.cisco.com/en/US/docs/ios-xml/ios/interface/command/ir-i1.html#GUID-8EC4EF91-F929-45F8-95CA-E4C9A9724FFF>

How to Define and Use Interface-Range Macros

You can define 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** command string, you must define the macro.

To define an interface-range macro, perform this task:

Command	Purpose
<pre>Router(config)# define interface-range macro_name {vlan vlan_ID - vlan_ID} {type slot/port - port} [, {type slot/port - port}]</pre>	Defines the interface-range macro and save it in NVRAM.

This example shows how to define an interface-range macro named `enet_list` to select Gigabit Ethernet ports 1/1 through 1/4:

```
Router(config)# define interface-range enet_list gigabitethernet 1/1 - 4
```

To show the defined interface-range macro configuration, perform this task:

Command	Purpose
Router# show running-config	Shows the defined interface-range macro configuration.

This example shows how to display the defined interface-range macro named `enet_list`:

```
Router# show running-config | include define  
define interface-range enet_list GigabitEthernet1/1 - 4  
Router#
```

To use an interface-range macro in the **interface range** command, perform this task:

Command	Purpose
Router(config)# interface range macro <i>macro_name</i>	Selects the interface range to be configured using the values saved in a named interface-range macro.

This example shows how to change to the interface-range configuration mode using the interface-range macro `enet_list`:

```
Router(config)# interface range macro enet_list  
Router(config-if)#
```

How to Configure Optional Interface Features

- [Configuring Ethernet Interface Speed and Duplex Mode, page 10-3](#)
- [Configuring Jumbo Frame Support, page 10-6](#)
- [Configuring IEEE 802.3x Flow Control, page 10-9](#)
- [Configuring the Port Debounce Timer, page 10-10](#)

Configuring Ethernet Interface Speed and Duplex Mode

- [Speed and Duplex Mode Configuration Guidelines, page 10-4](#)
- [Configuring the Ethernet Interface Speed, page 10-4](#)
- [Setting the Interface Duplex Mode, page 10-5](#)
- [Configuring Link Negotiation on Gigabit Ethernet Ports, page 10-5](#)
- [Displaying the Speed and Duplex Mode Configuration, page 10-6](#)

Speed and Duplex Mode Configuration Guidelines

You usually configure Ethernet port speed and duplex mode parameters to auto and allow ports to negotiate the speed and duplex mode. If you decide to configure the port speed and duplex modes manually, consider the following information:

- You cannot set the Ethernet port speed to auto (the **no speed** command) if the duplex mode is not set to auto (the **no duplex** command).
- If you configure an Ethernet port speed to a value other than auto (for example, 10, 100, or 1000 Mbps), configure the connecting port to match. Do not configure the connecting port to negotiate the speed.
- If you manually configure the Ethernet port speed to either 10 Mbps or 100 Mbps, the switch prompts you to also configure the duplex mode on the port.



Note

A LAN port cannot automatically negotiate Ethernet port speed and duplex mode if the connecting port is configured to a value other than auto.



Caution

Changing the Ethernet port speed and duplex mode configuration might shut down and reenable the interface during the reconfiguration.

Configuring the Ethernet Interface Speed



Note

If you configure the Ethernet port speed to **auto** on a 10/100/1000-Mbps Ethernet port, both speed and duplex are autonegotiated. 10-Gigabit Ethernet ports do not support autonegotiation.

To configure the port speed for a 10/100/1000-Mbps Ethernet port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface gigabitethernet slot/port	Selects the Ethernet port to be configured.
Step 2	Router(config-if)# speed {10 100 1000 {auto [10 100 [1000]]}}	Configures the speed of the Ethernet interface.

When configuring the port speed for a 10/100/1000-Mbps Ethernet port, note the following:

- Enter the **auto 10 100** keywords to restrict the negotiated speed to 10-Mbps or 100-Mbps.
- The **auto 10 100 1000** keywords have the same effect as the **auto** keyword by itself.

This example shows how to configure the speed to 100 Mbps on the Gigabit Ethernet port 1/4:

```
Router(config)# interface gigabitethernet 1/4
Router(config-if)# speed 100
```


Setting the Interface Duplex Mode



Note

- 10-Gigabit Ethernet and Gigabit Ethernet are full duplex only. You cannot change the duplex mode on 10-Gigabit Ethernet or Gigabit Ethernet ports or on a 10/100/1000-Mbps port configured for Gigabit Ethernet.
- If you set the port speed to auto on a 10/100/1000-Mbps Ethernet port, both speed and duplex are autonegotiated. You cannot change the duplex mode of autonegotiation ports.

To set the duplex mode of an Ethernet or Gigabit Ethernet port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface gigabitethernet <i>slot/port</i>	Selects the Ethernet port to be configured.
Step 2	Router(config-if)# duplex [auto full half]	Sets the duplex mode of the Ethernet port.

This example shows how to set the duplex mode to full on Gigabit Ethernet port 1/4:

```
Router(config)# interface gigabitethernet 1/4
Router(config-if)# duplex full
```

Configuring Link Negotiation on Gigabit Ethernet Ports



Note

Link negotiation does not negotiate port speed.

On Gigabit Ethernet ports, link negotiation exchanges flow-control parameters, remote fault information, and duplex information. Link negotiation is enabled by default.

The ports on both ends of a link must have the same setting. The link will not come up if the ports at each end of the link are set inconsistently (link negotiation enabled on one port and disabled on the other port).

[Table 10-1](#) shows the four possible link negotiation configurations and the resulting link status for each configuration.

Table 10-1 Link Negotiation Configuration and Possible Link Status

Link Negotiation State		Link Status	
Local Port	Remote Port	Local Port	Remote Port
Off	Off	Up	Up
On	On	Up	Up
Off	On	Up	Down
On	Off	Down	Up

To configure link negotiation on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface gigabitethernet <i>slot/port</i>	Selects the port to be configured.
Step 2	Router(config-if)# speed nonegotiate	Disables link negotiation.

This example shows how to enable link negotiation on Gigabit Ethernet port 1/4:

```
Router(config)# interface gigabitethernet 1/4
Router(config-if)# no speed nonegotiate
```

Displaying the Speed and Duplex Mode Configuration

To display the speed and duplex mode configuration for a port, perform this task:

Command	Purpose
Router# show interfaces <i>type slot/port</i> [transceiver properties]	Displays the speed and duplex mode configuration. To display autonegotiation status for speed and duplex, add the transceiver properties option.

Configuring Jumbo Frame Support

- [Information about Jumbo Frame Support, page 10-6](#)
- [Configuring MTU Sizes, page 10-8](#)

Information about Jumbo Frame Support

- [Jumbo Frame Support Overview, page 10-6](#)
- [Nondefault MTU Sizes on Ethernet Ports, page 10-7](#)
- [VLAN Interfaces, page 10-8](#)

Jumbo Frame Support Overview

A jumbo frame is a frame larger than the default Ethernet size. You enable jumbo frame support by configuring a larger-than-default maximum transmission unit (MTU) size on a port or VLAN interface and configuring the global LAN port MTU size.



Note

- Jumbo frame support fragments routed traffic in software on the route processor (RP).
- Jumbo frame support does not fragment bridged traffic.

Bridged and Routed Traffic Size Check at Ingress 10/100, and 100 Mbps Ethernet and 10-Gigabit Ethernet Ports

Jumbo frame support compares ingress traffic size with the global LAN port MTU size at ingress 10/100, and 100 Mbps Ethernet and 10-Gigabit Ethernet LAN ports that have a nondefault MTU size configured. The port drops traffic that is oversized. You can configure the global LAN port MTU size (see the

[“Configuring the Global Egress LAN Port MTU Size” section on page 10-9](#)).

Bridged and Routed Traffic Size Check at Ingress Gigabit Ethernet Ports

Gigabit Ethernet LAN ports configured with a nondefault MTU size accept frames containing packets of any size larger than 64 bytes. With a nondefault MTU size configured, Gigabit Ethernet LAN ports do not check for oversize ingress frames.

Routed Traffic Size Check on the PFC

For traffic that needs to be routed, Jumbo frame support on the PFC compares traffic sizes to the configured MTU sizes and provides Layer 3 switching for jumbo traffic between interfaces configured with MTU sizes large enough to accommodate the traffic. Between interfaces that are not configured with large enough MTU sizes, if the “do not fragment bit” is not set, the PFC sends the traffic to the RP to be fragmented and routed in software. If the “do not fragment bit” is set, the PFC drops the traffic.

Bridged and Routed Traffic Size Check at Egress 10, 10/100, and 100 Mbps Ethernet Ports

10, 10/100, and 100 Mbps Ethernet LAN ports configured with a nondefault MTU size transmit frames containing packets of any size larger than 64 bytes. With a nondefault MTU size configured, 10, 10/100, and 100 Mbps Ethernet LAN ports do not check for oversize egress frames.

Bridged and Routed Traffic Size Check at Egress Gigabit Ethernet and 10-Gigabit Ethernet Ports

Jumbo frame support compares egress traffic size with the global egress LAN port MTU size at egress Gigabit Ethernet and 10-Gigabit Ethernet LAN ports that have a nondefault MTU size configured. The port drops traffic that is oversized. You can configure the global LAN port MTU size (see the [“Configuring the Global Egress LAN Port MTU Size” section on page 10-9](#)).

Nondefault MTU Sizes on Ethernet Ports

- [Ethernet Port Overview, page 10-7](#)
- [Layer 3 Ethernet Ports, page 10-7](#)
- [Layer 2 Ethernet Ports, page 10-8](#)

Ethernet Port Overview

Configuring a nondefault MTU size on a 10, 10/100, or 100 Mbps Ethernet port limits ingress packets to the global LAN port MTU size and permits egress traffic of any size larger than 64 bytes.

Configuring a nondefault MTU size on a Gigabit Ethernet port permits ingress packets of any size larger than 64 bytes and limits egress traffic to the global LAN port MTU size.

Configuring a nondefault MTU size on a 10-Gigabit Ethernet port limits ingress and egress packets to the global LAN port MTU size.

You can configure the MTU size on any Ethernet port.

Layer 3 Ethernet Ports

On a Layer 3 port, you can configure an MTU size on each Layer 3 Ethernet port that is different than the global LAN port MTU size.



Note

Traffic through a Layer 3 Ethernet LAN port that is configured with a nondefault MTU size is also subject to the global LAN port MTU size (see the [“Configuring the Global Egress LAN Port MTU Size” section on page 10-9](#)).

Layer 2 Ethernet Ports

On a Layer 2 port, you can only configure an MTU size that matches the global LAN port MTU size (see the [“Configuring the Global Egress LAN Port MTU Size” section on page 10-9](#)).

VLAN Interfaces

You can configure a different MTU size on each Layer 3 VLAN interface. Configuring a nondefault MTU size on a VLAN interface limits traffic to the nondefault MTU size. You can configure the MTU size on VLAN interfaces to support jumbo frames.

Configuring MTU Sizes

- [Configuring the MTU Size, page 10-8](#)
- [Configuring the Global Egress LAN Port MTU Size, page 10-9](#)

Configuring the MTU Size

To configure the MTU size, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{ vlan <i>vlan_ID</i> } {{ type <i>slot/port</i> } {{ port-channel <i>port_channel_number</i> } <i>slot/port</i> }}	Selects the interface to configure.
Step 2	Router(config-if)# mtu <i>mtu_size</i>	Configures the MTU size.
Step 3	Router(config-if)# end	Exits configuration mode.

When configuring the MTU size, note the following information:

- For VLAN interfaces and Layer 3 Ethernet ports, supported MTU values are from 64 to 9216 bytes.



Note

For Cisco Catalyst C6840 Series switches the supported MTU values are from 64 to 9154 bytes.

- For Layer 2 Ethernet ports, you can configure only the global egress LAN port MTU size (see the [“Configuring the Global Egress LAN Port MTU Size” section on page 10-9](#)).
- MTU size on a Layer 2 interface is always set to maximum. If you reconfigure the MTU size to enable jumbo frame support, MTU size does not get updated in the MTU hardware table for L2 interface at the egress. The interface (L2) MTU size check happens only at the ingress port.

This example shows how to configure the MTU size on Gigabit Ethernet port 1/2:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/2
Router(config-if)# mtu 9216
Router(config-if)# end
```

This example shows how to verify the configuration:

```
Router# show interface gigabitethernet 1/2
GigabitEthernet1/2 is administratively down, line protocol is down
  Hardware is C6k 1000Mb 802.3, address is 0030.9629.9f88 (bia 0030.9629.9f88)
  MTU 9216 bytes, BW 1000000 Kbit, DLY 10 usec,
  <...Output Truncated...>
Router#
```

Configuring the Global Egress LAN Port MTU Size

To configure the global egress LAN port MTU size, perform this task:

	Command	Purpose
Step 1	Router(config)# system jumbomtu <i>mtu_size</i>	Configures the global egress LAN port MTU size. Note Because it would change all the interface MTU sizes to the default (1500), rather than to any configured nondefault interface MTU size, do not use the system jumbomtu command to set the MTU size to 1500. (CSCtq52016)
Step 2	Router(config)# end	Exits configuration mode.

Configuring IEEE 802.3x Flow Control

Gigabit Ethernet and 10-Gigabit Ethernet ports use flow control to stop the transmission of frames to the port for a specified time; other Ethernet ports use flow control to respond to flow-control requests.

If a Gigabit Ethernet or 10-Gigabit Ethernet port receive buffer becomes full, the port can be configured to transmit an IEEE 802.3x pause frame that requests the remote port to delay sending frames for a specified time. All Ethernet ports can be configured to respond to IEEE 802.3x pause frames from other devices.

To configure flow control on an Ethernet port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the port to configure.
Step 2	Router(config-if)# flowcontrol { receive send } { desired off on }	Configures a port to send or respond to pause frames.

When configuring flow control, note the following information:

- Because auto negotiation does not work on 10 Gigabit Ethernet fiber optic ports, they respond to pause frames by default. On 10 Gigabit Ethernet fiber optic ports, the flow-control operational mode is always the same as administrative mode.
- When configuring how a port responds to pause frames, note the following information:
 - For a Gigabit Ethernet port, when the configuration of a remote port is unknown, you can use the **receive desired** keywords to configure the Gigabit Ethernet port to respond to received pause frames. (Supported only on Gigabit Ethernet ports.)
 - Use the **receive on** keywords to configure a port to respond to received pause frames.
 - Use the **receive off** keywords to configure a port to ignore received pause frames.
- When configuring transmission of pause frames on a port, note the following information:
 - For a Gigabit Ethernet port, when the configuration of the remote ports is unknown, you can use the **send desired** keywords to configure the Gigabit Ethernet port to send pause frames. (Supported only on Gigabit Ethernet ports.)
 - Use the **send on** keywords to configure a port to send pause frames.
 - Use the **send off** keywords to configure a port not to send pause frames.

This example shows how to turn on receive flow control and how to verify the flow-control configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/2
Router(config-if)# flowcontrol receive on
Router(config-if)# end
Router# show interfaces flowcontrol

Interface Send      Receive
Gi1/1      Desired   OFF
Gi1/2      Desired   ON
<output truncated>
```

Configuring the Port Debounce Timer

The port debounce timer delays notification of a link change, which can decrease traffic loss due to network reconfiguration. You can configure the port debounce timer separately on each LAN port.



Caution

Enabling the port debounce timer causes link down detections to be delayed, resulting in loss of traffic during the debouncing period. This situation might affect the convergence and reconvergence of some Layer 2 and Layer 3 protocols.

To configure the debounce timer on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the port to configure.
Step 2	Router(config-if)# link debounce [time <i>debounce_time</i>]	Configures the debounce timer.

When configuring the debounce timer on a port, note the following information:

- The **time** keyword is supported only on fiber 1000 Mbps or faster Ethernet ports.
- You can increase the port debounce timer value in increments of 100 milliseconds up to 5000 milliseconds on ports operating at 1000 Mbps over copper media.
- The debounce timer recognizes 10-Gbps copper media and detects media-only changes.

Table 10-2 lists the time delay that occurs before notification of a link change.

Table 10-2 Default Port Debounce Timer Delay Times

Port Type	Debounce Timer Disabled	Debounce Timer Enabled
Ports operating at 10 Mbps or 100 Mbps:	300 milliseconds	3100 milliseconds
Ports operating at 1000 Mbps or 10 Gbps over copper media:	300 milliseconds	3100 milliseconds
Ports operating at 1000 Mbps or 10 Gbps over fiber media:	10 milliseconds	100 milliseconds
Note The show interfaces debounce command does not display the default value for 10-GigabitEthernet ports when the port debounce timer is disabled.		

**Note**

On all 10-Gigabit Ethernet ports, the Debounce Timer Disabled value is 10 milliseconds and the Debounce Timer Enabled value is 100 milliseconds.

This example shows how to enable the port debounce timer on Gigabit Ethernet port 1/12:

```
Router(config)# interface gigabitethernet 1/12
Router(config-if)# link debounce
Router(config-if)# end
```

This example shows how to display the port debounce timer settings:

```
Router# show interfaces debounce | include enable
Gi1/12  enable          3100
```

Information About Online Insertion and Removal

The online insertion and removal (OIR) feature allows you to remove and replace modules while the system is online. You can shut down the modules before removal and restart it after insertion without causing other software or interfaces to shut down.

**Note**

Do not remove or install more than one module at a time. After you remove or install a module, check the LEDs before continuing. For module LED descriptions, see the *Catalyst 6500 Series Switch Installation Guide*.

When a module has been removed or installed, the switch stops processing traffic for the module and scans the system for a configuration change. Each interface type is verified against the system configuration, and then the system runs diagnostics on the new module. There is no disruption to normal operation during module insertion or removal.

The switch can bring only an identical replacement module online. To support OIR of an identical module, the module configuration is not removed from the running-config file when you remove a module.

If the replacement module is different from the removed module, you must configure it before the switch can bring it online.

Layer 2 MAC addresses are stored in an EEPROM, which allows modules to be replaced online without requiring the system to update switching tables and data structures. Regardless of the types of modules installed, the Layer 2 MAC addresses do not change unless you replace the supervisor engine. If you do replace the supervisor engine, the Layer 2 MAC addresses of *all* ports change to those specified in the address allocator on the new supervisor engine.

How to Monitor and Maintain Interfaces

- [Monitoring Interface Status, page 10-12](#)
- [Clearing Counters on an Interface, page 10-12](#)
- [Resetting an Interface, page 10-13](#)
- [Shutting Down and Restarting an Interface, page 10-13](#)

Monitoring Interface Status

The software contains commands that you can enter at the EXEC prompt to display information about the interface including the version of the software and the hardware and statistics about interfaces. The following table lists some of the interface monitoring commands. (You can display the complete list of **show** commands by using the **show ?** command at the EXEC prompt.) These commands are described in the *Cisco IOS Interface Command Reference* publication.

To display information about the interface, perform these tasks:

Command	Purpose
Router# show ibc	Displays current internal status information.
Router# show eobc	Displays current internal out-of-band information.
Router# show interfaces [<i>type slot/port</i>]	Displays the status and configuration of all or a specific interface.
Router# show running-config	Displays the currently running configuration.
Router# show rif	Displays the current contents of the routing information field (RIF) cache.
Router# show protocols [<i>type slot/port</i>]	Displays the global (system-wide) and interface-specific status of any configured protocol.
Router# show version	Displays the hardware configuration, software version, the names and sources of configuration files, and the boot images.

Clearing Counters on an Interface

To clear the interface counters shown with the **show interfaces** command, perform this task:

Command	Purpose
Router# clear counters {{ vlan <i>vlan_ID</i> } { <i>type slot/port</i> } { port-channel <i>channel_ID</i> }}	Clears interface counters.

This example shows how to clear and reset the counters on Gigabit Ethernet port 1/5:

```
Router# clear counters gigabitethernet 1/5
Clear "show interface" counters on this interface [confirm] y
*Sep 30 08:42:55: %CLEAR-5-COUNTERS: Clear counter on interface GigabitEthernet1/5
```

The **clear counters** command clears all the current counters from the interface unless the optional arguments specify a specific interface.



Note

The **clear counters** command clears counters displayed with the EXEC **show interfaces** command, not counters retrieved using SNMP.

Resetting an Interface

To reset an interface, perform this task:

Command	Purpose
Router# clear interface <i>type slot/port</i>	Resets an interface.

This example shows how to reset Gigabit Ethernet port 1/5:

```
Router# clear interface gigabitethernet 1/5
```

Shutting Down and Restarting an Interface

You can shut down an interface, which disables all functions on the specified interface and shows 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 included in any routing updates.

To shut down an interface and then restart it, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{ vlan <i>vlan_ID</i> } { type slot/port } { port-channel <i>channel_ID</i> }}	Selects the interface to be configured.
Step 2	Router(config-if)# shutdown	Shuts down the interface.
Step 3	Router(config-if)# no shutdown	Reenables the interface.

This example shows how to shut down Gigabit Ethernet port 1/5:

```
Router(config)# interface gigabitethernet 1/5
Router(config-if)# shutdown
Router(config-if)#
```



Note

The link state messages (LINK-3-UPDOWN and LINEPROTO-5-UPDOWN) are disabled by default. Enter the **logging event link status** command on each interface where you want the messages enabled.

This example shows how to reenabale Gigabit Ethernet port 1/5:

```
Router(config-if)# no shutdown
Router(config-if)#
```

To check if an interface is disabled, enter the EXEC **show interfaces** command. An interface that has been shut down is shown as administratively down in the **show interfaces** command display.

How to Check Cable Status with the TDR

You can check the status of copper cables using the time domain reflectometer (TDR). The TDR detects a cable fault by sending a signal through the cable and reading the signal that is reflected back to it. All or part of the signal can be reflected back by any number of cable defects or by the end of the cable itself.

Use the TDR to determine if the cabling is at fault if you cannot establish a link. This test is especially important when replacing an existing switch, upgrading to Gigabit Ethernet, or installing new cables.

**Note**

- TDR can test cables up to a maximum length of 115 meters.
- TDR results are not meaningful for a link that is operating successfully.
- The port must be up before running the TDR test. If the port is down, you cannot enter the **test cable-diagnostics tdr** command successfully, and the following message is displayed:

```
Router# test cable-diagnostics tdr interface gigabitethernet2/12
% Interface Gi2/12 is administratively down
% Use 'no shutdown' to enable interface before TDR test start.
```

To start or stop the TDR test, perform this task:

Command	Purpose
test cable-diagnostics tdr interface {interface interface_number}	Starts or stops the TDR test.

This example shows how to run the TDR-cable diagnostics:

```
Router # test cable-diagnostics tdr interface gigabitethernet2/1
TDR test started on interface Gi2/1
A TDR test can take a few seconds to run on an interface
Use 'show cable-diagnostics tdr' to read the TDR results.
Router #
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

How to Convert Line Card to Operate in 40G Mode

You can configure C6800-x10G cards to operate in 10G mode and 40G mode. The ports are natively 10G but can support 40G by adding a CVR-4SFP10G-QSFP adapter. These cards need to be converted to 40G mode to support 40G interfaces. The cards can support both 10G and 40G interfaces on different port-groups.

**Note**

Converting the line card to 40G mode will reset the module.

To configure a module to operate in 40G mode, perform the following task:

Command	Purpose
Router(config)# hw-module { slot - slot-no } operation-mode { port-group - port-group } fortyGigabitEthernet	Configures the slot port-group to 40G Mode.

To show the module operation mode, run the following command:


Command	Purpose
Router# show hw-module { slot - slot-no } operation-mode	Shows the operation mode on slot per port-group.

Overview of Cisco CVR-4SFP10G-QSFP Adapter


The Cisco CVR-4SFP10G-QSFP Adapter converts four SFP+ ports into a QSFP+ port. With this adapter, you have the flexibility to use a QSFP+ module to connect to a 40 Gigabit Ethernet port on the other end of the link. This flexibility allows a cost-effective transition to 40 Gigabit Ethernet by maximizing the use of existing or new 10 Gigabit Ethernet SFP+ platforms.

For more information, see the Cisco 4 x SFP10G to QSFP Reverse Adapter Data Sheet.

Figure 10-1 Cisco CVR-4SFP10G-QSFP Adapter and QSFP+ Module



354671

Figure 10-2 Cisco CVR-4SFP10G-QSFP Adapter with QSFP+ Module Plugged In

354672

Compatible Line Cards

All Catalyst 6800 series line cards support CVR-4SFP10G-QSFP Adapter. Due to power limitations on line card, one port per port-group can be operated on 40G and connected to CVR-4SFP10G-QSFP Adapter.

Table 10-3 40G Ports supported on port groups

Line Card	Port-group Mapping	40G Ports – Physical Port Mapping	Supported 40G Ports per Port-Group
C6800-32P10G	PG1 – 1,3,5,7,9,11,13,15	Port 33 – 1,3,5,7 & Port 35 – 9,11,13,15	PG1 – Port 33
	PG2 – 2,4,6,8,10,12,14,16	Port 34 – 2,4,6,8 & Port 36 – 10,12,14,16	PG2 – Port 34
	PG3 – 17,19,21,23,25,27,29,31	Port 37 – 17,19,21,23 & Port 39 – 25,27,29,31	PG3 – Port 37
	PG4 – 18,20,22,24,26,28,30,32	Port 38 – 18,20,22,24 & Port 40 – 26,28,30,32	PG4 – Port 38
C6800-16P10G	PG1 – 1,2,3,4,5,6,7,8	Port 17 – 1,2,3,4 & Port 18 – 5,6,7,8	PG1 – Port 17
	PG2 – 9,10,11,12,13,14,15,16	Port 19 – 9,10,11,12 & Port 20 – 13,14,15,16	PG2 – Port 19
C6800-8P10G	PG1 – 1,2,3,4	Port 9 – 1,2,3,4	PG1 – Port 9
	PG2 – 5,6,7,8	Port 10 – 5,6,7,8	PG2 – Port 10

**Note**

40G mode conversion on line card is per port-group, ports in the non-converted port-groups will continue to operate in 10G mode and no ports are disabled.

11

UniDirectional Link Detection (UDLD)

- [Prerequisites for UDLD, page 11-1](#)
- [Restrictions for UDLD, page 11-1](#)
- [Information About UDLD, page 11-1](#)
- [Default Settings for UDLD, page 11-4](#)
- [How to Configure UDLD, page 11-4](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for UDLD

None.

Restrictions for UDLD

None.

Information About UDLD

- [UDLD Overview, page 11-2](#)
- [UDLD Aggressive Mode, page 11-3](#)

- [Fast UDLD, page 11-3](#)

UDLD Overview

The Cisco-proprietary UDLD protocol allows devices connected through fiber-optic or copper (for example, Category 5 cabling) Ethernet cables connected to LAN ports to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a unidirectional link is detected, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.

UDLD is a Layer 2 protocol that works with the Layer 1 protocols to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected LAN ports. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic transmitted by the local device over a link is received by the neighbor but traffic transmitted from the neighbor is not received by the local device. If one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both fibers are working normally at Layer 1, then UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation, because autonegotiation operates at Layer 1.

LAN ports with UDLD enabled periodically transmit UDLD packets to neighbor devices. If the packets are echoed back within a specific time frame and they are lacking a specific acknowledgment (echo), the link is flagged as unidirectional and the LAN port is shut down. Devices on both ends of the link must support UDLD in order for the protocol to successfully identify and disable unidirectional links.




Note

By default, UDLD is locally disabled on copper LAN ports to avoid sending unnecessary control traffic on this type of media since it is often used for access ports.

[Figure 11-1](#) shows an example of a unidirectional link condition. Switch B successfully receives traffic from Switch A on the port. However, Switch A does not receive traffic from Switch B on the same port. UDLD detects the problem and disables the port.

Figure 11-1 Unidirectional Link



UDLD Aggressive Mode

UDLD aggressive mode is disabled by default. Configure UDLD aggressive mode only on point-to-point links between network devices that support UDLD aggressive mode. With UDLD aggressive mode enabled, when a port on a bidirectional link that has a UDLD neighbor relationship established stops receiving UDLD packets, UDLD tries to reestablish the connection with the neighbor. After eight failed retries, the port is disabled.

To prevent spanning tree loops, nonaggressive UDLD with the default interval of 15 seconds is fast enough to shut down a unidirectional link before a blocking port transitions to the forwarding state (with default spanning tree parameters).

When you enable UDLD aggressive mode, you receive additional benefits in the following situations:

- One side of a link has a port stuck (both Tx and Rx)
- One side of a link remains up while the other side of the link has gone down

In these cases, UDLD aggressive mode disables one of the ports on the link, which prevents traffic from being discarding.

**Note**

In UDLD normal mode, when a unidirectional error is detected, the port is not disabled. In UDLD aggressive mode, when a unidirectional error is detected, the port is disabled.

Fast UDLD

Release 15.0(1)SY1 and later releases support fast UDLD.

Fast UDLD is a per-port configuration option that supports UDLD message time intervals between 200 and 1000 milliseconds. Fast UDLD can be configured to provide subsecond unidirectional link detection. (Without fast UDLD, the message time intervals are 7 through 90 seconds).

When configuring fast UDLD, note the following guidelines and restrictions:

- Fast UDLD is disabled by default.
- Normal and aggressive mode both support fast UDLD.
- Fast UDLD ports do not support the **link debounce** command.
- Fast UDLD supports only point-to-point with at least two links between network devices that support fast UDLD.
- Fast UDLD does not report a unidirectional link if the same error occurs simultaneously on more than one link to the same neighbor device.
- Fast UDLD cannot detect unidirectional links when the CPU utilization exceeds 60 percent.
- Fast UDLD is supported on 60 ports with a Supervisor Engine 6T.

Default Settings for UDLD

Feature	Default Value
UDLD global enable state	Globally disabled.
UDLD aggressive mode	Disabled.
UDLD per-port enable state for fiber-optic media	Enabled on all Ethernet fiber-optic LAN ports.
UDLD per-port enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX LAN ports.
Fast UDLD	Disabled.
Fast UDLD error reporting	Disabled.

How to Configure UDLD

- [Enabling UDLD Globally, page 11-4](#)
- [Enabling UDLD on LAN Interfaces, page 11-5](#)
- [Disabling UDLD on Nonfiber-Optic LAN Interfaces, page 11-5](#)
- [Disabling UDLD on Fiber-Optic LAN Interfaces, page 11-5](#)
- [Configuring the UDLD Probe Message Interval, page 11-6](#)
- [Configuring Fast UDLD, page 11-6](#)
- [Resetting Disabled LAN Interfaces, page 11-7](#)

Enabling UDLD Globally

To enable UDLD globally on all fiber-optic LAN ports, perform this task:

Command	Purpose
Router(config)# udld { enable aggressive }	Enables UDLD globally on fiber-optic LAN ports. Note This command only configures fiber-optic LAN ports. Individual LAN port configuration overrides the setting of this command.

Enabling UDLD on LAN Interfaces

To enable UDLD on a LAN port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# udld port [aggressive]	Enables UDLD on a LAN port. - Enter the aggressive keyword to enable aggressive mode. - On a fiber-optic LAN port, this command overrides the udld enable global configuration command setting. - On fiber-optic LAN ports, the no udld port command reverts the LAN port configuration to the udld enable global configuration command setting.

Disabling UDLD on Nonfiber-Optic LAN Interfaces

To disable UDLD on a nonfiber-optic LAN port., perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# no udld port [aggressive]	Disables UDLD on a nonfiber-optic LAN port.

Disabling UDLD on Fiber-Optic LAN Interfaces

To disable UDLD on individual fiber-optic LAN ports, perform this task:

Command	Purpose
Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Router(config-if)# udld port disable	Disables UDLD on a fiber-optic LAN port. Note The no form of this command, which reverts to the udld enable global configuration command setting, is only supported on fiber-optic LAN ports.

Configuring the UDLD Probe Message Interval

To configure the time between UDLD probe messages on ports that are in advertisement mode and are currently determined to be bidirectional, perform this task:

Command	Purpose
Router(config)# udld message time interval	Configures the time between UDLD probe messages on ports that are in advertisement mode and are currently determined to be bidirectional; valid values are from 7 to 90 seconds.

Configuring Fast UDLD

Release 15.0(1)SY1 and later releases support fast UDLD. These sections describe how to configure fast UDLD:

- [Configuring Fast UDLD on a Port, page 11-6](#)
- [Enabling Fast UDLD Error Reporting, page 11-7](#)



Note

You can configure fast UDLD on ports where UDLD is not enabled, but fast UDLD is active only when UDLD is enabled on the port.

Configuring Fast UDLD on a Port

To configure fast UDLD on a port, perform this task:

	Command	Purpose
Step 1	Router(config-if)# udld fast-hello interval	Configures the fast UDLD probe message interval on a port. • See the guidelines and restrictions in the “Fast UDLD” section on page 11-3. • When selecting the value, follow these guidelines: – Valid values are from 200 to 1000 milliseconds. – Adjust the fast UDLD probe message interval to the longest interval possible that will provide the desired link failure detection time. A shorter message interval increases the likelihood that UDLD will falsely report link failures under stressful conditions.
Step 2	Router# show udld fast-hello	Displays fast UDLD configuration and operational state.
Step 3	Router# show udld fast-hello type¹ slot/number	Verifies the per-port fast UDLD configuration and operational state.

1. *type* = fastethernet, gigabitethernet, tengigabitethernet, or fortygigabitethernet

Enabling Fast UDLD Error Reporting

By default, fast UDLD error-disables ports with unidirectional links. You can globally enable fast UDLD to report unidirectional links with a message displayed on the console instead of error-disabling ports with unidirectional links.

**Note**

When fast UDLD error reporting is enabled, you must manually take the action appropriate for the state of the link.

To globally enable fast UDLD error reporting, perform this task:

Command	Purpose
Router(config)# udld fast-hello error-reporting	Enables fast UDLD error reporting.

Resetting Disabled LAN Interfaces

To reset all LAN ports that have been shut down by UDLD, perform this task:

Command	Purpose
Router# udld reset	Resets all LAN ports that have been shut down by UDLD.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

12

Instant Access (IA)

- [Prerequisites for Instant Access, page 12-1](#)
- [Restrictions for Instant Access, page 12-2](#)
- [Information About Instant Access, page 12-6](#)
- [Default Settings for Instant Access, page 12-6](#)
- [How to Configure Instant Access, page 12-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

Prerequisites for Instant Access

- An IA parent—A VSS-mode Catalyst 6800 switch or a VSS-mode Cisco Catalyst 6880-X switch or a VSS-mode Catalyst 6500 switch equipped with a [Supervisor Engine 2T](#) and one or more [WS-X6904-40G-2T](#), [C6800-32P10G/XL](#), [C6800-16P10G/XL](#), and [C6800-8P10G/XL](#) switching modules, configured to support 1/10GE links.
- IA clients—[Catalyst 6800ia access switches](#)
See this publication for more information:
http://www.cisco.com/en/US/prod/collateral/switches/ps10902/ps715/ps13198/data_sheet_c78-728230.html
- See this publication for more information about Instant Access:
http://www.cisco.com/en/US/prod/collateral/switches/ps10902/ps715/ps13198/white_paper_c11-728265.html

- For image download during ISSU upgrade, remove **ip tftp source-interface** config on IA Parent, and add a static route to copy the image through **mgmt** interface.

Restrictions for Instant Access

- The IA parent must operate in VSS mode.



Note

- You can enable VSS mode on a single chassis to support IA clients.
- The VSS Quad-Sup SSO (VS4O) feature is supported with IA clients from Release 15.1(2)SY2.

- The IA parent-client connection is supported on links between WS-X6904-40G-2T, C6800-32P10G/XL, C6800-16P10G/XL, and C6800-8P10G/XL switching modules 1/10GE ports and Catalyst 6800ia access switches 10GE ports.
 - You can use up to 8 IA client 10GE ports in the IA parent-client link. See this document for information about the port configuration for WS-X6904-40G-2T, C6800-32P10G/XL, C6800-16P10G/XL, and C6800-8P10G/XL switching modules:
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11-696669.html
 - IA client 10-Gigabit Ethernet ports require no configuration.
 - UDLD, LLDP, and CDP are not supported on the IA parent-client link.
 - IA does not use STP on the IA parent-client connection.
 - Use only XL based modules for scale FEX QoS configuration to prevent issues with TCAM (ternary content-addressable memory) utilization. When QoS policy is configured on 1500 FEX host ports, the first 511 interfaces share the TCAM utilization. But, remaining ports will start using new TCAM entries for each interface and will exhaust non-XL TCAM utilization.
- IA client maximum values:

Value Description:	Maximum Value
Maximum IA client ports	2016 ports across 42 Catalyst 6800ia access switches Note The Supervisor Engine 2T module can support only up to 1200 FEX host ports with 15.2(1)SY and 15.2(1)SY0a and up to 1500 FEX host ports with 15.2(1)SY1 and later.
Maximum IA client switches	42 (defined by IA client FEX number 1–42 range.) Note The Supervisor Engine 2T module can support only up to 25 IA client switches with 15.2(1)SY and 15.2(1)SY0a and up to 32 IA client switches with 15.2(1)SY1 and later.

Value Description:	Maximum Value
Maximum Catalyst 6800ia access switches per IA client stack	5 • An IA client stack acts as single switch unit. • Instant access only supports connection with stacking cables to form a stack. • With an IA client that has multiple Catalyst 6800ia access switches, the switches in the stack assign incrementing switch numbers to themselves (automatic stacking capability). • If you add Catalyst 6800ia access switches to a configured IA client, the additional switches assign incrementing switch numbers to themselves. • The IA client configuration does not persist if the access switch number changes.
Maximum number of VLANs per IA client stack	You can set upto 1,000 VLANs, we recommend to set not more than 20 VLANs per FEX.
Maximum Number of Port Channels	512

- IA client ports do not support these features:
 - Configuring EtherChannels with a combination of FEX Ports using different FEX-IDs or a combination of FEX host port and an IA parent linecard port is not supported. However, FEX host port channel from the same FEX is supported.
 - FEX host port EtherChannel load balancing is not supported.
 - Port debounce timer
 - UDLR tunnel ARP and IGMP proxy
 - Uni-Directional Link Routing (UDLR)
 - IEEE 802.1Q tunneling
 - VLAN Mapping
 - VLAN Translation
 - IEEE 802.1Q custom ethertypes
 - L2PT - Layer 2 protocol tunneling
 - L2PT - Layer 2 protocol tunneling on trunk ports
 - 802.1ad tunneling
 - Port security on 802.1Q tunnel ports
 - Private VLANs (PVLAN)
 - VACL capture
 - Per-VLAN load balancing for Advanced QinQ service mapping
 - Cisco TrustSec NDAC (Network Device Admission Control)
 - Cisco TrustSec confidentiality and integrity with MACsec (IEEE 802.1AE)
 - Cisco TrustSec identity port mapping
 - Network edge authentication topology (NEAT)

- AutoQoS
- MQC queuing policy support
- QoS aggregated DSCP values for WRED
- QoS aggregated precedence values for WRED
- Class based weighted fair queuing (CBWFQ)
- Class-based shaping
- DiffServ-compliant dWRED
- Diffserv-compliant WRED
- Selective packet discard (SPD)
- Weighted fair queueing (WFQ)
- Weighted RED (WRED)
- QoS policer rate increase to 256G
- Ethernet over MPLS (EoMPLS) - IEEE 802.1q Tag Stacking
- H-VPLS N-PE redundancy for QinQ access
- Connectivity fault management (CFM)
- Ethernet connectivity fault management (E-CFM)
- Ethernet local management interface (LMI) at provider edge (PE)
- Ethernet operations, administration, and Maintenance (OAM)
- Ethernet-OAM 3.0: CFM over BD, Untagged
- IEEE 802.1ag - D8.1 standard Compliant CFM, Y.1731 multicast LBM / AIS / RDI / LCK, IP SLA for Ethernet
- IEEE 802.1ag Compliant CFM (D8.1)
- To use an IA client port as a SPAN destination, add the IA client port VLAN to the SPAN allowed VLAN list with the **switchport trunk allowed vlan** command.
- When FEX IA parent-client link portchannel is configured as SPAN source in Tx direction or both directions, the SPAN destination should not be on the same FEX. This is applicable for both stacked and standalone FEX.
- IA client port QoS:

- Configure ingress QoS on the IA parent port-channel interface.
- The egress QoS configuration on IA client ports is not configurable.
- Port architecture (Rx/Tx): 1p3q3t

1p3q3t strict-priority egress queue (queue 1)		CoS	Not supported
		DSCP	32, 33, 40–47
		Tail-drop	100% (nonconfigurable)
		WRED-drop	Not supported
1p3q3t standard egress queue 2 (high priority)	Threshold 1	CoS	Not supported
		DSCP	16–23, 26–31, 34–39
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported
	Thresholds 2	CoS	Not supported
		DSCP	24
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported
	Thresholds 3	CoS	Not supported
		DSCP	48–63
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported
1p3q3t standard egress queue 3 (medium priority)	Threshold 1	CoS	Not supported
		DSCP	25
		Tail-drop	Disabled; 70%
		WRED-drop	Not supported
	Threshold 2	CoS	Not supported
		DSCP	None.
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported
	Thresholds 3	CoS	Not supported
		DSCP	0–7
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported

1p3q3t standard egress queue 4 (lowest priority)	Threshold 1	CoS	Not supported
		DSCP	8, 9, 11, 13, 15
		Tail-drop	Disabled; 70%
		WRED-drop	Not supported
	Threshold 2	CoS	Not supported
		DSCP	10, 12, 14
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported
	Threshold 3	CoS	Not supported
		DSCP	None.
		Tail-drop	Disabled; 100%
		WRED-drop	Not supported

- Maximum dot1x hosts supported with IBNS 1.0 configuration is 1000.

Information About Instant Access

The Instant Access (IA) feature supports multiple Catalyst 6800ia access switches that function as clients of the IA parent switch. The IA parent and client switches form a single extended switch with a single management domain, managed by the IA parent.

The IA parent uses the Satellite Discovery Protocol (SDP) and the Satellite Registration Protocol (SRP) to automatically discover IA clients when they connect and monitor the IA client-parent link. The IA parent upgrades the IA client software image if it is not the same as the parent FEX Image Bundle.

The IA parent features are applied to IA client FEX ports. The IA clients do not perform any local packet forwarding. All traffic originating from IA client ports are sent to the IA parent, which makes all the switching and forwarding decisions.

The IA client switches support the following online diagnostic tests:

- [TestFexModeLoopback, page 1-4](#)
- [TestFexFabricLinkStatus, page 1-39](#)

Default Settings for Instant Access

None.

How to Configure Instant Access

- [Configure Instant Access Staggered Initialization Mode, page 12-7](#)
- [Enable FEX Auto-Config, page 12-7](#)
- [Enable IA Client Preprovisioning, page 12-8](#)
- [Configure Instant Access Port-Channel Interfaces, page 12-8](#)
- [Configure Instant Access Channel Groups, page 12-9](#)
- [Identify Connected IA Client Stack Modules, page 12-9](#)
- [Configure IA Clients, page 12-13](#)
- [Display or Clear SDP and SRP Traffic, page 12-13](#)
- [Upgrade FEX using In-Service Software Upgrade \(ISSU\) Process, page 12-13](#)
- [Configure Optional Parameters for an IA Client, page 12-14](#)

Configure Instant Access Staggered Initialization Mode

Instant Access staggered initialization mode avoids any excessively high CPU utilization that might occur if multiple IA clients attempt to initialize simultaneously. To configure Instant Access staggered initialization mode, perform this task:

Command	Purpose
Router(config)# fex stagger <i>delay_value</i>	Configures Instant Access staggered initialization mode. The <i>delay_value</i> can be 0 through 500. Note The recommended <i>delay_value</i> is 120.

This example shows how to configure Instant Access staggered mode:

```
Router# configure terminal  
Router(config)# fex stagger 120
```

Enable FEX Auto-Config

FEX Auto-config mode allows increased scalability with reduced configuration. Configuration using Auto-config includes FEX ID selection, Port-channel ID selection, converting RSL port to FEX-fabric mode, and RSL port bundling.

To enable FEX auto-config, perform this task:

Command	Purpose
Router(config)# fex auto-config	Enable FEX auto-config support globally.

This example shows how to configure FEX auto-config mode:

```
Router# configure terminal  
Router(config)# fex auto-config
```

Enable IA Client Preprovisioning

To allow IA client port configuration before the IA client is connected, perform this task:

Command	Purpose
Router# module provision create fex <i>fex_number</i> type {C6800IA-48FPD C6800IA-48TD C6800IA-48FPDR WS-C3560CX-12PD-S WS-C3560CX-8XPD-S} [slot <i>switch_number</i>]	Enables IA client port configuration before the IA client is connected. Enter the slot <i>switch_number</i> keyword and argument to enable configuration of a specific IA client stack member or an additional stack member before it is added to the IA client stack.

WS-C3560CX-8XPD-S and WS-C3560CX-12PD-S will undergo a number of reloads as FEX client before coming online based on the following conditions:

- Reload1: To convert itself from switch mode to FEX mode on receiving the SDP packet from controller
- Reload2: If a mismatch of IOS version occurs between FEX controller and FEX client, the FEX client downloads the image from controller and boots up.
- Reload3: If a mismatch of MTU occurs between FEX controller and FEX client, the FEX client reloads again to change the MTU of FEX host ports.

Configure Instant Access Port-Channel Interfaces

To create a port channel interface to support IA clients, perform this task:

	Command	Purpose
Step 1	Router(config)# interface port-channel <i>group_number</i>	Creates the port channel interface. The <i>group_number</i> can be 1 through 512, up to a maximum of 512 port-channel interfaces (42 port-channel interfaces can be used to support IA clients). Note If desired, you can configure the <i>group_number</i> to match the IA client FEX number.
Step 2	Router(config-if)# switchport	Configures the port channel interface for Layer 2 switching.
Step 3	Router(config-if)# switchport mode fex-fabric	Configures the port channel interface to support IA clients.
Step 4	Router(config-if)# fex associate <i>fex_number</i>	Configures the IA client FEX number. • The valid value range is 101–199. • Maximum of 42 IA client FEX numbers.

This example shows how to create port channel interface 1 and configure it to support IA FEX number 118:

```
Router# configure terminal
Router(config)# interface port-channel 118
Router(config-if)# switchport
```

```
Router(config-if)# switchport mode fex-fabric
Router(config-if)# fex associate 118
```

Configure Instant Access Channel Groups

To configure channel groups to support IA clients, perform this task for the 10 Gigabit Ethernet LAN ports that connect to IA clients:

	Command	Purpose
Step 1	Router(config)# interface range <i>first_10ge_port, last_10ge_port</i>	Selects the ports to configure.
Step 2	Router(config-if)# switchport	Configures the port channel interface for Layer 2 switching.
Step 3	Router(config-if)# switchport mode fex-fabric	Configures the port channel interface to support IA clients.
Step 4	Router(config-if)# channel-group <i>group_number</i> mode on	Configures the LAN port in an IA Client port channel and configures the mode as on .



Note

More links can be added to the channel group at any time.

This example shows how to configure 10 Gigabit Ethernet ports 1/2/5 and 2/2/5 into port channel 118 with mode **on**:

```
Router# configure terminal
Router(config)# interface range tengigabitethernet 1/2/5, 2/2/5
Router(config-if)# switchport
Router(config-if)# switchport mode fex-fabric
Router(config-if)# channel-group 118 mode on
Router(config-if)# end
```

This example shows how to verify the IA configuration when the IA client is connected:

```
Router# show fex 118 detail
FEX: 118 Description: FEX0118 state: online
FEX version: version_string
Extender Model: C6800IA-48TD, Extender Serial: serial_number
FCP ready: yes
Image Version Check: enforced
Fabric Portchannel Ports: 2
Fabric port for control traffic: Te1/2/5
Fabric interface state:
Po20 - Interface Up.
Te1/2/5 - Interface Up. state: bound
Te2/2/5 - Interface Up. state: bound
```

Identify Connected IA Client Stack Modules

- [Identify IA Client Stack Modules by Serial Number, page 12-10](#)
- [Identify IA Client Modules by Beacon LED, page 12-10](#)
- [Easy FEX, page 12-10](#)

Identify IA Client Stack Modules by Serial Number

This example shows how to identify IA client stack modules by serial number:

```
Router# show mod fex 101
```

```
Switch Number:   101   Role:                               FEX

Mod Ports Card Type                               Model                Serial No.
-----
 1   48  C6800IA 48GE                               C6800IA-48TD         FOC1737W0PF
 2   48  C6800IA 48GE POE                          C6800IA-48FPD        FOC1736Z036
 3   48  C6800IA 48GE                               C6800IA-48TD         FOC1737W0NP
 4   48  C6800IA 48GE POE                          C6800IA-48FPD        FOC1741S58N
 5   48  C6800IA 48GE POE                          C6800IA-48FPD        FOC1736Z03L
```

Identify IA Client Modules by Beacon LED

```
Router(config)# hw-module fex <> slot <> led beacon
```

This example shows how to activate the beacon LED on IA client 118, module1:

```
Router(config)# hw-module fex 118 module 1 led beacon
```

Easy FEX

Easy FEX allows interface naming of FEX host port interfaces. To create interface naming, perform this task:

Command	Purpose
Router(config-if)# alias <i>string</i>	Configures alias under native FEX interface sub-mode.
Router(config-fex)# alias <i>string</i>	Configures a common alias name for all the FEX interfaces under a particular FEX-id.
Router(config)# interface <i>alias string</i>	Configures interface with the help of alias set.

This example shows all aliases using ‘all’ keyword:

```
fex-vss# show interfaces alias all
Interface Name                               Alias
-----
TenGigabitEthernet1/4/5                     test
GigabitEthernet101/1/0/1                    my_fex_loc_A/1/0/1
GigabitEthernet101/1/0/2                    my_fex_loc_A/1/0/2
GigabitEthernet101/1/0/3                    my_fex_loc_A/1/0/3
GigabitEthernet101/1/0/4                    my_fex_loc_A/1/0/4
GigabitEthernet101/1/0/5                    my_fex_loc_A/1/0/5
```


Renumbering FEX Switch-ID

The renumbering of IA clients can be managed using switch-id allocation from controller, after stack boot up. Also, a priority can be assigned to the FEX members.

The following conditions must exist for successful execution of FEX switch-id allocation:

- For renumbering the source slot should be online and the target slot should be offline.
- If the source slot FEX type is different than target slot FEX type, you would lose the interface configurations if you proceed with renumbering.
- Same target slot cannot be used for multiple source slot renumbering.
- You can enter multiple renumbering entries along with different swapping scenarios.
- When priority is modified for a member IA, the whole stack will reload.

To renumber FEX switch-id, perform this task:

	Command	Purpose
Step 1	Router# module provision update fex <i>fex_id</i>	Enters into switch renumber sub-mode.
Step 2	Router(exec-fex-update)# renumber <i>source_slot</i> to <i>target_slot</i>	Renumbers <i>source_slot</i> to <i>target_slot</i>
Step 3	Router(exec-fex-update)# priority <i>source_slot</i> value <i>number</i>	Assigns the mentioned priority <i>number</i> to the <i>source_slot</i> .
Step 4	Router(exec-fex-update)# commit	Commits all entries entered under <i>exec-fex-update</i> sub-mode.

To renumber FEX switch-id when scale is set to maximum FEX slots, perform this task:

	Command	Purpose
Step 5	Router# module provision update fex <i>fex_id</i> temp-vslot-allow enable	When maximum vslots are already allocated, temp-vslot-allow will enable the temporary vslot module to come online.
Step 1	Router# module provision update fex <i>fex_id</i>	Enters into switch renumber sub-mode.
Step 2	Router(exec-fex-update)# renumber <i>source_slot</i> to <i>target_slot</i>	Renumbers <i>source_slot</i> to <i>target_slot</i>
Step 3	Router(exec-fex-update)# priority <i>source_slot</i> value <i>number</i>	Assigns the mentioned priority <i>number</i> to the <i>source_slot</i> .
Step 4	Router(exec-fex-update)# commit	Commits all entries entered under <i>exec-fex-update</i> sub-mode.
Step 5	Router# module provision update fex <i>fex_id</i> temp-vslot-allow disable	Disables the temporary vslot module to go offline.

This example shows how to identify when a particular temporary FEX vslot is online:

```
TED# show fex system platform usage
FEX id usage details
  Fex-ids inuse: 150
```

```

Fex-ids online: 150
Total      Used      Free
-----
42         1         41
FEX slot usage details
FEX-id     Switch-id   Vslot   Pslot   Status
-----
150        3         51      2      In-use
150        3         52      3      In-use
150        3         53      5      Reserved
150        3         55      4      In-use
150        3         92      1      Temp-In-use
Total      Used      Reserved  Temp-Use/Free  Free
-----
47         3         1         1/4         42
Current Temp vslot allowed FEXs: 150
Current Temp VSLot usage:
-----
FEX 150 module 1

FEX ports usage details
FEX-id     Switch-id   Ports
-----
150        3         192
Total      Used      Free
-----
2016       192     1824

Stack members usage details
FEX-id     Switch-id   Used      Free

```

This example verifies the active entries under sub-mode and also the FEX IDs on which "temp-vslot-allow" is enabled.

```

Router(exec-fex-update)#show
Current module renumber mappings for FEX 101
-----
renumber 1 to 2
Current module Priority mappings for FEX 101
-----
priority 1 value 15
Temp vslots allowed:YES
Current Temp vslot allowed FEXs:101
TERM-VSS-FEX(exec-fex-update)#

```

Configure IA Clients

The configuration for IA clients can be entered on the IA parent before or after the IA clients are connected. IA client 10-Gigabit Ethernet ports require no configuration. IA client Gigabit Ethernet ports use this format:

gigabitethernet/fex_number/access_switch_number/0/port_number

- *fex_number*—The IA client FEX number:
 - Maximum of 42 IA FEX numbers.
 - The valid value range is 101–199.
- *access_switch_number*—The access switch number:
 - The valid value range is 1-5.
 - Multiple-switch stacks assign incrementing switch numbers to themselves.
 - See the [“Identify Connected IA Client Stack Modules”](#) section on page 12-9.
- The third interface parameter is always zero.
- The *port_number* valid value range is 1–48.



Note

- IA client configuration does not persist if the access switch number changes.
- The interface-range configuration mode supports IA clients ports (see [“How to Configure a Range of Interfaces”](#) section on page 10-2)

Display or Clear SDP and SRP Traffic

To display the counters that record the SDP packet traffic on IA client 118, enter the following command:

```
Router# show fex 118 protocol | incl SDP
      130 SDP pkts sent
      129 SDP pkts received
      130 SDP pkts sent
      129 SDP pkts received
```



Note

The command displays a sent and received value for each link in the IA channel group.

To clear the protocol counters, enter the **clear fex fex_number {sdp | srp}** command.

Upgrade FEX using In-Service Software Upgrade (ISSU) Process

To upgrade the software image of the FEX hosts using ISSU process, perform the following steps:

	Command	Purpose
Step 1	Router# issu loadversion device:filename	Initiates the ISSU process.
Step 2	Router# issu runversion	Forces a switchover.

	Command	Purpose
Step 3	Router# issu acceptversion	Stops the rollback timer and ensures that the software upgrade is not automatically aborted during the ISSU process.
Step 4	Router# issu runversion fex < <i>fex-id</i> <i>all</i> >	Forces a switchover of the active to the standby processor and starts the image download and upgrade procedure on the FEX (6800IA). The FEX initiates the image download from the new software bundle of the Sup2T. If you are using FEX stacks, the stack master initiates the image download for its members.
Step 5	Router# issu commitversion	Upgrades the VSS standby chassis and completes the ISSU process.
Step 6	Router# show issu state detail	Displays the state and the current version of the RPs after the ISSU process.

Configure Optional Parameters for an IA Client

- [Enter the IA Client Configuration Mode, page 12-14](#)
- [Configure a Description, page 12-14](#)
- [Configure the Custom Location Type Feature, page 12-15](#)
- [Configure MTU, page 12-15](#)

Enter the IA Client Configuration Mode

To enter the IA client configuration mode, perform this task:

Command	Purpose
Router(config)# fex <i>fex_number</i>	Enters IA client configuration mode. Note Sets the IA client description to FEX <i>fex_number</i> .

Configure a Description

To configure a description for the IA client or for each module in the IA client stack, perform this task:

Command	Purpose
Router(config-fex)# [module <i>module_number</i>] description <i>description_string</i>	Configures a description for the IA FEX number or for a module in the IA client stack.

Configure the Custom Location Type Feature

You can configure the custom location type feature for the IA client in IA client configuration mode. See these publications for information about the **location** command:

<http://www.cisco.com/en/US/docs/ios-xml/ios/cether/command/ce-cr-book.html>

<http://www.cisco.com/en/US/docs/ios-xml/ios/cether/command/ce-e1.html>



Note

The **location** commands support the optional **fex-location** keyword for IA clients.

Configure MTU

You can configure MTU on the IA FEX using the **mtu** command in fex config mode. In an IA client stack, the configured MTU value is applied to all the host members in the stack.

To configure MTU for an IA client, perform this task:

Command	Purpose
Router# configure terminal	Enters the global configuration mode.
Router(config)# fex 110	Enters IA client configuration mode.
Router(config-fex)# mtu 2000	Resets the MTU value for the FEX system. All the FEX ports will be set to use the new MTU value. The default MTU value for a FEX host port is 9216. Note For Cisco Catalyst C6840 series switch, the maximum supported MTU for native interface is 9154. The default MTU value is 1500. If the MTU of peer (Catalyst 6880 or Sup2T) is configured higher than 9154, the packet will not be processed. When using existing configuration from a Catalyst 6880 or Sup2T, ensure that the MTU is reset or removed.
Router# end	Exits the global configuration mode. Note Reload the FEX for the MTU change to take effect.

If you want to avoid reloading the FEX, you can configure an explicit connectionless network service (CLNS) MTU size on the IA client and peer ports as shown in the following example:

```
Router# configure terminal
Router(config)# interface interface Gig118/1/0/1
Router(config-if)# ip router isis
Router(config-if)# clns mtu 1497
```


13

Configuring EnergyWise

EnergyWise is a Cisco energy management architecture that provides a common approach to configuring, reporting and managing the power consumed by Cisco switches and attached devices. With Cisco EnergyWise, power can be managed on a network, sub-network or network element level.

Release	Feature Modification
12.2(33)SX14	EnergyWise Phase 2 introduced on Catalyst 6500 Series switches
15.0(1)SY1	EnergyWise Phase 2.6 introduced on Catalyst 6500 Series switches
15.1(1)SY1	EnergyWise Phase 2.7 introduced on Catalyst 6500 Series switches

For hardware compatibility matrices, new feature information, and to understand EnergyWise release numbering, see the Cisco EnergyWise release notes:

http://www.cisco.com/en/US/products/ps10195/prod_release_notes_list.html

EnergyWise configuration guides and EnergyWise orchestrator configuration guides are located at the following URL:

http://www.cisco.com/en/US/products/ps10195/products_installation_and_configuration_guides_list.html

Additional Cisco EnergyWise documents, such as White Papers, Data Sheets, FAQs, and are located at the following URL:

<http://www.cisco.com/en/US/products/ps10195/>

14

Power Management

- [Power Management Overview, page 14-1](#)
- [How to Enable or Disable Power Redundancy, page 14-2](#)
- [How to Power Modules Off and On, page 14-3](#)
- [How to Display System Power Status, page 14-3](#)
- [How to Power Cycle Modules, page 14-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Power Management Overview

In systems with redundant power supplies, both power supplies must be of the same wattage. The Catalyst 6500 series switches allow you to use both AC-input and DC-input power supplies in the same chassis. For detailed information on supported power supply configurations, see the *Catalyst 6500 Series Switch Installation Guide*.

The modules have different power requirements, and some configurations require more power than a single power supply can provide. The power management feature allows you to power all installed modules with two power supplies. However, redundancy is not supported in this configuration because the total power drawn from both power supplies is at no time greater than the capability of one supply. Redundant and nonredundant power configurations are described in the following sections.

How to Enable or Disable Power Redundancy

To disable or enable redundancy (redundancy is enabled by default) from global configuration mode, enter the **power redundancy-mode combined | redundant** commands. You can change the configuration of the power supplies to redundant or nonredundant at any time.

To disable redundancy, use the **combined** keyword. In a nonredundant configuration, the power available to the system is the combined power capability of both power supplies. The system powers up as many modules as the combined capacity allows. However, if one power supply fails and there is not enough power for all of the previously powered-up modules, the system powers down those modules.

To enable redundancy, use the **redundant** keyword. In a redundant configuration, the total power drawn from both power supplies is not greater than the capability of one power supply. If one supply malfunctions, the other supply can take over the entire system load. When you install and power up two power supplies, each concurrently provides approximately half of the required power to the system. Load sharing and redundancy are enabled automatically; no software configuration is required.

To view the current state of modules and the total power available for modules, enter the **show power** command (see the “[How to Display System Power Status](#)” section on page 14-3).

Table 14-1 describes how the system responds to changes in the power supply configuration.

Table 14-1 Effects of Power Supply Configuration Changes

Configuration Change	Effect
Redundant to nonredundant	- System log and syslog messages are generated. - System power is increased to the combined power capability of both power supplies. - Modules marked <i>power-deny</i> in the show power oper state field are brought up if there is sufficient power.
Nonredundant to redundant (both power supplies must be of equal wattage)	- System log and syslog messages are generated. - System power is decreased to the power capability of one supply. - If there is not enough power for all previously powered-up modules, some modules are powered down and marked as <i>power-deny</i> in the show power oper state field.
Equal wattage power supply is inserted with redundancy enabled	- System log and syslog messages are generated. - System power equals the power capability of one supply. - No change in module status because the power capability is unchanged.
Equal wattage power supply is inserted with redundancy disabled	- System log and syslog messages are generated. - System power is increased to the combined power capability of both power supplies. - Modules marked <i>power-deny</i> in the show power oper state field are brought up if there is sufficient power.
Higher or lower wattage power supply is inserted with redundancy enabled	- System log and syslog messages are generated. - The system will operate in a non-redundant combined mode.
Higher or lower wattage power supply is inserted with redundancy disabled	- System log and syslog messages are generated. - System power is increased to the combined power capability of both power supplies. - Modules marked <i>power-deny</i> in the show power oper state field are brought up if there is sufficient power.

Table 14-1 *Effects of Power Supply Configuration Changes (continued)*

Configuration Change	Effect
Power supply is removed with redundancy enabled	- System log and syslog messages are generated. - No change in module status because the power capability is unchanged.
Power supply is removed with redundancy disabled	- System log and syslog messages are generated. - System power is decreased to the power capability of one supply. - If there is not enough power for all previously powered-up modules, some modules are powered down and marked as <i>power-deny</i> in the show power oper state field.
System is booted with power supplies of different wattage installed and redundancy enabled	- System log and syslog messages are generated. - The system does not allow you to have power supplies of different wattage installed in a redundant configuration. The lower wattage supply shuts down.
System is booted with power supplies of equal or different wattage installed and redundancy disabled	- System log and syslog messages are generated. - System power equals the combined power capability of both power supplies. - The system powers up as many modules as the combined capacity allows.

How to Power Modules Off and On

To power modules off and on from the CLI, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# power enable module <i>slot_number</i>	Powers a module on.
Step 3	Router(config)# no power enable module <i>slot_number</i>	Powers a module off.



Note

When you enter the **no power enable module** *slot* command to power down a module, the module's configuration is not saved.

This example shows how to power on the module in slot 3:

```
Router# configure terminal
Router(config)# power enable module 3
```

How to Display System Power Status

The **show power** command displays the current power status of system components:

```
Router# show power
system power redundancy mode = combined
system power redundancy operationally = combined(2+0)
system power total = 5699.72 Watts (109.61 Amps @ 52V)
system power used = 3930.16 Watts (75.58 Amps @ 52V)
system power available = 1769.56 Watts (34.03 Amps @ 52V)
Power-Capacity PS-Fan Output Oper
```

```

PS Type Watts A @52V Status Status State
-----
1 C6800-XL-3KW-AC 2999.88 57.69 OK OK on
2 C6800-XL-3KW-AC 2999.88 57.69 OK OK on
3 none
4 none
Pwr-Allocated Oper
Fan Type Watts A @52V State
-----
1 C6807-XL-FAN 260.00 5.00 OK
Pwr-Allocated Oper
PwrCon Type Watts A @52V State
-----
1 C6800-XL-PS-CONV 12.48 0.24 OK
2 C6800-XL-PS-CONV 12.48 0.24 OK

Pwr-Requested Pwr-Allocated Admin Oper
Slot Card-Type Watts A @52V Watts A @52V State State
-----
1 C6800-32P10G-XL 587.60 11.30 587.60 11.30 on on
2 C6800-32P10G 587.60 11.30 587.60 11.30 on on
3 C6800-SUP6T-XL 353.60 6.80 353.60 6.80 on on
4 C6800-SUP6T-XL 353.60 6.80 353.60 6.80 on on
5 C6800-32P10G-XL 587.60 11.30 587.60 11.30 on on
6 C6800-32P10G 587.60 11.30 587.60 11.30 on on
7 C6800-32P10G-XL 587.60 11.30 587.60 11.30 on on
system auxiliary power mode = off
system auxiliary power redundancy operationally = non-redundant
system primary connector power limit = 9100.00 Watts (175.00 Amps @ 52V)
system auxiliary connector power limit = 13000.00 Watts (250.00 Amps @ 52V)
system primary power used = 3930.16 Watts (75.58 Amps @ 52V)
system auxiliary power used = 0 Watt

Router#

```

The **show power** command displays the current power status of a specific power supply:

```

Router# show power status power-supply 1
Power-Capacity PS-Fan Output Oper
PS Type Watts A @52V Status Status State
-----
1 C6800-XL-3KW-AC 2999.88 57.69 OK OK on
Router#

```

You can display power supply input fields by specifying the power supply number in the command. A new power-output field with operating mode is displayed for power supplies with more than one output mode. Enter the **show environment status power-supply** command as follows:

```

Router# show environment status power-supply 1
switch 1 power-supply 1:
power-supply 1 fan-fail: OK
power-supply 1 power-input: AC high
power-supply 1 power-output-mode: high
power-supply 1 incompatible with fan: OK
power-supply 1 power-output-fail: OK
Router# show environment status power-supply 2
switch 1 power-supply 2:
power-supply 2 fan-fail: OK
power-supply 2 power-input: AC high
power-supply 2 power-output-mode: high
power-supply 2 incompatible with fan: OK
power-supply 2 power-output-fail: OK

```

How to Power Cycle Modules

You can power cycle (reset) a line card module by entering the **hw-module module {mod #} reset** command. The module powers off for 5 seconds, and then powers on.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

15

Environmental Monitoring

- [Environmental Monitoring Overview, page 15-1](#)
- [How to Determine Sensor Temperature Thresholds, page 15-1](#)
- [How to Monitor the System Environmental Status, page 15-5](#)
- [Information About LED Environmental Indications, page 15-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Environmental Monitoring Overview

Environmental monitoring of chassis components provides early-warning indications of possible component failures, which ensures a safe and reliable system operation and avoids network interruptions. This section describes the monitoring of these critical system components, which allows you to identify and rapidly correct hardware-related problems in your system.

How to Determine Sensor Temperature Thresholds

The system sensors set off alarms based on different temperature threshold settings. Use the **show environment alarm threshold** command to display the sensor temperature thresholds:

```
Router> show environment alarm threshold
environmental alarm thresholds:

power-supply 1 fan-fail: OK
threshold #1 for power-supply 1 fan-fail:
(sensor value != FALSE ) is system minor alarm
```

```

power-supply 1 power-output-fail: OK
  threshold #1 for power-supply 1 power-output-fail:
    (sensor value != FALSE ) is system minor alarm
power-supply 1 power-output-mode: high
  threshold #1 for power-supply 1 power-output-mode:
    (sensor value < 1 ) is system minor alarm
power-supply 1 incompatible with fan: OK
  threshold #1 for power-supply 1 incompatible with fan:
    (sensor value != FALSE ) is system minor alarm
power-supply 2 fan-fail: OK
  threshold #1 for power-supply 2 fan-fail:
    (sensor value != FALSE ) is system minor alarm
power-supply 2 power-output-fail: OK
  threshold #1 for power-supply 2 power-output-fail:
    (sensor value != FALSE ) is system minor alarm
power-supply 2 power-output-mode: high
  threshold #1 for power-supply 2 power-output-mode:
    (sensor value < 1 ) is system minor alarm
power-supply 2 incompatible with fan: OK
  threshold #1 for power-supply 2 incompatible with fan:
    (sensor value != FALSE ) is system minor alarm
fan-tray 1 fan-fail: OK
  threshold #1 for fan-tray 1 fan-fail:
    (sensor value == 1 ) is system minor alarm
operating clock count: 2
  threshold #1 for operating clock count:
    (sensor value < 2 ) is system minor alarm
  threshold #2 for operating clock count:
    (sensor value < 1 ) is system major alarm
operating VTT count: 3
  threshold #1 for operating VTT count:
    (sensor value < 3 ) is system minor alarm
  threshold #2 for operating VTT count:
    (sensor value < 2 ) is system major alarm
operating fan count: 1
  threshold #1 for operating fan count:
    (sensor value < 1 ) is system minor alarm
VTT 1 outlet temperature: 37C
  threshold #1 for VTT 1 outlet temperature:
    (sensor value >= 100C) is system minor alarm
  threshold #2 for VTT 1 outlet temperature:
    (sensor value >= 115C) is system major alarm
VTT 1 OK: OK
  threshold #1 for VTT 1 OK:
    (sensor value != FALSE ) is system minor alarm
VTT 2 outlet temperature: 39C
  threshold #1 for VTT 2 outlet temperature:
    (sensor value >= 100C) is system minor alarm
  threshold #2 for VTT 2 outlet temperature:
    (sensor value >= 115C) is system major alarm
VTT 2 OK: OK
  threshold #1 for VTT 2 OK:
    (sensor value != FALSE ) is system minor alarm
VTT 3 outlet temperature: 42C
  threshold #1 for VTT 3 outlet temperature:
    (sensor value >= 100C) is system minor alarm
  threshold #2 for VTT 3 outlet temperature:
    (sensor value >= 115C) is system major alarm
VTT 3 OK: OK
  threshold #1 for VTT 3 OK:
    (sensor value != FALSE ) is system minor alarm
clock 1 OK: OK
  threshold #1 for clock 1 OK:
    (sensor value != FALSE ) is system minor alarm

```



```

clock 2 OK: OK
  threshold #1 for clock 2 OK:
    (sensor value != FALSE ) is system minor alarm
module 1 power-output-fail: OK
  threshold #1 for module 1 power-output-fail:
    (sensor value != FALSE ) is system minor alarm
module 1 outlet temperature: 56C
  threshold #1 for module 1 outlet temperature:
    (sensor value >= 75C) is system minor alarm
  threshold #2 for module 1 outlet temperature:
    (sensor value >= 90C) is system major alarm
module 1 inlet temperature: 41C
  threshold #1 for module 1 inlet temperature:
    (sensor value >= 50C) is system minor alarm
  threshold #2 for module 1 inlet temperature:
    (sensor value >= 75C) is system major alarm
module 1 device-1 temperature: 63C
  threshold #1 for module 1 device-1 temperature:
    (sensor value >= 75C) is system minor alarm
  threshold #2 for module 1 device-1 temperature:
    (sensor value >= 90C) is system major alarm
module 1 device-2 temperature: 45C
  threshold #1 for module 1 device-2 temperature:
    (sensor value >= 50C) is system minor alarm
  threshold #2 for module 1 device-2 temperature:
    (sensor value >= 75C) is system major alarm
module 1 aux-1 temperature: 57C
  threshold #1 for module 1 aux-1 temperature:
    (sensor value >= 75C) is system minor alarm
  threshold #2 for module 1 aux-1 temperature:
    (sensor value >= 90C) is system major alarm
module 1 aux-2 temperature: 46C
  threshold #1 for module 1 aux-2 temperature:
    (sensor value >= 50C) is system minor alarm
  threshold #2 for module 1 aux-2 temperature:
    (sensor value >= 75C) is system major alarm
module 1 insufficient cooling: OK
  threshold #1 for module 1 insufficient cooling:
    (sensor value != FALSE ) is system major alarm
module 1 asic-1 temperature: 61C
  threshold #1 for module 1 asic-1 temperature:
    (sensor value >= 110C) is system minor alarm
  threshold #2 for module 1 asic-1 temperature:
    (sensor value >= 125C) is system major alarm
module 1 asic-2 temperature: 61C
  threshold #1 for module 1 asic-2 temperature:
    (sensor value >= 110C) is system minor alarm
  threshold #2 for module 1 asic-2 temperature:
    (sensor value >= 125C) is system major alarm
module 1 asic-3 temperature: 61C
  threshold #1 for module 1 asic-3 temperature:
    (sensor value >= 110C) is system minor alarm
  threshold #2 for module 1 asic-3 temperature:
    (sensor value >= 125C) is system major alarm
module 1 asic-4 temperature: 63C
  threshold #1 for module 1 asic-4 temperature:
    (sensor value >= 110C) is system minor alarm
  threshold #2 for module 1 asic-4 temperature:
    (sensor value >= 125C) is system major alarm
module 3 power-output-fail: OK
  threshold #1 for module 3 power-output-fail:
    (sensor value != FALSE ) is system minor alarm
module 3 module outlet-1 temperature: 39C
  threshold #1 for module 3 module outlet-1 temperature:

```

```

(sensor value >= 70C) is system minor alarm
threshold #2 for module 3 module outlet-1 temperature:
(sensor value >= 85C) is system major alarm
module 3 module inlet-1 temperature: 37C
threshold #1 for module 3 module inlet-1 temperature:
(sensor value >= 55C) is system minor alarm
threshold #2 for module 3 module inlet-1 temperature:
(sensor value >= 70C) is system major alarm
module 3 module outlet-2 temperature: 50C
threshold #1 for module 3 module outlet-2 temperature:
(sensor value >= 70C) is system minor alarm
threshold #2 for module 3 module outlet-2 temperature:
(sensor value >= 85C) is system major alarm
module 3 module inlet-2 temperature: 43C
threshold #1 for module 3 module inlet-2 temperature:
(sensor value >= 55C) is system minor alarm
threshold #2 for module 3 module inlet-2 temperature:
(sensor value >= 70C) is system major alarm
module 3 insufficient cooling: OK
threshold #1 for module 3 insufficient cooling:
(sensor value != FALSE ) is system major alarm
module 3 fan-upgrade required: OK
threshold #1 for module 3 fan-upgrade required:
(sensor value != FALSE ) is system major alarm
module 3 asic-1 temperature: 58C
threshold #1 for module 3 asic-1 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 3 asic-1 temperature:
(sensor value >= 125C) is system major alarm
module 3 asic-2 temperature: 56C
threshold #1 for module 3 asic-2 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 3 asic-2 temperature:
(sensor value >= 125C) is system major alarm
module 3 asic-3 temperature: 55C
threshold #1 for module 3 asic-3 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 3 asic-3 temperature:
(sensor value >= 125C) is system major alarm
module 1 EARL 1 device-1 temperature: 82C
threshold #1 for module 1 EARL 1 device-1 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 1 EARL 1 device-1 temperature:
(sensor value >= 125C) is system major alarm
module 1 EARL 1 device-2 temperature: 81C
threshold #1 for module 1 EARL 1 device-2 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 1 EARL 1 device-2 temperature:
(sensor value >= 125C) is system major alarm
module 1 EARL 2 device-1 temperature: 81C
threshold #1 for module 1 EARL 2 device-1 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 1 EARL 2 device-1 temperature:
(sensor value >= 125C) is system major alarm
module 1 EARL 2 device-2 temperature: 80C
threshold #1 for module 1 EARL 2 device-2 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 1 EARL 2 device-2 temperature:
(sensor value >= 125C) is system major alarm
module 3 EARL outlet-1 temperature: 73C
threshold #1 for module 3 EARL outlet-1 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 3 EARL outlet-1 temperature:
(sensor value >= 125C) is system major alarm

```

```

module 3 EARL inlet-1 temperature: 73C
threshold #1 for module 3 EARL inlet-1 temperature:
(sensor value >= 110C) is system minor alarm
threshold #2 for module 3 EARL inlet-1 temperature:
(sensor value >= 125C) is system major alarm
power-converter 1 : 3.3V Output: OK
threshold #1 for power-converter 1 : 3.3V Output:
(sensor value != 1 ) is system minor alarm
power-converter 2 : 3.3V Output: OK
threshold #1 for power-converter 2 : 3.3V Output:
(sensor value != 1 ) is system minor alarm

```

How to Monitor the System Environmental Status

To display system status information, enter the **show environment [alarm | cooling | status | temperature]** command. The keywords display the following information:

- **alarm**—Displays environmental alarms.
 - **status**—Displays alarm status.
 - **thresholds**—Displays alarm thresholds.
- **cooling**—Displays fan tray status, chassis cooling capacity, ambient temperature, and per-slot cooling capacity.
- **status**—Displays field-replaceable unit (FRU) operational status and power and temperature information.
- **temperature**—Displays FRU temperature information.

To view the system status information, enter the **show environment** command:

```

Router# show environment
environmental alarms:
no alarms

Router# show environment alarm
environmental alarms:
no alarms

Router# show environment cooling
fan-tray 1:
fan-tray 1 type: C6807-XL-FAN
fan-tray 1 mode: High-power
fan-tray 1 fan-fail: OK
chassis per slot cooling capacity: 97 cfm
module 1 cooling requirement: 84 cfm
module 3 cooling requirement: 84 cfm

Router# show environment status
switch 1 backplane:
operating clock count: 2
operating VTT count: 3
operating fan count: 1

switch 1 fan-tray 1:
switch 1 fan-tray 1 type: C6807-XL-FAN
switch 1 fan-tray 1 mode: High-power
switch 1 fan-tray 1 fan-fail: OK
switch 1 VTT 1:
switch 1 VTT 1 OK: OK
switch 1 VTT 1 outlet temperature: 36C

```

```

switch 1 VTT 2:
  switch 1 VTT 2 OK: OK
  switch 1 VTT 2 outlet temperature: 38C
switch 1 VTT 3:
  switch 1 VTT 3 OK: OK
  switch 1 VTT 3 outlet temperature: 41C
switch 1 clock 1:
  switch 1 clock 1 OK: OK, switch 1 clock 1 clock-inuse: in-use
switch 1 clock 2:
  switch 1 clock 2 OK: OK, switch 1 clock 2 clock-inuse: not-in-use
switch 1 power-supply 1:
  switch 1 power-supply 1 fan-fail: OK
  switch 1 power-supply 1 power-input: AC high
  switch 1 power-supply 1 power-output-mode: high
  switch 1 power-supply 1 incompatible with fan: OK
  switch 1 power-supply 1 power-output-fail: OK

switch 1 power-supply 2:
  switch 1 power-supply 2 fan-fail: OK
  switch 1 power-supply 2 power-input: AC high
  switch 1 power-supply 2 power-output-mode: high
  switch 1 power-supply 2 incompatible with fan: OK
  switch 1 power-supply 2 power-output-fail: OK

switch 1 power-converter 1:
  switch 1 power-converter 1 : 3.3V Output: OK
switch 1 power-converter 2:
  switch 1 power-converter 2 : 3.3V Output: OK
switch 1 module 1:
  switch 1 module 1 power-output-fail: OK
  switch 1 module 1 outlet temperature: 53C
  switch 1 module 1 inlet temperature: 40C
  switch 1 module 1 device-1 temperature: 60C
  switch 1 module 1 device-2 temperature: 45C
  switch 1 module 1 aux-1 temperature: 55C
  switch 1 module 1 aux-2 temperature: 45C
  switch 1 module 1 asic-1 temperature: 58C
  switch 1 module 1 asic-2 temperature: 59C
  switch 1 module 1 asic-3 temperature: 60C
  switch 1 module 1 asic-4 temperature: 61C
  switch 1 module 1 EARL 1 device-1 temperature: 78C
  switch 1 module 1 EARL 1 device-2 temperature: 77C
  switch 1 module 1 EARL 2 device-1 temperature: 77C
  switch 1 module 1 EARL 2 device-2 temperature: 77C
switch 1 module 3:
  switch 1 module 3 power-output-fail: OK
  switch 1 module 3 module outlet-1 temperature: 38C
  switch 1 module 3 module inlet-1 temperature: 36C
  switch 1 module 3 module outlet-2 temperature: 49C
  switch 1 module 3 module inlet-2 temperature: 43C
  switch 1 module 3 asic-1 temperature: 56C
  switch 1 module 3 asic-2 temperature: 56C
  switch 1 module 3 asic-3 temperature: 54C
  switch 1 module 3 EARL outlet-1 temperature: 72C
  switch 1 module 3 EARL inlet-1 temperature: 71C

```

Information About LED Environmental Indications

The LEDs can indicate two alarm types: major and minor. Major alarms indicate a critical problem that could lead to the system being shut down. Minor alarms are for informational purposes only, giving you notice of a problem that could turn critical if corrective action is not taken.

When the system has an alarm (major or minor), that indicates an overtemperature condition, the alarm is not canceled nor is any action taken (such as module reset or shutdown) for 5 minutes. If the temperature falls 5°C (41°F) below the alarm threshold during this period, the alarm is canceled.

[Table 15-1](#) lists the environmental indicators for the supervisor engine and switching modules.

**Note**

See the *Catalyst 6500 Series Switch Module Installation Guide* for additional information on LEDs, including the supervisor engine SYSTEM LED.

Table 15-1 *Environmental Monitoring for Supervisor Engine and Switching Modules*

Component	Alarm Type	LED Indication	Action
Supervisor engine temperature sensor exceeds major threshold	Major	STATUS LED red	Generates syslog message and an SNMP trap. If there is a redundancy situation, the system switches to a redundant supervisor engine and the active supervisor engine shuts down. If there is no redundancy situation and the overtemperature condition is not corrected, the system shuts down after 5 minutes.
Note - Temperature sensors monitor key supervisor engine components including daughter cards. - A STATUS LED is located on the supervisor engine front panel and all module front panels. - The STATUS LED is red on the failed supervisor engine. If there is no redundant supervisor, the SYSTEM LED is red also.			
Supervisor engine temperature sensor exceeds minor threshold	Minor	STATUS LED orange	Generates syslog message and an SNMP trap. Monitors the condition.
Redundant supervisor engine temperature sensor exceeds major or minor threshold	Major	STATUS LED red	Generates syslog message and an SNMP trap. If a major alarm is generated and the overtemperature condition is not corrected, the system shuts down after 5 minutes.
	Minor	STATUS LED orange	Monitors the condition if a minor alarm is generated.
Switching module temperature sensor exceeds major threshold	Major	STATUS LED red	Generates syslog message and SNMP. Powers down the module (see the “ How to Power Modules Off and On ” section on page 14-3 for instructions).
Switching module temperature sensor exceeds minor threshold	Minor	STATUS LED orange	Generates syslog message and an SNMP trap. Monitors the condition.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

16

Online Diagnostics

- [Prerequisites for Online Diagnostics, page 16-1](#)
- [Restrictions for Online Diagnostics, page 16-1](#)
- [Information About Online Diagnostics, page 16-2](#)
- [Default Settings for Online Diagnostics, page 16-2](#)
- [How to Configure Online Diagnostics, page 16-2](#)
- [How to Run Online Diagnostic Tests, page 16-7](#)
- [How to Perform Memory Tests, page 16-29](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Online Diagnostics

None.

Restrictions for Online Diagnostics

None.

Information About Online Diagnostics

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 control signals. Disruptive online diagnostic tests, such as the built-in self-test (BIST) and the disruptive loopback test, and nondisruptive online diagnostic tests, such as packet switching, run during bootup, module online insertion and removal (OIR), and system reset. The nondisruptive online diagnostic tests run as part of background health monitoring. Either disruptive or nondisruptive tests can be run at the user's request (on-demand).

The online diagnostics detect problems in the following areas:

- Hardware components
- Interfaces (GBICs, Ethernet ports, and so forth)
- Connectors (loose connectors, bent pins, and so forth)
- Solder joints
- Memory (failure over time)

Online diagnostics is one of the requirements for the high availability feature. High availability is a set of quality standards that seek to limit the impact of equipment failures on the network. A key part of high availability is detecting hardware failures and taking corrective action while the switch runs in a live network. Online diagnostics in high availability detect hardware failures and provide feedback to high availability software components to make switchover decisions.

Online diagnostics are categorized as bootup, on-demand, schedule, or health-monitoring diagnostics. Bootup diagnostics run during bootup; on-demand diagnostics run from the CLI; schedule diagnostics run at user-designated intervals or specified times when the switch is connected to a live network; and health-monitoring runs in the background.

Default Settings for Online Diagnostics

See the default information for each test in [Appendix 1, “Online Diagnostic Tests.”](#)

How to Configure Online Diagnostics

- [Setting Bootup Online Diagnostics Level, page 16-2](#)
- [Configuring On-Demand Online Diagnostics, page 16-4](#)
- [Scheduling Online Diagnostics, page 16-6](#)

Setting Bootup Online Diagnostics Level

You can set the bootup diagnostics level as minimal or complete or you can bypass the bootup diagnostics entirely. Enter the **complete** keyword to run all diagnostic tests; enter the **minimal** keyword to run only EARL tests and loopback tests for all ports in the switch. Enter the **no** form of the command to bypass all diagnostic tests. The default bootup diagnostics level is minimal.

To set the bootup diagnostic level, perform this task:

Command	Purpose
Router(config)# diagnostic bootup level {minimal complete}	Sets the bootup diagnostic level.

This example shows how to set the bootup online diagnostic level:

```
Router(config)# diagnostic bootup level complete
Router(config)#
```

This example shows how to display the bootup online diagnostic level:

```
Router(config)# show diagnostic bootup level
Current bootup diagnostic level: complete

Router(config)#
```

Combined Legacy Diagnostic Tests

The following legacy diagnostic tests have been combined:

TestAclRedirectRBacl

The legacy tests **TestRBacl** and **TestAclRedirect** have been merged into this test.

The combined test will verify both the ACL redirect feature and the Role based ACL feature of the Layer 3 forwarding engine.

TestQosDQUP

The legacy tests **TestQos** and **TestDQUP** have been merged into this test.

The combined test will verify both the QoS input/output TCAM functionality and that DQUP/PUP packets can be generated when diagnostic packets hit QoS entry.

TestIngressEgressSpan

The legacy tests **TestIngressSpan** and **TestEgressSpan** have been merged into this test.

This combined test will ensure the port ASIC will be able to tag packets for ingress span and verify the egress span replication functionality of the ingress module replication engine.

New Diagnostic Tests

TestEobcStressPing

This test is designed to stress test the EOBC link of a line card with the supervisor.

In this test a large number (20,000) of ICMP packets are sent from the Supervisor to the line card along with varying the size of the ICMP packet. The test is declared pass when all the ICMP packets are responded to within the timeout (2 seconds).

The test can be run on-demand or scheduled as it is a disruptive test. It is possible the test will fail if the EOBC was already under heavy load and all packets are not responded to. No corrective action is needed.

TestVSActiveToStandbyLoopback

This test is a VSS system wide test and only permitted in on-demand mode.

During this test, the standby VSS supervisor's uplink port is configured in loopback mode and a VLAN flood packet is sent from the Active VSS's inband port to the standby VSS's uplink port.

This is a disruptive test that can be run on-demand only. It doesn't require a VSL link from Active to Standby supervisor. The uplink port cannot be specified. The test cannot be run on the Active Supervisor.

TestVslBridgeLink

This test verifies the bridge link data path between two fabric channel and rewrite engine based data path complexes for VSL configurable Supervisor or line card.

The packet is sent across the bridge link from the inband port on one data path complex to a good loopback port on another data path complex.

TestVslLocalLoopback

This test verifies the hardware functionality of each port on a VSL card before any VSL member link in the VSL port-channel group is up.

A diagnostic packet is sent from the local inband port of the VSL card to each port to run a loopback test.

TestVslStatus

This test reports the status change of the link detected by VSLP protocol.

If any link problem is detected by VSLP, the link status and result are updated. The loopback test to check the hardware status, as requested by VSLP is also triggered.

TestFabricLinkStatus

This test monitors the FEX fabric link on the FEX controller.

It will report failure if the fabric link port has lost connection to the associated FEX.

Configuring On-Demand Online Diagnostics

You can run the on-demand online diagnostic tests from the CLI. You can set the execution action to either stop or continue the test when a failure is detected or to stop the test after a specific number of failures occur by using the failure count setting. You can configure a test to run multiple times using the iteration setting.

You should run packet-switching tests before memory tests.



Note

Do not use the **diagnostic start all** command until all of the following steps are completed.

Because some on-demand online diagnostic tests can affect the outcome of other tests, you should perform the tests in the following order:

1. Run the nondisruptive tests.
2. Run all tests in the relevant functional area.
3. Run the TestTrafficStress test.
4. Run the TestEobcStressPing test.
5. Run the exhaustive-memory tests.

To run on-demand online diagnostic tests, perform this task:

Step 1 Run the nondisruptive tests.

To display the available tests and their attributes, and determine which commands are in the nondisruptive category, enter the **show diagnostic content** command.

Step 2 Run all tests in the relevant functional area.

Packet-switching tests fall into specific functional areas. When a problem is suspected in a particular functional area, run all tests in that functional area. If you are unsure about which functional area you need to test, or if you want to run all available tests, enter the **complete** keyword.

Step 3 Run the TestTrafficStress test.

This is a disruptive packet-switching test. This test switches packets between pairs of ports at line rate for the purpose of stress testing. During this test all of the ports are shut down, and you may see link flaps. The link flaps will recover after the test is complete. The test takes several minutes to complete.

Disable all health-monitoring tests before running this test by using the **no diagnostic monitor module number test all** command.

Step 4 Run the TestEobcStressPing test.

This is a disruptive test and tests the Ethernet over backplane channel (EOBC) connection for the module. The test takes several minutes to complete. You cannot run any of the packet-switching tests described in previous steps after running this test. However, you can run tests described in subsequent steps after running this test.

Disable all health-monitoring tests before running this test by using the **no diagnostic monitor module number test all** command. The EOBC connection is disrupted during this test and will cause the health-monitoring tests to fail and take recovery action.

Step 5 Run the exhaustive-memory tests.

Before running the exhaustive-memory tests, all health-monitoring tests should be disabled because the tests will fail with health monitoring enabled and the switch will take recovery action. Disable the health-monitoring diagnostic tests by using the **no diagnostic monitor module number test all** command.

Perform the exhaustive-memory tests in the following order:

1. TestFibTcamSSRAM
2. TestAclQosTcam
3. TestNetFlowTcam
4. TestAsicMemory
5. TestAsicMemory

You must reboot the after running the exhaustive-memory tests before it is operational again. You cannot run any other tests on the switch after running the exhaustive-memory tests. Do not save the configuration when rebooting as it will have changed during the tests. After the reboot, reenable the health-monitoring tests using the **diagnostic monitor module number test all** command.

To set the bootup diagnostic level, perform this task:

Command	Purpose
Router# diagnostic ondemand {iteration iteration_count} {action-on-error {continue stop} [error_count]}	Configures on-demand diagnostic tests to run, how many times to run (iterations), and what action to take when errors are found.

This example shows how to set the on-demand testing iteration count:

```
Router# diagnostic ondemand iteration 3
Router#
```

This example shows how to set the execution action when an error is detected:

```
Router# diagnostic ondemand action-on-error continue 2
Router#
```

Scheduling Online Diagnostics

You can schedule online diagnostics to run at a designated time of day or on a daily, weekly, or monthly basis. You can schedule tests to run only once or to repeat at an interval. Use the **no** form of this command to remove the scheduling.

To schedule online diagnostics, perform this task:

Command	Purpose
Router(config)# diagnostic schedule module <i>number</i> test { <i>test_id</i> <i>test_id_range</i> all } [port { <i>num</i> <i>num_range</i> all }] { on <i>mm dd yyyy hh:mm</i> } { daily <i>hh:mm</i> } { weekly <i>day_of_week hh:mm</i> }	Schedules on-demand diagnostic tests on the specified module for a specific date and time, how many times to run (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 port on module 1:

```
Router(config)# diagnostic schedule module 1 test 1,2,5-9 port 3 on january 3 2003 23:32
Router(config)#
```

This example shows how to schedule diagnostic testing to occur daily at a certain time for a specific port:

```
Router(config)# diagnostic schedule module 1 test 1,2,5-9 port 3 daily 12:34
Router(config)#
```

This example shows how to schedule diagnostic testing to occur weekly on a certain day for a specific port:

```
Router(config)# diagnostic schedule module 1 test 1,2,5-9 port 3 weekly friday 09:23
Router(config)#
```

Configuring Health-Monitoring Diagnostics

You can configure health-monitoring diagnostic testing while the switch is connected to a live network. You can configure the execution interval for each health-monitoring test, the generation of a system message upon test failure, or the enabling or disabling an individual test. Use the **no** form of this command to disable testing.

To configure health-monitoring diagnostic testing, perform this task:

	Command	Purpose
Step 1	Router(config)# diagnostic monitor interval module <i>number test {test_id test_id_range all} [hour</i> <i>hh] [min mm] [second ss] [millisec ms] [day day]</i>	Configures the health-monitoring interval of the specified tests. The no form of this command will change the interval to the default interval, or zero.
Step 2	Router(config)# [no] diagnostic monitor module <i>number test {test_id test_id_range all}</i>	Enables or disables health-monitoring diagnostic tests.

This example shows how to configure the specified test to run every two minutes on module 1:

```
Router(config)# diagnostic monitor interval module 1 test 1 min 2
Router(config)#
```

This example shows how to run the test if health monitoring has not previously been enabled:

```
Router(config)# diagnostic monitor module 1 test 1
```

This example shows how to enable the generation of a syslog message when any health-monitoring test fails:

```
Router(config)# diagnostic monitor syslog
Router(config)#
```

How to Run Online Diagnostic Tests

- [Overview of Diagnostic Test Operation, page 16-8](#)
- [Starting and Stopping Online Diagnostic Tests, page 16-8](#)
- [Running All Online Diagnostic Tests, page 16-9](#)
- [Displaying Online Diagnostic Tests and Test Results, page 16-9](#)

Overview of Diagnostic Test Operation

After you configure online diagnostics, you can start or stop diagnostic tests or display the test results. You can also see which tests are configured and what diagnostic tests have already run.

- Enable the logging console/monitor to see all warning messages before you enable any online diagnostics tests.
- When you are running disruptive tests, run the tests when connected through the console. When disruptive tests are complete, a warning message on the console recommends that you reload the system to return to normal operation. Strictly follow this warning.
- While tests are running, all ports are shut down because a stress test is being performed with ports configured to loop internally; external traffic might alter the test results. The switch must be rebooted to bring the switch to normal operation. When you issue the command to reload the switch, the system will ask you if the configuration should be saved. Do not save the configuration.
- If you are running the tests on a supervisor engine, after the test is initiated and complete, you must reload or power down and then power up the entire system.
- If you are running the tests on a switching module, rather than the supervisor engine, after the test is initiated and complete, you must reset the switching module.

Starting and Stopping Online Diagnostic Tests

After you configure diagnostic tests to run, you can use the **start** and **stop** to begin or end a diagnostic test. To start or stop an online diagnostic command, perform one of these tasks:

Command	Purpose
Router# diagnostic start module number test {test_id test_id_range minimal complete basic per-port non-disruptive all} [port {num port#_range all}]	Starts a diagnostic test on a port or range of ports on the specified module.
Router# diagnostic stop module number	Stops a diagnostic test on the specified module.

This example shows how to start a diagnostic test on module 1:

```
Router# diagnostic start module 1 test 5
Module 1:Running test(s) 5 may disrupt normal system operation
Do you want to run disruptive tests? [no]yes
00:48:14:Running OnDemand Diagnostics [Iteration #1] ...
00:48:14:%DIAG-SP-6-TEST_RUNNING:Module 1:Running TestNewLearn{ID=5} ...
00:48:14:%DIAG-SP-6-TEST_OK:Module 1:TestNewLearn{ID=5} has completed successfully
00:48:14:Running OnDemand Diagnostics [Iteration #2] ...
00:48:14:%DIAG-SP-6-TEST_RUNNING:Module 1:Running TestNewLearn{ID=5} ...
00:48:14:%DIAG-SP-6-TEST_OK:Module 1:TestNewLearn{ID=5} has completed successfully
Router#
```

This example shows how to stop a diagnostic test:

```
Router# diagnostic stop module 1
Router#
```

Running All Online Diagnostic Tests

You can run all diagnostic tests, disruptive and nondisruptive, at once with a single command. In this case, all test dependencies will be handled automatically.



Note

- Running all online diagnostic tests will disrupt normal system operation. Reset the system after the **diagnostic start system test all** command has completed.
- Do not insert, remove, or power down modules or the supervisor while the system test is running.
- Do not issue any diagnostic command other than the **diagnostic stop system test all** command while the system test is running.
- Make sure no traffic is running in background.

To start or stop all online diagnostic tests, perform one of these tasks:

Command	Purpose
Router# diagnostic start system test all	Executes all online diagnostic tests.
Router# diagnostic stop system test all	Stops the execution of all online diagnostic tests.

This example shows how to start all online diagnostic tests:

```
Router# diagnostic start system test all
*****
* WARNING:
* 'diagnostic start system test all' will disrupt normal system
* operation. The system requires RESET after the command
* 'diagnostic start system test all' has completed prior to
* normal use.
*
* IMPORTANT:
* 1. DO NOT INSERT, OIR, or POWER DOWN Linecards or
* Supervisor while system test is running.
*
* 2. DO NOT ISSUE ANY DIAGNOSTIC COMMAND except
* "diagnostic stop system test all" while system test
* is running.
*
* 3. PLEASE MAKE SURE no traffic is running in background.
*****
Do you want to continue? [no]:
```

Displaying Online Diagnostic Tests and Test Results

You can display the online diagnostic tests that are configured and check the results of the tests using the following **show** commands:

- **show diagnostic content**
- **show diagnostic health**

To display the diagnostic tests that are configured, perform this task:

Command	Purpose
show diagnostic { bootup level content [module num] events [module num] [event-type <i>event-type</i>] health ondemand settings result [module num] [detail] schedule [module num] }	Displays the test results of online diagnostics and lists supported test suites.

This example shows how to display the online diagnostics that are configured on switch 1 module 5:

```
Router# show diagnostic content switch 1 module 5
```

```
Switch 1 Module 5: Supervisor Engine 6T 10GE w/ CTS (Active)
```

Diagnostics test suite attributes:

```

M/C/* - Minimal bootup level test / Complete bootup level test / NA
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/* - Power-down line cards and need reload supervisor / NA
K/* - Require resetting the line card after the test has completed / NA
T/* - Shut down all ports and need reload supervisor / NA
L/* - FEX switch will reload after test list completion / NA

```

ID	Test Name	Attributes	Test Interval day hh:mm:ss.ms	Thre- shold
1)	TestTransceiverIntegrity	****PD*X**I****	not configured	n/a
2)	TestLoopback	M****PD*X**I****	not configured	n/a
3)	TestFexModeLoopback	M****PD*X**I****	not configured	n/a
4)	TestVSActiveToStandbyLoopback	***DSX**I****	not configured	n/a
5)	TestL2CTSLoopback	M****PD*X**I****	not configured	n/a
6)	TestL3CTSLoopback	M****PD*X**I****	not configured	n/a
7)	TestScratchRegister	***N****I****	000 00:00:30.00	5
8)	TestNewIndexLearn	M***N****I****	000 00:00:15.00	10
9)	TestDontConditionalLearn	M***N****I****	000 00:00:15.00	10
10)	TestBpduTrap	M****D*X**I****	not configured	n/a
11)	TestMatchCapture	M****D*X**I****	not configured	n/a
12)	TestProtocolMatchChannel	M****D*X**I****	not configured	n/a
13)	TestPortSecurity	M****D*X**I****	not configured	n/a
14)	TestIPv4FibShortcut	M***N****I****	000 00:00:15.00	10
15)	TestL3Capture2	M****D*X**I****	not configured	n/a
16)	TestIPv6FibShortcut	M***N****I****	000 00:00:15.00	10
17)	TestMPLSFibShortcut	M***N****I****	000 00:00:15.00	10
18)	TestNATFibShortcut	M***N****I****	000 00:00:15.00	10
19)	TestAclPermit	M***N****I****	000 00:00:15.00	10
20)	TestAclDeny	M****D*X**I****	not configured	n/a
21)	TestAclRedirectRBacl	M***N****I****	not configured	n/a
22)	TestQosDQUP	M****D*X**I****	not configured	n/a
23)	TestL3VlanMet	M****D*X**I****	not configured	n/a
24)	TestIngressEgressSpan	M****D*X**I****	not configured	n/a
25)	TestNetflowShortcut	M****D*X**I****	not configured	n/a
26)	TestInbandEdit	M****D*X**I****	not configured	n/a
27)	TestFabricVlanLoopback	M****D*X**I****	not configured	n/a
28)	TestTrafficStress	***D*X**I**T*	not configured	n/a


```

29) TestL3TcamMonitoring -----> ***N***I*** 000 00:00:15.00 10
30) TestFibTcam -----> ***D*X**IR*** not configured n/a
31) TestAclQosTcam -----> ***D*X**IR*** not configured n/a
32) TestEarlMemOnBootup -----> M**N*X**I*** not configured n/a
33) TestAsicMemory -----> ***D*X**IR*** not configured n/a
34) TestEobcStressPing -----> ***D*X**I*** not configured n/a
35) TestFirmwareDiagStatus -----> M**N***I*** 000 00:00:15.00 10
36) TestAsicSync -----> ***N***I*** 000 00:00:15.00 10
37) TestUnusedPortLoopback -----> **PN***I*** 000 00:01:00.00 10
38) TestNonDisruptiveLoopback -----> **PN***I*** 000 00:00:10.00 10
39) TestFabricFlowControlStatus -----> ***N***I*** 000 00:00:15.00 10
40) TestPortTxMonitoring -----> **PN***I*** 000 00:01:15.00 5
41) TestOBFL -----> M**N***I*** 000 00:00:15.00 10
42) TestCFRW -----> M*VN*X**I*** not configured n/a
43) TestLtlFpoeMemoryConsistency -----> ***N***I*** 000 00:00:30.00 1
44) TestRwEngineOverSubscription -----> ***N***I*** 000 00:00:01.00 10
45) TestErrorCounterMonitor -----> ***N***I*** 000 00:00:30.00 10
46) TestVslBridgeLink -----> M**D*X**I*** not configured n/a
47) TestVslLocalLoopback -----> M*PD*X**I*** not configured n/a
48) TestVslStatus -----> M*PD*X**I*** not configured n/a
49) TestEARLInternalTables -----> ***N***I*** 000 00:05:00.00 1
50) TestNVRAMBatteryMonitor -----> ***N**F*I*** 000 01:00:00.00 72
51) TestMediaLoopback -----> **PD*X**I*** not configured n/a
52) TestMicroburst -----> **PN***I*** 000 00:00:05.00 10
53) TestFEXFabricLinkStatus -----> **PN***I*** 000 00:00:03.00 3

```

Router#

This example shows how to display the online diagnostic results for module 6:

```
Router# show diagnostic result switch 1 module 5
```

```
Current bootup diagnostic level: complete
```

```
Switch 1 Module 5: Supervisor Engine 6T 10GE w/ CTS (Active) SerialNo : SAL1927JELC
```

```
Overall Diagnostic Result for Switch 1 Module 5 : PASS
```

```
Diagnostic level at card bootup: complete
```

```
Test results: (. = Pass, F = Fail, U = Untested)
```

```
1) TestTransceiverIntegrity:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
U U U U . . . U . . U U U U U U U U U U
```

```
2) TestLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
. . . . . U U . . . U U U U U U U U U U
```

```
3) TestFexModeLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
. . . . . U U . U U U U U U U U U U U U
```

```
4) TestVSActiveToStandbyLoopback ---> U
```

```
5) TestL2CTSLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
. . . . . U U . . . U U U U U U U U U U
```

```
6) TestL3CTSLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
. . . . . U U . . . U U U U U U U U U U
```

```
7) TestScratchRegister -----> U
```

```
8) TestNewIndexLearn -----> .
```

```
9) TestDontConditionalLearn -----> .
```

```
10) TestBpduTrap -----> .
```

```
11) TestMatchCapture -----> .
```

```
12) TestProtocolMatchChannel -----> U
```

```
13) TestPortSecurity -----> .
```

```
14) TestIPv4FibShortcut -----> .
```

```
15) TestL3Capture2 -----> .
```

```
16) TestIPv6FibShortcut -----> .
```

```
17) TestMPLSFibShortcut -----> .
```

```
18) TestNATFibShortcut -----> .
```

```
19) TestAclPermit -----> .
```

```
20) TestAclDeny -----> .
```

```
21) TestAclRedirectRBacl -----> .
```

```

22) TestQosDQUP -----> .
23) TestL3VlanMet -----> .
24) TestIngressEgressSpan -----> .
25) TestNetflowShortcut -----> .
26) TestInbandEdit -----> .
27) TestFabricVlanLoopback -----> .
28) TestTrafficStress -----> U
29) TestL3TcamMonitoring -----> U
30) TestFibTcam -----> U
31) TestAclQosTcam -----> U
32) TestEarlMemOnBootup -----> .
33) TestAsicMemory -----> U
34) TestEobcStressPing -----> U
35) TestFirmwareDiagStatus -----> .
36) TestAsicSync -----> U
37) TestUnusedPortLoopback:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
U U U U U U U U U U U U U U U U U U U U

38) TestNonDisruptiveLoopback:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
U U U U U U U U U U U U U U U U U U U U

39) TestFabricFlowControlStatus -----> U
40) TestPortTxMonitoring:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
U U U U U U U U U U U U U U U U U U U U

41) TestOBFL -----> .
42) TestCFRW:

Device 1
-----
U

43) TestLtlFpoeMemoryConsistency ----> U
44) TestRwEngineOverSubscription ----> U
45) TestErrorCounterMonitor -----> U
46) TestVslBridgeLink -----> .
47) TestVslLocalLoopback:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
. . . . . U U U U U U U U U U

48) TestVslStatus:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
U U U U U . . U U U U U U U U U U U U U U

49) TestEARLInternalTables -----> U

```

```

50) TestNVRAMBatteryMonitor -----> U
51) TestMediaLoopback:

```

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----

```

```

U U U U U U U U U U U U U U U U U U U U U

```

```

52) TestMicroburst:

```

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----

```

```

U U U U U U U U U U U U U U U U U U U U U

```

```

53) TestFEXFabricLinkStatus:

```

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----

```

```

U U U U U U U U U U U U U U U U U U U U U

```

```

Router#

```

This example shows how to display the detailed online diagnostic results for module 6:

```
Router# show diagnostic result switch 1 module 5 detail
```

```
Current bootup diagnostic level: complete
```

```
Switch 1 Module 5: Supervisor Engine 6T 10GE w/ CTS (Active) SerialNo : SAL1927JELC
```

```
Overall Diagnostic Result for Switch 1 Module 5 : PASS
```

```
Diagnostic level at card bootup: complete
```

```
Test results: (. = Pass, F = Fail, U = Untested)
```

```
1) TestTransceiverIntegrity:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
U U U U . . . U . . U U U U U U U U U U
```

```
Error code -----> 0 (DIAG_SUCCESS)
```

```
Total run count -----> 5
```

```
Last test testing type -----> OnDemand
```

```
Last test execution time ----> Apr 07 2016 10:39:09
```

```
First test failure time -----> n/a
```

```
Last test failure time -----> n/a
```

```
Last test pass time -----> Apr 07 2016 10:39:09
```

```
Total failure count -----> 0
```

```
Consecutive failure count ----> 0
```

```
2) TestLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
. . . . . U U . . . U U U U U U U U U U
```

```
Error code -----> 0 (DIAG_SUCCESS)
```

```
Total run count -----> 1
```

```
Last test testing type -----> Bootup
```

```
Last test execution time ----> Apr 07 2016 10:38:42
```

```
First test failure time -----> n/a
```

```
Last test failure time -----> n/a
```

```
Last test pass time -----> Apr 07 2016 10:38:42
```

```
Total failure count -----> 0
```

```
Consecutive failure count ----> 0
```

```
3) TestFexModeLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----  
. . . . . U U . U U U U U U U U U U U
```

```
Error code -----> 0 (DIAG_SUCCESS)
```

```
Total run count -----> 1
```

```
Last test testing type -----> Bootup
```

```
Last test execution time ----> Apr 07 2016 10:38:47
```

```
First test failure time -----> n/a
```

```
Last test failure time -----> n/a
```

```

Last test pass time -----> Apr 07 2016 10:38:47
Total failure count -----> 0
Consecutive failure count ---> 0

```

4) TestVSActiveToStandbyLoopback ---> U

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

5) TestL2CTSLoopback:

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
. . . . . U U . . . U U U U U U U U U U

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:49
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:49
Total failure count -----> 0
Consecutive failure count ---> 0

```

6) TestL3CTSLoopback:

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
. . . . . U U . . . U U U U U U U U U U

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:52
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:52
Total failure count -----> 0
Consecutive failure count ---> 0

```

7) TestScratchRegister -----> U

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0

```

Consecutive failure count ---> 0

8) TestNewIndexLearn -----> .

Error code -----> 0 (DIAG_SUCCESS)
 Total run count -----> 1
 Last test testing type -----> Bootup
 Last test execution time -----> Apr 07 2016 10:38:56
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> Apr 07 2016 10:38:56
 Total failure count -----> 0
 Consecutive failure count ---> 0

9) TestDontConditionalLearn -----> .

Error code -----> 0 (DIAG_SUCCESS)
 Total run count -----> 1
 Last test testing type -----> Bootup
 Last test execution time -----> Apr 07 2016 10:38:56
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> Apr 07 2016 10:38:56
 Total failure count -----> 0
 Consecutive failure count ---> 0

10) TestBpduTrap -----> .

Error code -----> 0 (DIAG_SUCCESS)
 Total run count -----> 1
 Last test testing type -----> Bootup
 Last test execution time -----> Apr 07 2016 10:38:56
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> Apr 07 2016 10:38:56
 Total failure count -----> 0
 Consecutive failure count ---> 0

11) TestMatchCapture -----> .

Error code -----> 0 (DIAG_SUCCESS)
 Total run count -----> 1
 Last test testing type -----> Bootup
 Last test execution time -----> Apr 07 2016 10:38:56
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> Apr 07 2016 10:38:56
 Total failure count -----> 0
 Consecutive failure count ---> 0

12) TestProtocolMatchChannel -----> U

Error code -----> 358 (DIAG_SKIPPED_PORT_LOOPBACK_NOT_ALLOWED)
 Total run count -----> 0
 Last test testing type -----> n/a
 Last test execution time -----> n/a
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> n/a

```
Total failure count -----> 0
Consecutive failure count ---> 0
```

```
13) TestPortSecurity -----> .
```

```
Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ---> 0
```

```
14) TestIPv4FibShortcut -----> .
```

```
Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ---> 0
```

```
15) TestL3Capture2 -----> .
```

```
Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ---> 0
```

```
16) TestIPv6FibShortcut -----> .
```

```
Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ---> 0
```

```
17) TestMPLSFibShortcut -----> .
```

```
Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time -----> n/a
Last test failure time -----> n/a
```



```

Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ----> 0

```

```
18) TestNATFibShortcut -----> .
```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time ----> n/a
Last test failure time ----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ----> 0

```

```
19) TestAclPermit -----> .
```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time ----> n/a
Last test failure time ----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ----> 0

```

```
20) TestAclDeny -----> .
```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time ----> n/a
Last test failure time ----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ----> 0

```

```
21) TestAclRedirectRBacl -----> .
```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time ----> n/a
Last test failure time ----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ----> 0

```

```
22) TestQosDQUP -----> .
```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time ----> n/a

```

```

Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

23) TestL3VlanMet -----> .

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:57
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:57
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

24) TestIngressEgressSpan -----> .

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:58
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:58
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

25) TestNetflowShortcut -----> .

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

26) TestInbandEdit -----> .

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

27) TestFabricVlanLoopback -----> .

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59

```

```

First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

28) TestTrafficStress -----> U

```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

29) TestL3TcamMonitoring -----> U

```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

30) TestFibTcam -----> U

```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

31) TestAclQosTcam -----> U

```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```

32) TestEarlMemOnBootup -----> .

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup

```

```

Last test execution time ----> Apr 07 2016 10:38:59
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

```
33) TestAsicMemory -----> U
```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```
34) TestEobcStressPing -----> U
```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```
35) TestFirmwareDiagStatus -----> .
```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

```
36) TestAsicSync -----> U
```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

```
37) TestUnusedPortLoopback:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----
```

U U U U U U U U U U U U U U U U U

```
Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time -----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ----> 0
```

38) TestNonDisruptiveLoopback:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20

U U U U U U U U U U U U U U U U U

```
Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time -----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ----> 0
```

39) TestFabricFlowControlStatus -----> U

```
Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time -----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ----> 0
```

40) TestPortTxMonitoring:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20

U U U U U U U U U U U U U U U U U

```
Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time -----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ----> 0
```

41) TestOBFL -----> .

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

42) TestCFRW:

```

Device 1
-----
U

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ---> 0

```

43) TestLtlFpoeMemoryConsistency ----> U

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

44) TestRwEngineOverSubscription ----> U

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

45) TestErrorCounterMonitor -----> U

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a

```

```

Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ----> 0
Error Records -----> n/a

```

46) TestVslBridgeLink -----> .

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:38:59
First test failure time ----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:38:59
Total failure count -----> 0
Consecutive failure count ----> 0

```

47) TestVslLocalLoopback:

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
. . . . . U U U U U U U U U

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 1
Last test testing type -----> Bootup
Last test execution time ----> Apr 07 2016 10:39:00
First test failure time ----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 10:39:00
Total failure count -----> 0
Consecutive failure count ----> 0

```

48) TestVslStatus:

```

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
-----
U U U U U . . U U U U U U U U U U U

```

```

Error code -----> 0 (DIAG_SUCCESS)
Total run count -----> 2
Last test testing type -----> Health Monitoring
Last test execution time ----> Apr 07 2016 12:06:26
First test failure time ----> n/a
Last test failure time -----> n/a
Last test pass time -----> Apr 07 2016 12:06:26
Total failure count -----> 0
Consecutive failure count ----> 0

```

49) TestEARLInternalTables -----> U

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time ----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a

```

```

Total failure count -----> 0
Consecutive failure count ---> 0

AGE GROUP
Total CC run count -----> 0
Table size -----> 0
Total fail count -----> 0
Total correction count -----> 0
Last completion time -----> Apr 07 2016 10:33:49
Last result -----> UNKNOWN
Last fail count -----> 0
Last correction count -----> 0
Last entries checked -----> 0
Consistency checker -----> OFF

```

```

BUNDLE PORT MAP
Total CC run count -----> 0
Table size -----> 0
Total fail count -----> 0
Total correction count -----> 0
Last completion time -----> Apr 07 2016 10:33:49
Last result -----> UNKNOWN
Last fail count -----> 0
Last correction count -----> 0
Last entries checked -----> 0
Consistency checker -----> OFF

```

```

BUNDLE EXTENSION MAP
Total CC run count -----> 0
Table size -----> 0
Total fail count -----> 0
Total correction count -----> 0
Last completion time -----> Apr 07 2016 10:33:49
Last result -----> UNKNOWN
Last fail count -----> 0
Last correction count -----> 0
Last entries checked -----> 0
Consistency checker -----> OFF

```

```

VLAN ACCESS MODE MEMORY
Total CC run count -----> 0
Table size -----> 0
Total fail count -----> 0
Total correction count -----> 0
Last completion time -----> Apr 07 2016 10:33:49
Last result -----> UNKNOWN
Last fail count -----> 0
Last correction count -----> 0
Last entries checked -----> 0
Consistency checker -----> OFF

```

```

50) TestNVRAMBatteryMonitor -----> U

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time ----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ---> 0

```

51) TestMediaLoopback:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20

 U U U U U U U U U U U U U U U U U U

Error code -----> 3 (DIAG_SKIPPED)
 Total run count -----> 0
 Last test testing type -----> n/a
 Last test execution time -----> n/a
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> n/a
 Total failure count -----> 0
 Consecutive failure count ----> 0

52) TestMicroburst:

Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20

 U U U U U U U U U U U U U U U U U U

Error code -----> 3 (DIAG_SKIPPED)
 Total run count -----> 0
 Last test testing type -----> n/a
 Last test execution time -----> n/a
 First test failure time -----> n/a
 Last test failure time -----> n/a
 Last test pass time -----> n/a
 Total failure count -----> 0
 Consecutive failure count ----> 0

Consecutive microbursts count ----> 0
 Last microburst event time -----> n/a
 Last microburst ports -----> n/a
 Port 1:
 Ingress Drop Count -----> n/a
 Egress Drop Count -----> n/a

Port 2:
 Ingress Drop Count -----> n/a
 Egress Drop Count -----> n/a

Port 3:
 Ingress Drop Count -----> n/a
 Egress Drop Count -----> n/a

Port 4:
 Ingress Drop Count -----> n/a
 Egress Drop Count -----> n/a

Port 5:
 Ingress Drop Count -----> n/a
 Egress Drop Count -----> n/a

Port 6:
 Ingress Drop Count -----> n/a
 Egress Drop Count -----> n/a

Port 7:

```
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 8:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 9:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 10:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 11:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 12:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 13:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 14:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 15:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 16:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 17:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 18:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 19:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
Port 20:
Ingress Drop Count -----> n/a
Egress Drop Count -----> n/a
```

```
53) TestFEXFabricLinkStatus:
```

```
Port 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
-----
U U U U U U U U U U U U U U U U U U
```

```

Error code -----> 3 (DIAG_SKIPPED)
Total run count -----> 0
Last test testing type -----> n/a
Last test execution time -----> n/a
First test failure time -----> n/a
Last test failure time -----> n/a
Last test pass time -----> n/a
Total failure count -----> 0
Consecutive failure count ----> 0

```

Router#

This example shows how to display the output for the health checks performed:

```

Router# show diagnostic health
Non-zero port counters for 6/4 -
13.                               linkChange = 8530

Non-zero port counters for 6/5 -
13.                               linkChange = 8530

Router#

```

How to Perform Memory Tests

Most online diagnostic tests do not need any special setup or configuration. However, the memory tests, which include the TestFibTcamSSRAM and TestLinecardMemory tests, have some required tasks and some recommended tasks that you should complete before running them.

Before you run any of the online diagnostic memory tests, perform the following tasks:

- Required tasks
 - Isolate network traffic by disabling all connected ports.
 - Do not send test packets during a memory test.
 - Reset the system before returning the system to normal operating mode.
- Turn off all background health-monitoring tests using the **no diagnostic monitor module number test all** command.

How to Perform a Diagnostic Sanity Check

You can run the diagnostic sanity check in order to see potential problem areas in your network. The sanity check runs a set of predetermined checks on the configuration with a possible combination of certain system states to compile a list of warning conditions. The checks are designed to look for anything that seems out of place and are intended to serve as an aid for maintaining the system sanity.

To run the diagnostic sanity check, perform this task:

Command	Purpose
show diagnostic sanity	Runs a set of tests on the configuration and certain system states.

This example displays samples of the messages that could be displayed with the **show diagnostic sanity** command:

```
Router# show diagnostic sanity
Pinging default gateway 10.6.141.1 ....
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.6.141.1, timeout is 2 seconds:
..!!!.
Success rate is 0 percent (0/5)

IGMP snooping disabled please enable it for optimum config.

IGMP snooping disabled but RGMP enabled on the following interfaces,
please enable IGMP for proper config :
Vlan1, Vlan2, GigabitEthernet1/1

Multicast routing is enabled globally but not enabled on the following
interfaces:
GigabitEthernet1/1, GigabitEthernet1/2

A programming algorithm mismatch was found on the device bootflash:
Formatting the device is recommended.

The bootflash: does not have enough free space to accomodate the crashinfo file.

Please check your confreg value : 0x0.

Please check your confreg value on standby: 0x0.

The boot string is empty. Please enter a valid boot string .
Could not verify boot image "disk0:" specified in the boot string on the
slave.

Invalid boot image "bootflash:asdasd" specified in the boot string on the
slave.

Please check your boot string on the slave.

UDLD has been disabled globally - port-level UDLD sanity checks are
being bypassed.
OR
[
The following ports have UDLD disabled. Please enable UDLD for optimum
config:
Gi1/22

The following ports have an unknown UDLD link state. Please enable UDLD
on both sides of the link:
Gi1/22
]

The following ports have portfast enabled:
Gi1/20, Gi1/22

The following ports have trunk mode set to on:
Gi1/1, Gi1/13

The following trunks have mode set to auto:
Gi1/2, Gi1/3

The following ports with mode set to desirable are not trunking:
Gi1/3, Gi1/4
```

The following trunk ports have negotiated to half-duplex:
Gi1/3, Gi1/4

The following ports are configured for channel mode on:
Gi1/1, Gi1/2, Gi1/3, Gi1/4

The following ports, not channeling are configured for channel mode
desirable:
Gi1/14

The following vlan(s) have a spanning tree root of 32768:
1

The following vlan(s) have max age on the spanning tree root different from
the default:
1-2

The following vlan(s) have forward delay on the spanning tree root different
from the default:
1-2

The following vlan(s) have hello time on the spanning tree root different
from the default:
1-2

The following vlan(s) have max age on the bridge different from the
default:
1-2

The following vlan(s) have fwd delay on the bridge different from the
default:
1-2

The following vlan(s) have hello time on the bridge different from the
default:
1-2

The following vlan(s) have a different port priority than the default
on the port gigabitEthernet1/1
1-2

The following ports have receive flow control disabled:
Gi1/20, Gi1/22

The following inline power ports have power-deny/faulty status:
Gi1/1, Gi1/2

The following ports have negotiated to half-duplex:
Gi1/22

The following vlans have a duplex mismatch:
Gig 1/22

The following interfaces have a native vlan mismatch:
interface (native vlan - neighbor vlan)
Gig 1/22 (1 - 64)

The value for Community-Access on read-only operations for SNMP is the same
as default. Please verify that this is the best value from a security point
of view.

The value for Community-Access on write-only operations for SNMP is the same
as default. Please verify that this is the best value from a security point
of view.

The value for Community-Access on read-write operations for SNMP is the same
as default. Please verify that this is the best value from a security point
of view.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

17

Onboard Failure Logging (OBFL)

- [Prerequisites for OBFL, page 17-1](#)
- [Restrictions for OBFL, page 17-2](#)
- [Information About OBFL, page 17-2](#)
- [Default Settings for OBFL, page 17-9](#)
- [Enabling OBFL, page 17-10](#)
- [Configuration Examples for OBFL, page 17-11](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for OBFL

None.

Restrictions for OBFL

- **Software Restrictions**—If a device (router or switch) intends to use *linear* flash memory as its OBFL storage media, Cisco IOS software must reserve a minimum of two physical sectors (or physical blocks) for the OBFL feature. Because an erase operation for a linear flash device is done on per-sector (or per-block) basis, one extra physical sector is needed. Otherwise, the minimum amount of space reserved for the OBFL feature on any device must be at least 8 KB.
- **Firmware Restrictions**—If a line card or port adapter runs an operating system or firmware that is different from the Cisco IOS operating system, the line card or port adapter must provide device driver level support or an interprocess communications (IPC) layer that allows the OBFL file system to communicate to the line card or port adapter. This requirement is enforced to allow OBFL data to be recorded on a storage device attached to the line card or port adapter.
- **Hardware Restrictions**—To support the OBFL feature, a device must have at least 8 KB of nonvolatile memory space reserved for OBFL data logging.

Information About OBFL

- [Overview of OBFL, page 17-2](#)
- [Information about Data Collected by OBFL, page 17-2](#)

Overview of OBFL

The Onboard Failure Logging (OBFL) feature collects data such as operating temperatures, hardware uptime, interrupts, and other important events and messages from system hardware installed in a Cisco router or switch. The data is stored in nonvolatile memory and helps technical personnel diagnose hardware problems.

Information about Data Collected by OBFL

- [OBFL Data Overview, page 17-2](#)
- [Temperature, page 17-3](#)
- [Operational Uptime, page 17-4](#)
- [Interrupts, page 17-8](#)
- [Message Logging, page 17-9](#)

OBFL Data Overview

The OBFL feature records operating temperatures, hardware uptime, interrupts, and other important events and messages that can assist with diagnosing problems with hardware cards (or *modules*) installed in a Cisco router or switch. Data is logged to files stored in nonvolatile memory. When the onboard hardware is started up, a first record is made for each area monitored and becomes a base value for subsequent records. The OBFL feature provides a circular updating scheme for collecting continuous records and archiving older (historical) records, ensuring accurate data about the system. Data is recorded in one of two formats: continuous information that displays a snapshot of measurements and

samples in a continuous file, and summary information that provides details about the data being collected. The data is displayed using the **show logging onboard** command. The message “No historical data to display” is seen when historical data is not available.

Temperature

Temperatures surrounding hardware modules can exceed recommended safe operating ranges and cause system problems such as packet drops. Higher than recommended operating temperatures can also accelerate component degradation and affect device reliability. Monitoring temperatures is important for maintaining environmental control and system reliability. Once a temperature sample is logged, the sample becomes the base value for the next record. From that point on, temperatures are recorded either when there are changes from the previous record or if the maximum storage time is exceeded. Temperatures are measured and recorded in degrees Celsius.

Temperature Example

```
Router# show logging onboard switch 1 module 3 tem
```

TEMPERATURE SUMMARY INFORMATION

Number of sensors : 9

Sampling frequency : 5 minutes

Maximum time of storage : 120 minutes

Sensor | ID | Maximum Temperature 0C

MB-Out1 1190201 39

MB-In1 1190202 38

MB-Out2 1190203 52

MB-In2 1190204 44

Aludra0 1340301 78

Aludra1 1340301 76

Kunlun0-3 1340301 58

Kunlun1-3 1340301 58

SM15 1340301 58

Temp Sensor ID

0C 1 2 3 4 5 6 7 8 9

No historical data to display

TEMPERATURE CONTINUOUS INFORMATION

Sensor	ID
MB-Out	980201
MB-In	980202
MB	980203
MB	980204
EARL-Out	910201
EARL-In	910202
SSA 1	980301
SSA 2	980302
JANUS 1	980303
JANUS 2	980304
GEMINI 1	980305
GEMINI 2	980306

Time Stamp	Sensor Temperature 0C											
MM/DD/YYYY HH:MM:SS	1	2	3	4	5	6	7	8	9	10	11	12
03/06/2007 22:32:51	31	26	27	27	NA	NA	33	32	30	29	NA	NA
03/06/2007 22:37:51	43	28	29	38	NA	NA	38	36	36	35	NA	NA

To interpret this data:

- Number of sensors is the total number of temperature sensors that will be recorded. A column for each sensor is displayed with temperatures listed under the number of each sensor, as available.
- Sampling frequency is the time between measurements.
- Maximum time of storage determines the maximum amount of time, in minutes, that can pass when the temperature remains unchanged and the data is not saved to storage media. After this time, a temperature record will be saved even if the temperature has not changed.
- The Sensor column lists the name of the sensor.
- The ID column lists an assigned identifier for the sensor.
- Maximum Temperature 0C shows the highest recorded temperature per sensor.
- Temp indicates a recorded temperature in degrees Celsius in the historical record. Columns following show the total time each sensor has recorded that temperature.
- Sensor ID is an assigned number, so that temperatures for the same sensor can be stored together.

Operational Uptime

The operational uptime tracking begins when the module is powered on, and information is retained for the life of the module.

Operational Uptime Example

Router#show logging onboard switch 1 module 5 up

```
-----
UPTIME SUMMARY INFORMATION
-----
```

First customer power on : 09/04/2015 15:14:59

Total uptime : 0 years 50 weeks 5 days 12 hours 0 minutes

Total downtime : 0 years 8 weeks 1 days 17 hours 13 minutes

Number of resets : 859

Number of slot changes : 6

Current reset reason : 0x44

Current reset timestamp : 04/07/2016 10:38:34

Current slot : 5

Current uptime : 0 years 0 weeks 0 days 13 hours 0 minutes

Reset | | Reset |

Code | Count | Reason |

0X44 211 Cold Power-on

0X45 172 Reload

0XA1 476 Unknown

Last 5 Boot-up Records

Record | Boot-up Time Stamp | Slot | Reset | Software | Firmware |

Number | MM/DD/YYYY HH:MM:SS | Number | Reason | Version | Version |

#1 04/07/2016 10:38:34 5 0X44 15.3(1)SY 15.3(04

#2 04/07/2016 07:32:56 5 0X44 15.3(1)SY 15.3(04

#3 04/07/2016 06:12:41 5 0X44 15.3(1)SY 15.3(04

#4 04/06/2016 20:54:26 5 0X44 15.3(1)SY 15.3(04

#5 04/06/2016 20:06:27 5 0X44 15.3(1)IE10 15.3(04

The operational uptime application tracks the following events:

- Date and time the customer first powered on a component.
- Total uptime and downtime for the component in years, weeks, days, hours, and minutes.
- Total number of component resets.
- Total number of slot (module) changes.
- Current reset timestamp to include the date and time.
- Current slot (module) number of the component.
- Current uptime in years, weeks, days, hours, and minutes.
- Reset reason; see [Table 17-1](#) to translate the numbers displayed.
- Count is the number of resets that have occurred for each reset reason.

Table 17-1 **Reset Reason Codes and Explanations**

Reset Reason Code (in hex)	Component/Explanation
0x01	Chassis on
0x02	Line card hot plug in
0x03	Supervisor requests line card off or on
0x04	Supervisor requests hard reset on line card
0x05	Line card requests Supervisor off or on
0x06	Line card requests hard reset on Supervisor
0x07	Line card self reset using the internal system register
0x08	—
0x09	—
0x0A	Momentary power interruption on the line card
0x0B	—
0x0C	—
0x0D	—
0x0E	—
0x0F	—
0x10	—
0x11	Off or on after Supervisor non-maskable interrupts (NMI)
0x12	Hard reset after Supervisor NMI
0x13	Soft reset after Supervisor NMI
0x14	—
0x15	Off or on after line card asks Supervisor NMI
0x16	Hard reset after line card asks Supervisor NMI
0x17	Soft reset after line card asks Supervisor NMI
0x18	—
0x19	Off or on after line card self NMI
0x1A	Hard reset after line card self NMI
0x1B	Soft reset after line card self NMI
0x21	Off or on after spurious NMI
0x22	Hard reset after spurious NMI
0x23	Soft reset after spurious NMI
0x24	—
0x25	Off or on after watchdog NMI
0x26	Hard reset after watchdog NMI
0x27	Soft reset after watchdog NMI
0x28	—

Table 17-1 **Reset Reason Codes and Explanations**

Reset Reason Code (in hex)	Component/Explanation
0x29	Off or on after parity NMI
0x2A	Hard reset after parity NMI
0x2B	Soft reset after parity NMI
0x31	Off or on after system fatal interrupt
0x32	Hard reset after system fatal interrupt
0x33	Soft reset after system fatal interrupt
0x34	—
0x35	Off or on after application-specific integrated circuit (ASIC) interrupt
0x36	Hard reset after ASIC interrupt
0x37	Soft reset after ASIC interrupt
0x38	—
0x39	Off or on after unknown interrupt
0x3A	Hard reset after unknown interrupt
0x3B	Soft reset after unknown interrupt
0x41	Off or on after CPU exception
0x42	Hard reset after CPU exception
0x43	Soft reset after CPU exception
0xA1	Reset data converted to generic data

Interrupts

Interrupts are generated by system components that require attention from the CPU such as ASICs and NMIs. Interrupts are generally related to hardware limit conditions or errors that need to be corrected.

The continuous format records each time a component is interrupted, and this record is stored and used as base information for subsequent records. Each time the list is saved, a timestamp is added. Time differences from the previous interrupt are counted, so that technical personnel can gain a complete record of the component's operational history when an error occurs.

Interrupts Example

```
Router#show logging onboard switch 1 module 5 interrupt
```

```
-----  
INTERRUPT SUMMARY INFORMATION  
-----
```

```
Name | ID | Offset | BitList | Count  
-----
```

```
NA 0 0x0000 117857024  
NA 0 0x0000 29361408  
NA 0 0x0000 7932672  
NA 0 0x0100 5120 5968  
NA 0 0x0000 117857024  
-----
```

```
Router#
```

To interpret this data:

- Name is a description of the component including its position in the device.
- ID is an assigned field for data storage.
- Offset is the register offset from a component register's base address.
- Bit is the interrupt bit number recorded from the component's internal register.
- The timestamp shows the date and time that an interrupt occurred down to the millisecond.

Message Logging

The OBFL feature logs standard system messages. Instead of displaying the message to a terminal, the message is written to and stored in a file, so the message can be accessed and read at a later time. System messages range from level 1 alerts to level 7 debug messages, and these levels can be specified in the **hw module logging onboard** command.

Error Message Log Example

```
-----
ERROR MESSAGE SUMMARY INFORMATION
-----
Facility-Sev-Name      | Count | Persistence Flag
MM/DD/YYYY HH:MM:SS
-----
No historical data to display
-----
ERROR MESSAGE CONTINUOUS INFORMATION
-----
MM/DD/YYYY HH:MM:SS Facility-Sev-Name
-----
03/06/2007 22:33:35  %GOLD_OBFL-3-GOLD : Diagnostic OBFL: Diagnostic OBFL testing
```

To interpret this data:

- A timestamp shows the date and time the message was logged.
- Facility-Sev-Name is a coded naming scheme for a system message, as follows:
 - The Facility code consists of two or more uppercase letters that indicate the hardware device (facility) to which the message refers.
 - Sev is a single-digit code from 1 to 7 that reflects the severity of the message.
 - Name is one or two code names separated by a hyphen that describe the part of the system from where the message is coming.
- The error message follows the Facility-Sev-Name codes. For more information about system messages, see the [Cisco IOS System and Error Messages](#) guide.
- Count indicates the number of instances of this message that is allowed in the history file. Once that number of instances has been recorded, the oldest instance will be removed from the history file to make room for new ones.
- The Persistence Flag gives a message priority over others that do not have the flag set.

Default Settings for OBFL

The OBFL feature is enabled by default. Because of the valuable information this feature offers technical personnel, it should not be disabled.

Enabling OBFL

To enable OBFL, perform this task:

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# hw-module switch <i>switch-number</i> module <i>module-number</i> logging onboard [message level {1-7}]	Enables OBFL on the specified hardware module. Note By default, all system messages sent to a device are logged by the OBFL feature. You can define a specific message level (only level 1 messages, as an example) to be logged using the message level keywords.
Step 4	Router(config)# end	Ends global configuration mode.

Configuration Examples for OBFL

The important OBFL feature is the information that is displayed by the **show logging onboard module** privileged EXEC command. This section provides the following examples of how to enable and display OBFL records.

- [Enabling OBFL Message Logging: Example](#)
- [OBFL Message Log: Example](#)
- [OBFL Component Uptime Report: Example](#)
- [OBFL Report for a Specific Time: Example](#)

Enabling OBFL Message Logging: Example

The following example shows how to configure OBFL message logging at level 3:

```
Router(config)# hw-module switch 2 module 1 logging onboard message level 3
```

OBFL Message Log: Example

The following example shows how to display the system messages that are being logged for module 2:

```
Router# show logging onboard module 2 message continuous
```

```
-----
ERROR MESSAGE CONTINUOUS INFORMATION
-----
```

```
MM/DD/YYYY HH:MM:SS Facility-Sev-Name
-----
```

```
03/06/2007 22:33:35 %SWITCH_IF-3-CAMERR : [chars], for VCI [dec] VPI [dec] in stdby data
path check, status: [dec]
-----
```

OBFL Component Uptime Report: Example

The following example shows how to display a summary report for component uptimes for module 2:

```
Router# show logging onboard module 2 uptime
```

```
-----
UPTIME SUMMARY INFORMATION
-----
```

```
First customer power on : 03/06/2007 22:32:51
Total uptime           : 0 years 0 weeks 0 days 0 hours 35 minutes
Total downtime         : 0 years 0 weeks 0 days 0 hours 0 minutes
Number of resets        : 1
Number of slot changes  : 0
Current reset reason     : 0xA1
Current reset timestamp : 03/06/2007 22:31:34
Current slot            : 2
Current uptime          : 0 years 0 weeks 0 days 0 hours 35 minutes
-----
```

```
Reset |      |
Reason | Count |
-----
```

```
No historical data to display
-----
```

OBFL Report for a Specific Time: Example

The following example shows how to display continuous reports for all components during a specific time period:

Router# **show logging onboard module 3 continuous start 15:01:57 1 Mar 2007 end 15:04:57 3 Mar 2007**

PID: WS-X6748-GE-TX , VID: , SN: SAL09063B85

----- UPTIME CONTINUOUS INFORMATION

Time Stamp	Reset	Uptime
MM/DD/YYYY HH:MM:SS	Reason	years weeks days hours minutes
03/01/2007 15:01:57	0xA1	0 0 0 10 0
03/03/2007 02:29:29	0xA1	0 0 0 5 0

----- TEMPERATURE CONTINUOUS INFORMATION

Sensor	ID
MB-Out	930201
MB-In	930202
MB	930203
MB	930204
EARL-Out	910201
EARL-In	910202
SSA 1	930301
SSA 2	930302
JANUS 1	930303
JANUS 2	930304
GEMINI 1	930305
GEMINI 2	930306

Time Stamp	Sensor	Temperature	0C
MM/DD/YYYY HH:MM:SS	1	2	3 4 5 6 7 8 9 10 11 12
03/01/2007 15:01:57	26	26	NA NA NA NA 0 0 0 0 0 0
03/01/2007 15:06:57	39	27	NA NA NA NA 39 37 36 29 32 32
03/01/2007 15:11:02	40	27	NA NA NA NA 40 38 37 30 32 32
03/01/2007 17:06:06	40	27	NA NA NA NA 40 38 37 30 32 32
03/01/2007 19:01:09	40	27	NA NA NA NA 40 38 37 30 32 32
03/03/2007 02:29:30	25	26	NA NA NA NA 0 0 0 0 0 0
03/03/2007 02:34:30	38	26	NA NA NA NA 39 37 36 29 31 31
03/03/2007 04:29:33	40	27	NA NA NA NA 40 38 36 30 32 32
03/03/2007 06:24:37	40	27	NA NA NA NA 40 38 36 29 32 32
03/03/2007 08:19:40	40	27	NA NA NA NA 40 38 36 29 32 32
03/03/2007 10:14:44	40	27	NA NA NA NA 40 38 36 30 32 32
03/03/2007 12:09:47	40	27	NA NA NA NA 40 38 36 30 32 32
03/03/2007 14:04:51	40	27	NA NA NA NA 40 38 36 30 32 32

----- CONTINUOUS INTERRUPT INFORMATION

MM/DD/YYYY HH:MM:SS mmm	Name	ID	Offset	Bit
03/01/2007 15:01:59 350	Port-ASIC #0	7	0x00E7	6
03/03/2007 02:29:34 650	Port-ASIC #0	7	0x00E7	6

```
-----  
ERROR MESSAGE CONTINUOUS INFORMATION  
-----  
MM/DD/YYYY HH:MM:SS Facility-Sev-Name  
-----  
03/01/2007 15:02:15 %GOLD_OBFL-3-GOLD : Diagnostic OBFL: Diagnostic OBFL testing  
03/03/2007 02:29:51 %GOLD_OBFL-3-GOLD : Diagnostic OBFL: Diagnostic OBFL testing  
-----
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

18

Switch Fabric Functionality

- [Prerequisites for Switch Fabric Functionality, page 18-1](#)
- [Restrictions for Switch Fabric Functionality, page 18-1](#)
- [Information About the Switch Fabric Functionality, page 18-1](#)
- [Default Settings for Switch Fabric Functionality, page 18-2](#)
-
- [Monitoring the Switch Fabric Functionality, page 18-2](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Switch Fabric Functionality

None.

Restrictions for Switch Fabric Functionality

None.

Information About the Switch Fabric Functionality

- [Switch Fabric Functionality Overview, page 18-2](#)
- [Forwarding Decisions for Layer 3-Switched Traffic, page 18-2](#)

Switch Fabric Functionality Overview

The switch fabric functionality is built into the supervisor engine and creates a dedicated connection between fabric-enabled modules and provides uninterrupted transmission of frames between these modules. In addition to the direct connection between fabric-enabled modules provided by the switch fabric functionality, fabric-enabled modules also have a direct connection to the forwarding bus.

Forwarding Decisions for Layer 3-Switched Traffic

Either a PFC or a Distributed Feature Card makes the forwarding decision for Layer 3-switched traffic as follows:

- A PFC makes all forwarding decisions for each packet that enters the switch through a module without a DFC.
- A DFC makes all forwarding decisions for each packet that enters the switch on a DFC-equipped module in these situations:
 - If the egress port is on the same module as the ingress port, the DFC forwards the packet locally (the packet never leaves the module).
 - If the egress port is on a different fabric-enabled module, the DFC sends the packet to the egress module, which sends it out the egress port.

Default Settings for Switch Fabric Functionality

Traffic is forwarded to and from modules in one of the following modes:

- Compact mode—The switch uses this mode for all traffic when only fabric-enabled modules are installed. In this mode, a compact version of the DBus header is forwarded over the switch fabric channel, which provides the best possible performance.
- Truncated mode—The switch uses this mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the switch sends a truncated version of the traffic (the first 64 bytes of the frame) over the switch fabric channel.

Monitoring the Switch Fabric Functionality

- [Displaying the Switch Fabric Redundancy Status, page 18-3](#)
- [Displaying Fabric Channel Switching Modes, page 18-3](#)
- [Displaying the Fabric Status, page 18-3](#)
- [Displaying the Fabric Utilization, page 18-4](#)
- [Displaying Fabric Errors, page 18-4](#)

Displaying the Switch Fabric Redundancy Status

To display the switch fabric redundancy status, perform this task:

Command	Purpose
Router# show fabric active	Displays switch fabric redundancy status.

```
Router# show fabric active
Active fabric card in slot 6
Backup fabric card in slot 5
Router#
```

Displaying Fabric Channel Switching Modes

To display the fabric channel switching mode of one or all modules, perform this task:

Command	Purpose
Router# show fabric switching-mode [module {slot_number all}]	Displays fabric channel switching mode of one or all modules.

This example shows how to display the fabric channel switching mode of all modules:

```
Router# show fabric switching-mode module all
Module Slot      Switching Mode
1          dCEF
4          dCEF
5          dCEF
6          dCEF
Router#
```

Displaying the Fabric Status

To display the fabric status of one or all switching modules, perform this task:

Command	Purpose
Router# show fabric status [slot_number all]	Displays fabric status.

This example shows how to display the fabric status of all modules:

```
Router# show fabric status
slot      channel      speed      module      fabric
status      status
1          0          8G          OK          OK
5          0          8G          OK          Up- Timeout
6          0          20G         OK          Up- BufError
8          0          8G          OK          OK
8          1          8G          OK          OK
9          0          8G          Down- DDRsync OK
Router#
```

Displaying the Fabric Utilization

To display the fabric utilization of one or all modules, perform this task:

Command	Purpose
Router# show fabric utilization [<i>slot_number</i> all]	Displays fabric utilization.

This example shows how to display the fabric utilization of all modules:

```
Router# show fabric utilization all
slot channel speed Ingress % Egress %
1 0 40G 0 0
1 1 40G 0 0
2 0 40G 0 0
2 1 40G 0 0
3 0 40G 0 0
3 1 40G 0 0
4 0 40G 0 0
4 1 40G 0 0
5 0 40G 0 0
5 2 40G 0 0
6 0 20G 0 0
6 1 20G 0 0
6 2 20G 0 0
6 3 20G 0 0
Router#
```

Displaying Fabric Errors

To display fabric errors of one or all modules, perform this task:

Command	Purpose
Router# show fabric errors [<i>slot_number</i> all]	Displays fabric errors.

This example shows how to display fabric errors on all modules:

```
Router# show fabric errors

Module errors:
slot channel crc hbeat sync DDR sync
1 0 0 0 0 0
1 1 0 0 0 0
4 0 0 0 0 0
4 1 0 0 0 0
5 0 0 0 0 0
5 2 0 0 0 0
6 0 0 0 0 0
6 1 0 0 0 0
6 2 0 0 0 0
6 3 0 0 0 0

Fabric errors:
slot channel sync buffer timeout
1 0 0 0 0
1 1 0 0 0
4 0 0 0 0
```



```
4 1 0 0 0
5 0 0 0 0
5 2 0 0 0
6 0 0 0 0
6 1 0 0 0
6 2 0 0 0
6 3 0 0 0
Router#
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

19

Cisco IP Phone Support

- [Prerequisites for Cisco IP Phone Support, page 19-1](#)
- [Restrictions for Cisco IP Phone Support, page 19-1](#)
- [Information About Cisco IP Phone Support, page 19-2](#)
- [Default Setting for Cisco IP Phone Support, page 19-4](#)
- [How to Configure Cisco IP Phone Support, page 19-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Cisco IP Phone Support

None.

Restrictions for Cisco IP Phone Support

- The information in this publication may be helpful in configuring support for non-Cisco IP phones, but we recommend that you see the manufacturer's documentation for those devices.
- You must enable the Cisco Discovery Protocol (CDP) on the port connected to the Cisco IP phone to send configuration information to the Cisco IP phone.
- You can configure a voice VLAN only on a Layer 2 LAN port.

- The following conditions indicate that the Cisco IP phone and a device attached to the Cisco IP phone are in the same VLAN and must be in the same IP subnet:
 - If they both use 802.1p or untagged frames
 - If the Cisco IP phone uses 802.1p frames and the device uses untagged frames
 - If the Cisco IP phone uses untagged frames and the device uses 802.1p frames
 - If the Cisco IP phone uses 802.1Q frames and the voice VLAN is the same as the access VLAN
- The Cisco IP phone and a device attached to the Cisco IP phone cannot communicate if they are in the same VLAN and subnet but use different frame types, because traffic between devices in the same subnet is not routed (routing would eliminate the frame type difference).
- You cannot use Cisco IOS software commands to configure the frame type used by traffic sent from a device attached to the access port on the Cisco IP phone.
- If you enable port security on a port configured with a voice VLAN and if there is a PC connected to the Cisco IP phone, set the maximum allowed secure addresses on the port to at least 2.
- You cannot configure static secure MAC addresses in the voice VLAN.
- Ports configured with a voice VLAN can be secure ports (see [Chapter 56, “Port Security”](#)).
- 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).

Information About Cisco IP Phone Support


- [Cisco IP Phone Connections, page 19-2](#)
- [Cisco IP Phone Voice Traffic, page 19-3](#)
- [Cisco IP Phone Data Traffic, page 19-4](#)
- [Other Cisco IP Phone Features, page 19-4](#)

Cisco IP Phone Connections

The Cisco IP phone contains an integrated 3-port 10/100 switch. The ports are dedicated connections to these devices:

- Port 1 connects to the switch.
- Port 2 is an internal 10/100 interface that carries the Cisco IP phone traffic.
- Port 3 connects to a PC or other device.

[Figure 19-1](#) shows a Cisco IP phone connected between a switch and a PC.

Figure 19-1 Cisco IP Phone Connected to a Switch

Cisco IP Phone Voice Traffic

The Cisco IP phone transmits voice traffic with Layer 3 IP precedence and Layer 2 CoS values, which are both set to 5 by default. The sound quality of a Cisco IP phone call can deteriorate if the voice traffic is transmitted unevenly.

You can configure Layer 2 access ports on the switch to send Cisco Discovery Protocol (CDP) packets that configure an attached Cisco IP phone to transmit voice traffic to the switch in any of the following 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).

To provide more predictable voice traffic flow, you can configure QoS on the switch to trust the Layer 3 IP precedence or Layer 2 CoS value in the received traffic (see [Chapter 32, “PFC QoS Overview”](#)).

The trusted boundary device verification feature configures ports on the switch to apply configured [QoS port trust commands](#) only when the Cisco Discovery Protocol (CDP) verifies that the device attached to the port is a Cisco IP phone. See the [“How to Configure Trusted Boundary with Cisco Device Verification” section on page 36-10](#).

You can configure a Layer 2 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 Cisco IP phone.

Cisco IP Phone Data Traffic

**Note**

- The ability to either trust or mark tagged data traffic from the device attached to the access port on the Cisco IP phone is called the “trusted boundary (extended trust for CDP devices)” feature.
- You cannot use Cisco IOS software commands to configure the frame type used by data traffic sent from a device attached to the access port on the Cisco IP phone.
- Untagged data traffic from the device attached to the Cisco IP phone passes through the Cisco IP phone unchanged, regardless of the trust state of the access port on the Cisco IP phone.

To process tagged data traffic (traffic in 802.1Q or 802.1p frame types) from the device attached to the access port on the Cisco IP phone (see [Figure 19-1](#)), you can configure Layer 2 access ports on the switch to send CDP packets that instruct an attached Cisco IP phone to configure the access port on the Cisco IP phone to either of these two modes:

- Trusted mode—All traffic received through the access port on the Cisco IP phone passes through the Cisco IP phone unchanged.
- Untrusted mode—All traffic in 802.1Q or 802.1p frames received through the access port on the Cisco IP phone is marked with a configured Layer 2 CoS value. The default Layer 2 CoS value is 0. Untrusted mode is the default.

Most IP phones have no ability to notify the switch of link state changes on the IP phone’s access port. When a device attached to the access port is disconnected or disabled administratively, the switch is unaware of the change. Some Cisco IP phones can send a CDP message containing a host presence type length value (TLV) indicating the changed state of the access port link.

Other Cisco IP Phone Features

The switch provides support for authentication, authorization, and accounting (AAA) for Cisco IP phones, as described in [Chapter 54, “IEEE 802.1X Port-Based Authentication.”](#)

The switch also supports automatic tracking for Cisco Emergency Responder (Cisco ER) to help you manage emergency calls in your telephony network. For further information, see this URL:

http://www.cisco.com/en/US/products/sw/voicesw/ps842/tsd_products_support_series_home.html

Default Setting for Cisco IP Phone Support

- Cisco IP phone support is disabled by default.
- When the voice VLAN feature is enabled, all untagged traffic is sent with the default CoS priority of the port.
- CoS values are not trusted for 802.1P or 802.1Q tagged traffic.

How to Configure Cisco IP Phone Support

- [Configuring Voice Traffic Support, page 19-5](#)
- [Configuring Data Traffic Support, page 19-6](#)

Configuring Voice Traffic Support

To configure the way in which the Cisco IP phone transmits voice traffic, perform this task:

	Command	Purpose
Step 1	Router(config)# interface gigabitethernet <i>slot/port</i>	Selects the port to configure.
Step 2	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. Note You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 port before you can enter additional switchport commands with keywords.
Step 3	Router(config-if)# switchport voice vlan { <i>voice_vlan_ID</i> dot1p none untagged }	Configures the way in which the Cisco IP phone transmits voice traffic.
Step 4	Router(config)# end	Exits configuration mode.

When configuring the way in which the Cisco IP phone transmits voice traffic, note the following information:

- Enter a voice VLAN ID to send CDP packets that configure the Cisco IP phone to transmit voice traffic in 802.1Q frames, tagged with the voice VLAN ID and a Layer 2 CoS value (the default is 5). Valid VLAN IDs are from 1 to 4094. The switch puts the 802.1Q voice traffic into the voice VLAN.
- Enter the **dot1p** keyword to send CDP packets that configure the Cisco IP phone to transmit voice traffic in 802.1p frames, tagged with VLAN ID 0 and a Layer 2 CoS value (the default is 5 for voice traffic and 3 for voice control traffic). The switch puts the 802.1p voice traffic into the access VLAN.
- Enter the **untagged** keyword to send CDP packets that configure the Cisco IP phone to transmit untagged voice traffic. The switch puts the untagged voice traffic into the access VLAN.
- Enter the **none** keyword to allow the Cisco IP phone to use its own configuration and transmit untagged voice traffic. The switch puts the untagged voice traffic into the access VLAN.
- In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5).
- See [Chapter 32, “PFC QoS Overview,”](#) for information about how to configure QoS.
- See the [“Configuring a LAN Interface as a Layer 2 Access Port”](#) section on page 11-13 for information about how to configure the port as a Layer 2 access port and configure the access VLAN.

This example shows how to configure Gigabit Ethernet port 5/1 to send CDP packets that tell the Cisco IP phone to use VLAN 101 as the voice VLAN:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/1
Router(config-if)# switchport voice vlan 101
Router(config-if)# exit
```

This example shows how to verify the configuration of Gigabit Ethernet port 5/1:

```
Router# show interfaces gigabitethernet 5/1 switchport
Name: Gi5/1
Switchport: Enabled
Administrative Mode: access
Operational Mode: access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: off
Access Mode VLAN: 100
Voice VLAN: 101
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: 900 ((Inactive)) 901 ((Inactive))
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
```

Configuring Data Traffic Support



Note

The trusted boundary feature is implemented with the **platform qos trust extend** command.

To configure the way in which an attached Cisco IP phone transmits data traffic, perform this task:

	Command	Purpose
Step 1	Router(config)# interface gigabitethernet slot/port	Selects the port to configure.
Step 2	Router(config-if)# platform qos trust extend [cos cos_value]	Configures the way in which an attached Cisco IP phone transmits data traffic.
Step 3	Router(config)# end	Exits configuration mode.

When configuring the way in which an attached Cisco IP phone transmits data traffic, note the following information:

- To send CDP packets that configure an attached Cisco IP phone to trust tagged traffic received from a device connected to the access port on the Cisco IP phone, do not enter the **cos** keyword and CoS value.
- To send CDP packets that configure an attached Cisco IP phone to mark tagged ingress traffic received from a device connected to the access port on the Cisco IP phone, enter the **cos** keyword and CoS value (valid values are 0 through 7).
- You cannot use Cisco IOS software commands to configure whether or not traffic sent from a device attached to the access port on the Cisco IP phone is tagged.

This example shows how to configure Gigabit Ethernet port 5/1 to send CDP packets that tell the Cisco IP phone to configure its access port as untrusted and to mark all tagged traffic received from a device connected to the access port on the Cisco IP phone with CoS 3:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/1
Router(config-if)# platform qos trust extend cos 3
```


This example shows how to configure Gigabit Ethernet port 5/1 to send CDP packets that tell the Cisco IP phone to configure its access port as trusted:

```
Router# configure terminal  
Router(config)# interface gigabitethernet 5/1  
Router(config-if)# platform qos trust extend
```

This example shows how to verify the configuration on Gigabit Ethernet port 5/1:

```
Router# show queueing interface gigabitethernet 5/1 | include Extend  
Extend trust state: trusted
```

**Note**

In addition to CDP, you can also use LLDP MED on the port connecting to the Cisco IP Phone to detect and configure it.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

20

Power over Ethernet (PoE) Support

- [Prerequisites for PoE, page 20-1](#)
- [Restrictions for PoE, page 20-1](#)
- [Information About PoE, page 20-1](#)
- [How to Configure PoE Support, page 20-4](#)

**Note**

- For information about switching modules that support PoE, see the *Release Notes for Cisco IOS Release 15.2SY* publication at this URL:
http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.1SY/release_notes.html
- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for PoE

None.

Restrictions for PoE

PoE is supported only on Layer 2 switchports.

Information About PoE

- [Device Roles, page 20-2](#)
- [PoE Overview, page 20-2](#)

- [CPD-Based PoE Management](#), page 20-3
- [Inline Power IEEE Power Classification Override](#), page 20-4
- [LLDP Inline Power Negotiation for PoE+ \(IEEE 802.3at\)](#), page 20-4

Device Roles

- Power sourcing equipment (PSE)—A device that provides power through a twisted-pair Ethernet connection. The switch, through switching modules equipped with Power over Ethernet (PoE) daughtercards, functions in the PSE role.
- Powered device (PD)—A device powered by a PSE (for example, IP phones, IP cameras, and wireless access points).



Note

Not all PoE-capable devices are powered from the switch. There are two sources of local power for PoE-capable devices:

- A power supply connected to the device.
- A power supply through a patch panel over the Ethernet connection to the device.

When a locally powered PoE-capable device is present on a switching module port, the switching module itself cannot detect its presence. If the device supports CDP, the supervisor engine can discover a locally powered PoE-capable device through CDP messaging with the device. If a locally powered PoE-capable device loses local power, the switching module can discover and supply power to the IP phone if the inline power mode is set to **auto**.

PoE Overview

Cisco PoE daughtercards support one or more PoE implementation:

- IEEE 802.3at standard, shown in Cisco Feature Navigator as “PoE Plus (PoE+, PoEP) support”.
 - Supported only with the PoE daughtercard on the WS-X6148E-GE-45AT switching module.
 - These features are supported for IEEE 802.3at-compliant class 4 PDs:
 - Class 4: 30.00 W at the PSE (12.95 W to 25.50 W at the PD).
 - Optionally, [LLDP Inline Power Negotiation for PoE+](#).
 - With releases earlier than Release 15.1(1)SY, maximum 16.8 W at the PSE (ePoE for 45 ports maximum).
- IEEE 802.3af standard.
 - Supported with the WS-F6K-48-AF PoE daughtercard and the PoE daughtercard on the WS-X6148E-GE-45AT switching module.
 - Maximum 16.80 W at the PSE.
 - The IEEE 802.3af PoE standard defines a method to sense a PD and to immediately classify the power requirement of the PD into these per port power ranges at the PSE:
 - Class 0: Up to 15.4 W (0.44–12.95 W at the PD; default classification)
 - Class 1: Up to 4 W (0.44–3.84 W at the PD)

- Class 2: Up to 7 W (3.84–6.49 W at the PD)
- Class 3: Up to 15.4 W (6.49–12.95 W at the PD)
- Cisco prestandard inline power—10 W at the PSE.

With a PoE daughtercard installed, a switching module can automatically detect and provision a PoE-capable device that adheres to a PoE implementation supported by the PoE daughtercard. The switching module can supply power to devices supporting other PoE implementations only through manual configuration.

Only a PD connected directly to the switch port can be powered from the switch. If a second PD is daisy-chained from the PD that is connected to the switch port, the second PD cannot be powered by the switch.

Each PD requires power to be allocated from the chassis power budget. Because each PD can have unique power requirements, more devices can be supported if the system's power management software can intelligently allocate the necessary power on a per-port basis.

You can configure ports to allocate power at a level based on the following:

- If a PD is detected, with auto mode configured:
 - Information sensed from the device
 - A default level
 - A configured maximum level
- Whether or not a PD is present on the port, with static mode configured:
 - A default level
 - A configured level

CPD-Based PoE Management

When a switching module port detects an unpowered PD, the default-allocated power is provided to the port. When the correct amount of power is determined through CDP messaging with the PD, the supervisor engine reduces or increases the allocated power, up to the hardware limit of the installed PoE daughtercard.



Caution

When a PD cable is plugged into a port and the power is turned on, the supervisor engine has a 4-second timeout waiting for the link to go up on the line. During those 4 seconds, if the IP phone cable is unplugged and a network device is plugged in, the network device could be damaged. We recommend that you wait at least 10 seconds between unplugging a network device and plugging in another network device.

Inline Power IEEE Power Classification Override

The IEEE 802.3af standard contains no provision for adjustment of the power allocation. 802.3af-compliant PDs that support CDP can use CDP to override the IEEE 802.3af power classification.

The WS-F6K-48-AF PoE daughtercard or the PoE daughtercard on the WS-X6148E-GE-45AT switching module support these inline power IEEE 802.3af power classification override features:

- Power use measurement—The ability to accurately measure the power provided by the port to the powered device.
- Power policing—The ability to monitor power usage on a port.

With power measurement and policing, you can safely override the IEEE 802.3af power classification of a device that requires a power level at the lower end of its IEEE power classification range.

PoE monitoring and policing compares the power consumption on ports with the administrative maximum value (either a configured maximum value or the port's default value). If the power consumption on a monitored port exceeds the administrative maximum value, the following actions occur:

- A syslog message is issued.
- The monitored port is shut down and error-disabled.
- The allocated power is freed.

LLDP Inline Power Negotiation for PoE+ (IEEE 802.3at)

The PoE daughtercard on the WS-X6148E-GE-45AT switching module supports [IEEE 802.3at](#)-compliant LLDP PoE power negotiation, which supports additional negotiation that can reduce power usage.



Note

- Enabled by default.
- The LLDP TLV used is DTE Power-via-MDI TLV.
- When a PD that performs power negotiation using multiple protocols (CDP and LLDP 802.3at) is connected to a switch, the switch locks to the first protocol packet (CDP or LLDP) that contains the power negotiation TLV. If you need to use any single protocol for power negotiation each time, you must administratively disable the other power negotiation protocols on the switch interface.
- See this publication for other the Link Layer Discovery Protocol (LLDP) configuration procedures: http://www.cisco.com/en/US/docs/ios/cether/configuration/guide/ce_lldp-med.html

How to Configure PoE Support

- [Displaying PoE Status, page 20-5](#)
- [Configuring Per-Port PoE Support, page 20-5](#)
- [Configuring PoE Power Priority, page 20-6](#)
- [Configuring PoE Monitoring and Policing, page 20-8](#)
- [Disabling LLDP Power Negotiation \(IEEE 802.3at\), page 20-8](#)

Displaying PoE Status

This example shows how to display the PoE status on switch:

```
Router# show power auxiliary
system auxiliary power mode = on
system auxiliary power redundancy operationally = redundant
system primary connector power limit = 7266.00 Watts (173.00 Amps @ 42V)
system auxiliary connector power limit = 10500.00 Watts (250.00 Amps @ 42V)
system primary power used = 1407.00 Watts (33.50 Amps @ 42V)
system auxiliary power used = 22.68 Watts ( 0.54 Amps @ 42V)
```

Slot	Card-Type	Inline Pwr-Limit Watts	A @42V	Inline-Pwr Used-Thru-Pri Watts	A @42V	Inline-Pwr Used-Thru-Aux Watts	A @42V	VDB Aux-Pwr Capable
2	WS-F6K-48-AT	1600.20	38.10	23.10	0.55	11.34	0.27	Yes
4	WS-F6K-48-AT	1600.20	38.10	23.10	0.55	11.34	0.27	Yes
Totals:				46.20	1.10	22.68	0.54	

Configuring Per-Port PoE Support

To configure per-port PoE support, perform this task:

	Command	Purpose
Step 1	Router(config-if)# power inline { auto static never } [max <i>milliwatts</i>]	Configures per-port PoE support and optionally specifies a maximum inline power level in milliwatts for the port.
Step 2	Router# show power inline { <i>type slot/port</i> <i>module slot</i> } [detail]	Verifies the configuration.

When configuring inline power support with the **power inline** command, note the following information:

- To configure auto-detection of a PD and PoE auto-allocation, enter the **auto** keyword.
- To configure auto-detection of a PD but reserve a fixed PoE allocation, enter the **static** keyword.
- To specify the maximum power to allocate to a port, enter either the **auto** or **static** keyword followed by the **max** keyword and the power level in milliwatts.
- When the **auto** keyword is entered and CDP is enabled on the port, a PD that supports CDP can negotiate a different power level.
- To disable auto-detection of a PD, enter the **never** keyword.

- With a WS-F6K-GE48-AF, WS-F6K-48-AF, or the PoE daughtercard on the WS-X6148E-GE-45AT switching module:
 - The configurable range of maximum power using the **max** keyword is 4000 to 16800 milliwatts. If no maximum power level is configured, the default maximum power is 15400 milliwatts.

**Note**

To support a large number of inline-powered ports using power levels above 15400 milliwatts on an inline power card, we recommend using the **static** keyword so that the power budget is deterministic.

- When the **auto** keyword is entered and CDP is enabled on the port, an inline-powered device that supports CDP can negotiate a power level up to 16800 milliwatts unless a lower maximum power level is configured.

This example shows how to disable inline power on GigabitEthernet port 2/10:

```
Router# configure terminal
Router(config)# interface gigabitethernet 2/10
Router(config-if)# power inline never
```

This example shows how to enable inline power on GigabitEthernet port 2/10:

```
Router# configure terminal
Router(config)# interface gigabitethernet 2/10
Router(config-if)# power inline auto
```

This example shows how to verify the inline power configuration on GigabitEthernet port 2/10:

```
Router# show power inline gigabitethernet 2/10
```

Interface	Admin	Priority (enabled)	Oper	Power(Watts) From PS To PD		Device	Class
Gi2/10	auto	low	on	14.5	13.1	Cisco IP Phone 9971	4


```
Router# show power inline gigabitethernet 2/10
```

Interface	AdminPowerMax (Watts)	Police	ActConsumption
Gi2/10	30.0	on	6.7

Configuring PoE Power Priority

You can configure how the switch responds if a power shortage occurs by setting the priority of ports providing PoE. The priority determines the order in which PoE is removed from ports if a power shortage occurs: low-priority, then high-priority, with power maintained for critical-priority ports as long as possible. These sections describe how to configure PoE power priority:

- [Setting the PoE Power Priority Global Enable State, page 20-7](#)
- [Configuring PoE Port Power Priority, page 20-7](#)

Setting the PoE Power Priority Global Enable State

To disable PoE power priority globally, perform this task:

	Command	Purpose
Step 1	Router(config)# no power inline priority enable	Disables PoE power priority globally (default: enabled).
Step 2	Router# show power inline	Verifies the configuration.

This example shows how to disable PoE power priority globally:

```
Router(config)# no power inline priority enable
```

The column heading of any **show power inline** command displays the PoE power priority global state (“disabled” in this example):

```
Router# show power inline
Interface Admin  Priority      Oper      Power(Watts)  Device      Class
                (disabled)    From PS To PD
-----
...
```

Configuring PoE Port Power Priority

To configure PoE port power priority, perform this task:

	Command	Purpose
Step 1	Router(config-if)# power inline auto priority {critical high low}	Enables PoE port power priority (default: low priority when power priority is enabled globally). If a power shortage occurs, PoE is removed from ports in the following order: - Low priority ports - High priority ports PoE is maintained for critical priority ports as long as possible.
Step 2	Router# show power inline type slot/port [detail]	Verifies the configuration.

This example shows how to configure the PoE port power priority of GigabitEthernet port 2/10 as high:

```
Router# configure terminal
Router(config)# interface gigabitethernet 2/10
Router(config-if)# power inline auto priority high
```

This example shows how to verify the PoE port power priority configuration of GigabitEthernet port 2/10:

```
Router# show power inline gigabitethernet 2/10 detail | include Priority
Priority: high
```

Configuring PoE Monitoring and Policing

With the WS-F6K-48-AF PoE daughtercard or the PoE daughtercard on the WS-X6148E-GE-45AT switching module, to configure PoE monitoring and policing, perform this task:

	Command	Purpose
Step 1	Router(config-if)# power inline police	Enables PoE monitoring and policing.
Step 2	Router# show power inline {type slot/port module slot}[detail]	Verifies the configuration.

This example shows how to enable monitoring and policing on GigabitEthernet port 1/9:

```
Router# configure terminal
Router(config)# interface gigabitethernet 2/10
Router(config-if)# power inline police
```

These examples shows how to verify the power monitoring and policing configuration on GigabitEthernet port 2/10:

```
Router# show power inline gigabitethernet 2/10 detail | include Police
Police: on
```

```
Router#
Router# show power inline gigabitethernet 2/10
Interface  Admin  Oper  Power (Watts)  Device  Class
          -----
          From PS To Device
-----
Gi2/10      auto   on    17.3    15.4        Ieee PD   3

Interface  AdminPowerMax (Watts)  Police  ActualConsumption
-----
Gi2/10      15.4                on      5.7
Router#
```

Disabling LLDP Power Negotiation (IEEE 802.3at)

With the WS-X6148E-GE-45AT switching module, LLDP power negotiation is enabled by default. To disable LLDP power negotiation, perform this task:

Command	Purpose
Router(config-if)# no lldp tlv-select power-management	Disables LLDP power negotiation (default: enabled).

This example shows how to display the LLDP power negotiation configuration on interface GigabitEthernet 3/1 when LLDP power negotiation is enabled:

```
Router# show power inline gigabitethernet 2/10 detail | begin LLDP
LLDP Power Classification    -- Sent to PD --  -- Rcvd from PD --
Power Type :                 type 2 PSE                type 2 PD
Power Source :                primary                   PSE
Power Priority :               low                       high
Requested Power (watts):      11.2                     11.2
Allocated Power (watts):      11.2                     11.2
Power class :                 4                         4
```

```
LLDP Legacy MDI TLV          -- Rcvd from PD --
MDI power support :          0
pse power pair :             0
MDI power class :            0
```

This example shows how to disable LLDP power negotiation on interface GigabitEthernet 2/10:

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# interface gigabitethernet 2/10
Router(config-if)# no lldp tlv-select power-management
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

21

LAN Ports for Layer 2 Switching

- [Prerequisites for Layer 2 LAN Interfaces, page 21-1](#)
- [Restrictions for Layer 2 LAN Interfaces, page 21-2](#)
- [Information About Layer 2 Switching, page 21-2](#)
- [Default Settings for Layer 2 LAN Interfaces, page 21-5](#)
- [How to Configure LAN Interfaces for Layer 2 Switching, page 21-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - To configure Layer 3 interfaces, see [Chapter 35, “Layer 3 Interfaces.”](#)
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Layer 2 LAN Interfaces

None.

Restrictions for Layer 2 LAN Interfaces

- When connecting Cisco switches through an 802.1q trunk, make sure the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning tree loops might result.
- Disabling spanning tree on the native VLAN of an 802.1Q trunk without disabling spanning tree on every VLAN in the network can cause spanning tree loops. We recommend that you leave spanning tree enabled on the native VLAN of an 802.1Q trunk. If this is not possible, disable spanning tree on every VLAN in the network. Make sure your network is free of physical loops before disabling spanning tree.
- When you connect two Cisco switches through 802.1Q trunks, the switches exchange spanning tree BPDUs on each VLAN allowed on the trunks. The BPDUs on the native VLAN of the trunk are sent untagged to the reserved IEEE 802.1d spanning tree multicast MAC address (01-80-C2-00-00-00). The BPDUs on all other VLANs on the trunk are sent tagged to the reserved Cisco Shared Spanning Tree (SSTP) multicast MAC address (01-00-0c-cc-cc-cd).
- Non-Cisco 802.1Q switches maintain only a single instance of spanning tree (the Mono Spanning Tree, or MST) that defines the spanning tree topology for all VLANs. When you connect a Cisco switch to a non-Cisco switch through an 802.1Q trunk, the MST of the non-Cisco switch and the native VLAN spanning tree of the Cisco switch combine to form a single spanning tree topology known as the Common Spanning Tree (CST).
- Because Cisco switches transmit BPDUs to the SSTP multicast MAC address on VLANs other than the native VLAN of the trunk, non-Cisco switches do not recognize these frames as BPDUs and flood them on all ports in the corresponding VLAN. Other Cisco switches connected to the non-Cisco 802.1q cloud receive these flooded BPDUs. This allows Cisco switches to maintain a per-VLAN spanning tree topology across a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single broadcast segment between all switches connected to the non-Cisco 802.1q cloud through 802.1q trunks.
- Make certain that the native VLAN is the same on all of the 802.1q trunks connecting the Cisco switches to the non-Cisco 802.1q cloud.
- If you are connecting multiple Cisco switches to a non-Cisco 802.1q cloud, all of the connections must be through 802.1q trunks. You cannot connect Cisco switches to a non-Cisco 802.1q cloud through access ports. Doing so causes the switch to place the access port into the spanning tree “port inconsistent” state and no traffic will pass through the port.

Information About Layer 2 Switching

- [Information about Layer 2 Ethernet Switching, page 21-2](#)
- [Information about VLAN Trunks, page 21-4](#)
- [Layer 2 LAN Port Modes, page 21-4](#)

Information about Layer 2 Ethernet Switching

- [Layer 2 Ethernet Switching Overview, page 21-3](#)
- [Building the MAC Address Table, page 21-3](#)

Layer 2 Ethernet Switching Overview

Layer 2 Ethernet ports on Cisco switches support simultaneous, parallel connections between Layer 2 Ethernet segments. Switched connections between Ethernet segments last only for the duration of the packet. New connections can be made between different segments for the next packet.

Layer 2 LAN switching (hardware-supported bridging) avoids congestion by assigning each connected device to its own collision domain. Because each LAN port connects to a separate Ethernet collision domain, attached devices in a properly configured switched environment achieve full access to network bandwidth.

Building the MAC Address Table

- [Overview of the MAC Address Table, page 21-3](#)
- [Synchronization and Sharing of the Address Table, page 21-3](#)
- [Notification of Address Table Changes, page 21-3](#)

Overview of the MAC Address Table

When stations connected to different LAN ports need to communicate, the switch forwards frames from one LAN port to the other at wire speed to ensure that each session receives full bandwidth.

To switch frames between LAN ports efficiently, the switch maintains a MAC address table. When a frame enters the switch, it associates the MAC address of the sending network device with the LAN port on which it was received.

The MAC address table is built by using the source MAC address of the frames received. When the switch receives a frame for a destination MAC address not listed in its MAC address table, it floods the frame to all LAN ports of the same VLAN except the port that received the frame. When the destination station replies, the switch adds its relevant source MAC address and port ID to the MAC address table. The switch then forwards subsequent frames to a single LAN port without flooding to all LAN ports.

The MAC address table can store at least 128,000 address entries without flooding any entries. The switch uses an aging mechanism, configured by the **mac address-table aging-time** command, so if an address remains inactive for a specified number of seconds, it is removed from the address table.

Synchronization and Sharing of the Address Table

With distributed switching, each DFC-equipped switching module learns MAC addresses, maintains an address table, and ages table entries. MAC address table synchronization over the Ethernet Out of Band Channel (EOBC) synchronizes address tables among the PFC and all DFCs, eliminating the need for flooding by a DFC for an address that is active on another module. MAC synchronization is enabled by default.

Notification of Address Table Changes

You can configure the switch to maintain a history of dynamic additions and removals of address table entries associated with a particular LAN port. The change history can be sent as an SNMP trap notification or it can be read manually from the SNMP MIB.

Information about VLAN Trunks



Note

For information about VLANs, see [Chapter 26, “Virtual Local Area Networks \(VLANs\).”](#)

A trunk is a point-to-point link between the switch and another networking device. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network. 802.1Q, an industry-standard trunking encapsulation, is available on all Ethernet ports.

You can configure a trunk on a single Ethernet port or on an EtherChannel. For more information about EtherChannel, see [Chapter 23, “EtherChannels.”](#)

Ethernet trunk ports support several trunking modes (see [Table 21-1 on page 21-4](#)).

The Dynamic Trunking Protocol (DTP) manages trunk autonegotiation on LAN ports.

To autonegotiate trunking, the LAN ports must be in the same VTP domain. Use the **trunk** or **nonegotiate** keywords to force LAN ports in different domains to trunk. For more information on VTP domains, see [Chapter 25, “VLAN Trunking Protocol \(VTP\).”](#)

Layer 2 LAN Port Modes

Table 21-1 *Layer 2 LAN Port Modes*

Mode	Function
switchport mode access	Puts the LAN port into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The LAN port becomes a nontrunk port even if the neighboring LAN port does not agree to the change.
switchport mode dynamic desirable	Makes the LAN port actively attempt to convert the link to a trunk link. The LAN port becomes a trunk port if the neighboring LAN port is set to trunk , desirable , or auto mode. This is the default mode for all LAN ports.
switchport mode dynamic auto	Makes the LAN port willing to convert the link to a trunk link. The LAN port becomes a trunk port if the neighboring LAN port is set to trunk or desirable mode.
switchport mode trunk	Puts the LAN port into permanent trunking mode and negotiates to convert the link into a trunk link. The LAN port becomes a trunk port even if the neighboring port does not agree to the change.
switchport nonegotiate	Puts the LAN port into permanent trunking mode but prevents the port from generating DTP frames. You must configure the neighboring port manually as a trunk port to establish a trunk link.



Note

DTP is a point-to-point protocol. However, some internetworking devices might forward DTP frames improperly. To avoid this problem, ensure that LAN ports connected to devices that do not support DTP are configured with the **access** keyword if you do not intend to trunk across those links. To enable trunking to a device that does not support DTP, use the **nonegotiate** keyword to cause the LAN port to become a trunk but not generate DTP frames.

Default Settings for Layer 2 LAN Interfaces

Feature	Default
Interface mode: - Before entering the switchport command - After entering the switchport command	Layer 3 (unconfigured) switchport mode dynamic desirable
Allowed VLAN range	VLANs 1 to 4094, except reserved VLANs (see Table 26-1 on page 26-3)
VLAN range eligible for pruning	VLANs 2 to 1001
Default access VLAN	VLAN 1
Native VLAN (for 802.1Q trunks)	VLAN 1
Spanning Tree Protocol (STP)	Enabled for all VLANs
STP port priority	128
STP port cost	100 for 10-Mbps Ethernet LAN ports 19 for 10/100-Mbps Fast Ethernet LAN ports 19 for 100-Mbps Fast Ethernet LAN ports 4 for 1,000-Mbps Gigabit Ethernet LAN ports 2 for 10,000-Mbps 10-Gigabit Ethernet LAN ports

How to Configure LAN Interfaces for Layer 2 Switching

- [Configuring a LAN Port for Layer 2 Switching, page 21-6](#)
- [Enabling Out-of-Band MAC Address Table Synchronization, page 21-6](#)
- [Configuring MAC Address Table Notification, page 21-7](#)
- [Configuring a Layer 2 Switching Port as a Trunk, page 21-8](#)
- [Configuring a LAN Interface as a Layer 2 Access Port, page 21-14](#)
- [Configuring a Custom IEEE 802.1Q EtherType Field Value, page 21-15](#)



Note

Use the **default interface** {**fastethernet** | **gigabitethernet** | **tengigabitethernet**} *slot/port* command to revert an interface to its default configuration.

Configuring a LAN Port for Layer 2 Switching

To configure a LAN port for Layer 2 switching, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# shutdown	(Optional) Shuts down the interface to prevent traffic flow until configuration is complete.
Step 3	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. Note You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 port before you can enter additional switchport commands with keywords.
Step 4	Router(config-if)# no shutdown	Activates the interface. (Required only if you shut down the interface.)
Step 5	Router(config-if)# end	Exits configuration mode.

After you enter the **switchport** command, the default mode is **switchport mode dynamic desirable**. If the neighboring port supports trunking and is configured to allow trunking, the link becomes a Layer 2 trunk when you enter the **switchport** command.



Note

When using the **switchport** command, if a port configured for Layer 3 is now configured for Layer 2, the configuration for Layer 3 is retained in the memory but not in the running configuration and is applied to the port whenever the port switches back to Layer 3. Also, if a port configured for Layer 2 is now configured for Layer 3, the configuration for Layer 2 is retained in the memory but not in the running configuration and is applied to the port whenever the port switches back to Layer 2. To restore the default configuration of the port in the memory and in the running configuration, use the **default interface** command. To avoid potential issues while changing the role of a port using the **switchport** command, shut down the interface before applying the **switchport** command.

Enabling Out-of-Band MAC Address Table Synchronization

To enable the out-of-band MAC address table synchronization feature, perform this task:

Command	Purpose
Router(config)# mac address-table synchronize [activity-time <i>seconds</i>]	Enables out-of-band synchronization of MAC address tables among DFC-equipped switching modules. activity-time <i>seconds</i>—(Optional) Specifies the activity timer interval.

When configuring out-of-band MAC address table synchronization, note the following information:

- By default, out-of-band MAC address table synchronization is disabled.
- Out-of-band MAC address table synchronization is enabled automatically if a WS-6708-10G switching module is installed in the switch.
- The activity timer interval can be configured as 160, 320, and 640 seconds. The default is 160 seconds.

This example shows how to enable out-of-band MAC address table synchronization:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# mac address-table synchronize activity-time 320
```

Configuring MAC Address Table Notification



Note

- Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching” section on page 21-6](#) before performing the tasks in this section.
- To send SNMP trap notifications using this feature, you must also enable the global MAC trap flag, using the **snmp-server enable mac-notification change** command.

To configure the MAC address table notification feature, perform this task:

	Command	Purpose
Step 1	Router(config)# mac address-table notification change [interval value]	Enables sending notification of dynamic changes to MAC address table. (Optional) Sets the minimum change-sending interval in seconds. Note The no form of the command reverts to the default without sending any change information.
Step 2	Router(config)# mac address-table notification change [history size]	Enables sending notification of dynamic changes to MAC address table. (Optional) Sets the number of entries in the history buffer. Note The no form of the command reverts to the default without sending any change information.
Step 3	Router(config)# interface type slot/port	Selects the LAN port to configure.
Step 4	Router(config-if)# snmp trap mac-notification change [added removed]	For MAC addresses that are associated with this LAN port, enable SNMP trap notification when MAC addresses are added to or removed from the address table. (Optional) To notify only when a MAC address is added to the table, use the added option. To notify only when a MAC address is removed, use the removed option.
Step 5	Router(config-if)# end	Exits interface configuration mode.

When configuring the notification parameters, note the following information:

- The **interval** *value* parameter can be configured from 0 seconds (immediate) to 2,147,483,647 seconds. The default is 1 second.
- The **history** *size* parameter can be configured from 0 entries to 500 entries. The default is 1 entry.

This example shows how to configure the SNMP notification of dynamic additions to the MAC address table of addresses on the Gigabit Ethernet ports 5/7 and 5/8. Notifications of changes will be sent no more frequently than 5 seconds, and up to 25 changes can be stored and sent in that interval:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# mac address-table notification change interval 5
Router(config)# mac address-table notification change history 25
Router(config)# interface gigabitethernet 5/7
Router(config-if)# snmp trap mac-notification change added
Router(config-if)# end
Router(config)# interface gigabitethernet 5/8
Router(config-if)# snmp trap mac-notification change added
Router(config-if)# end
Router# exit
```

Configuring a Layer 2 Switching Port as a Trunk

- [Configuring the Layer 2 Switching Port as an 802.1Q Trunk, page 21-8](#)
- [Configuring the Layer 2 Trunk to Use DTP, page 21-9](#)
- [Configuring the Layer 2 Trunk Not to Use DTP, page 21-9](#)
- [Configuring the Access VLAN, page 21-10](#)
- [Configuring the 802.1Q Native VLAN, page 21-11](#)
- [Configuring the List of VLANs Allowed on a Trunk, page 21-11](#)
- [Configuring the List of Prune-Eligible VLANs, page 21-12](#)
- [Completing Trunk Configuration, page 21-13](#)
- [Verifying Layer 2 Trunk Configuration, page 21-13](#)
- [Configuration and Verification Examples, page 21-13](#)

Configuring the Layer 2 Switching Port as an 802.1Q Trunk



Note

- Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching”](#) section on page 21-6 before performing the tasks in this section.
- When you enter the **switchport** command with no other keywords ([Step 3](#) in the previous section), the default mode is **switchport mode dynamic desirable**.

To configure the Layer 2 switching port as an 802.1Q trunk, perform this task:

Command	Purpose
Router(config-if) # switchport trunk encapsulation dot1q	(Optional) Configures the encapsulation as 802.1Q.

To support the **switchport mode trunk** command, you must configure the encapsulation as 802.1Q.

**Note**

Complete the steps in the [“Completing Trunk Configuration” section on page 21-13](#) after performing the tasks in this section.

Configuring the Layer 2 Trunk to Use DTP

**Note**

Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching” section on page 21-6](#) before performing the tasks in this section.

To configure the Layer 2 trunk to use DTP, perform this task:

Command	Purpose
Router(config-if) # switchport mode dynamic {auto desirable}	(Optional) Configures the trunk to use DTP. Note The no form of the command reverts to the default trunk trunking mode (switchport mode dynamic desirable).

When configuring the Layer 2 trunk to use DTP, note the following information:

- Required only if the interface is a Layer 2 access port or to specify the trunking mode.
- See [Table 21-1 on page 21-4](#) for information about trunking modes.

**Note**

Complete the steps in the [“Completing Trunk Configuration” section on page 21-13](#) after performing the tasks in this section.

Configuring the Layer 2 Trunk Not to Use DTP

**Note**

Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching” section on page 21-6](#) before performing the tasks in this section.

To configure the Layer 2 trunk not to use DTP, perform this task:

	Command	Purpose
Step 1	Router(config-if)# switchport mode trunk	(Optional) Configures the port to trunk unconditionally.
Step 2	Router(config-if)# switchport nonegotiate	(Optional) Configures the trunk not to use DTP. Note The no form of the command enables DTP on the port.

When configuring the Layer 2 trunk not to use DTP, note the following information:

- Before entering the **switchport mode trunk** command, you must configure the encapsulation (see the “[Configuring the Layer 2 Switching Port as an 802.1Q Trunk](#)” section on page 21-8).
- To support the **switchport nonegotiate** command, you must enter the **switchport mode trunk** command.
- Enter the **switchport mode dynamic trunk** command. See [Table 21-1 on page 21-4](#) for information about trunking modes.
- Before entering the **switchport nonegotiate** command, you must configure the encapsulation (see the “[Configuring the Layer 2 Switching Port as an 802.1Q Trunk](#)” section on page 21-8) and configure the port to trunk unconditionally with the **switchport mode trunk** command (see the “[Configuring the Layer 2 Trunk to Use DTP](#)” section on page 21-9).


Note

Complete the steps in the “[Completing Trunk Configuration](#)” section on page 21-13 after performing the tasks in this section.

Configuring the Access VLAN


Note

Complete the steps in the “[Configuring a LAN Port for Layer 2 Switching](#)” section on page 21-6 before performing the tasks in this section.

To configure the access VLAN, perform this task:

Command	Purpose
Router(config-if)# switchport access vlan <i>vlan_ID</i>	(Optional) Configures the access VLAN, which is used if the interface stops trunking. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3). Note • If VLAN locking is enabled, enter the VLAN name instead of the VLAN number. For more information, see the “VLAN Locking” section on page 26-4. • The no form of the command reverts to the default VLAN (VLAN 1).

**Note**

Complete the steps in the [“Completing Trunk Configuration” section on page 21-13](#) after performing the tasks in this section.

Configuring the 802.1Q Native VLAN

**Note**

Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching” section on page 21-6](#) before performing the tasks in this section.

To configure the 802.1Q native VLAN, perform this task:

Command	Purpose
Router(config-if)# switchport trunk native vlan <i>vlan_ID</i>	(Optional) Configures the 802.1Q native VLAN. Note If VLAN locking is enabled, enter the VLAN name instead of the VLAN number. For more information, see the “VLAN Locking” section on page 26-4 .

When configuring the native VLAN, note the following information:

- The *vlan_ID* value can be 1 through 4094, except reserved VLANs (see [Table 26-1 on page 26-3](#)).
- The access VLAN is not automatically used as the native VLAN.

**Note**

Complete the steps in the [“Completing Trunk Configuration” section on page 21-13](#) after performing the tasks in this section.

Configuring the List of VLANs Allowed on a Trunk

**Note**

Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching” section on page 21-6](#) before performing the tasks in this section.

To configure the list of VLANs allowed on a trunk, perform this task:

Command	Purpose
Router(config-if)# switchport trunk allowed vlan [add except none remove] <i>vlan</i> [, <i>vlan</i> [, <i>vlan</i> [, ...]]]	(Optional) Configures the list of VLANs allowed on the trunk. Note • If VLAN locking is enabled, enter VLAN names instead of VLAN numbers. For more information, see the “VLAN Locking” section on page 26-4. • The no form of the command reverts to the default value (all VLANs allowed).

When configuring the list of VLANs allowed on a trunk, note the following information:

- The *vlan* parameter is either a single VLAN number from 1 through 4094, or a range of VLANs described by two VLAN numbers, the lesser one first, separated by a dash. Do not enter any spaces between comma-separated *vlan* parameters or in dash-specified ranges.
- If VLAN locking is enabled, enter VLAN names instead of VLAN numbers. When entering a range of VLAN names, you must leave spaces between the VLAN names and the dash.
- All VLANs are allowed by default.
- You can remove VLAN 1. If you remove VLAN 1 from a trunk, the trunk interface continues to send and receive management traffic, for example, Cisco Discovery Protocol (CDP), VLAN Trunking Protocol (VTP), Port Aggregation Protocol (PAgP), and DTP in VLAN 1.


Note

Complete the steps in the [“Completing Trunk Configuration” section on page 21-13](#) after performing the tasks in this section.

Configuring the List of Prune-Eligible VLANs


Note

Complete the steps in the [“Configuring a LAN Port for Layer 2 Switching” section on page 21-6](#) before performing the tasks in this section.

To configure the list of prune-eligible VLANs on the Layer 2 trunk, perform this task:

Command	Purpose
Router(config-if)# switchport trunk pruning vlan {none {{add except remove} vlan[,vlan[,vlan[,...]]}}	(Optional) Configures the list of prune-eligible VLANs on the trunk (see the “VTP Pruning” section on page 25-7). Note The no form of the command reverts to the default value (all VLANs prune-eligible).

When configuring the list of prune-eligible VLANs on a trunk, note the following information:

- The *vlan* parameter is either a single VLAN number from 1 through 4094, except reserved VLANs (see [Table 26-1 on page 26-3](#)), or a range of VLANs described by two VLAN numbers, the lesser one first, separated by a dash. Do not enter any spaces between comma-separated *vlan* parameters or in dash-specified ranges.
- The default list of VLANs allowed to be pruned contains all VLANs.
- Network devices in VTP transparent mode do not send VTP Join messages. On trunk connections to network devices in VTP transparent mode, configure the VLANs used by the transparent-mode network devices or that need to be carried across the transparent-mode network devices as pruning ineligible.


Note

Complete the steps in the [“Completing Trunk Configuration” section on page 21-13](#) after performing the tasks in this section.

Completing Trunk Configuration

To complete Layer 2 trunk configuration, perform this task:

	Command	Purpose
Step 1	Router(config-if)# no shutdown	Activates the interface. (Required only if you shut down the interface.)
Step 2	Router(config-if)# end	Exits configuration mode.

Verifying Layer 2 Trunk Configuration

To verify Layer 2 trunk configuration, perform this task:

	Command	Purpose
Step 1	Router# show running-config interface type slot/port	Displays the running configuration of the interface.
Step 2	Router# show interfaces [type slot/port] switchport	Displays the switch port configuration of the interface.
Step 3	Router# show interfaces [type slot/port] trunk	Displays the trunk configuration of the interface.

Configuration and Verification Examples

This example shows how to configure the Gigabit Ethernet port 5/8 as an 802.1Q trunk. This example assumes that the neighbor port is configured to support 802.1Q trunking:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/8
Router(config-if)# shutdown
Router(config-if)# switchport
Router(config-if)# switchport mode dynamic desirable
Router(config-if)# switchport trunk encapsulation dot1q
Router(config-if)# no shutdown
Router(config-if)# end
Router# exit
```

This example shows how to verify the configuration:

```
Router# show running-config interface gigabitethernet 5/8
Building configuration...
Current configuration:
!
interface GigabitEthernet5/8
 no ip address
 switchport
 switchport trunk encapsulation dot1q
end
```

```
Router# show interfaces gigabitethernet 5/8 switchport
Name: Gi5/8
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: trunk
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: Enabled
Access Mode VLAN: 1 (default)
```

```

Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: ALL

Router# show interfaces gigabitethernet 5/8 trunk

Port      Mode           Encapsulation  Status        Native vlan
Gi5/8     desirable     n-802.1q       trunking      1

Port      Vlans allowed on trunk
Gi5/8 1-1005

Port      Vlans allowed and active in management domain
Gi5/8 1-6,10,20,50,100,152,200,300,303-305,349-351,400,500,521,524,570,801-802,850,917,999,1002-1005

Port      Vlans in spanning tree forwarding state and not pruned
Gi5/8 1-6,10,20,50,100,152,200,300,303-305,349-351,400,500,521,524,570,801-802,850,917,999,1002-1005

Router#

```

Configuring a LAN Interface as a Layer 2 Access Port



Note

If you assign a LAN port to a VLAN that does not exist, the port is shut down until you create the VLAN in the VLAN database (see the [“Creating or Modifying an Ethernet VLAN”](#) section on page 26-5).

To configure a LAN port as a Layer 2 access port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# shutdown	(Optional) Shuts down the interface to prevent traffic flow until configuration is complete.
Step 3	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. Note You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 port before you can enter additional switchport commands with keywords.
Step 4	Router(config-if)# switchport mode access	Configures the LAN port as a Layer 2 access port.
Step 5	Router(config-if)# switchport access vlan <i>vlan_ID</i>	Places the LAN port in a VLAN. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3). Note If VLAN locking is enabled, enter the VLAN name instead of the VLAN number. For more information, see the “VLAN Locking” section on page 26-4.
Step 6	Router(config-if)# no shutdown	Activates the interface. (Required only if you shut down the interface.)
Step 7	Router(config-if)# end	Exits configuration mode.

This example shows how to configure the Gigabit Ethernet port 5/6 as an access port in VLAN 200:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/6
Router(config-if)# shutdown
Router(config-if)# switchport
Router(config-if)# switchport mode access
Router(config-if)# switchport access vlan 200
Router(config-if)# no shutdown
Router(config-if)# end
Router# exit
```

This example shows how to verify the configuration:

```
Router# show running-config interface gigabitethernet 5/6
Building configuration...
!
Current configuration:
interface GigabitEthernet5/6
  no ip address
  switchport access vlan 200
  switchport mode access
end

Router# show interfaces gigabitethernet 5/6 switchport
Name: Gi5/6
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Enabled
Access Mode VLAN: 200 (VLAN0200)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: ALL

Router#
```

Configuring a Custom IEEE 802.1Q EtherType Field Value

You can configure a custom EtherType field value on a port to support network devices that do not use the standard 0x8100 EtherType field value on 802.1Q-tagged or 802.1p-tagged frames.

To configure a custom value for the EtherType field, perform this task:

Command	Purpose
Router(config-if)# switchport dot1q ethertype value	Configures the 802.1Q EtherType field value for the port.

When configuring a custom EtherType field value, note the following information:

- To use a custom EtherType field value, all network devices in the traffic path across the network must support the custom EtherType field value.
- You can configure a custom EtherType field value on trunk ports, access ports, and tunnel ports.
- You can configure a custom EtherType field value on the member ports of an EtherChannel.
- You cannot configure a custom EtherType field value on a port-channel interface.

- Each port supports only one EtherType field value. A port that is configured with a custom EtherType field value does not recognize frames that have any other EtherType field value as tagged frames. For example, a trunk port that is configured with a custom EtherType field value does not recognize the standard 0x8100 EtherType field value on 802.1Q-tagged frames and cannot put the frames into the VLAN to which they belong.

**Caution**

A port that is configured with a custom EtherType field value considers frames that have any other EtherType field value to be untagged frames. A trunk port with a custom EtherType field value places frames with any other EtherType field value into the native VLAN. An access port or tunnel port with a custom EtherType field value places frames that are tagged with any other EtherType field value into the access VLAN. If you misconfigure a custom EtherType field value, frames might be placed into the wrong VLAN.

- See the *Release Notes for Cisco IOS Release 15.2SY* for a list of the modules that support custom IEEE 802.1Q EtherType field values:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.1SY/release_notes.html

This example shows how to configure the EtherType field value to 0x1234:

```
Router (config-if)# switchport dot1q ethertype 1234
Router (config-if)#
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

22

Flex Links

- [Prerequisites for Flex Links, page 22-1](#)
- [Restrictions for Flex Links, page 22-1](#)
- [Information About Flex Links, page 22-2](#)
- [Default Settings for Flex Links, page 22-4](#)
- [How to Configure Flex Links, page 22-4](#)
- [Monitoring Flex Links, page 22-6](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Flex Links

None.

Restrictions for Flex Links

- You can configure only one Flex Links 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 Links pair. An interface can be a backup link for only one active link. An active link cannot belong to another Flex Links 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 as the active link (Fast Ethernet, Gigabit Ethernet, or port channel). However, you should configure both Flex Links with similar characteristics so that there are no loops or changes in operation if the standby link becomes active.
- STP is disabled on Flex Links ports. If STP is disabled on the switch, be sure that there are no Layer 2 loops in the network topology.
- Do not configure any STP features (for example, PortFast, BPDU Guard, and so forth) on Flex Links ports or the ports to which the links connect.
- Local administrative shut down or a link that starts forwarding again due to preemption is not considered a link failure. In those cases, the feature flushes the dynamic hosts and does not move them.
- Static MAC addresses that are configured on the primary link are not moved to the standby link.
- Static MAC addresses configured on a flex links port are restored when it starts forwarding again.


Information About Flex Links

Flex Links are a pair of Layer 2 interfaces (ports or port channels), where one interface is configured to act as a backup to the other. Flex Links are typically configured in service-provider or enterprise networks where customers do not want to run STP. Flex Links provide link-level redundancy that is an alternative to Spanning Tree Protocol (STP). STP is automatically disabled on Flex Links interfaces.

Release 15.2SY supports a maximum of 16 Flex Links. Flex Links are supported only on Layer 2 ports and port channels, not on VLANs or on Layer 3 ports.

To configure the Flex Links feature, you configure one Layer 2 interface as the standby link for the link that you want to be primary. With Flex Links configured for a pair of interfaces, only one of the interfaces is in the linkup state and is forwarding traffic. If the primary link shuts down, the standby link starts forwarding traffic. When the inactive link comes back up, it goes into standby mode.

In [Figure 22-1](#), 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 and the other one 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 to forward traffic.

Figure 22-1 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. When a primary link fails, the feature takes these actions:


- Detects the failure.
- Moves any dynamic unicast MAC addresses that are learned on the primary link to the standby link.
- Moves the standby link to a forwarding state.
- Transmits dummy multicast packets over the new active interface. The dummy multicast packet format is:
 - Destination: 01:00:0c:cd:cd:cd
 - Source: MAC address of the hosts or ports on the newly active Flex Link port.

In [Figure 22-2](#), ports 1 and 2 on switch A are connected to switches B and D through a Flex Link pair. Port 1 is forwarding traffic, and port 2 is in the blocking 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 port 1 shuts down, port 2 starts forwarding traffic. If there is no traffic from the PC to the server after failover to port 2, switch C does not learn the MAC address of the PC on port 4, and because of that, switch C keeps forwarding traffic from the server to the PC out of port 3. There is traffic loss from the server to the PC because port 1 is down. To alleviate this problem, the feature sends out a dummy multicast packet with the source MAC address of the PC over port 2. Switch C learns the PC MAC address on port 4 and start forwarding traffic from the server to the PC out of port 4. One dummy multicast packet is sent out for every MAC address.

Flex links interface preemption specifies one of the ports in a flex links pair as preferred for traffic forwarding. The preference can be unconditional or it can be based on bandwidth availability. See the [“How to Configure Flex Links”](#) section on page 22-4.

Figure 22-2 Flexlink Dummy Multicast Packets Example



Default Settings for Flex Links

- Flex links: not configured.
- Flex links interface preemption: not configured.
- Flex links interface preemption delay: 35 seconds.

How to Configure Flex Links

To configure Flex Links, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(conf)# interface {{type slot/port} {port-channel number}}	Specifies a Layer 2 interface.

	Command	Purpose
Step 3	Router(conf-if)# switchport backup interface <i>interface_id</i>	Configures the interface as part of a flex links pair. - The <i>interface_id</i> specifies can be a physical port or a port-channel interface. - The backup interface must already be configured as a Layer 2 port.
Step 4	Router(conf-if)# switchport backup interface <i>interface_id</i> preemption mode [forced bandwidth off]	(Optional) Configures the preferred flex links interface preemption port in the flex links pair. - The <i>interface_id</i> can be a physical port or a port-channel interface. 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.
Step 5	Router(conf-if)# switchport backup interface <i>interface_id</i> preemption delay <i>delay_time</i>	(Optional) Configures the flex links interface preemption delay time for the flex links pair. - The <i>interface_id</i> can be a physical port or a port-channel interface. - The range for <i>delay_time</i> is 1 through 300 seconds.
Step 6	Router(conf-if)# exit	Exits configuration mode.

This example shows how to configure an interface with a flex links backup interface and how to verify the configuration:

```
Router# configure terminal
Router(conf)# interface tengigabitethernet 2/9
Router(conf-if)# switchport backup interface tengigabitethernet 2/12
Router(conf-if)# switchport backup interface tengigabitethernet 2/12 preemption mode
[forced | bandwidth | off]
Router(conf-if)# switchport backup interface tengigabitethernet 2/12 preemption delay 35
Router(conf-if)# exit
Router# show interface switchport backup detail
```

Switch Backup Interface Pairs:

```
Active Interface      Backup Interface      State
-----
Te2/9                 Te2/12                Active Up/Backup Standby
Interface Pair       : Te2/9, Te2/12
Preemption Mode      : forced
Preemption Delay     : 35 seconds (default)
Bandwidth            : 10000000 Kbit (Te2/9), 10000000 Kbit (Te2/12)
```

Monitoring Flex Links

To monitor the Flex Links configuration, perform this task:

Command	Purpose
<code>show interface [{type slot/port} {port-channel number}] switchport backup</code>	Displays the Flex Links backup interface configured for an interface, or displays all Flex Links configured on the switch and the state of each active and backup interface (up or standby mode).

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

23

EtherChannels

- [Prerequisites for EtherChannels, page 23-1](#)
- [Restrictions for EtherChannels, page 23-2](#)
- [Information About EtherChannels, page 23-3](#)
- [Default Settings for EtherChannels, page 23-7](#)
- [How to Configure EtherChannels, page 23-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for EtherChannels

None.

Restrictions for EtherChannels

- LACP EtherChannels and the [802.1ad provider-bridge mode](#) are mutually exclusive. LACP EtherChannels cannot transmit traffic when the 802.1ad provider-bridge mode is enabled.
- LACP does not support half-duplex links. Half-duplex links in an LACP EtherChannel are put in the suspended state.



Caution

Serious traffic problems can result from mixing manual mode with LACP mode, or by connecting an EtherChannel member port to a port not configured as part of the EtherChannel. For example, if a port configured in **on** mode is connected to another port configured in **desirable** mode, or to a port not configured as a member of the EtherChannel, a bridge loop is created and a broadcast storm can occur. If one end uses the **on** mode, the other end must also.

- When EtherChannel interfaces are configured improperly, they are disabled automatically to avoid network loops and other problems.
- Frames with SAP/SNAP encapsulation are load-balanced as Layer 2 traffic.
- The commands in this chapter can be used on all Layer 2 Ethernet ports, including the ports on the supervisor engine and a redundant supervisor engine.
- All Layer 2 Ethernet ports on all modules, including those on a redundant supervisor engine, support EtherChannels (maximum of eight LAN ports) with no requirement that the LAN ports be physically contiguous or on the same module.
- Configure all LAN ports in an EtherChannel to use the same EtherChannel protocol; you cannot run two EtherChannel protocols in one EtherChannel.
- Configure all LAN ports in an EtherChannel to operate at the same speed and in the same duplex mode.
- Enable all LAN ports in an EtherChannel. If you shut down a LAN port in an EtherChannel, it is treated as a link failure and its traffic is transferred to one of the remaining ports in the EtherChannel.
- An EtherChannel will not form if any of the LAN ports is a Switched Port Analyzer (SPAN) destination port.
- For Layer 3 EtherChannels, assign Layer 3 addresses to the port channel logical interface, not to the LAN ports in the channel.
- For Layer 2 EtherChannels:
 - Assign all LAN ports in the EtherChannel to the same VLAN or configure them as trunks.
 - If you configure an EtherChannel from trunking LAN ports, verify that the trunking mode is the same on all the trunks. LAN ports in an EtherChannel with different trunk modes can operate unpredictably.
 - An EtherChannel supports the same allowed range of VLANs on all the LAN ports in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the LAN ports do not form an EtherChannel.
 - LAN ports with different STP port path costs can form an EtherChannel as long they are compatibly configured with each other. If you set different STP port path costs, the LAN ports are not incompatible for the formation of an EtherChannel.
 - An EtherChannel will not form if protocol filtering is set differently on the LAN ports.

- Configure static MAC addresses on the EtherChannel only and not on physical member ports of the EtherChannel.
- After you configure an EtherChannel, the configuration that you apply to the port channel interface affects the EtherChannel. The configuration that you apply to the LAN ports affects only the LAN port where you apply the configuration.
- Cisco IOS Release 15.4SY does not support ISL trunk encapsulation. If a non-trunking Layer 2 EtherChannel includes member ports that are not capable of ISL trunk encapsulation, the **switchport trunk encapsulation dot1q** command is added to the port-channel interface. The command has no affect when the switchport mode is “access” (CSCta45114).
- When QoS is enabled, enter the **no platform qos channel-consistency** port-channel interface command to support EtherChannels that have ports with and without strict-priority queues.

Information About EtherChannels

- [EtherChannel Feature Overview, page 23-3](#)
- [Information about EtherChannel Configuration, page 23-3](#)
- [Information about Port Channel Interfaces, page 23-7](#)
- [Information about LACP 1:1 Redundancy, page 23-6](#)
- [Information about Load Balancing, page 23-7](#)

EtherChannel Feature Overview

An EtherChannel bundles individual Ethernet links into a single logical link that provides the aggregate bandwidth of up to eight physical links.

Cisco IOS Release 15.4SY supports a maximum of 512 EtherChannels. You can form an EtherChannel with up to eight compatibly configured LAN ports on any switching module. All LAN ports in each EtherChannel must be the same speed and must all be configured as either Layer 2 or Layer 3 LAN ports.



Note

The network device to which a switch is connected may impose its own limits on the number of ports in an EtherChannel.

If a segment within an EtherChannel fails, traffic previously carried over the failed link switches to the remaining segments within the EtherChannel. When a failure occurs, the EtherChannel feature sends a trap that identifies the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one segment in an EtherChannel are blocked from returning on any other segment of the EtherChannel.

Information about EtherChannel Configuration

- [EtherChannel Configuration Overview, page 23-4](#)
- [Information about Manual EtherChannel Configuration, page 23-5](#)
- [Information about PAGP EtherChannel Configuration, page 23-5](#)
- [Information about IEEE 802.3ad LACP EtherChannel Configuration, page 23-5](#)

EtherChannel Configuration Overview

You can configure EtherChannels manually or you can use the Port Aggregation Control Protocol (PAgP) or the Link Aggregation Control Protocol (LACP) to form EtherChannels. The EtherChannel protocols allow ports with similar characteristics to form an EtherChannel through dynamic negotiation with connected network devices. PAgP is a Cisco-proprietary protocol and LACP is defined in IEEE 802.3ad.

PAgP and LACP do not interoperate with each other. Ports configured to use PAgP cannot form EtherChannels with ports configured to use LACP. Ports configured to use LACP cannot form EtherChannels with ports configured to use PAgP. Neither interoperates with ports configured manually.

[Table 23-1](#) lists the user-configurable EtherChannel modes.

[Table 23-2](#) lists the EtherChannel member port states.

Table 23-1 EtherChannel Modes

Mode	Description
on	Mode that forces the LAN port to channel unconditionally. In the on mode, a usable EtherChannel exists only when a LAN port group in the on mode is connected to another LAN port group in the on mode. Because ports configured in the on mode do not negotiate, there is no negotiation traffic between the ports. You cannot configure the on mode with an EtherChannel protocol. If one end uses the on mode, the other end must also.
auto	PAgP mode that places a LAN port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not initiate PAgP negotiation. (Default)
desirable	PAgP mode that places a LAN port into an active negotiating state, in which the port initiates negotiations with other LAN ports by sending PAgP packets.
passive	LACP mode that places a port into a passive negotiating state, in which the port responds to LACP packets it receives but does not initiate LACP negotiation. (Default)
active	LACP mode that places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets.

Table 23-2 EtherChannel Member Port States

Port States	Description
bundled	The port is part of an EtherChannel and can send and receive BPDUs and data traffic.
suspended	The port is not part of an EtherChannel. The port can receive BPDUs but cannot send them. Data traffic is blocked.
standalone	The port is not bundled in an EtherChannel. The port functions as a standalone data port. The port can send and receive BPDUs and data traffic. Note When one end of an EtherChannel has more members than the other, the unmatched ports enter the standalone state. In a topology that is not protected from Layer 2 loops by the spanning tree protocol (STP), a port in the standalone state can cause significant network errors. You can enter the port-channel standalone-disable interface configuration mode command to put ports into the suspended state instead of the standalone state. See the “Configuring LACP Port-Channel Standalone Disable” section on page 23-16.

Information about Manual EtherChannel Configuration

Manually configured EtherChannel ports do not exchange EtherChannel protocol packets. A manually configured EtherChannel forms only when you configure all ports in the EtherChannel compatibly.

Information about PAgP EtherChannel Configuration

PAgP supports the automatic creation of EtherChannels by exchanging PAgP packets between LAN ports. PAgP packets are exchanged only between ports in **auto** and **desirable** modes.

The protocol learns the capabilities of LAN port groups dynamically and informs the other LAN ports. Once PAgP identifies correctly matched Ethernet links, it facilitates grouping the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.

Both the **auto** and **desirable** modes allow PAgP to negotiate between LAN ports to determine if they can form an EtherChannel, based on criteria such as port speed and trunking state. Layer 2 EtherChannels also use VLAN numbers.

LAN ports can form an EtherChannel when they are in different PAgP modes if the modes are compatible. For example:

- A LAN port in **desirable** mode can form an EtherChannel successfully with another LAN port that is in **desirable** mode.
- A LAN port in **desirable** mode can form an EtherChannel with another LAN port in **auto** mode.
- A LAN port in **auto** mode cannot form an EtherChannel with another LAN port that is also in **auto** mode, because neither port will initiate negotiation.

Information about IEEE 802.3ad LACP EtherChannel Configuration

LACP supports the automatic creation of EtherChannels by exchanging LACP packets between LAN ports. LACP packets are exchanged only between ports in **passive** and **active** modes.

The protocol learns the capabilities of LAN port groups dynamically and informs the other LAN ports. Once LACP identifies correctly matched Ethernet links, it facilitates grouping the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.

Both the **passive** and **active** modes allow LACP to negotiate between LAN ports to determine if they can form an EtherChannel, based on criteria such as port speed and trunking state. Layer 2 EtherChannels also use VLAN numbers.

LAN ports can form an EtherChannel when they are in different LACP modes as long as the modes are compatible. For example:

- A LAN port in **active** mode can form an EtherChannel successfully with another LAN port that is in **active** mode.
- A LAN port in **active** mode can form an EtherChannel with another LAN port in **passive** mode.
- A LAN port in **passive** mode cannot form an EtherChannel with another LAN port that is also in **passive** mode, because neither port will initiate negotiation.

LACP uses the following parameters:

- LACP system priority—You must configure an LACP system priority on each switch running LACP. The system priority can be configured automatically or through the CLI (see the [“Configuring the LACP System Priority and System ID” section on page 23-11](#)). LACP uses the system priority with the switch MAC address to form the system ID and also during negotiation with other systems.

**Note**

The LACP system ID is the combination of the LACP system priority value and the MAC address of the switch.

- LACP port priority—You must configure an LACP port priority on each port configured to use LACP. The port priority can be configured automatically or through the CLI (see the [“Configuring Channel Groups” section on page 23-9](#)). LACP uses the port priority with the port number to form the port identifier. LACP uses the port priority to decide which ports should be put in standby mode when there is a hardware limitation that prevents all compatible ports from aggregating.
- LACP administrative key—LACP automatically configures an administrative key value equal to the channel group identification number on each port configured to use LACP. The administrative key defines the ability of a port to aggregate with other ports. A port’s ability to aggregate with other ports is determined by these factors:
 - Port physical characteristics, such as data rate, duplex capability, and point-to-point or shared medium
 - Configuration restrictions that you establish

On ports configured to use LACP, LACP tries to configure the maximum number of compatible ports in an EtherChannel, up to the maximum allowed by the hardware (eight ports). If LACP cannot aggregate all the ports that are compatible (for example, the remote system might have more restrictive hardware limitations), then all the ports that cannot be actively included in the channel are put in hot standby state and are used only if one of the channeled ports fails. You can configure an additional 8 standby ports (total of 16 ports associated with the EtherChannel).

Information about LACP 1:1 Redundancy

The LACP 1:1 redundancy feature supports an EtherChannel configuration with one active link and fast switchover to a hot standby link. The link connected to the port with the lower port priority number (and therefore a higher priority) will be the active link, and the other link will be in a hot standby state. If the active link goes down, LACP performs fast switchover to the hot standby link to keep the EtherChannel up. When the failed link becomes operational again, LACP performs another fast switchover to revert to the original active link.

To allow the higher priority port to stabilize when it becomes active again after a higher-priority to lower-priority switchover, the LACP 1:1 hot standby dampening feature configures a timer that delays switchover back to the higher priority port after it becomes active.

See the [“Configuring LACP 1:1 Redundancy” section on page 23-14](#).

Information about Port Channel Interfaces

Each EtherChannel has a numbered port channel interface. You can configure a maximum of 128 port-channel interfaces, numbered from 1 to 256. The configuration that you apply to the port channel interface affects all LAN ports assigned to the port channel interface.

After you configure an EtherChannel, the configuration that you apply to the port channel interface affects the EtherChannel; the configuration that you apply to the LAN ports affects only the LAN port where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply the configuration commands to the port channel interface, for example, Spanning Tree Protocol (STP) commands or commands to configure a Layer 2 EtherChannel as a trunk.

Information about Load Balancing

An EtherChannel balances the traffic load across the links in an EtherChannel 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. EtherChannel load balancing can also use Layer 4 port numbers. EtherChannel load balancing can use either source or destination or both source and destination addresses or ports. The selected mode applies to all EtherChannels configured on the switch. EtherChannel load balancing can use MPLS Layer 2 information.

Use the option that provides the balance criteria with the greatest variety in your configuration. For example, if the traffic on an EtherChannel is going only to a single MAC address and you use the destination MAC address as the basis of EtherChannel load balancing, the EtherChannel always chooses the same link in the EtherChannel; using source addresses or IP addresses might result in better load balancing.

Default Settings for EtherChannels

None.

How to Configure EtherChannels

- [Configuring Port Channel Logical Interfaces, page 23-8](#)
- [Configuring Channel Groups, page 23-9](#)
- [Configuring the LACP System Priority and System ID, page 23-11](#)
- [Configuring EtherChannel Load Balancing, page 23-11](#)
- [Configuring the EtherChannel Hash-Distribution Algorithm, page 23-12](#)
- [Configuring the EtherChannel Min-Links Feature, page 23-13](#)
- [Configuring LACP 1:1 Redundancy, page 23-14](#)

- [Configuring Auto Interleaved Port Priority For LACP Port Channels, page 23-15](#)
- [Configuring LACP Port-Channel Standalone Disable, page 23-16](#)

**Note**

Make sure that the LAN ports are configured correctly (see the “[Restrictions for EtherChannels](#)” section on page 23-2).

Configuring Port Channel Logical Interfaces

**Note**

To move an IP address from a Layer 3 LAN port to an EtherChannel, you must delete the IP address from the Layer 3 LAN port before configuring it on the port channel logical interface.

To create a port channel interface for a Layer 3 EtherChannel, perform this task:

	Command	Purpose
Step 1	Router(config)# interface port-channel <i>group_number</i>	Creates the port channel interface.
Step 2	Router(config-if)# ip address <i>ip_address mask</i>	Assigns an IP address and subnet mask to the EtherChannel.
Step 3	Router(config-if)# end	Exits configuration mode.

The *group_number* can be 1 through 256, up to a maximum of 256 port-channel interfaces in standalone mode. In VSS-mode the group number can be through 1 through 512, up to a maximum of 512 port channel-interfaces. This example shows how to create port channel interface 1:

```
Router# configure terminal
Router(config)# interface port-channel 1
Router(config-if)# ip address 172.32.52.10 255.255.255.0
Router(config-if)# end
```

This example shows how to verify the configuration of port channel interface 1:

```
Router# show running-config interface port-channel 1
Building configuration...

Current configuration:
!
interface Port-channel1
 ip address 172.32.52.10 255.255.255.0
 no ip directed-broadcast
end
Router#
```

Configuring Channel Groups



Note

For Cisco IOS to create port channel interfaces for Layer 2 EtherChannels, the Layer 2 LAN ports must be connected and functioning.

To configure channel groups, perform this task for each LAN port:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects a LAN port to configure.
Step 2	Router(config-if)# no ip address	Ensures that there is no IP address assigned to the LAN port.
Step 3	Router(config-if)# channel-protocol { lACP pagp }	(Optional) On the selected LAN port, restricts the channel-group command to the EtherChannel protocol configured with the channel-protocol command.
Step 4	Router(config-if)# channel-group <i>group_number</i> mode { active auto desirable on passive }	Configures the LAN port in a port channel and specifies the mode (see Table 23-1 on page 23-4). PAgP supports only the auto and desirable modes. LACP supports only the active and passive modes.
Step 5	Router(config-if)# lACP port-priority <i>priority_value</i>	(Optional for LACP) Valid values are 1 through 65535. Higher numbers have lower priority. The default is 32768.
Step 6	Router(config-if)# do show interface port-channel <i>group_number</i>	(Optional) Displays the interface configuration information of the specified port channel. Note: The output of this command shows that the newly created port channel interface will be in shutdown state.
Step 7	Router(config-if)# no shutdown	Brings the port channel and its members up and all the configuration changes that you apply to the port channel is applied to every member interface of that port channel.
Step 8	Router(config-if)# end	Exits configuration mode.

This example shows how to configure Gigabit Ethernet ports 5/6 and 5/7 into port channel 2 with PAgP mode **desirable**:

```
Router# configure terminal
Router(config)# interface range gigabitethernet 5/6 -7
Router(config-if)# channel-group 2 mode desirable
Router(config-if)# no shutdown
Router(config-if)# end
```



Note

See the “[How to Configure a Range of Interfaces](#)” section on [page 10-2](#) for information about the **range** keyword.

This example shows how to verify the configuration of port channel interface 2:

```
Router# show running-config interface port-channel 2
Building configuration...

Current configuration:
!
interface Port-channel2
```

```

no ip address
switchport
switchport access vlan 10
switchport mode access
end
Router#

```

This example shows how to verify the configuration of Gigabit Ethernet port 5/6:

```
Router# show running-config interface gigabitethernet 5/6
```

```
Building configuration...
```

```
Current configuration:
```

```

!
interface GigabitEthernet5/6
no ip address
switchport
switchport access vlan 10
switchport mode access
channel-group 2 mode desirable
end

```

```
Router# show interfaces gigabitethernet 5/6 etherchannel
```

```

Port state      = Down Not-in-Bndl
Channel group   = 12           Mode = Desirable-S1      Gcchange = 0
Port-channel    = null        GC   = 0x00000000        Pseudo port-channel = Po1
2
Port index      = 0           Load = 0x00           Protocol =   PAgP

```

```

Flags:  S - Device is sending Slow hello.  C - Device is in Consistent state.
        A - Device is in Auto mode.        P - Device learns on physical port.
        d - PAgP is down.

```

```

Timers: H - Hello timer is running.        Q - Quit timer is running.
        S - Switching timer is running.    I - Interface timer is running.

```

```
Local information:
```

Port	Flags	State	Timers	Hello Interval	Partner Count	PAgP Priority	Learning Method	Group Ifindex
Gi5/2	d	U1/S1		1s	0	128	Any	0

```
Age of the port in the current state: 04d:18h:57m:19s
```

This example shows how to verify the configuration of port channel interface 2 after the LAN ports have been configured:

```
Router# show etherchannel 12 port-channel
```

```
Port-channels in the group:
```

```
-----
```

```
Port-channel: Po12
```

```
-----
```

```

Age of the Port-channel    = 04d:18h:58m:50s
Logical slot/port          = 14/1           Number of ports = 0
GC                         = 0x00000000     HotStandBy port = null
Port state                 = Port-channel Ag-Not-Inuse
Protocol                   =   PAgP

```

```
Router#
```

Configuring the LACP System Priority and System ID

The LACP system ID is the combination of the LACP system priority value and the MAC address of the switch.

To configure the LACP system priority and system ID, perform this task:

	Command	Purpose
Step 1	Router(config)# lacp system-priority <i>value</i>	(Optional for LACP) Valid values are 1 through 65535. Higher numbers have lower priority. The default is 32768.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the LACP system priority:

```
Router# configure terminal
Router(config)# lacp system-priority 23456
Router(config)# end
Router(config)#
```

This example shows how to verify the configuration:

```
Router# show lacp sys-id
23456,0050.3e8d.6400
Router#
```

The system priority is displayed first, followed by the MAC address of the switch.

Configuring EtherChannel Load Balancing

To configure EtherChannel load balancing, perform this task:

	Command	Purpose
Step 1	Router(config)# port-channel per-module load-balance	(Optional) Enables the ability to specify the load-balancing method on a per-module basis.
Step 2	Router(config)# port-channel load-balance [<i>all</i> <i>fex</i> <i>fex_number</i>] <i>method</i> [<i>module</i> <i>slot</i>]	Configures the EtherChannel load-balancing method. The <i>method</i> is globally applied to all port channels. Optionally, you can configure the load-balancing method for a specific module. The default method is src-dst-ip . A configurable load balancing method is only applicable for unicast traffic. For multicast traffic the load balancing method used is always src-dst-ip . Note - If you configure EtherChannel load balancing on a switch that is not in VSS mode, traffic is disrupted while the EtherChannel member ports transition through the shutdown and then no shutdown states. - There is no disruption on switches in VSS mode.
Step 3	Router(config)# end	Exits configuration mode.

The load-balancing *method* keywords indicate the following information:

- **dst-ip**—Destination IP addresses
- **dst-mac**—Destination MAC addresses
- **dst-mixed-ip-port**—Destination IP address and TCP/UDP port
- **dst-port**—Destination Layer 4 port
- **mpls**—Load balancing for MPLS packets
- **src-dst-ip**—(Default) Source and destination IP addresses
- **src-dst-mac**—Source and destination MAC addresses
- **src-dst-mixed-ip-port**—Source and destination IP address and TCP/UDP port
- **src-dst-port**—Source and destination Layer 4 port
- **src-ip**—Source IP addresses
- **src-mac**—Source MAC addresses
- **src-mixed-ip-port**—Source IP address and TCP/UDP port
- **src-port**—Source Layer 4 port
- **vlan-dst-ip**—VLAN number and destination IP address
- **vlan-dst-mixed-ip-port**—VLAN number and destination IP address and TCP/UDP port
- **vlan-src-dst-ip**—VLAN number and source and destination IP address
- **vlan-src-dst-mixed-ip-port**—VLAN number and source and destination IP address and TCP/UDP port
- **vlan-src-ip**—VLAN number and source IP address
- **vlan-src-mixed-ip-port**—VLAN number and source IP address and TCP/UDP port

The optional **module** keyword allows you to specify the load-balancing method for a specific module. This capability is supported only on DFC-equipped switching modules. You must enable per-module load balancing globally before configuring the feature on a module.

This example shows how to configure EtherChannel to use source and destination IP addresses:

```
Router# configure terminal
Router(config)# port-channel load-balance src-dst-ip
Router(config)# end
Router(config)#
```

This example shows how to verify the configuration:

```
Router# show etherchannel load-balance
Source XOR Destination IP address
```

Configuring the EtherChannel Hash-Distribution Algorithm

When you add a port to an EtherChannel or delete a port from an EtherChannel, the fixed algorithm updates the port ASIC for each port in the EtherChannel, which causes a short outage on each port.

The default adaptive algorithm does not need to update the port ASIC for existing member ports. You can configure a global value for the adaptive algorithm. You can also specify the algorithm for individual port channels.

When you change the algorithm, the change is applied at the next member link event (link down, link up, addition, deletion, no shutdown, and shutdown). When you enter the command to change the algorithm, the command console issues a warning that the command does not take effect until the next member link event.

**Note**

- Some external devices require the fixed algorithm. For example, the service control engine (SCE) requires incoming and outgoing packets to use the same port.
- If you change the load-balancing method, EtherChannel ports on DFC-equipped switching modules or on an active supervisor engine in a dual supervisor engine configuration will flap.

Configuring the Hash-Distribution Algorithm Globally

To configure the load-sharing algorithm globally, perform this task:

	Command	Purpose
Step 1	Router(config)# port-channel hash-distribution {adaptive fixed}	Sets the hash distribution algorithm to adaptive or fixed.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to globally set the hash distribution to adaptive:

```
Router(config)# port-channel hash-distribution adaptive
```

Configuring the Hash-Distribution Algorithm for a Port Channel

To configure the hash-distribution algorithm for a specific port channel, perform this task:

	Command	Purpose
Step 1	Router(config)# interface port-channel channel-num	Enters interface configuration mode for the port channel.
Step 2	Router(config-if)# port-channel port hash-distribution {adaptive fixed}	Sets the hash distribution algorithm for this interface
Step 3	Router(config-if)# end	Exits interface configuration mode.

This example shows how to set the hash distribution algorithm to adaptive on port channel 10:

```
Router (config)# interface port-channel 10
Router (config-if)# port-channel port hash-distribution adaptive
```

Configuring the EtherChannel Min-Links Feature

The EtherChannel min-links feature is supported on [LACP](#) EtherChannels. This feature allows you to configure the minimum number of member ports that must be in the link-up state and bundled in the EtherChannel for the port channel interface to transition to the link-up state. You can use the EtherChannel min-links feature to prevent low-bandwidth LACP EtherChannels from becoming active. This feature also causes LACP EtherChannels to become inactive if they have too few active member

ports to supply your required minimum bandwidth. In addition, when LACP max-bundle values are specified in conjunction with min-links, the configuration is verified and an error message is returned if the min-links value is not compatible with (equal to or less than) the max-bundle value.

To configure the EtherChannel min-links feature, perform this task:

	Command	Purpose
Step 1	Router(config)# interface port-channel <i>group_number</i>	Selects an LACP port channel interface.
Step 2	Router(config-if)# port-channel min-links <i>number</i>	Configures the minimum number of member ports that must be in the link-up state and bundled in the EtherChannel for the port channel interface to transition to the link-up state.
Step 3	Router(config-if)# end	Exits configuration mode.

**Note**

Although the EtherChannel min-links feature works correctly when configured only on one end of an EtherChannel, for best results, configure the same number of minimum links on both ends of the EtherChannel.

This example shows how to configure port channel interface 1 to be inactive if fewer than two member ports are active in the EtherChannel:

```
Router# configure terminal
Router(config)# interface port-channel 1
Router(config-if)# port-channel min-links 2
Router(config-if)# end
```

Configuring LACP 1:1 Redundancy

To configure the [LACP 1:1 redundancy feature](#), perform this task:

	Command	Purpose
Step 1	Router(config)# interface port-channel <i>group_number</i>	Selects an LACP port channel interface.
Step 2	Router(config-if)# lacp fast-switchover	Enables the LACP 1:1 redundancy feature on the EtherChannel.
Step 3	Router(config-if)# lacp max-bundle 1	Sets the maximum number of active member ports to be one. The only value supported with LACP 1:1 redundancy is “1”.
Step 4	Router(config-if)# lacp fast-switchover dampening <i>seconds</i>	(Optional) Enables the LACP 1:1 hot standby dampening feature for this EtherChannel. The range for the time parameter is 35 through 180 seconds.
Step 5	Router(config-if)# end	Exits configuration mode.

**Note**

LACP 1:1 redundancy must be enabled at both ends of the LACP EtherChannel.

This example shows how to configure an LACP EtherChannel with 1:1 redundancy. Because Gigabit Ethernet port 5/6 is configured with a higher port priority number (and therefore a lower priority) than the default of 32768, it will be the standby port.

```
Router# configure terminal
Router(config)# lacp system-priority 33000
Router(config)# interface range gigabitethernet 5/6 -7
Router(config-if)# channel-protocol lacp
Router(config-if)# channel-group 1 mode active
Router(config)# interface gigabitethernet 5/6
Router(config-if)# lacp port-priority 33000
Router(config)# interface port-channel 1
Router(config-if)# lacp fast-switchover
Router(config-if)# lacp max-bundle 1
Router(config-if)# lacp fast-switchover dampening 30
Router(config-if)# end
```

Configuring Auto Interleaved Port Priority For LACP Port Channels

To configure auto interleaved port priority for LACP on a port channel, perform this task on the port channel interface:

	Command	Purpose
Step 1	Router(config)# interface port-channel <i>channel-group</i>	Selects a port channel interface to configure.
Step 2	Router(config-if)# lacp active-port distribution automatic	Configures the port channel to use interleaved port priority.
		 Note You need to perform a shutdown and no shutdown for interleaved port priority to be enabled.
Step 3	Router(config-if)# shutdown	Disables the interface.
Step 4	Router(config-if)# no shutdown	Enables the interface.
Step 5	Router(config-if)# end	Exits configuration mode.
Step 6	Router# show etherchannel <i>channel-group</i> port-channel Router# show etherchannel <i>channel-group</i> summary	Verifies the configuration.

This example shows how to configure auto interleaved port priority on a port channel:

```
Router(config)# interface port-channel23
Router(config-if)# lacp active-port distribution automatic
Please shut/no shut the port-channel for configuration to take effect immediately.
Router(config-if)# shutdown
Router(config-if)# no shutdown

Router(config-if)# end
```

This example shows how to verify the configuration of port channel interface 23:

```
Router# show running interfaces port-channel23
Building configuration...

Current configuration : 81 bytes
!
```

```

interface Port-channel23
  no switchport
  no ip address
  lacp max-bundle 4
  lacp active-port distribution automatic
end

```

This example shows how to verify the configuration of EtherChannel 23:

Router# **show etherchannel 23 summary**

```

Flags:  D - down          P - bundled in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3       S - Layer2
        U - in use       N - not in use, no aggregation
        f - failed to allocate aggregator

        M - not in use, no aggregation due to minimum links not met
        m - not in use, port not aggregated due to minimum links not met
        u - unsuitable for bundling
        d - default port

        w - waiting to be aggregated
Number of channel-groups in use: 9
Number of aggregators:          9

Group  Port-channel  Protocol    Ports
-----+-----+-----+-----
23     Po23(RU)        LACP        Gi1/1/21(P)  Gi1/1/22(P)  Gi1/1/23(H)
                   Gi1/1/24(H)  Gi2/1/17(P)  Gi2/1/18(P)
                   Gi2/1/19(H)  Gi2/1/20(H)

Last applied Hash Distribution Algorithm: Fixed

```



Note

The above example shows that the four bundled ports are distributed 2 per chassis and slot.

Configuring LACP Port-Channel Standalone Disable

To disable the standalone EtherChannel member port state on a port channel (see [Table 23-2 on page 23-4](#)), perform this task on the port channel interface:

	Command	Purpose
Step 1	Router(config)# interface port-channel <i>channel-group</i>	Selects a port channel interface to configure.
Step 2	Router(config-if)# port-channel standalone-disable	Disables the standalone mode on the port-channel interface.
Step 3	Router(config-if)# end	Exits configuration mode.
Step 4	Router# show etherchannel <i>channel-group</i> port-channel Router# show etherchannel <i>channel-group</i> detail	Verifies the configuration.

This example shows how to disable the standalone EtherChannel member port state on port channel 42:

```

Router(config)# interface port-channel channel-group
Router(config-if)# port-channel standalone-disable

```

This example shows how to verify the configuration:

```
Router# show etherchannel 42 port-channel | include Standalone
Standalone Disable = enabled
Router# show etherchannel 42 detail | include Standalone
Standalone Disable = enabled
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

24

IEEE 802.1ak MVRP and MRP

- [Prerequisites for IEEE 802.1ak MVRP and MRP, page 24-1](#)
- [Restrictions for IEEE 802.1ak MVRP and MRP, page 24-1](#)
- [Information About IEEE 802.1ak MVRP and MRP, page 24-2](#)
- [Default Settings for IEEE 802.1ak MVRP and MRP, page 24-8](#)
- [How to Configure IEEE 802.1ak MVRP and MRP, page 24-8](#)
- [Troubleshooting the MVRP Configuration, page 24-10](#)
- [Configuration Examples for IEEE 802.1ak MVRP and MRP, page 24-11](#)

**Note**

- This feature appears in Cisco Feature navigator as “IEEE 802.1ak - MVRP and MRP.”
- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IEEE 802.1ak MVRP and MRP

None.

Restrictions for IEEE 802.1ak MVRP and MRP

- In releases where CSCta96338 is not resolved, a physical port with an MVRP configuration and enable state that differs from what is configured on a port-channel interface cannot become an active member of that EtherChannel.

- In releases where CSCta96338 is resolved, a physical port with an MVRP configuration and enable state that differs from what is configured on a port-channel interface can become an active member of the EtherChannel because the physical port will use the port-channel interface MVRP configuration and enable state.
- A non-Cisco device can interoperate with a Cisco device only through 802.1Q trunks.
- MVRP runs on ports where it is enabled. VTP pruning can run on ports where MVRP is not enabled.
- MVRP can be configured on both physical interfaces and EtherChannel interfaces, but is not supported on EtherChannel member ports.
- MVRP dynamic VLAN creation is not supported when the device is running in VTP server or client mode.
- MVRP and Connectivity Fault Management (CFM) can coexist but if the module does not have enough MAC address match registers to support both protocols, the MVRP ports on those modules are put in the error-disabled state. To use the ports that have been shut down, disable MVRP on the ports, and then enter **shutdown** and **no shutdown** commands.
- 802.1X authentication and authorization takes place after the port becomes active and before the Dynamic Trunking Protocol (DTP) negotiations start prior to MVRP running on the port.
- Do not enable MVRP automatic MAC address learning on edge switches that are configured with access ports. Enable MVRP automatic MAC address learning only on core switches where all the trunk interfaces are running MVRP.
- MVRP is supported only on Layer 2 trunks. MVRP is not supported on subinterfaces.

Information About IEEE 802.1ak MVRP and MRP

- [Overview, page 24-2](#)
- [Dynamic VLAN Creation, page 24-4](#)
- [MVRP Interoperability with VTP, page 24-4](#)
- [MVRP Interoperation with Non-Cisco Devices, page 24-6](#)
- [MVRP Interoperability with Other Software Features and Protocols, page 24-6](#)

Overview


The IEEE 802.1ak Multiple VLAN Registration Protocol (MVRP) supports dynamic registration and deregistration of VLANs on ports in a VLAN bridged network. IEEE 802.1ak uses more efficient Protocol Data Units (PDUs) and protocol design to provide better performance than the Generic VLAN Registration Protocol (GARP) VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP) protocols.

A VLAN-bridged network usually restricts unknown unicast, multicast, and broadcast traffic to those links that the traffic uses to access the appropriate network devices. In a large network, localized topology changes can affect the service over a much larger portion of the network. IEEE 802.1ak replaces GARP with the Multiple Registration Protocol (MRP), which provides improved resource utilization and bandwidth conservation.

With the 802.1ak MRP attribute encoding scheme, MVRP only needs to send one PDU that includes the state of all 4094 VLANs on a port. MVRP also transmits Topology Change Notifications (TCNs) for individual VLANs. This is an important feature for service providers because it allows them to localize

topology changes. Figure 24-1 illustrates MVRP deployed in a provider network on provider and customer bridges.

Figure 24-1 MVRP Deployed on Provider and Customer Bridges




274481

Because most providers do not wish to filter traffic by destination MAC addresses, a pruning protocol like MVRP is important in a Metro Ethernet provider network, which often uses thousands of VLANs.

Figure 24-2 displays redundant links that are configured between the access switch and two distribution switches on the cloud. When the link with VLAN 104 fails over, MVRP needs to send only one TCN for VLAN 104. Without MVRP, an STP TCN would need to be sent out for the whole MST region (VLANs 1-1000), which could cause unnecessary network interruption.

STP sets the tcDetected variable to signal MVRP that MVRP must decide whether to send an MVRP TCN. MVRP can flush filtering database entries rapidly on a per-VLAN basis following a topology change because when a port receives an attribute declaration marked as new, any entries in the filtering database for that port and for that VLAN are removed.

Figure 24-2 MVRP TCN Application



Dynamic VLAN Creation

Virtual Trunking Protocol (VTP) is a Cisco proprietary protocol that distributes VLAN configuration information across multiple devices within a VTP domain. When VTP is running on MVRP-aware devices, all of the VLANs allowed on the Cisco bridged LAN segments are determined by VTP.

Only the VTP transparent mode supports MVRP dynamic VLAN creation. When dynamic VLAN creation is disabled, the MVRP trunk ports can register and propagate the VLAN messages only for existing VLANs. MVRP PDUs and MVRP messages for the nonexistent VLANs are discarded.

For a switch to be configured in full compliance with the MVRP standard, the switch VTP mode must be transparent and MVRP dynamic VLAN creation must be enabled.

MVRP Interoperability with VTP

- [Overview, page 24-5](#)
- [VTP in Transparent or Off Mode, page 24-5](#)
- [VTP in Server or Client Mode and VTP Pruning is Disabled, page 24-5](#)
- [VTP in Server or Client Mode and VTP Pruning is Enabled, page 24-5](#)

Overview

The VLAN Trunking Protocol (VTP) is a Cisco proprietary protocol that distributes VLAN configuration information across multiple devices within a VTP domain. VTP pruning is an extension of VTP. It has its own Join message that can be exchanged with VTP PDUs. VTP PDUs can be transmitted on both 802.1Q trunks and ISL trunks. A VTP-capable device is in one of the VTP modes: server, client, transparent, or off.

When VTP Pruning and MVRP are both enabled globally, MVRP runs on trunks where it is enabled and VTP Pruning runs on other trunks. MVRP or VTP pruning can be enabled on a trunk, but not both.

VTP in Transparent or Off Mode

When VTP is in transparent or off mode, VTP pruning is not supported and VTP PDUs are not processed.

When a port receives an MVRP join message for a VLAN, the port transmits broadcast, multicast, and unknown unicast frames in that VLAN and adds the traffic definition to the MRP Attribute Propagation (MAP) port configured for that VLAN. The mapping is removed when the VLAN is no longer registered on the port.

For each interface that is forwarding in each VLAN, MVRP issues a join request to each MRP Attribute Declaration (MAD) instance and an MVRP Join message is sent out on each corresponding MVRP port.

MVRP dynamic VLAN creation can be enabled in VTP transparent or off mode. If it is enabled and the VLAN registered by a join message does not exist in the VLAN database in the device, then the VLAN will be created.

VTP in Server or Client Mode and VTP Pruning is Disabled

MVRP functions like VTP in transparent or off mode, except that MVRP dynamic VLAN creation is not allowed.

VTP in Server or Client Mode and VTP Pruning is Enabled

MVRP and VTP with pruning disabled can be supported on the same port and these two protocols need to communicate and exchange pruning information.

When VTP receives a VTP join message on a VTP trunk, MVRP is notified so that join request can be posted to the MVRP port MAD instances, and MVRP join messages are out on the MVRP ports to the MVRP network.

When VTP pruning removes a VLAN from a VTP trunk, MVRP sends a leave request to all the MAD instances and the MAD instances send a leave or empty message from the MVRP ports to indicate that the VLAN is not configured on the device.

When an MVRP port received an MVRP join message, MVRP propagates the event to other MVRP ports in the same MAP context, and notifies VTP so that VTP pruning can send a VTP join message from the VTP trunk ports.

If MVRP learns that a VLAN is no longer declared by the neighboring devices, MVRP sends a withdrawal event to VTP and then VTP pruning verifies that it should continue sending VTP join messages.

For VLANs that are configured as VTP pruning non-eligible on the VTP trunks, the VTP pruning state variables are set to joined for the VLANs. MVRP join requests are sent to those VLANs through the MVRP ports.

MVRP Interoperation with Non-Cisco Devices

Non-Cisco devices can interoperate with a Cisco device only through 802.1q trunks.

MVRP Interoperability with Other Software Features and Protocols

- [802.1x and Port Security, page 24-6](#)
- [DTP, page 24-6](#)
- [EtherChannel, page 24-7](#)
- [Flex Links, page 24-7](#)
- [High Availability, page 24-7](#)
- [ISSU and eFSU, page 24-7](#)
- [L2PT, page 24-7](#)
- [SPAN, page 24-7](#)
- [Unknown Unicast and Multicast Flood Control, page 24-7](#)
- [STP, page 24-7](#)
- [UDLR, page 24-7](#)
- [VLANs with MVRP, page 24-8](#)

802.1x and Port Security

802.1x authenticates and authorizes a port after it transitions to the link-up state, but before DTP negotiation occurs and MVRP runs on a port. Port security works independently of MVRP.

**Note**

When MVRP is globally enabled, the MVRP MAC address auto detect and provision feature is disabled by default (**mvrp mac-learning auto**). In some situations, MVRP MAC address auto detect and provision can disable MAC address learning and prevent correct port security operation. For example, on ports where port security is configured, when the number of streams exceeds the configured maximum number of MAC addresses, no port security violation occurs because MAC address learning is disabled, which prevents updates to port security about the streams coming into the port. To avoid incorrect port security operation, use caution when enabling the MVRP MAC address auto detect and provision feature on ports where port security is configured.

DTP

DTP negotiation occurs after ports transition to the link-up state and before transition to the forwarding state. If MVRP is administratively enabled globally and enabled on a port, it becomes operational when the port starts trunking.

EtherChannel

An EtherChannel port-channel interface can be configured as an MVRP participant. The EtherChannel member ports cannot be MVRP participants. MVRP learns the STP state of EtherChannel port-channel interfaces. The MAP context applies to the EtherChannel port-channel interfaces, but not to the EtherChannel member ports.

Flex Links

MVRP declares VLANs on STP forwarding ports but not on ports in the blocking state. On flex links ports, MVRP declares VLANs on the active ports but not on the standby ports. When a standby port takes over and an active port transitions to the link-down state, MVRP declares the VLANs on the newly active port.

High Availability

State Switchover (SSO) and ISSU supports MVRP.

ISSU and eFSU

Enhanced Fast Software Upgrade (EFSU) is an enhanced software upgrade procedure. MVRP is serviced by the ISSU client identified as ISSU_MVRP_CLIENT_ID.

L2PT

Layer 2 Protocol Tunneling (L2PT) does not support MVRP PDUs on 802.1Q tunnel ports.

SPAN

MVRP ports can be configured as either Switched Port Analyzer (SPAN) sources or destinations.

Unknown Unicast and Multicast Flood Control

MVRP and the Unknown Unicast and Multicast Flood Control feature, configured with the **switchport block** command, cannot be configured on the same port.

STP

An STP mode change causes forwarding ports to leave the forwarding state until STP reconverges in the newly configured mode. The reconvergence might cause an MVRP topology change because join messages might be received on different forwarding ports, and leave timers might expire on other ports.

UDLR

MVRP and unidirectional link routing (UDLR) cannot be configured on the same port.

VLANs with MVRP

- [VLAN Translation, page 24-8](#)
- [802.1Q Native VLAN Tagging, page 24-8](#)
- [Private VLANs, page 24-8](#)

VLAN Translation

VLAN translation and MVRP cannot be configured on the same port.

802.1Q Native VLAN Tagging

Other MVRP participants might not be able to accept tagged MVRP PDUs in the 802.1Q native VLAN. Compatibility between MVRP and 802.1Q native VLAN tagging depends on the specific network configuration.

Private VLANs

Private VLAN ports cannot support MVRP.

Default Settings for IEEE 802.1ak MVRP and MRP

None.

How to Configure IEEE 802.1ak MVRP and MRP

- [Enabling MVRP, page 24-8](#)
- [Enabling Automatic Detection of MAC Addresses, page 24-9](#)
- [Enabling MVRP Dynamic VLAN Creation, page 24-9](#)
- [Changing the MVRP Registrar State, page 24-10](#)

Enabling MVRP

MVRP must be enabled globally and on trunk ports. To enable MVRP, perform this task:

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# mvrp global	Globally enables MVRP.

	Command or Action	Purpose
Step 4	Router(config)# interface <i>type number</i>	Specifies a trunk port and enters interface configuration mode.
Step 5	Router(config-if)# mvrp	Enables MVRP on the interface. Note If MVRP is not successfully enabled on the port, the port is put in the errdisabled state. Enter the no mvrp command on the interface or the no mvrp global command to clear the errdisabled state.

This example shows how to enable MVRP globally and on an interface:

```
Router> enable
Router# configure terminal
Router(config)# mvrp global
Router(config)# interface FastEthernet 2/1
Router(config-if)# mvrp
```

Enabling Automatic Detection of MAC Addresses

MVRP automatic detection of MAC addresses is disabled by default. To enable MVRP automatic detection of MAC addresses on VLANs, perform this task:

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# mvrp mac-learning auto	Enables MAC address learning.

This example shows how to enable automatic MAC address learning:

```
Router> enable
Router# configure terminal
Router(config)# mvrp mac-learning auto
```

Enabling MVRP Dynamic VLAN Creation

To enable MVRP dynamic VLAN creation, perform this task:

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# vtp mode transparent	Sets VTP mode to transparent. Note Required for MVRP dynamic VLAN creation.
Step 4	Router(config)# mvrp vlan creation	Enables MVRP dynamic VLAN creation.

This example shows how to enable MVRP dynamic VLAN creation:

```
Router> enable
Router# configure terminal
Router(config)# vtp mode transparent
Router(config)# mvrp vlan create
```

Changing the MVRP Registrar State

The MRP protocol allows one participant per application in an end station, and one per application per port in a bridge. To set the MVRP registrar state, perform this task:

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# interface <i>type number</i>	Specifies an interface and enters interface configuration mode.
Step 4	Router(config-if)# mvrp registration [normal fixed forbidden]	Registers MVRP with the MAD instance.

This example shows how to set the MVRP registrar state to normal:

```
Router> enable
Router# configure terminal
Router(config)# interface FastEthernet 2/1
Router(config-if)# mvrp registration normal
```

Troubleshooting the MVRP Configuration

Use the **show mvrp summary** and **show mvrp interface** commands to display configuration information and interface states, and the **debug mvrp** command to enable all or a limited set of output messages related to an interface.

To troubleshoot the MVRP configuration, perform this task:

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# show mvrp summary	Displays the MVRP configuration.
Step 3	Router# show mvrp interface <i>interface-type port/slot</i>	Displays the MVRP interface states for the specified interface.
Step 4	Router# debug mvrp	Displays MVRP debugging information.
Step 5	Router# clear mvrp statistics	Clears MVRP statistics on all interfaces.

The following is sample output from the **show mvrp summary** command. This command can be used to display the MVRP configuration at the device level.

```
Router# show mvrp summary
```

```

MVRP global state           : enabled
MVRP VLAN creation          : disabled
VLANs created via MVRP      : 20-45, 3001-3050
Learning disabled on VLANs  : none

```

The following is sample output from the **show mvrp interface** command. This command can be used to display MVRP interface details of the administrative and operational MVRP states of all or one particular trunk port in the device.

```
Router# show mvrp interface
```

```

Port      Status   Registrar State
Fa3/1     off      normal

Port      Join Timeout  Leave Timeout  Leaveall Timeout
Fa3/1     201 600       700           1000

Port      Vlans Declared
Fa3/1     none

Port      Vlans Registered
Fa3/1     none

Port      Vlans Registered and in Spanning Tree Forwarding State
Fa3/1     none

```

Configuration Examples for IEEE 802.1ak MVRP and MRP

- [Enabling MVRP, page 24-11](#)
- [Enabling MVRP Automatic Detection of MAC Addresses, page 24-11](#)
- [Enabling Dynamic VLAN Creation, page 24-12](#)
- [Changing the MVRP Registrar State, page 24-12](#)

Enabling MVRP

The following example shows how to enable MVRP:

```

Router> enable
Router# configure terminal
Router(config)# mvrp global
Router(config)# interface fastethernet2/1
Router(config-if)# mvrp

```

Enabling MVRP Automatic Detection of MAC Addresses

The following example shows how to enable MAC address learning:

```

Router> enable
Router# configure terminal
Router(config)# mvrp mac-learning auto

```

Enabling Dynamic VLAN Creation

The following example shows how to enable dynamic VLAN creation:

```
Router> enable
Router# configure terminal
Router(config)# vtp mode transparent
Router(config)# mvrp vlan create
```

Changing the MVRP Registrar State

The following example shows how to change the MVRP registrar state:

```
Router> enable
Router# configure terminal
Router(config)# mvrp registration normal
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

25

VLAN Trunking Protocol (VTP)

- [Prerequisites for VTP, page 25-1](#)
- [Restrictions for VTP, page 25-1](#)
- [Information About VTP, page 25-2](#)
- [Default Settings for VTP, page 25-9](#)
- [How to Configure VTP, page 25-10](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for VTP

None.

Restrictions for VTP

- Supervisor engine redundancy does not support nondefault VLAN data filenames or locations. Do not enter the **vtp file *file_name*** command on a switch that has a redundant supervisor engine.
- Before installing a redundant supervisor engine, enter the **no vtp file** command to return to the default configuration.
- All network devices in a VTP domain must run the same VTP version.
- You must configure a password on each network device in the management domain when in secure mode.

**Caution**

If you configure VTP in secure mode, the management domain will not function properly if you do not assign a management domain password to each network device in the domain.

- A VTP version 2-capable network device can operate in the same VTP domain as a network device that runs VTP version 1 if VTP version 2 is disabled on the VTP version 2-capable network device (VTP version 2 is disabled by default).
- Do not enable VTP version 2 on a network device unless all of the network devices in the same VTP domain are version 2-capable. When you enable VTP version 2 on a network device, all of the version 2-capable network devices in the domain enable VTP version 2.
- In a Token Ring environment, you must enable VTP version 2 for Token Ring VLAN switching to function properly.
- When you enable or disable VTP pruning on a VTP server, VTP pruning for the entire management domain is enabled or disabled.
- The pruning-eligibility configuration applies globally to all trunks on the switch. You cannot configure pruning eligibility separately for each trunk.
- When you configure VLANs as pruning eligible or pruning ineligible, pruning eligibility for those VLANs is affected on that switch only, not on all network devices in the VTP domain.
- VTP version 1 and VTP version 2 do not propagate configuration information for extended-range VLANs (VLAN numbers 1006 to 4094). You must configure extended-range VLANs manually on each network device.
- VTP version 3 supports extended-range VLANs (VLAN numbers 1006 to 4094). If you convert from VTP version 3 to VTP version 2, the VLANs in the range 1006 to 4094 are removed from VTP control.
- VTP version 3 supports propagation of any database in a domain by allowing you to configure a primary and secondary server.
- The network administrator has to manually configure VTP version 3 on the switches that need to run VTP version 3.
- VTP version 3 is not supported on private VLAN (PVLAN) ports.
- Prior to configuring VTP version 3, you must ensure that the **spanning-tree extend system-id** command has been enabled.
- If there is insufficient DRAM available for use by VTP, the VTP mode changes to transparent.
- Network devices in VTP transparent mode do not send VTP Join messages. On trunk connections to network devices in VTP transparent mode, configure the VLANs that are used by the transparent-mode network devices or that need to be carried across trunks as pruning ineligible. For information about configuring prune eligibility, see the [“Configuring the List of Prune-Eligible VLANs” section on page 21-12](#).

Information About VTP

- [VTP Overview, page 25-3](#)
- [VTP Domains, page 25-3](#)
- [VTP Modes, page 25-4](#)
- [VTP Advertisements, page 25-4](#)
- [VTP Authentication, page 25-5](#)

- [VTP Version 2, page 25-5](#)
- [VTP Version 3, page 25-6](#)
- [VTP Pruning, page 25-7](#)
- [VLAN Interaction, page 25-9](#)

**Note**

For complete information on configuring VLANs, see [Chapter 26, “Virtual Local Area Networks \(VLANs\).”](#)

VTP Overview

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs within a VTP domain. A VTP domain (also called a VLAN management domain) is made up of one or more network devices that share the same VTP domain name and that are interconnected with trunks. VTP minimizes misconfigurations and configuration inconsistencies that can result in a number of 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. With VTP, you can make configuration changes centrally on one or more network devices and have those changes automatically communicated to all the other network devices in the network.

VTP Domains

A VTP domain (also called a VLAN management domain) is made up of one or more interconnected network devices that share the same VTP domain name. A network device can be configured to be in one and only one VTP domain. You make global VLAN configuration changes for the domain using either the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

VTP server mode is the default and the switch is in the no-management domain state until it receives an advertisement for a domain over a trunk link or you configure a management domain.

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 ignores advertisements with a different management domain name or an earlier configuration revision number.

If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch. The valid VLAN ranges are as follows:

- VTP version 1 and version 2 support VLANs 1 to 1000 only.
- VTP version 3 supports the entire VLAN range (VLANs 1 to 4094).
- The pruning of VLANs still applies to VLANs 1 to 1000 only.
- Extended-range VLANs are supported only in VTP version 3. If converting from VTP version 3 to VTP version 2, VLANs in the range 1006 to 4094 are removed from VTP control.

By default, all devices come up as secondary servers. You can enter the **vtp primary** privileged EXEC mode command to specify a primary server.

When using VTP version 1 and version 2, a VTP server is used to back up the database to the NVRAM and allows you to change the database information.

In VTP version 3, there is a VTP-primary server and a VTP-secondary server. A primary server allows you to alter the database information and the database updates sent out are honored by all the devices in the system. A secondary server can only back up the updated VTP configuration received from the primary server in the NVRAMs. The status of the primary and secondary servers is a runtime status and is not configurable.

VTP maps VLANs dynamically across multiple LAN types with unique names and internal index associations. Mapping eliminates excessive device administration required from network administrators.

VTP Modes

You can configure any one of these VTP modes:

- **Server**—In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as VTP version and VTP pruning) for the entire VTP domain. VTP servers advertise their VLAN configuration to other network devices in the same VTP domain and synchronize their VLAN configuration with other network devices based on advertisements received over trunk links. VTP server is the default mode.
- **Client**—VTP clients behave the same way as VTP servers, but you cannot create, change, or delete VLANs on a VTP client.
- **Transparent**—VTP transparent network devices do not participate in VTP. A VTP transparent network device does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, a transparent network device will forward received VTP advertisements from its trunking LAN ports. In VTP version 3, a transparent network device is specific to an instance.
- **Off**—In VTP off mode, a network device functions in the same manner as a VTP transparent device except that it does not forward VTP advertisements.



Note

The VTP server mode automatically changes from VTP server mode to VTP client mode if the switch detects a failure while writing configuration to NVRAM. If this happens, the switch cannot be returned to VTP server mode until the NVRAM is functioning.

VTP Advertisements

Each network device in the VTP domain sends periodic advertisements out each trunking LAN port to a reserved multicast address. VTP advertisements are received by neighboring network devices, which update their VTP and VLAN configurations as necessary.

The following global configuration information is distributed in VTP version 1 and version 2 advertisements:

- VLAN IDs.
- Emulated LAN names (for ATM LANE).
- 802.10 SAID values (FDDI).
- VTP domain name.
- VTP configuration revision number.
- VLAN configuration, including the maximum transmission unit (MTU) size for each VLAN.
- Frame format.

In VTP version 3, the information distributed in VTP version 1 and version 2 advertisements are supported, as well as the following information:

- A primary server ID.
- An instance number.
- A start index.
- An advertisement request is sent by a Client or a Server in these situations:
 - On a trunk coming up on a switch with an invalid database.
 - On all trunks when the database of a switch becomes invalid as a result of a configuration change or a takeover message.
 - On a specific trunk where a superior database has been advertised.
- VTP version 3 adds the following fields to the subset advertisement request:
 - A primary server ID.
 - An instance number.
 - A window size.
 - A start index.

VTP Authentication

When VTP authentication is not configured, the secret that is used to validate the received VTP updates is visible in plain text in the **show** commands and the NVRAM file, `const_nvram:vlan.dat`. In the event that a device in a VTP domain is compromised, the administrator must change the VTP secret across all the devices in the VTP domain.

With VTP version 3, you can configure the authentication password to be hidden using the **vtp password** command. When you configure the authentication password to be hidden, it does not appear in plain text in the configuration. Instead, the secret associated with the password is saved in hexadecimal format in the running configuration. The *password-string* argument is an ASCII string from 8 to 64 characters identifying the administrative domain for the device.

VTP Version 2

VTP version 2 supports the following features not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring LAN switching and VLANs (Token Ring Bridge Relay Function [TrBRF] and Token Ring Concentrator Relay Function [TrCRF]). For more information about Token Ring VLANs, see the [“Information About VLANs” section on page 26-2](#).
- Unrecognized Type-Length-Value (TLV) Support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs that it is not able to parse. The unrecognized TLV is saved in NVRAM.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent network device inspects VTP messages for the domain name and version and forwards a message only if the version and domain name match. Because only one domain is supported, VTP version 2 forwards VTP messages in transparent mode without checking the version.

- **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 digest on a received VTP message is correct, its information is accepted without consistency checks.

VTP Version 3

VTP version 3 supports all the features in version 1 and version 2. VTP version 3 also supports the following features not supported in version 1 and version 2:

- **Enhanced authentication**—In VTP version 3, you can configure the authentication password to be hidden using the **vtp password** command. When you configure the authentication password to be hidden, it does not appear in plain text in the configuration. Instead, the secret associated with the password is saved in hexadecimal format in the running configuration. The *password-string* argument is an ASCII string from 8 to 64 characters identifying the administrative domain for the device.

The **hidden** and **secret** keywords for VTP password are supported only in VTP version 3. If converting to VTP version 2 from VTP version 3, you must remove the **hidden** or **secret** keyword prior to the conversion.
- **Support for extended range VLAN database propagation**—VTP version 1 and version 2 support VLANs 1 to 1000 only. In VTP version 3, the entire VLAN range is supported (VLANs 1 to 4094). The pruning of VLANs still applies to VLANs 1 to 1000 only. Extended-range VLANs are supported in VTP version 3 only. Private VLANs are supported in VTP version 3. If you convert from VTP version 3 to VTP version 2, the VLANs in the range 1006 to 4094 are removed from VTP control.
- VLANs 1002 to 1005 are reserved VLANs in VTP version 1, version 2, and version 3.
- **Support for propagation of any database in a domain**—In VTP version 1 and version 2, a VTP server is used to back up the database to the NVRAM and allows you to change the database information.



Note

VTP version 3 supports Multiple Spanning Tree (802.1s) (MST) database propagation separate from the VLAN database only. In the MST database propagation, there is a VTP primary server and a VTP secondary server. A primary server allows you to alter the database information, and the database updates sent out are honored by all the devices in the system. A secondary server can only back up the updated VTP configuration received from the primary server in the NVRAMs. The status of the primary and secondary servers is a runtime status and is not configurable.

By default, all devices come up as secondary servers. You can enter the **vtp primary** privileged EXEC mode command to specify a primary server.

The primary-server status is needed only when database changes have to be performed and is obtained when the administrator issues a takeover message in the domain. The primary-server status is lost when you reload, switch over, or the domain parameters change. The secondary servers back up the configuration and continue to propagate the database. You can have a working VTP domain without any primary servers. Primary and secondary servers may exist on an instance in the domain.

In VTP version 3, there is no longer a restriction to propagate only VLAN database information. You can use VTP version 3 to propagate any database information across the VTP domain. A separate instance of the protocol is running for each application that uses VTP.

Two VTP version 3 regions can only communicate over a VTP version 1 or VTP version 2 region in transparent mode.

- CLI to turn off/on VTP on a per-trunk basis—You can enable VTP on a per-trunk basis using the **vtp** interface configuration mode command. You can disable VTP on a per-trunk basis using the **no** form of this command. When you disable VTP on the trunking port, all the VTP instances for that port are disabled. You will not be provided with the option of setting VTP to OFF for the MST database and ON for the VLAN database.

VTP on a global basis—When you set VTP mode to OFF globally, this applies to all the trunking ports in the system. Unlike the per-port configuration, you can specify the OFF option on a per-VTP instance basis. For example, the system could be configured as VTP-server for the VLAN database and as VTP-off for the MST database. In this case, VLAN databases are propagated by VTP, MST updates are sent out on the trunk ports in the system, and the MST updates received by the system are discarded.

VTP Pruning

VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. By default, VTP pruning is disabled.

In VTP versions 1 and 2, when you enable or disable pruning, it is propagated to the entire domain and accepted by all the devices in that domain. In VTP version 3, the domain administrator must manually enable or disable VTP pruning explicitly on each device.

For VTP pruning to be effective, all devices in the management domain must support VTP pruning. On devices that do not support VTP pruning, you must manually configure the VLANs allowed on trunks.

[Figure 25-1](#) shows a switched network without VTP pruning enabled. Interface 1 on network Switch 1 and port 2 on Switch 4 are assigned to the Red VLAN. A broadcast is sent from the host connected to Switch 1. Switch 1 floods the broadcast, and every network device in the network receives it, even though Switches 3, 5, and 6 have no ports in the Red VLAN.

You enable pruning globally on the switch (see the [“Enabling VTP Pruning”](#) section on page 25-12). You configure pruning on Layer 2 trunking LAN ports (see the [“Configuring a Layer 2 Switching Port as a Trunk”](#) section on page 21-8).



Figure 25-1 *Flooding Traffic without VTP Pruning*

Figure 25-2 shows the same switched network with VTP pruning enabled. The broadcast traffic from Switch 1 is not forwarded to Switches 3, 5, and 6 because traffic for the Red VLAN has been pruned on the links indicated (port 5 on Switch 2 and port 4 on Switch 4).

Figure 25-2 *Flooding Traffic with VTP Pruning*

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning eligible. VTP pruning does not prune traffic from pruning-ineligible VLANs. VLAN 1 is always pruning ineligible; traffic from VLAN 1 cannot be pruned.

To configure VTP pruning on a trunking LAN port, use the **switchport trunk pruning vlan** command (see the “[Configuring a Layer 2 Switching Port as a Trunk](#)” section on page 21-8). VTP pruning operates when a LAN port is trunking. You can set VLAN pruning eligibility when VTP pruning is enabled or disabled for the VTP domain, when any given VLAN exists or not, and when the LAN port is currently trunking or not.

VLAN Interaction

This section describes the VLAN interaction between devices with different VTP versions:

- [Interaction Between VTP Version 3 and VTP Version 2 Devices, page 25-9](#)
- [Interaction Between VTP Version 3 and VTP Version 1 Devices, page 25-9](#)

Interaction Between VTP Version 3 and VTP Version 2 Devices

When a VTP version 3 device on a trunk port receives messages from a VTP version 2 device, the VTP version 3 device sends a scaled-down version of the VLAN database on that particular trunk in a VTP version 2 format. A VTP version 3 device does not send out VTP version 2-formatted packets on a trunk port unless it first receives VTP version 2 packets on that trunk. If the VTP version 3 device does not receive VTP version 2 packets for an interval of time on the trunk port, the VTP version 3 device stops transmitting VTP version 2 packets on that trunk port.

Even when a VTP version 3 device detects a VTP version 2 device on a trunk port, the VTP version 3 device continues to send VTP version 3 packets in addition to VTP version 2 packets, to allow two kinds of neighbors to coexist on the trunk. VTP version 3 sends VTP version 3 and VTP version 2 updates on VTP version 2-detected trunks.

A VTP version 3 device does not accept configuration from a VTP version 2 (or VTP version 1) device.

Unlike in VTP version 2, when you configure the VTP version to be version 3, version 3 does not configure all the VTP version 3-capable devices in the domain to start behaving as VTP version 3 systems.

Interaction Between VTP Version 3 and VTP Version 1 Devices

When a VTP version 1 device that is capable of VTP version 2 or VTP version 3 receives a VTP version 3 packet, it will be configured as a VTP version 2 device if VTP version 2 conflicts do not exist.

VTP version 1-only capable devices cannot interoperate with VTP version 3 devices.

Default Settings for VTP

Feature	Default Value
VTP domain name	Null
VTP version 1 and version 2 mode	Server
VTP version 3 mode	The VTP version 3 VLAN database mode is the same as the VLAN database mode in VTP version 1 or 2 after the conversion from VTP version 1 or 2 to VTP version 3. For example, the VTP version 1 or 2 VLAN database mode is carried over to VTP version 3 VLAN database mode.
MST database mode	Transparent
VTP version 3 server type	Secondary
VTP version 2 state	Version 2 is disabled

Feature	Default Value
VTP password	None
VTP pruning	Disabled

How to Configure VTP

- [Configuring VTP Global Parameters, page 25-10](#)
- [Configuring the VTP Mode, page 25-15](#)
- [Configuring VTP Mode on a Per-Port Basis, page 25-16](#)
- [Displaying VTP Statistics, page 25-17](#)

Configuring VTP Global Parameters

- [Configuring VTP Version 1 and Version 2 Passwords, page 25-10](#)
- [Configuring VTP Version 3 Password, page 25-11](#)
- [Enabling VTP Pruning, page 25-12](#)
- [Enabling VTP Version 2, page 25-13](#)
- [Enabling VTP Version 3, page 25-13](#)



Note

You can enter the VTP global parameters in either global configuration mode or in EXEC mode.

Configuring VTP Version 1 and Version 2 Passwords

To configure the VTP version 1 and version 2 global parameters, perform this task:

Command	Purpose
Router(config)# vtp password <i>password-string</i>	Sets a password, which can be from 8 to 64 characters long, for the VTP domain.
Router(config)# no vtp password	Clears the password.

This example shows one way to configure a VTP password in global configuration mode:

```
Router# configure terminal
Router(config)# vtp password WATER
Setting device VLAN database password to WATER.
Router#
```

This example shows how to configure a VTP password in EXEC mode:

```
Router# vtp password WATER
Setting device VLAN database password to WATER.
Router#
```



Note

The password is not stored in the running-config file.

Configuring VTP Version 3 Password

To configure the VTP version 3 password, perform this task:

Command	Purpose
Router(config)# vtp password <i>password-string</i> [hidden secret]	Configures a password, which can be from 8 to 64 characters long or in 32-digit hexadecimal format, for the VTP domain. Note When entering the secret keyword, the <i>password-string</i> must be entered in 32-digit hexadecimal format.
Router(config)# no vtp password	Clears the password.

This example shows one way to configure a VTP password in global configuration mode:

```
Router# configure terminal
Router(config)# vtp password water
Setting device VTP database password to water.
Router#
```



Note

If you configure a VTP password in EXEC mode, the password is not stored in the running-config file.

This example shows one way to configure the password with a hidden key saved in hexadecimal format in the running configuration:

```
Router# configure terminal
Router(config)# vtp password 82214640C5D90868B6A0D8103657A721 hidden
Setting device VTP password
Router#
```

This example shows how you configure the password secret key in hexadecimal format:

```
Router# configure terminal
Router(config)# vtp password 300F060A2B0601035301020107010201 secret
Setting device VTP password
Router#
```

Configuring VTP Version 3 Server Type

To specify a primary server, perform this task:

Command	Purpose
Router# vtp primary [vlan mst] [force]	Configure this device as the primary server.

The **vtp primary** command does not have a **no** form. To return to the secondary server status, one of the following conditions must be met:

- System reload.
- Switchover between redundant supervisors.
- Takeover from another server.

- Change in the mode configuration.
- Any domain configuration change (version, domain name, domain password).

This example shows how to configure this device as the primary server if the password feature is disabled:

```
Router# vtp primary
This system is becoming primary server for feature vlan
No conflicting VTP version 3 devices found.
Do you want to continue? [confirm]y
Router#
```

This example shows how to configure this device as the primary server for the VTP VLAN feature if the password feature is disabled:

```
Router# vtp primary vlan
This system is becoming primary server for feature vlan
No conflicting VTP version 3 devices found.
Do you want to continue? [confirm]y
Router#
```

This example shows how to force this device to be the primary server for the VTP MST feature if the password feature is disabled:

```
Router# vtp primary mst force
This system is becoming primary server for feature MST
No conflicting VTP version 3 devices found.
Do you want to continue? [confirm]y
Router#
```

This example shows how to force this device to be the primary server for the VTP MST feature when the domain VTP password is set with the **hidden** or **secret** keyword:

```
Router# vtp primary mst force
Enter VTP password: water1
This switch is becoming Primary server for mst 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
Router#
```

Enabling VTP Pruning

To enable VTP pruning in the management domain, perform this task:

Command	Purpose
Router(config)# vtp pruning	Enables VTP pruning in the management domain.

This example shows one way to enable VTP pruning in the management domain:

```
Router# configure terminal
Router(config)# vtp pruning
Pruning switched ON
```

This example shows how to enable VTP pruning in the management domain with any release:

```
Router# vtp pruning
Pruning switched ON
```

This example shows how to verify the configuration:

```
Router# show vtp status | include Pruning
VTP Pruning Mode: Enabled
Router#
```

For information about configuring prune eligibility, see the [“Configuring the List of Prune-Eligible VLANs”](#) section on page 21-12.

Enabling VTP Version 2

VTP version 2 is disabled by default on VTP version 2-capable network devices. When you enable VTP version 2 on a network device, every VTP version 2-capable network device in the VTP domain enables version 2.



Caution

VTP version 1 and VTP version 2 are not interoperable on network devices in the same VTP domain. Every network device in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every network device in the VTP domain supports version 2.



Note

In a Token Ring environment, you must enable VTP version 2 for Token Ring VLAN switching to function properly on devices that support Token Ring interfaces.

To enable VTP version 2, perform this task:

Command	Purpose
Router(config)# vtp version 2	Enables VTP version 2.

This example shows one way to enable VTP version 2:

```
Router# configure terminal
Router(config)# vtp version 2
V2 mode enabled.
Router(config)#
```

This example shows how to enable VTP version 2 with any release:

```
Router# vtp version 2
V2 mode enabled.
Router#
```

This example shows how to verify the configuration:

```
Router# show vtp status | include V2
VTP V2 Mode: Enabled
Router#
```

Enabling VTP Version 3

VTP version 3 is disabled by default. You can enable version 3 in global configuration mode only. The network administrator has to manually configure VTP version 3 on the switches that need to run VTP version 3.

**Note**

Prior to configuring VTP version 3, you must ensure that the **spanning-tree extend system-id** command has been enabled.

**Caution**

In VTP version 3, both the primary and secondary servers may exist on an instance in the domain.

To enable VTP version 3, perform this task:

Command	Purpose
Router(config)# vtp version 3	Enables VTP version 3.

This example shows one way to enable VTP version 3:

```
Router# configure terminal
Router(config)# vtp version 3
Router(config)#
```

This example shows how to verify the configuration:

```
Router# show vtp status
VTP Version capable          : 1 to 3
VTP version running          : 3
VTP Domain Name              : lab_switch
VTP Pruning Mode             : Disabled
VTP Traps Generation         : Disabled
Device ID                    : 0015.c724.0040

Feature VLAN:
-----
VTP Operating Mode           : Server
Number of existing VLANs     : 6
Number of existing extended VLANs : 0
Configuration Revision       : 0
Primary ID                   : 0000.0000.0000
Primary Description          :
MD5 digest                   : 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00
                               0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00

Feature MST:
-----
VTP Operating Mode           : Transparent

Feature UNKNOWN:
-----
VTP Operating Mode           : Transparent
Router#
```

Configuring the VTP Mode

To configure the VTP mode, perform this task:

	Command	Purpose
Step 1	Router(config)# vtp mode { client server transparent off } { vlan mst unknown }	Configures the VTP mode.
Step 2	Router(config)# vtp domain <i>domain-name</i>	(Optional for server mode) Defines the VTP domain name, which can be up to 32 characters long. 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. Note You cannot clear the domain name.
Step 3	Router(config)# end	Exits VLAN configuration mode.



Note

When VTP is disabled, you can enter VLAN configuration commands in configuration mode instead of the VLAN database mode and the VLAN configuration is stored in the startup configuration file.

This example shows how to configure the switch as a VTP server:

```
Router# configuration terminal
Router(config)# vtp mode server
Setting device to VTP SERVER mode.
Router(config)# vtp domain lab_network
Setting VTP domain name to lab_network
Router(config)# end
Router#
```

This example shows how to configure the switch as a VTP client:

```
Router# configuration terminal
Router(config)# vtp mode client
Setting device to VTP CLIENT mode.
Router(config)# exit
Router#
```

This example shows how to disable VTP on the switch:

```
Router# configuration terminal
Router(config)# vtp mode transparent
Setting device to VTP TRANSPARENT mode.
Router(config)# end
Router#
```

This example shows how to disable VTP on the switch and to disable VTP advertisement forwarding:

```
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# vtp mode off
Setting device to VTP OFF mode.
Router(config)# exit
Router#
```

This example shows how to verify the configuration:

```
Router# show vtp status
VTP Version capable      : 1 to 3
VTP version running      : 3
VTP Domain Name          : lab_network
VTP Pruning Mode         : Disabled
VTP Traps Generation     : Disabled
Device ID                : 0015.c724.0040

Feature VLAN:
-----
VTP Operating Mode       : Server
Number of existing VLANs : 6
Number of existing extended VLANs : 0
Configuration Revision   : 0
Primary ID               : 0000.0000.0000
Primary Description      :
MD5 digest               : 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00
                        : 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00

Feature MST:
-----
VTP Operating Mode       : Transparent

Feature UNKNOWN:
-----
VTP Operating Mode       : Transparent

Router#
```

Configuring VTP Mode on a Per-Port Basis

You can configure VTP mode on a per-port basis. The VTP enable value will be applied only when a port becomes switched port in trunk mode. Incoming and outgoing vtp pdus are blocked; *not* forwarded. In VTP version 3, you can also configure VTP mode on a per-trunk basis. To configure VTP mode, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 2	Router(config-if)# vtp	Enables VTP on the specified port.
Step 3	Router(config-if)# end	Exits interface configuration mode.

This example shows how to configure VTP mode on a port:

```
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 3/5
Router(config-if)# vtp
Router(config-if)# end
Router#
```


This example shows how to disable VTP mode on a port:

```
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 3/5
Router(config-if)# no vtp
Router(config-if)# end
Router#
```

This example shows how to verify the configuration change:

```
Router# show vtp interface gigabitethernet 3/5
```

Interface	VTP Status
GigabitEthernet3/5	disabled

```
Router#
```

This example shows how to verify the interface:

```
Router# show vtp interface
```

Interface	VTP Status
GigabitEthernet3/1	enabled
GigabitEthernet3/2	enabled
GigabitEthernet3/3	enabled
GigabitEthernet3/4	enabled
GigabitEthernet3/5	disabled
GigabitEthernet3/6	enabled
...	

Displaying VTP Statistics

To display VTP statistics, including VTP advertisements sent and received and VTP errors, perform this task:

Command	Purpose
Router# show vtp counters	Displays VTP statistics.

This example shows how to display VTP statistics:

```
Router# show vtp counters
VTP statistics:
Summary advertisements received      : 7
Subset advertisements received      : 5
Request advertisements received     : 0
Summary advertisements transmitted : 997
Subset advertisements transmitted   : 13
Request advertisements transmitted  : 3
Number of config revision errors    : 0
Number of config digest errors      : 0
Number of V1 summary errors         : 0
```

VTP pruning statistics:

Trunk	Join Transmitted	Join Received	Summary advts received from non-pruning-capable device
-----	-----	-----	-----
Gi5/8	43071	42766	5



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

26

Virtual Local Area Networks (VLANs)

- [Prerequisites for VLANs, page 26-1](#)
- [Restrictions for VLANs, page 26-2](#)
- [Information About VLANs, page 26-2](#)
- [Default Settings for VLANs, page 26-3](#)
- [How to Configure VLANs, page 26-4](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for VLANs

The following recommendations apply to Fabric Extender (FEX) VLANs:

- FEX-configurable VLAN range—1000 VLANs of any 4K VLANs
- Total number of VLANs—Not to exceed 20 VLANs per FEX
- Maximum number of trunks—Not to exceed 4 trunks

Restrictions for VLANs

- If the switch is in VTP server or transparent mode (see the [“How to Configure VTP”](#) section on [page 25-10](#)), you can configure VLANs in global and config-vlan configuration modes. When you configure VLANs in global and config-vlan configuration modes, the VLAN configuration is saved in the vlan.dat files. To display the VLAN configuration, enter the **show vlan** command.

If the switch is in VLAN transparent mode, use the copy **running-config startup-config** command to save the VLAN configuration to the startup-config file. After you save the running configuration as the startup configuration, use the **show running-config** and **show startup-config** commands to display the VLAN configuration.

- When the switch boots, if the VTP domain name and the VTP mode in the startup-config file and vlan.dat files do not match, the switch uses the configuration in the vlan.dat file.
- You can configure extended-range VLANs only in global configuration mode.
- Supervisor engine redundancy does not support nondefault VLAN data file names or locations. Do not enter the **vtp file file_name** command on a switch that has a redundant supervisor engine.
- Before installing a redundant supervisor engine, enter the **no vtp file** command to return to the default configuration.
- Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. For information on configuring VTP, see [Chapter 25, “VLAN Trunking Protocol \(VTP\)”](#).
- The VLAN configuration is stored in the vlan.dat file, which is stored in nonvolatile memory. You can cause inconsistency in the VLAN database if you manually delete the vlan.dat file.
- To do a complete backup of your configuration, include the vlan.dat file in the backup.

Information About VLANs

- [VLAN Overview, page 26-2](#)
- [VLAN Ranges, page 26-2](#)

VLAN Overview

A VLAN is a group of end stations with a common set of requirements, independent of physical location. VLANs have the same attributes as a physical LAN but allow you to group end stations even if they are not located physically on the same LAN segment.

VLANs are usually associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. Traffic between VLANs must be routed. LAN port VLAN membership is assigned manually on an port-by-port basis.

VLAN Ranges



Note

You must enable the extended system ID to use 4094 VLANs (see the [“Information about the Bridge ID”](#) section on [page 31-3](#)).

Cisco IOS Release 15.4SY supports 4094 VLANs in accordance with the IEEE 802.1Q standard. These VLANs are organized into several ranges; you use each range slightly differently. Some of these VLANs are propagated to other switches in the network when you use the VLAN Trunking Protocol (VTP). The extended-range VLANs are not propagated, so you must configure extended-range VLANs manually on each network device.

Table 26-1 describes the VLAN ranges.

Table 26-1 **VLAN Ranges**

VLANs	Range	Usage	Propagated by VTP
0, 4095	Reserved	For system use only. You cannot see or use these VLANs.	—
1	Normal	Cisco default. You can use this VLAN but you cannot delete it.	Yes
2–1001	Normal	For Ethernet VLANs; you can create, use, and delete these VLANs.	Yes
1002–1005	Normal	Cisco defaults for FDDI and Token Ring. You cannot delete VLANs 1002–1005.	Yes
1006–4094	Extended	For Ethernet VLANs only.	No
3968–4031	MET VLANs	Reserved for internal usage.	—

The following information applies to VLAN ranges:

- Layer 3 LAN ports, WAN interfaces and subinterfaces, and some software features use internal VLANs in the extended range. You cannot use an extended range VLAN that has been allocated for internal use.
- To display the VLANs used internally, enter the **show vlan internal usage** command. With earlier releases, enter the **show vlan internal usage** and **show cwan vlans** commands.
- You can configure ascending internal VLAN allocation (from 1006 and up) or descending internal VLAN allocation (from 4094 and down).
- You must enable the extended system ID to use extended range VLANs (see the [“Information about the Bridge ID”](#) section on page 31-3).

Default Settings for VLANs

- VLAN ID: 1; range: 1–4094
- VLAN name:
 - VLAN 1: “default”
 - Other VLANs: “VLANvlan_ID”
- 802.1Q SAID: 10vlan_ID; range: 100001–104094
- MTU size: 1500; range: 1500–18190
- Translational bridge 1: 0; range: 0–1005
- Translational bridge 2: 0; range: 0–1005
- VLAN state: active: active, suspend

- Pruning eligibility:
 - VLANs 2–1001 are pruning eligible
 - VLANs 1006–4094 are pruning eligible

How to Configure VLANs

- [Configurable VLAN Parameters, page 26-4](#)
- [VLAN Locking, page 26-4](#)
- [Creating or Modifying an Ethernet VLAN, page 26-5](#)
- [Assigning a Layer 2 LAN Interface to a VLAN, page 26-6](#)
- [Configuring the Internal VLAN Allocation Policy, page 26-6](#)
- [Configuring VLAN Translation, page 26-7](#)
- [Saving VLAN Information, page 26-9](#)

Configurable VLAN Parameters



Note

- Ethernet VLAN 1 uses only default values.
- Except for the VLAN name, Ethernet VLANs 1006 through 4094 use only default values.
- You can configure the VLAN name for Ethernet VLANs 1006 through 4094.

You can configure the following parameters for VLANs 2 through 1001:

- VLAN name
- VLAN type (Ethernet, FDDI, FDDI network entity title [NET], TrBRF, or TrCRF)
- VLAN state (active or suspended)
- 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 Locking

The VLAN locking feature provides an extra level of verification to ensure that you have configured the intended VLAN. When VLAN locking is enabled, you need to specify the VLAN name when you change a port from one VLAN to another. This feature affects **switchport** commands (in interface configuration mode) that specify the VLANs or private VLANs for access and trunk ports.

For additional information about how to configure access and trunk ports with VLAN locking enabled, see the [“How to Configure LAN Interfaces for Layer 2 Switching” section on page 21-5](#).

For additional information about how to configure ports in private VLANs with VLAN locking enabled, see the [“How to Configure Private VLANs” section on page 27-10](#).

By default, the VLAN locking is disabled. To enable VLAN locking, perform this task:

Command	Purpose
Router(config)# vlan port provisioning	Enables VLAN locking.

Creating or Modifying an Ethernet VLAN

User-configured VLANs have unique IDs from 1 to 4094, except for reserved VLANs (see [Table 26-1 on page 26-3](#)). Enter the **vlan** command with an unused ID to create a VLAN. Enter the **vlan** command for an existing VLAN to modify the VLAN (you cannot modify an existing VLAN that is being used by a Layer 3 port or a software feature).

See the [“Default Settings for VLANs” section on page 26-3](#) for the list of default parameters that are assigned when you create a VLAN. If you do not specify the VLAN type with the **media** keyword, the VLAN is an Ethernet VLAN.

To create or modify a VLAN, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters configuration mode.
Step 2	Router(config)# vlan <i>vlan_ID</i> { [- <i>vlan_ID</i>] [, <i>vlan_ID</i>] }	Creates or modifies an Ethernet VLAN, a range of Ethernet VLANs, or several Ethernet VLANs specified in a comma-separated list (do not enter space characters).
Step 3	Router(config-vlan)# end	Updates the VLAN database and returns to privileged EXEC mode.

When you create or modify an Ethernet VLAN, note the following information:

- Because Layer 3 ports and some software features require internal VLANs allocated from 1006 and up, configure extended-range VLANs starting with 4094.
- You can configure extended-range VLANs only in global configuration mode. You cannot configure extended-range VLANs in VLAN database mode.
- Layer 3 ports and some software features use extended-range VLANs. If the VLAN you are trying to create or modify is being used by a Layer 3 port or a software feature, the switch displays a message and does not modify the VLAN configuration.

When deleting VLANs, note the following information:

- You cannot delete the default VLANs for the different media types: Ethernet VLAN 1 and FDDI or Token Ring VLANs 1002 to 1005.
- When you delete a VLAN, any LAN ports configured as access ports assigned to that VLAN become inactive. The ports remain associated with the VLAN (and inactive) until you assign them to a new VLAN.

This example shows how to create an Ethernet VLAN in global configuration mode and verify the configuration:

```
Router# configure terminal
```

```
Router(config)# vlan 3
Router(config-vlan)# end
Router# show vlan id 3
```

VLAN	Name	Status	Ports
3	VLAN0003	active	

VLAN	Type	SAID	MTU	Parent	RingNo	BridgeNo	Stp	BrdgMode	Trans1	Trans2
3	enet	100003	1500	-	-	-	-	-	0	0

Primary	Secondary	Type	Interfaces

Assigning a Layer 2 LAN Interface to a VLAN

A VLAN created in a management domain remains unused until you assign one or more LAN ports to the VLAN.



Note

Make sure you assign LAN ports to a VLAN of the appropriate type. Assign Ethernet ports to Ethernet-type VLANs.

To assign one or more LAN ports to a VLAN, complete the procedures in the [“How to Configure LAN Interfaces for Layer 2 Switching”](#) section on page 21-5.

Configuring the Internal VLAN Allocation Policy

For more information about VLAN allocation, see the [“VLAN Ranges”](#) section on page 26-2.



Note

The internal VLAN allocation policy is applied only following a reload.

To configure the internal VLAN allocation policy, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan internal allocation policy {ascending descending}	Configures the internal VLAN allocation policy.
Step 2	Router(config)# end	Exits configuration mode.
Step 3	Router# reload	Applies the new internal VLAN allocation policy.
		 Caution You do not need to enter the reload command immediately. Enter the reload command during a planned maintenance window.

When you configure the internal VLAN allocation policy, note the following information:

- Enter the **ascending** keyword to allocate internal VLANs from 1006 and up.
- Enter the **descending** keyword to allocate internal VLAN from 4094 and down.

This example shows how to configure descending as the internal VLAN allocation policy:

```
Router# configure terminal
Router(config)# vlan internal allocation policy descending
```

Configuring VLAN Translation

- [VLAN Translation Guidelines and Restrictions, page 26-7](#)
- [Configuring VLAN Translation on a Trunk Port, page 26-8](#)
- [Enabling VLAN Translation on Other Ports in a Port Group, page 26-9](#)



Note

- To avoid spanning tree loops, be careful not to misconfigure the VLAN translation feature.
- On trunk ports, you can translate one VLAN number to another VLAN number, which transfers all traffic received in one VLAN to the other VLAN.

VLAN Translation Guidelines and Restrictions

When translating VLANs, follow these guidelines and restrictions:

- A VLAN translation configuration is inactive if it is applied to ports that are not Layer 2 trunks.
- Do not configure translation of ingress native VLAN traffic on an 802.1Q trunk. Because 802.1Q native VLAN traffic is untagged, it cannot be recognized for translation. You can translate traffic from other VLANs to the native VLAN of an 802.1Q trunk.
- Do not remove the VLAN to which you are translating from the trunk.
- The VLAN translation configuration applies to all ports in a port group. VLAN translation is disabled by default on all ports in a port group. Enable VLAN translation on ports as needed.
- Cisco IOS Release 15.4SY supports only IEEE 802.1Q trunking.

Table 26-2 **Module Support for VLAN Translation**

Product Number	Number of Ports	Number of Port Groups	Port Ranges per Port Group	Translations per Port Group
C6800-SUP6T-XL C6800-SUP6T	Depends on the operational port mode of the Sup6T.*	2	Depends on the operational port mode of the Sup6T..*	16
WS-X6904-40G-2T	See the Release Notes .			16
WS-X6908-10GE	8	8	1 port in each group	16
WS-X6816-10T-2T, WS-X6716-10T	16	16	1 port in each group	16
WS-X6816-10G-2T, WS-X6716-10GE	16	16	1 port in each group	16
WS-X6704-10GE	4	4	1 port in each group	128

Table 26-2 **Module Support for VLAN Translation (continued)**

Product Number	Number of Ports	Number of Port Groups	Port Ranges per Port Group	Translations per Port Group
WS-X6848-TX-2T, WS-X6748-GE-TX	48	4	1–12 13–24 25–36 37–48	128
WS-X6848-SFP-2T, WS-X6748-SFP	48	4	1–23, odd 2–24, even 25–47, odd 26–48, even	128
WS-X6824-SFP-2T, WS-X6724-SFP	24	2	1–12 13–24	128

**Note**

To configure a port as a trunk, see the [“Configuring a Layer 2 Switching Port as a Trunk”](#) section on [page 21-8](#).

*Following are the different combinations based on the operational port mode on Supervisor Engine 6T:

- Default mode: port-group1 had 4 ports and port-group2 has 2 ports
- 16x10G Over-subscription mode: both port-group1 & port-group2 has 8 ports each.
- 8x10G Performance mode: both port-group1 & port-group2 has 4 ports each.
- 2x40G Performance mode: both port-group1 & port-group2 has 1 port each.

Configuring VLAN Translation on a Trunk Port

To translate VLANs on a trunk port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface type slot/port	Selects the Layer 2 trunk port to configure.
Step 2	Router(config-if)# switchport vlan mapping enable	Enables VLAN translation.
Step 3	Router(config-if)# switchport vlan mapping original_vlan_ID translated_vlan_ID	Translates a VLAN to another VLAN. The valid range is 1 to 4094. When you configure a VLAN mapping from the original VLAN to the translated VLAN on a port, traffic arriving on the original VLAN gets mapped or translated to the translated VLAN at the ingress of the switch port, and the traffic internally tagged with the translated VLAN gets mapped to the original VLAN before leaving the switch port. This method of VLAN mapping is a two-way mapping.
Step 4	Router(config-if)# end	Exits configuration mode.

This example shows how to map VLAN 1649 to VLAN 755 Gigabit Ethernet port 5/2:

```
Router# configure terminal
```

```
Router(config)# interface gigabitethernet 5/2
Router(config-if)# switchport vlan mapping 1649 755
Router(config-if)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show interface gigabitethernet 5/2 vlan mapping
State: enabled
Original VLAN Translated VLAN
-----
1649                755
```

Enabling VLAN Translation on Other Ports in a Port Group

To enable VLAN translation on other ports in a port group, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport vlan mapping enable	Enables VLAN translation.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to enable VLAN translation on a port:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/2
Router(config-if)# switchport vlan mapping enable
Router(config-if)# end
```

Saving VLAN Information

The VLAN database is stored in the vlan.dat file. You should create a backup of the vlan.dat file in addition to backing up the running-config and startup-config files. If you replace the existing supervisor engine, copy the startup-config file as well as the vlan.dat file to restore the system. The vlan.dat file is read on bootup and you will have to reload the supervisor engine after uploading the file. To view the file location, use the **dir vlan.dat** command. To copy the file (binary), use the **copy vlan.dat tftp** command.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

27

Private VLANs

- [Prerequisites for Private VLANs, page 27-1](#)
- [Restrictions for Private VLANs, page 27-1](#)
- [Information About Private VLANs, page 27-5](#)
- [Default Settings for Private VLANs, page 27-10](#)
- [How to Configure Private VLANs, page 27-10](#)
- [Monitoring Private VLANs, page 27-16](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Private VLANs

None.

Restrictions for Private VLANs

- [Secondary and Primary VLANs, page 27-2](#)
- [Private VLAN Ports, page 27-4](#)
- [Limitations with Other Features, page 27-4](#)

Secondary and Primary VLANs

- After you configure a private VLAN and set VTP to transparent mode, you are not allowed to change the VTP mode to client or server. For information about VTP, see [Chapter 25, “VLAN Trunking Protocol \(VTP\).”](#)
- After you have configured private VLANs, use the **copy running-config startup config** privileged EXEC command to save the VTP transparent mode configuration and private VLAN configuration in the startup-config file. If the switch resets it must default to VTP transparent mode to support private VLANs.
- In VTP versions 1 and 2, VTP does not propagate a private VLAN configuration and you must configure private VLANs on each device where you want private VLAN ports. In VTP version 3, VTP does propagate private VLAN configurations automatically.
- You cannot configure VLAN 1 or VLANs 1002 to 1005 as primary or secondary VLANs. Extended VLANs (VLAN IDs 1006 to 4094) cannot belong to private VLANs. Only Ethernet VLANs can be private VLANs.
- A primary VLAN can have one isolated VLAN and multiple community VLANs associated with it. An isolated or community VLAN can have only one primary VLAN associated with it.
- When a secondary VLAN is associated with the primary VLAN, the STP parameters of the primary VLAN, such as bridge priorities, are propagated to the secondary VLAN. However, STP parameters do not necessarily propagate to other devices. You should manually check the STP configuration to ensure that the primary, isolated, and community VLANs’ spanning tree topologies match so that the VLANs can properly share the same forwarding database.
- If you enable MAC address reduction on the switch, we recommend that you enable MAC address reduction on all the devices in your network to ensure that the STP topologies of the private VLANs match.
- In a network where private VLANs are configured, if you enable MAC address reduction on some devices and disable it on others (mixed environment), use the default bridge priorities to make sure that the root bridge is common to the primary VLAN and to all its associated isolated and community VLANs. Be consistent with the ranges employed by the MAC address reduction feature regardless of whether it is enabled on the system. MAC address reduction allows only discrete levels and uses all intermediate values internally as a range. You should disable a root bridge with private VLANs and MAC address reduction, and configure the root bridge with any priority higher than the highest priority range used by any nonroot bridge.
- You cannot apply VACLs to secondary VLANs. (See [Chapter 36, “VLAN ACLs \(VACLs\).”](#))
- You can enable DHCP snooping on private VLANs. When you enable DHCP snooping on the primary VLAN, it is propagated to the secondary VLANs. If you configure DHCP on a secondary VLAN, the configuration does not take effect if the primary VLAN is already configured.
- We recommend that you prune the private VLANs from the trunks on devices that carry no traffic in the private VLANs.
- You can apply different quality of service (QoS) configurations to primary, isolated, and community VLANs. (See [Chapter 24, “PFC QoS Overview.”](#))
- When you configure private VLANs, sticky Address Resolution Protocol (ARP) is enabled by default, and ARP entries learned on Layer 3 private VLAN interfaces are sticky ARP entries. For security reasons, private VLAN port sticky ARP entries do not age out. For information about configuring sticky ARP, see the [“Configuring Sticky ARP” section on page 38-9](#).
- We recommend that you display and verify private VLAN interface ARP entries.

- Sticky ARP prevents MAC address spoofing by ensuring that ARP entries (IP address, MAC address, and source VLAN) do not age out. You can configure sticky ARP on a per-interface basis. For information about configuring sticky ARP, see the [“Configuring Sticky ARP” section on page 38-9](#). The following guidelines and restrictions apply to private VLAN sticky ARP:

- ARP entries learned on Layer 3 private VLAN interfaces are sticky ARP entries.
- Connecting a device with a different MAC address but with the same IP address generates a message and the ARP entry is not created.
- Because the private VLAN port sticky ARP entries do not age out, you must manually remove private VLAN port ARP entries if a MAC address changes. You can add or remove private VLAN ARP entries manually as follows:

```
Router(config)# no arp 11.1.3.30
IP ARP:Deleting Sticky ARP entry 11.1.3.30

Router(config)# arp 11.1.3.30 0000.5403.2356 arpa
IP ARP:Overwriting Sticky ARP entry 11.1.3.30, hw:00d0.bb09.266e by
hw:0000.5403.2356
```

- You can configure VLAN maps on primary and secondary VLANs. (See the [“Applying a VLAN Access Map” section on page 36-4](#).) However, we recommend that you configure the same VLAN maps on private VLAN primary and secondary VLANs.
- When a frame is Layer 2 forwarded within a private VLAN, the same VLAN map is applied at the ingress side and at the egress side. When a frame is routed from inside a private VLAN to an external port, the private VLAN map is applied at the ingress side.
 - For frames going upstream from a host port to a promiscuous port, the VLAN map configured on the secondary VLAN is applied.
 - For frames going downstream from a promiscuous port to a host port, the VLAN map configured on the primary VLAN is applied.

To filter out specific IP traffic for a private VLAN, you should apply the VLAN map to both the primary and secondary VLANs.

- To apply Cisco IOS output ACLs to all outgoing private VLAN traffic, configure them on the Layer 3 VLAN interface of the primary VLAN. (See [Chapter 34, “MAC Address-Based Traffic Blocking”](#).)
- Cisco IOS ACLs applied to the Layer 3 VLAN interface of a primary VLAN automatically apply to the associated isolated and community VLANs.
- Do not apply Cisco IOS ACLs to isolated or community VLANs. Cisco IOS ACL configuration applied to isolated and community VLANs is inactive while the VLANs are part of the private VLAN configuration.
- Although private VLANs provide host isolation at Layer 2, hosts can communicate with each other at Layer 3.
- Private VLANs support these Switched Port Analyzer (SPAN) features:
 - You can configure a private VLAN port as a SPAN source port.
 - You can use VLAN-based SPAN (VSPAN) on primary, isolated, and community VLANs or use SPAN on only one VLAN to separately monitor egress or ingress traffic.
 - For more information about SPAN, see [Chapter 18, “Local SPAN, RSPAN, and ERSPAN.”](#)

Private VLAN Ports

- Use only the private VLAN configuration commands to assign ports to primary, isolated, or community VLANs. Layer 2 access ports assigned to the VLANs that you configure as primary, isolated, or community VLANs are inactive while the VLAN is part of the private VLAN configuration. Layer 2 trunk interfaces remain in the STP forwarding state.
- Do not configure ports that belong to a PAGP or LACP EtherChannel as private VLAN ports. While a port is part of the private VLAN configuration, any EtherChannel configuration for it is inactive.
- Enable PortFast and BPDU guard on isolated and community host ports to prevent STP loops due to misconfigurations and to speed up STP convergence. (See [Chapter 32, “Optional STP Features”](#).) When enabled, STP applies the BPDU guard feature to all PortFast-configured Layer 2 LAN ports. Do not enable PortFast and BPDU guard on promiscuous ports.
- If you delete a VLAN used in the private VLAN configuration, the private VLAN ports associated with the VLAN become inactive.
- Private VLAN ports can be on different network devices if the devices are trunk-connected and the primary and secondary VLANs have not been removed from the trunk.
- All primary, isolated, and community VLANs associated within a private VLAN must maintain the same topology across trunks. You are highly recommended to configure the same STP bridge parameters and trunk port parameters on all associated VLANs in order to maintain the same topology.

Limitations with Other Features

- VTP version 3 is not supported on private VLAN (PVLAN) ports.
- In some cases, the configuration is accepted with no error messages, but the commands have no effect.
- Do not configure fallback bridging on switches with private VLANs.
- A port is only affected by the private VLAN feature if it is currently in private VLAN mode and its private VLAN configuration indicates that it is a primary, isolated, or community port. If a port is in any other mode, such as Dynamic Trunking Protocol (DTP), it does not function as a private port.
- Do not configure private VLAN ports on interfaces configured for these other features:
 - Port Aggregation Protocol (PAgP)
 - Link Aggregation Control Protocol (LACP)
 - Voice VLAN
- You can configure IEEE 802.1x port-based authentication on a private VLAN port, but do not configure 802.1x with port security, voice VLAN, or per-user ACL on private VLAN ports.
- Do not configure a remote SPAN (RSPAN) VLAN as a private VLAN primary or secondary VLAN. For more information about SPAN, see [Chapter 18, “Local SPAN, RSPAN, and ERSPAN.”](#)

- A private VLAN host or promiscuous port cannot be a SPAN destination port. If you configure a SPAN destination port as a private VLAN port, the port becomes inactive.
- A destination SPAN port should not be an isolated port. (However, a source SPAN port can be an isolated port.) VSPAN could be configured to span both primary and secondary VLANs or, alternatively, to span either one if the user is interested only in ingress or egress traffic.
- If using the shortcuts between different VLANs (if any of these VLANs is private) consider both primary and isolated and community VLANs. The primary VLAN should be used both as the destination and as the virtual source, because the secondary VLAN (the real source) is always remapped to the primary VLAN in the Layer 2 FID table.
- If you configure a static MAC address on a promiscuous port in the primary VLAN, you must add the same static address to all associated secondary VLANs. If you configure a static MAC address on a host port in a secondary VLAN, you must add the same static MAC address to the associated primary VLAN. When you delete a static MAC address from a private VLAN port, you must remove all instances of the configured MAC address from the private VLAN.

**Note**

Dynamic MAC addresses learned in one VLAN of a private VLAN are replicated in the associated VLANs. For example, a MAC address learned in a secondary VLAN is replicated in the primary VLAN. When the original dynamic MAC address is deleted or aged out, the replicated addresses are removed from the MAC address table.

- Do not configure private VLAN ports as EtherChannels. A port can be part of the private VLAN configuration, but any EtherChannel configuration for the port is inactive.

Information About Private VLANs

- [Private VLAN Domains, page 27-5](#)
- [Private VLAN Ports, page 27-7](#)
- [Primary, Isolated, and Community VLANs, page 27-7](#)
- [Private VLAN Port Isolation, page 27-8](#)
- [IP Addressing Scheme with Private VLANs, page 27-8](#)
- [Private VLANs Across Multiple Switches, page 27-9](#)
- [Private VLAN Interaction with Other Features, page 27-9](#)

Private VLAN Domains


The private VLAN feature addresses two problems that service providers encounter when using VLANs:

- The switch supports up to 4096 VLANs. If a service provider assigns one VLAN per customer, the number of customers that service provider can support is limited.
- To enable IP routing, each VLAN is assigned a subnet address space or a block of addresses, which can result in wasting the unused IP addresses and creating IP address management problems.

Using private VLANs solves the scalability problem and provides IP address management benefits for service providers and Layer 2 security for customers.

The private VLAN feature partitions the Layer 2 broadcast domain of a VLAN into subdomains. A subdomain is represented by a pair of private VLANs: a primary VLAN and a secondary VLAN. A private VLAN domain can have multiple private VLAN pairs, one pair for each subdomain. All VLAN pairs in a private VLAN domain share the same primary VLAN. The secondary VLAN ID differentiates one subdomain from another (see [Figure 27-1](#)).

Figure 27-1 Private VLAN Domain



A private VLAN domain has only one primary VLAN. Every port in a private VLAN domain is a member of the primary VLAN. In other words, the primary VLAN is the entire private VLAN domain. Secondary VLANs provide Layer 2 isolation between ports within the same private VLAN domain. There are two types of secondary VLANs:

- Isolated VLANs—Ports within an isolated VLAN cannot communicate with each other at the Layer 2 level.
- Community VLANs—Ports within a community VLAN can communicate with each other but cannot communicate with ports in other communities at the Layer 2 level.

Private VLAN Ports

There are three types of private VLAN ports:

- **Promiscuous**—A promiscuous port belongs to the primary VLAN and can communicate with all interfaces, including the community and isolated host ports that belong to the secondary VLANs that are associated with the primary VLAN.
- **Isolated**—An isolated port is a host port that belongs to an isolated secondary VLAN. This port has complete Layer 2 isolation from other ports within the same private VLAN domain, except for the promiscuous ports. Private VLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports.
- **Community**—A community port is a host port that belongs to a community secondary VLAN. Community ports communicate with other ports in the same community VLAN and with promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities and from isolated ports within their private VLAN domain.

**Note**

Because trunks can support the VLANs carrying traffic between isolated, community, and promiscuous ports, isolated and community port traffic might enter or leave the switch through a trunk interface.

Primary, Isolated, and Community VLANs

Primary VLANs and the two types of secondary VLANs, isolated VLANs and community VLANs, have these characteristics:

- **Primary VLAN**— The primary VLAN carries unidirectional traffic downstream from the promiscuous ports to the (isolated and community) host ports and to other promiscuous ports.
- **Isolated VLAN** —A private VLAN domain has only one isolated VLAN. An isolated VLAN is a secondary VLAN that carries unidirectional traffic upstream from the hosts toward the promiscuous ports and the gateway.
- **Community VLAN**—A community VLAN is a secondary VLAN that carries upstream traffic from the community ports to the promiscuous port gateways and to other host ports in the same community. You can configure multiple community VLANs in a private VLAN domain.

A promiscuous port can serve only one primary VLAN, one isolated VLAN, and multiple community VLANs. Layer 3 gateways are connected typically to the switch through a promiscuous port. With a promiscuous port, you can connect a wide range of devices as access points to a private VLAN. For example, you can use a promiscuous port to monitor or back up all the private VLAN servers from an administration workstation.

In a switched environment, you can assign an individual private VLAN and associated IP subnet to each individual or common group of end stations. The end stations need to communicate only with a default gateway to communicate outside the private VLAN.

Private VLAN Port Isolation

You can use private VLANs to control access to end stations in these ways:

- Configure selected interfaces connected to end stations as isolated ports to prevent any communication at Layer 2. For example, if the end stations are servers, this configuration prevents Layer 2 communication between the servers.
- Configure interfaces connected to default gateways and selected end stations (for example, backup servers) as promiscuous ports to allow all end stations access to a default gateway.

You can extend private VLANs across multiple devices by trunking the primary, isolated, and community VLANs to other devices that support private VLANs. To maintain the security of your private VLAN configuration and to avoid other use of the VLANs configured as private VLANs, configure private VLANs on all intermediate devices, including devices that have no private VLAN ports.

IP Addressing Scheme with Private VLANs

When you assign a separate VLAN to each customer, an inefficient IP addressing scheme is created as follows:


- Assigning a block of addresses to a customer VLAN can result in unused IP addresses.
- If the number of devices in the VLAN increases, the number of assigned addresses might not be large enough to accommodate them.

These problems are reduced by using private VLANs, where all members in the private VLAN share a common address space, which is allocated to the primary VLAN. Hosts are connected to secondary VLANs, and the DHCP server assigns them IP addresses from the block of addresses allocated to the primary VLAN. Subsequent IP addresses can be assigned to customer devices in different secondary VLANs, but in the same primary VLAN. When new devices are added, the DHCP server assigns them the next available address from a large pool of subnet addresses.

Private VLANs Across Multiple Switches

As with regular VLANs, private VLANs can span multiple switches. A trunk port carries the primary VLAN and secondary VLANs to a neighboring switch. The trunk port deals with the private VLAN as any other VLAN. A feature of private VLANs across multiple switches is that traffic from an isolated port in switch A does not reach an isolated port on Switch B. (See [Figure 27-2](#).)

Figure 27-2 Private VLANs Across Switches



VLAN 100 = Primary VLAN
 VLAN 201 = Secondary isolated VLAN
 VLAN 202 = Secondary community VLAN

Because VTP versions 1 and 2 do not support private VLANs, you must manually configure private VLANs on all switches in the Layer 2 network. If you do not configure the primary and secondary VLAN association in some switches in the network, the Layer 2 databases in these switches are not merged. This situation can result in unnecessary flooding of private VLAN traffic on those switches.

VTP version 3 does support private VLANs, so you do not need to manually configure private VLANs on all switches in the Layer 2 network.

Private VLAN Interaction with Other Features

- [Private VLANs and Unicast, Broadcast, and Multicast Traffic, page 27-10](#)
- [Private VLANs and SVIs, page 27-10](#)

Private VLANs and Unicast, Broadcast, and Multicast Traffic

In regular VLANs, devices in the same VLAN can communicate with each other at the Layer 2 level, but devices connected to interfaces in different VLANs must communicate at the Layer 3 level. In private VLANs, the promiscuous ports are members of the primary VLAN, while the host ports belong to secondary VLANs. Because the secondary VLAN is associated to the primary VLAN, members of the these VLANs can communicate with each other at the Layer 2 level.

In a regular VLAN, broadcasts are forwarded to all ports in that VLAN. Private VLAN broadcast forwarding depends on the port sending the broadcast:

- An isolated port sends a broadcast only to the promiscuous ports or trunk ports.
- A community port sends a broadcast to all promiscuous ports, trunk ports, and ports in the same community VLAN.
- A promiscuous port sends a broadcast to all ports in the private VLAN (other promiscuous ports, trunk ports, isolated ports, and community ports).

Multicast traffic is routed or bridged across private VLAN boundaries and within a single community VLAN. Multicast traffic is not forwarded between ports in the same isolated VLAN or between ports in different secondary VLANs.

Private VLANs and SVIs

A switch virtual interface (SVI) is the Layer 3 interface of a Layer 2 VLAN. Layer 3 devices communicate with a private VLAN only through the primary VLAN and not through secondary VLANs. Configure Layer 3 VLAN SVIs only for primary VLANs. Do not configure Layer 3 VLAN interfaces for secondary VLANs. SVIs for secondary VLANs are inactive while the VLAN is configured as a secondary VLAN.

- If you try to configure a VLAN with an active SVI as a secondary VLAN, the configuration is not allowed until you disable the SVI.
- If you try to create an SVI on a VLAN that is configured as a secondary VLAN, and the secondary VLAN is already mapped at Layer 3, the SVI is not created, and an error is returned. If the SVI is not mapped at Layer 3, the SVI is created, but it is automatically shut down.

When the primary VLAN is associated with and mapped to the secondary VLAN, any configuration on the primary VLAN is propagated to the secondary VLAN SVIs. For example, if you assign an IP subnet to the primary VLAN SVI, this subnet is the IP subnet address of the entire private VLAN.

Default Settings for Private VLANs

None.

How to Configure Private VLANs

- [Configuring a VLAN as a Private VLAN, page 27-11](#)
- [Associating Secondary VLANs with a Primary VLAN, page 27-12](#)
- [Mapping Secondary VLANs to the Layer 3 VLAN Interface of a Primary VLAN, page 27-13](#)
- [Configuring a Layer 2 Interface as a Private VLAN Host Port, page 27-14](#)

- [Configuring a Layer 2 Interface as a Private VLAN Promiscuous Port, page 27-15](#)

**Note**

If the VLAN is not defined already, the private VLAN configuration process defines it.

Configuring a VLAN as a Private VLAN

To configure a VLAN as a private VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan <i>vlan_ID</i>	Enters VLAN configuration submode.
Step 2	Router(config-vlan)# private-vlan { community isolated primary }	Configures a VLAN as a private VLAN. Note These commands do not take effect until you exit VLAN configuration submode.
Step 3	Router(config-vlan)# end	Exits configuration mode.

This example shows how to configure VLAN 202 as a primary VLAN and verify the configuration:

```
Router# configure terminal
Router(config)# vlan 202
Router(config-vlan)# private-vlan primary
Router(config-vlan)# end
Router# show vlan private-vlan
```

```
Primary Secondary Type Interfaces
-----
202                primary
```

This example shows how to configure VLAN 303 as a community VLAN and verify the configuration:

```
Router# configure terminal
Router(config)# vlan 303
Router(config-vlan)# private-vlan community
Router(config-vlan)# end
Router# show vlan private-vlan
```

```
Primary Secondary Type Interfaces
-----
202                primary
303                community
```

This example shows how to configure VLAN 440 as an isolated VLAN and verify the configuration:

```
Router# configure terminal
Router(config)# vlan 440
Router(config-vlan)# private-vlan isolated
Router(config-vlan)# end
Router# show vlan private-vlan
```

```
Primary Secondary Type Interfaces
-----
202                primary
303                community
440                isolated
```

Associating Secondary VLANs with a Primary VLAN

To associate secondary VLANs with a primary VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan <i>primary_vlan_ID</i>	Enters VLAN configuration submode for the primary VLAN.
Step 2	Router(config-vlan)# private-vlan association { <i>secondary_vlan_list</i> add <i>secondary_vlan_list</i> remove <i>secondary_vlan_list</i> }	Associates the secondary VLANs with the primary VLAN.
Step 3	Router(config-vlan)# end	Exits VLAN configuration mode.

When you associate secondary VLANs with a primary VLAN, note the following information:

- The *secondary_vlan_list* parameter cannot contain spaces. It can contain multiple comma-separated items. Each item can be a single private VLAN ID or a hyphenated range of private VLAN IDs.
- The *secondary_vlan_list* parameter can contain multiple community VLAN IDs.
- The *secondary_vlan_list* parameter can contain only one isolated VLAN ID.
- Enter a *secondary_vlan_list* or use the **add** keyword with a *secondary_vlan_list* to associate secondary VLANs with a primary VLAN.
- Use the **remove** keyword with a *secondary_vlan_list* to clear the association between secondary VLANs and a primary VLAN.
- The command does not take effect until you exit VLAN configuration submode.

This example shows how to associate community VLANs 303 through 307 and 309 and isolated VLAN 440 with primary VLAN 202 and verify the configuration:

```
Router# configure terminal
Router(config)# vlan 202
Router(config-vlan)# private-vlan association 303-307,309,440
Router(config-vlan)# end
Router# show vlan private-vlan
```

Primary	Secondary	Type	Interfaces
202	303	community	
202	304	community	
202	305	community	
202	306	community	
202	307	community	
202	309	community	
202	440	isolated	
	308	community	

Mapping Secondary VLANs to the Layer 3 VLAN Interface of a Primary VLAN


Note

Isolated and community VLANs are both called secondary VLANs.

To map secondary VLANs to the Layer 3 VLAN interface of a primary VLAN to allow Layer 3 switching of private VLAN ingress traffic, perform this task:

	Command	Purpose
Step 1	Router(config)# interface vlan <i>primary_vlan_ID</i>	Enters interface configuration mode for the primary VLAN.
Step 2	Router(config-if)# private-vlan mapping { <i>secondary_vlan_list</i> add <i>secondary_vlan_list</i> remove <i>secondary_vlan_list</i> }	Maps the secondary VLANs to the Layer 3 VLAN interface of a primary VLAN to allow Layer 3 switching of private VLAN ingress traffic.
	Router(config-if)# [no] private-vlan mapping	Clears the mapping between the secondary VLANs and the primary VLAN.
Step 3	Router(config-if)# end	Exits configuration mode.

When you map secondary VLANs to the Layer 3 VLAN interface of a primary VLAN, note the following information:

- The **private-vlan mapping** interface configuration command only affects private VLAN ingress traffic that is Layer 3-switched.
- The *secondary_vlan_list* parameter cannot contain spaces. It can contain multiple comma-separated items. Each item can be a single private VLAN ID or a hyphenated range of private VLAN IDs.
- Enter a *secondary_vlan_list* parameter or use the **add** keyword with a *secondary_vlan_list* parameter to map the secondary VLANs to the primary VLAN.
- Use the **remove** keyword with a *secondary_vlan_list* parameter to clear the mapping between secondary VLANs and the primary VLAN.

This example shows how to permit routing of secondary VLAN ingress traffic from private VLANs 303 through 307, 309, and 440 and verify the configuration:

```
Router# configure terminal
Router(config)# interface vlan 202
Router(config-if)# private-vlan mapping add 303-307,309,440
Router(config-if)# end
Router# show interfaces private-vlan mapping
Interface Secondary VLAN Type
-----
vlan202    303          community
vlan202    304          community
vlan202    305          community
vlan202    306          community
vlan202    307          community
vlan202    309          community
vlan202    440          isolated
Router#
```

Configuring a Layer 2 Interface as a Private VLAN Host Port

To configure a Layer 2 interface as a private VLAN host port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. - You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 interface before you can enter additional switchport commands with keywords. - Required only if you have not entered the switchport command already for the interface.
Step 3	Router(config-if)# switchport mode private-vlan { host promiscuous }	Configures the Layer 2 port as a private VLAN host port.
Step 4	Router(config-if)# switchport private-vlan host-association <i>primary_vlan_ID</i> <i>secondary_vlan_ID</i>	Associates the Layer 2 port with a private VLAN. Note If VLAN locking is enabled, enter the VLAN name instead of the VLAN number. For more information, see the “VLAN Locking” section on page 26-4.
Step 5	Router(config-if)# end	Exits configuration mode.

This example shows how to configure interface GigabitEthernet 5/1 as a private VLAN host port and verify the configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/1
Router(config-if)# switchport mode private-vlan host
Router(config-if)# switchport private-vlan host-association 202 303
Router(config-if)# end
Router# show interfaces gigabitethernet 5/1 switchport | include private-vlan
Administrative Mode: private-vlan host
Administrative private-vlan host-association: 202 (VLAN0202) 303 (VLAN0303)
Administrative private-vlan mapping: none
Operational private-vlan: none
```

Configuring a Layer 2 Interface as a Private VLAN Promiscuous Port

To configure a Layer 2 interface as a private VLAN promiscuous port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN interface to configure.
Step 2	Router(config-if)# switchport	Configures the LAN interface for Layer 2 switching: - You must enter the switchport command once without any keywords to configure the LAN interface as a Layer 2 interface before you can enter additional switchport commands with keywords. - Required only if you have not entered the switchport command already for the interface.
Step 3	Router(config-if)# switchport mode private-vlan { host promiscuous }	Configures the Layer 2 port as a private VLAN promiscuous port.
Step 4	Router(config-if)# switchport private-vlan mapping <i>primary_vlan_ID</i> { <i>secondary_vlan_list</i> add <i>secondary_vlan_list</i> remove <i>secondary_vlan_list</i> }	Maps the private VLAN promiscuous port to a primary VLAN and to selected secondary VLANs. Note If VLAN locking is enabled, enter the VLAN name instead of the VLAN number. For more information, see the “VLAN Locking” section on page 26-4.
	Router(config-if)# no switchport private-vlan mapping	Clears all mapping between the private VLAN promiscuous port and the primary VLAN and any secondary VLANs.
Step 5	Router(config-if)# end	Exits configuration mode.

When you configure a Layer 2 interface as a private VLAN promiscuous port, note the following information:

- The *secondary_vlan_list* parameter cannot contain spaces. It can contain multiple comma-separated items. Each item can be a single private VLAN ID or a hyphenated range of private VLAN IDs.
- If VLAN locking is enabled, enter VLAN names instead of VLAN numbers in the *secondary_vlan_list*. When entering a range of VLAN names, you must leave spaces between the VLAN names and the dash.
- Enter a *secondary_vlan_list* value or use the **add** keyword with a *secondary_vlan_list* value to map the secondary VLANs to the private VLAN promiscuous port.
- Use the **remove** keyword with a *secondary_vlan_list* value to clear the mapping between secondary VLANs and the private VLAN promiscuous port.

This example shows how to configure interface GigabitEthernet 5/2 as a private VLAN promiscuous port and map it to a private VLAN:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/2
Router(config-if)# switchport mode private-vlan promiscuous
Router(config-if)# switchport private-vlan mapping 202 303,440
Router(config-if)# end
```

This example shows how to verify the configuration:

```
Router# show interfaces gigabitethernet 5/2 switchport | include private-vlan
Administrative Mode: private-vlan promiscuous
Administrative private-vlan host-association: none ((Inactive))
Administrative private-vlan mapping: 202 (VLAN0202) 303 (VLAN0303) 440 (VLAN0440)
Operational private-vlan: none
```

Monitoring Private VLANs

Table 27-1 shows the privileged EXEC commands for monitoring private VLAN activity.

Table 27-1 Private VLAN Monitoring Commands

Command	Purpose
show interfaces status	Displays the status of interfaces, including the VLANs to which they belong.
show vlan private-vlan [type]	Displays the private VLAN information for the switch.
show interface switchport	Displays private VLAN configuration on interfaces.
show interface private-vlan mapping	Displays information about the private VLAN mapping for VLAN SVIs.

This is an example of the output from the **show vlan private-vlan** command:

```
Switch(config)# show vlan private-vlan
Primary Secondary Type Ports
-----
10      501      isolated      Gi2/1, Gi3/1, Gi3/2
10      502      community    Gi2/11, Gi3/1, Gi3/4
10      503      non-operational
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

28

Private Hosts

- [Prerequisites for Private Hosts, page 28-1](#)
- [Restrictions for Private Hosts, page 28-1](#)
- [Information About Private Hosts, page 28-3](#)
- [How to Configure Private Hosts, page 28-8](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Private Hosts

None.

Restrictions for Private Hosts

- [General Private Host Restrictions, page 28-2](#)
- [Private Host ACL Restrictions, page 28-2](#)
- [Private Host VLAN on Trunk Port Restrictions, page 28-2](#)
- [Private Host Interaction with Other Features, page 28-3](#)
- [Private Host Spoofing Protection, page 28-3](#)
- [Private Host Multicast Operation, page 28-3](#)

General Private Host Restrictions

- Private hosts and private VLANs cannot both be configured on the same port (interface). Both features can coexist on the switch, but each feature must be configured on different ports.
- Private hosts is an end-to-end feature. You must enable the feature on all of the switches between the DSLAMs and an upstream device such as a BRAS or a multicast server.
- Only trusted ports can be configured as isolated ports.
- The private hosts feature is supported on Layer 2 interfaces that are configured as trunking switch ports.
- The private hosts feature is supported on port-channel interfaces (EtherChannel, Fast EtherChannel, and Gigabit EtherChannel). You must enable private hosts on the port-channel interface; you cannot enable the feature on member ports.
- DAI and DHCP snooping cannot be enabled on a private hosts port unless all of the VLANs on the port are configured for snooping.

Private Host ACL Restrictions

- This release of the private hosts feature uses protocol-independent MAC ACLs.
Do not apply IP-based ACLs to any port configured for private hosts or you will defeat the private hosts feature (because the switch will not be able to apply a private hosts MAC ACL to the port).
- You can configure the following interface types for protocol-independent MAC ACL filtering:
 - VLAN interfaces with no IP address
 - Physical LAN ports that support EoMPLS
 - Logical LAN subinterfaces that support EoMPLS
- Protocol-independent MAC ACL filtering applies MAC ACLs to all ingress traffic types (for example, IPv4 traffic, IPv6 traffic, and MPLS traffic, in addition to MAC-layer traffic).
- Ingress traffic that is permitted or denied by a protocol-independent MAC ACL is processed by egress interfaces as MAC-layer traffic. You cannot apply egress IP ACLs to traffic permitted or denied by a MAC ACL on an interface configured for protocol-independent MAC ACL filtering.
- Do not configure protocol-independent MAC ACL filtering on VLAN interfaces where you have configured an IP address.
- Do not configure protocol-independent MAC ACL filtering with microflow policing when the permitted traffic would be bridged or Layer 3 switched in hardware by the PFC or a DFC.
- Protocol-independent MAC ACL filtering supports microflow policing when the permitted traffic is routed in software.
- To prevent any existing VLAN ACLs (VACLs) and routing ACLs (RACLs) from interfering with the PACL on the trunk port, configure the access group mode of the trunk port interface as prefer port mode. Do not apply any VACLs or RACLs to a port configured for private hosts.

Private Host VLAN on Trunk Port Restrictions

- You can enable IGMP snooping on VLANs that use trunk ports configured for private hosts.
- You cannot enable IP multicast on a VLAN that uses a trunk port that is configured for private hosts.

- Because PACLs operate in override mode on trunk ports, you cannot apply VLAN-based features to switch ports.
- The Multicast VLAN Registration (MVR) feature can coexist with private hosts as long as the multicast source exists on a promiscuous port.

Private Host Interaction with Other Features

- Private hosts do not affect Layer 2-based services such as MAC limiting, unicast flood protection (UFP), or unknown unicast flood blocking (UUFB).
- The private hosts features does not affect IGMP snooping. However, if IGMP snooping is globally disabled, IGMP control packets will be subject to ACL checks. To permit IGMP control packets, the private hosts software adds a multicast permit statement to the PACLs for isolated hosts. This operation occurs automatically and no user intervention is required.
- Port security can be enabled on isolated ports to provide added security to those ports.
- When enabled on promiscuous or mixed-mode ports, the port security feature may restrict a change in source port for an upstream device (such as a BRAS or a multicast server).
- When enabled on an access port, 802.1X is not affected by the private hosts feature.

Private Host Spoofing Protection

The private hosts feature prevents MAC address spoofing but does not validate the customer MAC or IP address. To prevent MAC address spoofing, the private hosts feature does the following:

- Uses a static MAC address for a BRAS or a multicast server.
- Disables learning in the Layer 2 forwarding table.
- Alerts the switch software when a BRAS or multicast server moves from one source port to another. The software then validates the move and updates the Layer 2 forwarding table.

Private Host Multicast Operation

Multicast traffic that originates from an upstream device (such as a BRAS or a multicast server) is always permitted. In addition, the private hosts PACLs are not applied to multicast control packets (such as IGMP query and join requests). This operation allows isolated hosts to participate in multicast groups, respond to IGMP queries, and receive traffic from any groups of interest.

Multicast traffic that originates from a host is dropped by the private hosts PACLs. However, if other hosts need to receive multicast traffic originating from a host, the private hosts feature adds a *multicast permit* entry to the PACLs.

Information About Private Hosts

- [Private Hosts Overview](#), page 28-4
- [Isolating Hosts in a VLAN](#), page 28-4
- [Restricting Traffic Flow \(Using Private Hosts Port Mode and PACLs\)](#), page 28-5
- [Port ACLs](#), page 28-7

Private Hosts Overview

Service providers typically deliver triple-play services (voice, video, and data) using three different VLANs over a single physical interface for each end user. The service infrastructure would be simpler and more scalable if the service provider could deploy a single set of VLANs to multiple end users for the same set of services, but the service provider must be able to isolate traffic between the users (hosts) at Layer 2. The private hosts feature provides this isolation, allowing VLAN sharing among multiple end users.

The private hosts feature provides these key benefits:

- Isolates traffic among hosts (subscribers) that share the same VLAN ID.
- Reuses VLAN IDs across different subscribers, which improves VLAN scalability by making better use of the 4096 VLANs allowed.
- Prevents media access control (MAC) address spoofing to prevent denial of service (DOS) attacks.

The private hosts feature uses protocol-independent port-based access control lists (PACLs) to provide Layer 2 isolation between hosts on trusted ports within a strictly Layer 2 domain. The PACLs isolate the hosts by imposing Layer 2 forwarding constraints on the switch ports.

Isolating Hosts in a VLAN

By isolating the hosts, a service provider can use a single set of VLANs to deliver the same set of broadband or metro Ethernet services to multiple end users while ensuring that none of the hosts in the VLAN can communicate directly with each other. For example, VLAN 10 can be used for voice traffic, VLAN 20 for video traffic, and VLAN 30 for data traffic.

When the switch is used as a Digital Subscriber Line Access Multiplexer (DSLAM) gigabit Ethernet aggregator, the DSLAM is connected to the switch through a trunk port that can carry data for multiple VLANs. The service provider uses a single physical port and a single set of VLANs to deliver the same set of services to different end users (isolated hosts). A separate VLAN is used for each service (voice, video, and data).


[Figure 28-1](#) shows an example of triple-play services being delivered from the switch to multiple end users attached to a DSLAM. In the figure, note the following:

- A single trunk link between the switch and the DSLAM carries traffic for all three VLANs.
- Virtual circuits (VCs) deliver the VLAN traffic from the DSLAM to individual end users.

- All unicast traffic between access switches (switches connected to each other) is blocked except for traffic directed toward a BRAS or a multicast server.
- The unknown unicast flood blocking (UUFB) feature is used to block unknown unicast traffic on DSLAM-facing ports.

Figure 28-2 shows the different types of port modes (isolated, promiscuous, and mixed) used in a private hosts configuration.

Figure 28-2 Private Hosts Port Types (Modes)



1	Promiscuous ports	Permit all traffic from a BRAS to hosts.

2	Mixed ports	Permit broadcast traffic from a BRAS. Redirect broadcast traffic from hosts to promiscuous and mixed-mode ports. Permit traffic from a BRAS to hosts and from hosts to a BRAS. Deny all other host to host traffic.
3	Isolated ports	Permit unicast traffic from host to a BRAS only; block unicast traffic between ports. Redirect all broadcast traffic from host to a BRAS. Deny traffic from a BRAS (to prevent spoofing). Permit multicast traffic (IPv4 and IPv6).
Note In this description of port types, the term BRAS represents an upstream devices such as a BRAS, a multicast server (such as a video server), or a core network device that provides access to these devices.		

Port ACLs

The private hosts feature creates several types of port ACLs (PACLs) to impose Layer 2 forwarding constraints on switch ports. The software creates PACLs for the different types of private hosts ports based on the MAC addresses of the content servers providing broadband services and the VLAN IDs of the isolated hosts to deliver those services to. You specify the mode in which each private hosts port is to operate and the software applies the appropriate PACL to the port based on the port's mode (isolated, promiscuous, or mixed).

The following are examples of the different types of PACLs that are used by the private hosts feature.

Isolated Hosts PACL

This example shows a PACL for isolated ports:

```
deny host BRAS_MAC any
permit any host BRAS_MAC
redirect any host FFFF.FFFF.FFFF to LTLIndex 6
permit any 0100.5E00.0000/0000.007F.FFFF
permit any 3333.0000.0000/000.FFFF.FFFF
deny any any
```

Promiscuous Port PACL

This example shows a PACL for promiscuous ports:

```
permit host BRAS_MAC any
deny any any
```

Mixed Port PACL

This example shows a PACL for mixed ports:

```
permit host BRAS_MAC ffff.ffff.ffff
redirect any host FFFF.FFFF.FFFF to LTLIndex 6
permit host BRAS_MAC any
permit any host BRAS_MAC
deny any any
```

Default Settings for Private Hosts

None.

How to Configure Private Hosts

- [Configuration Summary, page 28-8](#)
- [Detailed Configuration Steps, page 28-9](#)
- [Configuration Examples, page 28-10](#)

Configuration Summary

1. Determine which switch ports (interfaces) to use for the private hosts feature. You can configure the feature on trunking switch ports or port-channel interfaces. Private hosts must be enabled on the port-channel interface; you cannot enable the feature on member ports.
2. Configure each port (interface) for normal, non-private hosts service. Configure the access group mode of the port as prefer port mode. You can configure the VLANs at this step or later.
3. Determine which VLAN or set of VLANs will be used to deliver broadband services to end users. The private hosts feature will provide Layer 2 isolation among the hosts in these VLANs.
4. Identify the MAC addresses of all of the BRASs and multicast servers that are being used to provide broadband services to end users (isolated hosts).

**Note**

If a server is not connected directly to the switch, determine the MAC address of the core network device that provides access to the server.

5. (Optional) If you plan to offer different types of broadband service to different sets of isolated hosts, create multiple MAC and VLAN lists.
 - Each MAC address list identifies a server or set of servers providing a particular type of service.
 - Each VLAN list identifies the isolated hosts to deliver that service to.
6. Configure promiscuous ports and specify a MAC and VLAN list to identify the server and receiving hosts for a particular type of service.

**Note**

You can specify multiple MAC and VLAN combinations to allow different types of services to be delivered to different sets of hosts. For example, the BRAS at xxxx.xxxx.xxxx could be used to deliver a basic set of services over VLANs 20, 25, and 30, and the BRAS at yyyy.yyyy.yyyy could be used to deliver a premium set of services over VLANs 5, 10, and 15.

7. Globally enable private hosts.
8. Enable private hosts on individual ports (interfaces) and specify the mode in which the port is to operate. To determine port mode, you need to know if the port faces upstream (toward content servers or core network), faces downstream (toward DSLAM and isolated hosts), or is connected to another switch (typically, in a ring topology). See the [“Restricting Traffic Flow \(Using Private Hosts Port Mode and PACLs\)”](#) section on page 28-5.

After you enable the feature on individual ports, the switch is ready to run the private hosts feature. The private hosts software uses the MAC and VLAN lists you defined to create the isolated, promiscuous, and mixed-mode PACLs for your configuration. The software then applies the appropriate PACL to each private hosts port based on the port's mode.

Detailed Configuration Steps

To configure the private hosts feature, perform the following steps. These steps assume that you have already configured the Layer 2 interfaces that you plan for private hosts.



Note

You can configure private hosts only on trunking switch ports or EtherChannel ports. In addition, you must enable private hosts on all of the switches between the DSLAMs and upstream devices.

	Command or Action	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# private-hosts mac-list <i>mac_list_name</i> <i>mac_address</i> [remark <i>device-name</i> <i>comment</i>]	Creates a list of MAC addresses that identify the BRAS and multicast servers providing broadband services. • <i>mac_list_name</i> specifies a name to assign to this list of content servers. • <i>mac_address</i> identifies the BRAS or multicast server (or set of servers) providing a particular broadband service or set of services. • remark allows you to specify an optional device name or comment to assign to this MAC list. Specify the MAC address of every content server being used to deliver services. If you plan to offer different types of services to different sets of hosts, create a separate MAC list for each server or set of servers providing a particular service. Note If a server is not directly connected to the switch, specify the MAC address of the core network device that provides access to the server.
Step 3	Router(config)# private-hosts vlan-list <i>vlan-IDs</i>	Creates a list of the VLANs (<i>vlan-IDs</i>) whose hosts need to be isolated so that the hosts can receive broadband services. Create separate VLAN lists if you plan to offer particular services to different sets of hosts. Otherwise, all of the broadband services will be delivered to all isolated hosts.

	Command or Action	Purpose
Step 4	Router(config)# private-hosts promiscuous <i>mac-list-name</i> [vlan-list <i>vlan-IDs</i>]	Identifies the content servers for broadband services and the end users (isolated hosts) to which to deliver the services. <i>mac-list-name</i> specifies the name of the MAC address lists that identifies the BRAS or multicast server (or set of servers) providing a particular type of broadband service or set of services. <i>vlan-IDs</i> identifies the VLAN or set of VLANs whose hosts are to receive services from the above servers. If no VLAN list is specified, the software uses the global VLAN list (configured in Step 3). Note You can enter this command multiple times to configure multiple MAC and VLAN combinations, each defining the server and receiving hosts for a particular type of service.
Step 5	Router(config)# private-hosts	Globally enables private hosts on the switch.
Step 6	Router(config)# interface <i>interface</i>	Selects the trunking switch port or EtherChannel to enable for private hosts.
Step 7	Router(config-if)# access-group mode prefer port	Specifies that any existing VACLs or RACLs on the trunk port will be ignored.
Step 8	Router(config-if)# private-hosts mode { promiscuous isolated mixed }	Enables private hosts on the port. Use one of the following keywords to define the mode that the port is to operate in: promiscuous—Upstream-facing ports that connect to broadband servers (BRAS, multicast, or video) or to core network devices providing access to the servers. isolated—Ports that connect to DSLAMs. mixed—Ports that connect to other switches, typically in a ring topology. Note You must perform this step on each port being used for private hosts.
Step 9	Router(config-if)# end	Exits interface and global configuration modes and returns to privileged EXEC mode. Private Hosts configuration is complete.

Configuration Examples

The following example creates a MAC address list and a VLAN list and isolates the hosts in VLANs 10, 12, 15, and 200 through 300. The BRAS-facing port is made promiscuous and two host-connected ports are made isolated:

```
Router# configure terminal
Router(config)# private-hosts mac-list BRAS_list 0000.1111.1111 remark BRAS_SanJose
Router(config)# private-hosts vlan-list 10,12,15,200-300
Router(config)# private-hosts promiscuous BRAS_list vlan-list 10,12,15,200-300
Router(config)# private-hosts
Router(config)# interface gig 4/2
Router(config-if)# private-hosts mode promiscuous
Router(config-if)# exit
```

```
Router(config)# interface gig 5/2
Router(config-if)# private-hosts mode isolated
Router(config-if)# exit
Router(config)# interface gig 5/3
Router(config-if)# private-hosts mode isolated
Router(config-if)# end
Router#
```

The following example shows the interface configuration of a private hosts isolated port:

```
Router# show run interface gig 5/2
Building configuration...

Current configuration : 200 bytes
!
interface GigabitEthernet5/2
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
  access-group mode prefer port
  private-hosts mode isolated
end
```

The following example shows the interface configuration of a private hosts promiscuous port:

```
Router# show run interface gig 4/2
Building configuration...

Current configuration : 189 bytes
!
interface GigabitEthernet4/2
  switchport
  switchport access vlan 200
  switchport mode access
  private-hosts mode promiscuous
end

private-hosts
private-hosts vlan-list 200
private-hosts promiscuous bras-list
private-hosts mac-list bras-list 0000.1111.1111 remark BRAS-SERVER
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

29

IEEE 802.1Q Tunneling

- [Prerequisites for 802.1Q Tunneling, page 29-1](#)
- [Restrictions for 802.1Q Tunneling, page 29-1](#)
- [Information About 802.1Q Tunneling, page 29-3](#)
- [Default Settings for 802.1Q Tunneling, page 29-5](#)
- [How to Configure 802.1Q Tunneling, page 29-6](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for 802.1Q Tunneling

None.

Restrictions for 802.1Q Tunneling

- Use asymmetrical links to put traffic into a tunnel or to remove traffic from a tunnel.
- Configure tunnel ports only to form an asymmetrical link.
- Dedicate one VLAN for each tunnel.
- Assign only tunnel ports to VLANs used for tunneling.
- Trunks require no special configuration to carry tunnel VLANs.
- Tunnel ports are not trunks. Any commands to configure trunking are inactive while the port is configured as a tunnel port.

- Tunnel ports learn customer MAC addresses.
- We recommend that you use ISL trunks to carry tunnel traffic between devices that do not have tunnel ports. Because of the 802.1Q native VLAN feature, using 802.1Q trunks requires that you be very careful when you configure tunneling: a mistake might direct tunnel traffic to a non-tunnel port.
- By default, the native VLAN traffic of a dot1q trunk is sent untagged, which cannot be double-tagged in the service provider network. Because of this situation, the native VLAN traffic might not be tunneled correctly. Be sure that the native VLAN traffic is always sent tagged in an asymmetrical link. To tag the native VLAN egress traffic and drop all untagged ingress traffic, enter the global **vlan dot1q tag native** command.
- Configure jumbo frame support on tunnel ports:
 - See the “[Configuring Jumbo Frame Support](#)” section on page 10-6.
 - Take note of the modules listed in the “Configuring Jumbo Frame Support” section that do not support jumbo frames.
- Jumbo frames can be tunneled as long as the jumbo frame length combined with the 802.1Q tag does not exceed the maximum frame size.
- Because tunnel traffic has the added ethertype and length field and retains the 802.1Q tag within the switch, the following restrictions exist:
 - The Layer 3 packet within the Layer 2 frame cannot be identified in tunnel traffic.
 - Layer 3 and higher parameters cannot be identified in tunnel traffic (for example, Layer 3 destination and source addresses).
 - Because the Layer 3 addresses cannot be identified within the packet, tunnel traffic cannot be routed.
 - The switch can provide only MAC-layer filtering for tunnel traffic (VLAN IDs and source and destination MAC addresses).
 - The switch can provide only MAC-layer access control and QoS for tunnel traffic.
 - QoS cannot detect the received CoS value in the 802.1Q 2-byte Tag Control Information field.
- On an asymmetrical link, the Cisco Discovery Protocol (CDP) reports a native VLAN mismatch if the VLAN of the tunnel port does not match the native VLAN of the 802.1Q trunk. The 802.1Q tunnel feature does not require that the VLANs match. Ignore the messages if your configuration requires nonmatching VLANs.
- Asymmetrical links do not support the Dynamic Trunking Protocol (DTP) because only one port on the link is a trunk. Configure the 802.1Q trunk port on an asymmetrical link to trunk unconditionally.
- The 802.1Q tunneling feature cannot be configured on ports configured to support private VLANs.
- The 802.1Q tunneling feature cannot be configured on ports configured to support EVCs.
- The following Layer 2 protocols work between devices connected by an asymmetrical link:
 - CDP
 - UniDirectional Link Detection (UDLD)
 - Port Aggregation Protocol (PAgP)
 - Link Aggregation Control Protocol (LACP)
- PortFast BPDU filtering is enabled automatically on tunnel ports.
- CDP is automatically disabled on tunnel ports.

- VLAN Trunk Protocol (VTP) does not work between the following devices:
 - Devices connected by an asymmetrical link
 - Devices communicating through a tunnel



Note VTP works between tunneled devices if Layer 2 protocol tunneling is enabled. See [Chapter 30, “Layer 2 Protocol Tunneling,”](#) for configuration details.

- To configure an EtherChannel as an asymmetrical link, all ports in the EtherChannel must have the same tunneling configuration. Because the Layer 3 packet within the Layer 2 frame cannot be identified, you must configure the EtherChannel to use MAC-address-based frame distribution.

The following configuration guidelines are *required* for your Layer 2 protocol tunneling configuration:

- On all the service provider edge switches, PortFast BPDU filtering must be enabled on the 802.1Q tunnel ports as follows:

```
Router(config-if)# spanning-tree bpdufilter enable
Router(config-if)# spanning-tree portfast
```



Note PortFast BPDU filtering is enabled automatically on tunnel ports.

- At least one VLAN must be available for native VLAN tagging (**vlan dot1q tag native** option). If you use all the available VLANs and then try to enable the **vlan dot1q tag native** option, the option will not be enabled.
- On all the service provider core switches, tag native VLAN egress traffic and drop untagged native VLAN ingress traffic by entering the following command:

```
Router(config)# vlan dot1q tag native
```

- On all the customer switches, *either* enable or disable the global **vlan dot1q tag native** option.



Note If this option is enabled on one switch and disabled on another switch, all traffic is dropped; all customer switches must have this option configured the same on each switch.

The following configuration guidelines are *optional* for your Layer 2 protocol tunneling configuration:

- Because all the BPDUs are being dropped, spanning tree PortFast can be enabled on Layer 2 protocol tunnel ports as follows:

```
Router(config-if)# spanning-tree portfast trunk
```

- If the service provider does not want the customer to see its switches, CDP should be disabled on the 802.1Q tunnel port as follows:

```
Router(config-if)# no cdp enable
```

Information About 802.1Q Tunneling

802.1Q tunneling enables service providers to use a single VLAN to support customers who have multiple VLANs, while preserving customer VLAN IDs and keeping traffic in different customer VLANs segregated.

A port configured to support 802.1Q tunneling is called a tunnel port. When you configure tunneling, you assign a tunnel port to a VLAN that you dedicate to tunneling, which then becomes a tunnel VLAN. To keep customer traffic segregated, each customer requires a separate tunnel VLAN, but that one tunnel VLAN supports all of the customer's VLANs.

802.1Q tunneling is not restricted to point-to-point tunnel configurations. Any tunnel port in a tunnel VLAN is a tunnel entry and exit point. An 802.1Q tunnel can have as many tunnel ports as are needed to connect customer switches.

The customer switches are trunk connected, but with 802.1Q tunneling, the service provider switches only use one service provider VLAN to carry all the customer VLANs, instead of directly carrying all the customer VLANs.

With 802.1Q tunneling, tagged customer traffic comes from an 802.1Q trunk port on a customer device and enters the service-provider edge switch through a tunnel port. The link between the 802.1Q trunk port on a customer device and the tunnel port is called an asymmetrical link because one end is configured as an 802.1Q trunk port and the other end is configured as a tunnel port. You assign the tunnel port to an access VLAN ID unique to each customer. See [Figure 29-1 on page 29-4](#) and [Figure 29-2 on page 29-5](#).

Figure 29-1 IEEE 802.1Q Tunnel Ports in a Service-Provider Network





Figure 29-2 Untagged, 802.1Q-Tagged, and Double-Tagged Ethernet Frames

When a tunnel port receives tagged customer traffic from an 802.1Q trunk port, it does not strip the received 802.1Q tag from the frame header; instead, the tunnel port leaves the 802.1Q tag intact, adds a 2-byte Ethertype field (0x8100) followed by a 2-byte field containing the priority (CoS) and the VLAN. The received customer traffic is then put into the VLAN to which the tunnel port is assigned. This Ethertype 0x8100 traffic, with the received 802.1Q tag intact, is called tunnel traffic.

A VLAN carrying tunnel traffic is an 802.1Q tunnel. The tunnel ports in the VLAN are the tunnel's ingress and egress points.

The tunnel ports do not have to be on the same network device. The tunnel can cross other network links and other network devices before reaching the egress tunnel port. A tunnel can have as many tunnel ports as required to support the customer devices that need to communicate through the tunnel.

An egress tunnel port strips the 2-byte Ethertype field (0x8100) and the 2-byte length field and transmits the traffic with the 802.1Q tag still intact to an 802.1Q trunk port on a customer device. The 802.1Q trunk port on the customer device strips the 802.1Q tag and puts the traffic into the appropriate customer VLAN.

**Note**

Tunnel traffic carries a second 802.1Q tag only when it is on a trunk link between service-provider network devices, with the outer tag containing the service-provider-assigned VLAN ID and the inner tag containing the customer-assigned VLAN IDs.

Default Settings for 802.1Q Tunneling

None.

How to Configure 802.1Q Tunneling

- [Configuring 802.1Q Tunnel Ports, page 29-6](#)
- [Configuring the Switch to Tag Native VLAN Traffic, page 29-6](#)



Caution

Ensure that only the appropriate tunnel ports are in any VLAN used for tunneling and that one VLAN is used for each tunnel. Incorrect assignment of tunnel ports to VLANs can forward traffic inappropriately.

Configuring 802.1Q Tunnel Ports

To configure 802.1Q tunneling on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. • You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 interface before you can enter additional switchport commands with keywords. • Required only if you have not entered the switchport command already for the interface.
Step 3	Router(config-if)# switchport mode dot1q-tunnel	Configures the Layer 2 port as a tunnel port.
Step 4	Router(config-if)# no lldp transmit	(Required on PE ports) Disables LLDP. Note CDP is automatically disabled.
Step 5	Router(config-if)# end	Exits configuration mode.

This example shows how to configure tunneling on port 4/1 and verify the configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 4/1
Router(config-if)# switchport mode dot1q-tunnel
Router(config-if)# no lldp transmit
Router(config-if)# end
Router# show dot1q-tunnel interface
```

Configuring the Switch to Tag Native VLAN Traffic

- [Configuring the Switch to Tag Native VLAN Traffic Globally, page 29-7](#)
- [Configuring Ports Not to Tag Native VLAN Traffic, page 29-7](#)

Configuring the Switch to Tag Native VLAN Traffic Globally

To configure the switch to tag traffic in the native VLAN globally, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan dot1q tag native	Configures the switch to tag native VLAN traffic globally, and admit only 802.1Q tagged frames on 802.1Q trunks, dropping any untagged traffic, including untagged traffic in the native VLAN. Note On ports where you enter the no switchport trunk native vlan tag interface command, the function of the vlan dot1q tag native global command is disabled.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the switch to tag native VLAN traffic and verify the configuration:

```
Router# configure terminal
Router(config)# vlan dot1q tag native
Router(config)# end
Router# show vlan dot1q tag native | include globally
dot1q native vlan tagging is enabled globally
Router(config)#
```

Configuring Ports Not to Tag Native VLAN Traffic

To configure a port not to tag traffic in the native VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. - You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 interface before you can enter additional switchport commands with keywords. - Required only if you have not entered the switchport command already for the port.
Step 3	Router(config-if)# no switchport trunk native vlan tag	When the switch is configured to tag native VLAN traffic globally, configures the Layer 2 port not to tag native VLAN traffic.
Step 4	Router(config-if)# end	Exits configuration mode.



Note

The **switchport trunk native vlan tag** interface command does not enable native VLAN tagging unless the switch is configured to tag native VLAN traffic globally.

This example shows how to configure Gigabit Ethernet port 1/4 to tag traffic in the native VLAN and verify the configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/4
Router(config-if)# switchport trunk native vlan tag
Router(config-if)# end
Router# show interface gigabitethernet 1/4 switchport | include tagging
Administrative Native VLAN tagging: enabled
Operational Native VLAN tagging: disabled
Router#
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

30

Layer 2 Protocol Tunneling

- [Prerequisites for Layer 2 Protocol Tunneling, page 30-1](#)
- [Restrictions for Layer 2 Protocol Tunneling, page 30-1](#)
- [Information About Layer 2 Protocol Tunneling, page 30-1](#)
- [Default Settings for Layer 2 Protocol Tunneling, page 30-2](#)
- [How to Configure Layer 2 Protocol Tunneling, page 30-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Layer 2 Protocol Tunneling

None.

Restrictions for Layer 2 Protocol Tunneling

None.


Information About Layer 2 Protocol Tunneling

Layer 2 protocol tunneling allows Layer 2 protocol data units (PDUs) (CDP, STP, and VTP) to be tunneled through a network. This section uses the following terminology:

- Edge switch—The switch connected to the customer switch and placed on the boundary of the service provider network (see [Figure 30-1](#)).
- Layer 2 protocol tunnel port—A port on the edge switch on which a specific tunneled protocol can be encapsulated or deencapsulated. The Layer 2 protocol tunnel port is configured through CLI commands.
- Tunneled PDU—A CDP, STP, or VTP PDU.

Without Layer 2 protocol tunneling, tunnel ports drop STP and VTP packets and process CDP packets. This handling of the PDUs creates different spanning tree domains (different spanning tree roots) for the customer switches. For example, STP for a VLAN on switch 1 (see [Figure 30-1](#)) builds a spanning tree topology on switches 1, 2, and 3 without considering convergence parameters based on switches 4 and 5. To provide a single spanning tree domain for the customer, a generic scheme to tunnel BPDUs was created for control protocol PDUs (CDP, STP, and VTP). This process is referred to as Generic Bridge PDU Tunneling (GBPT).

Figure 30-1 Layer 2 Protocol Tunneling Network Configuration



GBPT provides a scalable approach to PDU tunneling by software encapsulating the PDUs in the ingress edge switches and then multicasting them in hardware. All switches inside the service provider network treat these encapsulated frames as data packets and forward them to the other end. The egress edge switch listens for these special encapsulated frames and deencapsulates them; they are then forwarded out of the tunnel.

The encapsulation involves rewriting the destination media access control (MAC) address in the PDU. An ingress edge switch rewrites the destination MAC address of the PDUs received on a Layer 2 tunnel port with the Cisco proprietary multicast address (01-00-0c-cd-cd-d0). The PDU is then flooded to the native VLAN of the Layer 2 tunnel port. If you enable Layer 2 protocol tunneling on a port, PDUs of an enabled protocol are not sent out. If you disable Layer 2 protocol tunneling on a port, the disabled protocols function the same way they were functioning before Layer 2 protocol tunneling was enabled on the port.

Default Settings for Layer 2 Protocol Tunneling

None.

How to Configure Layer 2 Protocol Tunneling



Note

- Encapsulated PDUs received by an 802.1Q tunnel port are transmitted from other tunnel ports in the same VLAN on the switch.
- Configure jumbo frame support on Layer 2 protocol tunneling ports:
 - See the [“Configuring Jumbo Frame Support” section on page 10-6](#).
 - Take note of the modules listed in the “Configuring Jumbo Frame Support” section that do not support jumbo frames.

To configure Layer 2 protocol tunneling on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport	Configures the LAN port for Layer 2 switching. - You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 interface before you can enter additional switchport commands with keywords. - Required only if you have not entered the switchport command already for the interface.
Step 3	Router(config-if)# l2protocol-tunnel [cdp lldp stp vtp]	Configures the Layer 2 port as a Layer 2 protocol tunnel port for all protocols or only the specified protocol.
Step 4	Router(config-if)# l2protocol-tunnel drop-threshold {[cdp lldp stp vtp] <i>packets</i> }	(Optional) Configures the port as a Layer 2 protocol tunnel port and sets a drop threshold for all protocols or only the specified protocol.
Step 5	Router(config-if)# l2protocol-tunnel shutdown-threshold {[cdp lldp stp vtp] <i>packets</i> }	(Optional) Configures the port as a Layer 2 protocol tunnel port and sets a shutdown threshold for all protocols or only the specified protocol.
Step 6	Router(config-if)# no lldp transmit	(Required on PE ports) Disables LLDP. Note CDP is automatically disabled.
Step 7	Router(config)# end	Exits configuration mode.

When you configure a Layer 2 port as a Layer 2 protocol tunnel port, note the following information:

- Optionally, you may specify a drop threshold for the port. The drop threshold value, from 1 to 4096, determines the number of packets to be processed for that protocol on that interface in one second. When the drop threshold is exceeded, PDUs for the specified protocol are dropped for the remainder of the one-second period. If a drop threshold is not specified, the value is 0 (drop threshold disabled).
- Optionally, you may specify a shutdown threshold for the port. The shutdown threshold value, from 1 to 4096, determines the number of packets to be processed for that protocol on that interface in one second. When the shutdown threshold is exceeded, the port is put in errdisable state. If a shutdown threshold is not specified, the value is 0 (shutdown threshold disabled).

- If you specify both a drop threshold and a shutdown threshold for the port, packets exceeding the drop threshold will not be forwarded but will be counted toward the shutdown threshold.

**Note**

The commands support the **l2ptguard** keyword:

- **errdisable detect cause**
- **errdisable recovery**

This example shows how to configure Layer 2 protocol tunneling and drop and shutdown thresholds on port 5/1 for CDP, STP, and VTP, and verify the configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/1
Router(config-if)# switchport
Router(config-if)# l2protocol-tunnel shutdown-threshold cdp 400
Router(config-if)# l2protocol-tunnel shutdown-threshold stp 400
Router(config-if)# l2protocol-tunnel shutdown-threshold vtp 400
Router(config-if)# l2protocol-tunnel drop-threshold vtp 200
Router(config-if)# no lldp transmit
Router(config-if)# end
Router# show l2protocol-tunnel summary
COS for Encapsulated Packets: 5
Drop Threshold for Encapsulated Packets: 0
```

Port	Protocol	Shutdown Threshold	Drop Threshold	Status
		(cdp/lldp/stp/vtp)	(cdp/lldp/stp/vtp)	
Gi5/1	-- -- -- --	400/----/ 400/ 400	----/----/----/ 200	down(trunk)

```
Router#
```

This example shows how to display counter information for port 5/1:

```
Router# show l2protocol-tunnel interface gigabitethernet 5/1
COS for Encapsulated Packets: 5
```

Port	Protocol	Thresholds	Counters
		Shutdown Drop	Encap Decap Drop
-----	-----	-----	-----

```
Router#
```

This example shows how to clear the Layer 2 protocol tunneling configuration from port 5/1:

```
Router(config-if)# no l2protocol-tunnel shutdown-threshold cdp 400
Router(config-if)# no l2protocol-tunnel shutdown-threshold stp 400
Router(config-if)# no l2protocol-tunnel shutdown-threshold vtp 400
Router(config-if)# no l2protocol-tunnel drop-threshold vtp 200
Router(config-if)# no l2protocol-tunnel cdp
Router(config-if)# no l2protocol-tunnel stp
Router(config-if)# no l2protocol-tunnel vtp
Router(config-if)# lldp transmit
Router(config-if)# end
Router# show l2protocol-tunnel summary
COS for Encapsulated Packets: 5
Drop Threshold for Encapsulated Packets: 0
```

Port	Protocol	Shutdown Threshold	Drop Threshold	Status

```
(cdp/lldp/stp/vtp) (cdp/lldp/stp/vtp)
```

```
Router#
```

This example shows how to clear Layer 2 protocol tunneling port counters:

```
Router# clear l2protocol-tunnel counters  
Router#
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

31

Spanning Tree Protocols

- [Prerequisites for Spanning Tree Protocols, page 31-1](#)
- [Restrictions for Spanning Tree Protocols, page 31-1](#)
- [Information About Spanning Tree Protocols, page 31-2](#)
- [Default Settings for Spanning Tree Protocols, page 31-25](#)
- [How to Configure Spanning Tree Protocols, page 31-26](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - This chapter describes the Spanning Tree Protocol (STP) and Multiple Spanning Tree (MST) protocol. For information on configuring the PortFast, UplinkFast, and BackboneFast STP enhancements, see [Chapter 32, “Optional STP Features.”](#)
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Spanning Tree Protocols

None.

Restrictions for Spanning Tree Protocols

- The 802.1s MST standard allows up to 65 MST instances. You can map an unlimited number of VLANs to an MST instance.
- Rapid PVST+ and MST are supported, but only one version can be active at any time.

- VTP does not propagate the MST configuration. You must manually configure the MST configuration (region name, revision number, and VLAN-to-instance mapping) on each switch within the MST region through the command-line interface (CLI) or SNMP.
- For load balancing across redundant paths in the network to work, all VLAN-to-instance mapping assignments must match; otherwise, all traffic flows on a single link.
- All MST boundary ports must be forwarding for load balancing between a PVST+ and an MST cloud or between a rapid-PVST+ and an MST cloud. For this to occur, the CIST regional root of the MST cloud must 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 rapid-PVST+ cloud.
- 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 non-Layer 2 devices.
- Adding or removing VLANs to an existing MST instance will trigger spanning tree recalculation for that MST instance, and the traffic of all the VLANs for that MST instance will be disrupted.
- The **show vlan virtual-port** command displays the number of virtual interfaces. The maximum numbers of virtual interfaces supported on stand-alone switch/VSS system limits are as follows:
 - MST: 100,000 total
 - RPVST+: 12,000 total
 - PVST+: 15,000 total

**Note**

Cisco IOS software displays a message if you exceed the maximum number of virtual interfaces.

Information About Spanning Tree Protocols

- [Information about STP, page 31-2](#)
- [Information about IEEE 802.1w RSTP, page 31-13](#)
- [Information about MST, page 31-18](#)
- [Detecting Unidirectional Link Failure, page 31-25](#)

Information about STP

- [STP Overview, page 31-3](#)
- [Information about the Bridge ID, page 31-3](#)
- [Information about Bridge Protocol Data Units, page 31-4](#)
- [Election of the Root Bridge, page 31-5](#)
- [STP Protocol Timers, page 31-5](#)
- [Creating the Spanning Tree Topology, page 31-5](#)
- [STP Port States, page 31-6](#)
- [STP and IEEE 802.1Q Trunks, page 31-12](#)

STP Overview

STP, the IEEE 802.1D bridge protocol, is a Layer 2 link-management protocol that provides path redundancy while preventing undesirable loops in the network. For a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. STP 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.

In an extension known as per-VLAN spanning tree (PVST), Layer 2 Ethernet ports can use STP on all VLANs. Rapid-PVST+, enabled by default on each configured VLAN (provided you do not manually disable it), uses RSTP to provide faster convergence. Independent VLANs run their own RSTP instance.

With rapid-PVST+, dynamic entries are flushed immediately on a per-port basis upon receiving a topology change. UplinkFast and BackboneFast configurations are ignored in rapid-PVST+ mode; both features are included in RSTP. Rapid-PVST+ mode supports unidirectional link failure detection as described in the [“Detecting Unidirectional Link Failure” section on page 31-25](#).

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. The STP algorithm calculates the best loop-free path throughout a switched Layer 2 network. Layer 2 LAN ports send and receive STP frames at regular intervals. Network devices do not forward these frames, but use the frames to construct a loop-free path.

Multiple active paths between end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages and network devices might learn end station MAC addresses on multiple Layer 2 LAN ports. These conditions result in an unstable network.

STP defines a tree with a root bridge and a loop-free path from the root to all network devices in the Layer 2 network. STP forces redundant data paths into a standby (blocked) state. If a network segment in the spanning tree fails and a redundant path exists, the STP algorithm recalculates the spanning tree topology and activates the standby path.

When two Layer 2 LAN ports on a network device are part of a loop, the STP port priority and port path cost setting determine which port is put in the forwarding state and which port is put in the blocking state. The STP port priority value represents the location of a port in the network topology and how efficiently that location allows the port to pass traffic. The STP port path cost value represents media speed.

Information about the Bridge ID

- [Bridge Priority Value, page 31-3](#)
- [Extended System ID, page 31-3](#)
- [STP MAC Address Allocation, page 31-4](#)

Bridge Priority Value

Each VLAN on each network device has a unique 64-bit bridge ID consisting of a bridge priority value, an extended system ID, and an STP MAC address allocation. The bridge priority is a 4-bit value when the extended system ID is enabled (see [Table 31-1 on page 31-4](#) and the [“Configuring the Bridge Priority of a VLAN” section on page 31-34](#)).

Extended System ID

A 12-bit extended system ID field is part of the bridge ID (see [Table 31-1 on page 31-4](#)). Chassis that support only 64 MAC addresses always use the 12-bit extended system ID. On chassis that support 1,024 MAC addresses, you can enable use of the extended system ID.

The extended system ID is enabled by default under the following conditions:

- Chassis that support only 64 MAC addresses
- When the STP mode is MST

STP uses the VLAN ID as the extended system ID. See the [“Enabling the Extended System ID” section on page 31-28](#).

Table 31-1 Bridge Priority Value and Extended System ID with the Extended System ID Enabled

Bridge Priority Value				Extended System ID (Set Equal to the VLAN/MST Instance 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

STP MAC Address Allocation

The chassis has either 64 or 1024 MAC addresses available to support software features such as STP. To view the MAC address range on your chassis, enter the **show catalyst6000 chassis-mac-address** command.

For chassis with 64 MAC addresses, STP uses the extended system ID plus a MAC address to make the bridge ID unique for each VLAN.

When the extended system ID is not enabled, STP uses one MAC address per VLAN to make the bridge ID unique for each VLAN.

If you have a network device in your network with the extended system ID enabled, you should also enable the extended system ID on all other Layer 2 connected network devices to avoid undesirable root bridge election and spanning tree topology issues.

When the extended system ID is enabled, the root bridge priority becomes a multiple of 4096 plus the VLAN ID. With the extended system ID enabled, a switch bridge ID (used by the spanning tree algorithm to determine the identity of the root bridge, the lowest being preferred) can only be specified as a multiple of 4096. Only the following values are possible: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440.

If another bridge in the same spanning tree domain does not have the extended system ID enabled, it could win root bridge ownership because of the finer granularity in the selection of its bridge ID.

Information about Bridge Protocol Data Units

Bridge protocol data units (BPDUs) are transmitted in one direction from the root bridge. Each network device sends configuration BPDUs to communicate and compute the spanning tree topology. Each configuration BPDU contains the following minimal information:

- The unique bridge ID of the network device that the transmitting network device believes to be the root bridge
- The STP path cost to the root
- The bridge ID of the transmitting bridge
- Message age
- The identifier of the transmitting port
- Values for the hello, forward delay, and max-age protocol timers

When a network device transmits a BPDU frame, all network devices connected to the LAN on which the frame is transmitted receive the BPDU. When a network device receives a BPDU, it does not forward the frame but instead uses the information in the frame to calculate a BPDU, and, if the topology changes, initiate a BPDU transmission.

A BPDU exchange results in the following:

- One network device is elected as the root bridge.
- The shortest distance to the root bridge is calculated for each network device based on the path cost.
- A designated bridge for each LAN segment is selected. This is the network device closest to the root bridge through which frames are forwarded to the root.
- A root port is selected. This is the port providing the best path from the bridge to the root bridge.
- Ports included in the spanning tree are selected.

Election of the Root Bridge

For each VLAN, the network device with the highest-priority bridge ID (the lowest numerical ID value) is elected as the root bridge. If all network devices are configured with the default priority (32768), the network device with the lowest MAC address in the VLAN becomes the root bridge. The bridge priority value occupies the most significant bits of the bridge ID.

When you change the bridge priority value, you change the probability that the switch will be elected as the root bridge. Configuring a higher-priority value increases the probability; a lower-priority value decreases the probability.

The STP root bridge is the logical center of the spanning tree topology in a Layer 2 network. All paths that are not needed to reach the root bridge from anywhere in the Layer 2 network are placed in STP blocking mode.

BPDU contains information about the transmitting bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and path cost. STP uses this information to elect the root bridge for the Layer 2 network, to elect the root port leading to the root bridge, and to determine the designated port for each Layer 2 segment.

STP Protocol Timers


Variable	Description
Hello timer	Determines how often the network device broadcasts hello messages to other network devices.
Forward delay timer	Determines how long each of the listening and learning states last before the port begins forwarding.
Maximum age timer	Determines the amount of time protocol information received on an port is stored by the network device.

Creating the Spanning Tree Topology

In [Figure 31-1](#), Switch A is elected as the root bridge because the bridge priority of all the network devices is set to the default (32768) and Switch A has the lowest MAC address. However, due to traffic patterns, number of forwarding ports, or link types, Switch A might not be the ideal root bridge. By

increasing the priority (lowering the numerical value) of the ideal network device so that it becomes the root bridge, you force an STP recalculation to form a new spanning tree topology with the ideal network device as the root.

Figure 31-1 *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 a port that has a higher number than the current 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 fiber-optic link, and another port on Switch B (an unshielded twisted-pair [UTP] link) is the root port. Network traffic might be more efficient over the high-speed fiber-optic link. By changing the STP port priority on the fiber-optic port to a higher priority (lower numerical value) than the root port, the fiber-optic port becomes the new root port.

STP Port States

- [STP Port State Overview, page 31-6](#)
- [Blocking State, page 31-8](#)
- [Listening State, page 31-9](#)
- [Learning State, page 31-10](#)
- [Forwarding State, page 31-11](#)
- [Disabled State, page 31-12](#)

STP Port State Overview

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 a Layer 2 LAN port transitions directly from nonparticipation in the spanning tree topology to the forwarding state, it can create temporary data loops. Ports 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 frames that have been forwarded using the old topology.

Each Layer 2 LAN port using STP exists in one of the following five states:

- **Blocking**—The Layer 2 LAN port does not participate in frame forwarding.
- **Listening**—First transitional state after the blocking state when STP determines that the Layer 2 LAN port should participate in frame forwarding.


- Learning—The Layer 2 LAN port prepares to participate in frame forwarding.
- Forwarding—The Layer 2 LAN port forwards frames.
- Disabled—The Layer 2 LAN port does not participate in STP and is not forwarding frames.

A Layer 2 LAN port moves through these five states as follows:

- 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 31-2 illustrates how a Layer 2 LAN port moves through the five states.

Figure 31-2 STP Layer 2 LAN Interface States



When you enable STP, every port, VLAN, and network goes through the blocking state and the transitory states of listening and learning at power up. If properly configured, each Layer 2 LAN port stabilizes to the forwarding or blocking state.


When the STP algorithm places a Layer 2 LAN port in the forwarding state, the following process occurs:

1. The Layer 2 LAN port is put into the listening state while it waits for protocol information that suggests it should go to the blocking state.
2. The Layer 2 LAN port waits for the forward delay timer to expire, moves the Layer 2 LAN port to the learning state, and resets the forward delay timer.
3. In the learning state, the Layer 2 LAN port continues to block frame forwarding as it learns end station location information for the forwarding database.
4. The Layer 2 LAN port waits for the forward delay timer to expire and then moves the Layer 2 LAN port to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A Layer 2 LAN port in the blocking state does not participate in frame forwarding, as shown in Figure 31-3. After initialization, a BPDU is sent out to each Layer 2 LAN port. A network device initially assumes it is the root until it exchanges BPDUs with other network devices. This exchange establishes which network device in the network is the root or root bridge. If only one network device is in the network, no exchange occurs, the forward delay timer expires, and the ports move to the listening state. A port always enters the blocking state following initialization.

Figure 31-3 Interface 2 in Blocking State




A Layer 2 LAN port in the blocking state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Does not incorporate end station location into its address database. (There is no learning on a blocking Layer 2 LAN port, so there is no address database update.)
- Receives BPDUs and directs them to the system module.
- Does not transmit BPDUs received from the system module.
- Receives and responds to network management messages.

Listening State

The listening state is the first transitional state a Layer 2 LAN port enters after the blocking state. The Layer 2 LAN port enters this state when STP determines that the Layer 2 LAN port should participate in frame forwarding. [Figure 31-4](#) shows a Layer 2 LAN port in the listening state.

Figure 31-4 Interface 2 in Listening State




A Layer 2 LAN port in the listening state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another LAN port for forwarding.
- Does not incorporate end station location into its address database. (There is no learning at this point, so there is no address database update.)
- Receives BPDUs and directs them to the system module.
- Receives, processes, and transmits BPDUs received from the system module.
- Receives and responds to network management messages.

Learning State

A Layer 2 LAN port in the learning state prepares to participate in frame forwarding. The Layer 2 LAN port enters the learning state from the listening state. [Figure 31-5](#) shows a Layer 2 LAN port in the learning state.

Figure 31-5 *Interface 2 in Learning State*




A Layer 2 LAN port in the learning state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Incorporates end station location into its address database.
- Receives BPDUs and directs them to the system module.
- Receives, processes, and transmits BPDUs received from the system module.
- Receives and responds to network management messages.

Forwarding State

A Layer 2 LAN port in the forwarding state forwards frames, as shown in [Figure 31-6](#). The Layer 2 LAN port enters the forwarding state from the learning state.

Figure 31-6 **Interface 2 in Forwarding State**




A Layer 2 LAN port in the forwarding state performs as follows:

- Forwards frames received from the attached segment.
- Forwards frames switched from another port for forwarding.
- Incorporates end station location information into its address database.
- Receives BPDUs and directs them to the system module.
- Processes BPDUs received from the system module.
- Receives and responds to network management messages.

Disabled State

A Layer 2 LAN port in the disabled state does not participate in frame forwarding or STP, as shown in Figure 31-7. A Layer 2 LAN port in the disabled state is virtually nonoperational.

Figure 31-7 *Interface 2 in Disabled State*



A disabled Layer 2 LAN port performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Does not incorporate end station location into its address database. (There is no learning, so there is no address database update.)
- Does not receive BPDUs.
- Does not receive BPDUs for transmission from the system module.

STP and IEEE 802.1Q Trunks

802.1Q trunks impose some limitations on the STP strategy for a network. In a network of Cisco network devices connected through 802.1Q trunks, the network devices maintain one instance of STP for each VLAN allowed on the trunks. However, non-Cisco 802.1Q network devices maintain only one instance of STP for all VLANs allowed on the trunks.

When you connect a Cisco network device to a non-Cisco device through an 802.1Q trunk, the Cisco network device combines the STP instance of the 802.1Q VLAN of the trunk with the STP instance of the non-Cisco 802.1Q network device. However, all per-VLAN STP information is maintained by Cisco network devices separated by a cloud of non-Cisco 802.1Q network devices. The non-Cisco 802.1Q cloud separating the Cisco network devices is treated as a single trunk link between the network devices.

For more information on 802.1Q trunks, see [Chapter 21, “LAN Ports for Layer 2 Switching.”](#)

Information about IEEE 802.1w RSTP

- [RSTP Overview, page 31-13](#)
- [Port Roles and the Active Topology, page 31-13](#)
- [Rapid Convergence, page 31-14](#)
- [Synchronization of Port Roles, page 31-15](#)
- [Bridge Protocol Data Unit Format and Processing, page 31-16](#)
- [Topology Changes, page 31-17](#)

RSTP Overview

RSTP, which is enabled by default, takes advantage of point-to-point wiring and provides rapid convergence of the spanning tree. Reconfiguration of the spanning tree can occur in less than 1 second (in contrast to 50 seconds with the default settings in the 802.1D spanning tree).

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 802.1D STP to select the switch with the highest switch priority (lowest numerical priority value) as the root bridge as described in the [“Election of the Root Bridge” section on page 31-5](#). The RSTP then 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 bridge.
- Designated port—Connects to the designated switch, which incurs the lowest path cost when forwarding packets from that LAN to the root bridge. 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 bridge 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 802.1D). The port state controls the operation of the forwarding and learning processes. [Table 31-2](#) provides a comparison of 802.1D and RSTP port states.

Table 31-2 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 31-8](#), switch A is connected to switch B through a point-to-point link, and all of the ports are in the blocking state. Assume that the priority of switch A is a smaller numerical value than the priority of switch B. Switch A sends a proposal message (a configuration BPDU with the proposal flag set) to switch B, proposing itself as the designated switch.

After receiving the proposal message, switch B selects as its new root port the port from which the proposal message was received, forces all nonedge ports to the blocking state, and sends an agreement message (a BPDU with the agreement flag set) through its new root port.

After receiving switch B's agreement message, switch A also immediately transitions its designated port to the forwarding state. No loops in the network are formed because switch B blocked all of its nonedge ports and because there is a point-to-point link between switches A and B.

When switch C is connected to switch B, a similar set of handshaking messages are exchanged. Switch C selects the port connected to switch B as its root port, and both ends immediately transition to the forwarding state. With each iteration of this handshaking process, one more switch joins the active topology. As the network converges, this proposal-agreement handshaking progresses from the root toward the leaves of the spanning tree.

The switch learns the link type from the port duplex mode: a full-duplex port is considered to have a point-to-point connection and 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 31-8 Proposal and Agreement Handshaking for Rapid Convergence

88760

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 31-9](#).

Figure 31-9 Sequence of Events During Rapid Convergence

Bridge Protocol Data Unit Format and Processing

- [BPDU Format and Processing Overview, page 31-16](#)
- [Processing Superior BPDU Information, page 31-17](#)
- [Processing Inferior BPDU Information, page 31-17](#)

BPDU Format and Processing Overview

The RSTP BPDU format is the same as the 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 31-3](#) describes the RSTP flag fields.

Table 31-3 RSTP BPDU Flags

Bit	Function
0	Topology change (TC)
1	Proposal
2–3:	Port role:
00	Unknown
01	Alternate port or backup 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 TCN BPDU. It uses the topology change (TC) flag to show the topology changes. However, for interoperability with 802.1D switches, the RSTP switch processes and generates TCN BPDUs.

The learning and forwarding flags are set according to the state of the sending port.

Processing Superior BPDU Information

A superior BPDU is a BPDU with root information (such as lower switch ID or lower path cost) that is superior to what is currently stored for the port.

If a port receives a superior BPDU, the RSTP triggers a reconfiguration. If the port is proposed and is selected as the new root port, RSTP forces all the other ports to synchronize.

If the BPDU received is an RSTP BPDU with the proposal flag set, the switch sends an agreement message after all of the other ports are synchronized. If the BPDU is an 802.1D BPDU, the switch does not set the proposal flag and starts the forward-delay timer for the port. The new root port requires twice the forward-delay time to transition to the forwarding state.

If the superior information received on the port causes the port to become a backup port or an alternate port, RSTP sets the port to the blocking state and sends an 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

An inferior BPDU is a BPDU with root information (such as higher switch ID or higher path cost) that is inferior to what is currently stored for the port.

If a designated port receives an inferior BPDU, it immediately replies with its own information.

Topology Changes

These are the differences between the RSTP and the 802.1D in handling spanning tree topology changes:

- **Detection**—Unlike 802.1D in which *any* transition between the blocking and the forwarding state causes a topology change, *only* transitions from the blocking to the forwarding state cause a topology change with RSTP (only an increase in connectivity is considered a topology change). State changes on an edge port do not cause a topology change. When an RSTP switch detects a topology change, it deletes the learned information on all of its nonedge ports except on those from which it received the TC notification.
- **Notification**—The RSTP does not use TCN BPDUs, unlike 802.1D. However, for 802.1D interoperability, an RSTP switch processes and generates TCN BPDUs.
- **Acknowledgement**—When an RSTP switch receives a TCN message on a designated port from an 802.1D switch, it replies with an 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the TC timer in 802.1D) is active on a root port connected to an 802.1D switch and a configuration BPDU with the TCA set is received, the TC-while timer is reset.

This method of operation is only required to support 802.1D switches. The RSTP BPDUs never have the TCA bit set.

- **Propagation**—When an RSTP switch receives a TC message from another switch through a designated or root port, it propagates the change to all of its nonedge, designated ports and to the root port (excluding the port on which it is received). The switch starts the TC-while timer for all such ports and flushes the information learned on them.
- **Protocol migration**—For backward compatibility with 802.1D switches, RSTP selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.

When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which RSTP BPDUs are sent), and RSTP BPDUs are sent. While this timer is active, the switch processes all BPDUs received on that port and ignores the protocol type.

If the switch receives an 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an 802.1D switch and starts using only 802.1D BPDUs. However, if the RSTP switch is using 802.1D BPDUs on a port and receives an RSTP BPDU after the timer has expired, it restarts the timer and starts using RSTP BPDUs on that port.

Information about MST

- [MST Overview, page 31-18](#)
- [MST Regions, page 31-19](#)
- [IST, CIST, and CST, page 31-19](#)
- [Hop Count, page 31-22](#)
- [Boundary Ports, page 31-22](#)
- [Standard-Compliant MST Implementation, page 31-23](#)
- [Interoperability with IEEE 802.1D-1998 STP, page 31-24](#)

MST Overview

MST maps multiple VLANs 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. MST 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 MST is in the backbone and distribution layers of a Layer 2 switched network. This deployment provides the kind of highly available network that is required in a service-provider environment.

MST provides rapid spanning tree convergence through explicit handshaking, which eliminates the 802.1D forwarding delay and quickly transitions root bridge ports and designated ports to the forwarding state.

MST improves spanning tree operation and maintains backward compatibility with these STP versions:

- Original 802.1D spanning tree
- Existing Cisco-proprietary Multiple Instance STP (MISTP)
- Existing Cisco per-VLAN spanning tree plus (PVST+)
- Rapid per-VLAN spanning tree plus (rapid PVST+)

**Note**

- IEEE 802.1w defined the Rapid Spanning Tree Protocol (RSTP) and was incorporated into IEEE 802.1D.
- IEEE 802.1s defined MST and was incorporated into IEEE 802.1Q.

MST Regions

For switches to participate in 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 31-10 on page 31-21](#).

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.

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

- [IST, CIST, and CST Overview, page 31-19](#)
- [Spanning Tree Operation Within an MST Region, page 31-20](#)
- [Spanning Tree Operations Between MST Regions, page 31-20](#)
- [IEEE 802.1s Terminology, page 31-21](#)

IST, CIST, and CST Overview

Unlike other spanning tree protocols, in which all the spanning tree instances are independent, MST establishes and maintains internal spanning tree (IST), common and internal spanning tree (CIST), and common spanning tree (CST) instances:

- An IST is the spanning tree that runs in an MST region.

Within each MST region, MST maintains multiple spanning tree instances. Instance 0 is a special instance for a region, known as the 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 MSTP records (M-records), which are encapsulated within MST BPDUs. Because the MST 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 bridge 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 CIST is a collection of the ISTs in each MST region.
- The CST 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 802.1w, 802.1s, and 802.1D standards. The CIST inside an MST region is the same as the CST outside a region.

For more information, see the [“Spanning Tree Operation Within an MST Region” section on page 31-20](#) and the [“Spanning Tree Operations Between MST Regions” section on page 31-20](#).

Spanning Tree Operation Within an MST Region

The IST connects all the MST 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 802.1s standard) as shown in [Figure 31-10 on page 31-21](#). 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 MST switches at the boundary of the region is selected as the CIST regional root.

When an MST switch initializes, it sends BPDUs that identify 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 from a neighbor in the same region, they leave their old subregions and join the new subregion that contains the true CIST regional root, which causes all subregions to 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.

Spanning Tree Operations Between MST Regions

If there are multiple regions or 802.1D switches within the network, MST establishes and maintains the CST, which includes all MST regions and all 802.1D 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 MST 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 31-10](#) shows a network with three MST regions and an 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.

Figure 31-10 MST Regions, CIST Regional Roots, 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, 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.

MST switches use Version 3 BPDUs or 802.1D STP BPDUs to communicate with 802.1D switches. MST switches use MST BPDUs to communicate with MST switches.

IEEE 802.1s Terminology

Some MST naming conventions used in the prestandard implementation have been changed to include identification of some *internal* and *regional* parameters. These parameters are used only within an MST region, compared to external parameters that are used throughout the whole network. Because the CIST is the only spanning tree instance that spans the whole network, only the CIST parameters require the external qualifiers and not the internal or regional qualifiers.

- The CIST root is the root bridge for the CIST, which is the unique instance that spans the whole network.
- 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 to 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 bridge 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 31-4 compares the IEEE standard and the Cisco prestandard terminology.

Table 31-4 *Prestandard and Standard Terminology*

IEEE Standard Definition	Cisco Prestandard Implementation	Cisco Standard Implementation
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

MST does 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 bridge 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 one of these STP regions:

- A single spanning tree region running RSTP
- A single spanning tree region running PVST+ or rapid PVST+
- 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 802.1s standard. The 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, which 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 from 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 an 802.1D STP switch on the segment, messages are always considered external.

The other change from the prestandard implementation is that the CIST regional root bridge ID field is now inserted where an RSTP or legacy 802.1s 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.

Standard-Compliant MST Implementation

- [Changes in Port-Role Naming, page 31-23](#)
- [Spanning Tree Interoperation Between Legacy and Standard-Compliant Switches, page 31-23](#)

**Note**

The standard-compliant MST implementation 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.

Changes in Port-Role Naming

The boundary role was deleted from the final MST standard, but this boundary concept is maintained in the standard-compliant implementation. However, an MST instance (MSTI) port at a boundary of the region might not follow the state of the corresponding CIST port. The following two situations currently exist:

- The boundary port is the root port of the CIST regional root—When the CIST instance port is proposed and is synchronized, it can send back an agreement and move to the forwarding state only after all the corresponding MSTI ports are synchronized (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 (M-records). In this situation, although the boundary role no longer exists, when you enter the **show** commands, they identify a port as boundary in the *type* column of the output.


Spanning Tree Interoperation Between Legacy and Standard-Compliant 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 before using the CIST. Only the capability of load balancing over different instances is lost in this specific situation. The CLI displays different flags depending on the

port configuration when the 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 31-11 illustrates a standard-compliant switch connected to a prestandard switch. Assume that A is the standard-compliant switch and B is a prestandard switch, both configured to be in the same region. A is the root bridge for the CIST, and so 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 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 31-11 Standard-Compliant and Prestandard Switch Interoperation



Note

We recommend that you minimize the interaction between standard and prestandard MST implementations.

Interoperability with IEEE 802.1D-1998 STP

A switch running MST supports a built-in protocol migration feature that enables it to interoperate with 802.1D switches. If this switch receives an 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MST switch also can detect that a port is at the boundary of a region when it receives an 802.1D BPDU, an MST BPDU (Version 3) associated with a different region, or an RSTP BPDU (Version 2).

However, the switch does not automatically revert to the MST mode if it no longer receives 802.1D BPDUs because it cannot detect whether the 802.1D switch has been removed from the link unless the 802.1D 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 802.1D switches on the link are RSTP switches, they can process MST BPDUs as if they are RSTP BPDUs. Therefore, MST switches send either a Version 0 configuration and topology change notification (TCN) BPDUs or Version 3 MST 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.


Detecting Unidirectional Link Failure

Using the dispute mechanism included in the IEEE 802.1D-2004 RSTP and IEEE 802.1Q-2005 MSTP standard, the switch 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 a discarding (blocking) state because disrupting connectivity in case of inconsistency is preferable to opening a bridging loop.

Figure 31-12 illustrates a unidirectional link failure that typically creates a bridging loop. Switch A is the root bridge, 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 bridge. As a result, switch A blocks (or keeps blocking) its port, thus preventing the bridging loop.

Figure 31-12 Detecting Unidirectional Link Failure



Default Settings for Spanning Tree Protocols

- [Default STP Configuration, page 31-25](#)
- [Default MST Configuration, page 31-26](#)

Default STP Configuration

Feature	Default Value	
Mode	Rapid PVST+	
Enable state	Enabled for all VLANs	
Bridge priority	32768	
STP port priority (configurable on a per-port basis—used on LAN ports configured as Layer 2 access ports)	128	
STP port cost (configurable on a per-port basis—used on LAN ports configured as Layer 2 access ports)	10-Gigabit Ethernet:	2
	Gigabit Ethernet:	4
	Fast Ethernet:	19
	Ethernet:	100
STP VLAN port priority (configurable on a per-VLAN basis—used on LAN ports configured as Layer 2 trunk ports)	128	

Feature	Default Value	
STP VLAN port cost (configurable on a per-VLAN basis—used on LAN ports configured as Layer 2 trunk ports)	10-Gigabit Ethernet:	2
	Gigabit Ethernet:	4
	Fast Ethernet:	19
	Ethernet:	100
Hello time	2 seconds	
Forward delay time	15 seconds	
Maximum aging time	20 seconds	

Default MST Configuration

Feature	Default Setting	
Spanning tree mode	Rapid PVST+ (MST is disabled)	
Switch priority (configurable on a per-MST port basis)	32768	
Spanning tree port priority (configurable on a per-MST instance port basis)	128	
Spanning tree port cost (configurable on a per-MST instance port basis)	10-Gigabit Ethernet:	2,000
	Gigabit Ethernet:	20,000
	Fast Ethernet:	200,000
	Ethernet:	2,000,000
Hello time	2 seconds	
Forward-delay time	15 seconds	
Maximum-aging time	20 seconds	
Maximum hop count	20 hops	

How to Configure Spanning Tree Protocols

- [Configuring STP, page 31-26](#)
- [Configuring MST, page 31-37](#)

Configuring STP

- [Enabling STP, page 31-27](#)
- [Enabling the Extended System ID, page 31-28](#)
- [Configuring the Root Bridge, page 31-29](#)
- [Configuring a Secondary Root Bridge, page 31-30](#)

- [Configuring STP Port Priority, page 31-31](#)
- [Configuring STP Port Cost, page 31-32](#)
- [Configuring the Bridge Priority of a VLAN, page 31-34](#)
- [Configuring the Hello Time, page 31-35](#)
- [Configuring the Forward-Delay Time for a VLAN, page 31-35](#)
- [Configuring the Maximum Aging Time for a VLAN, page 31-36](#)
- [Enabling Rapid-PVST+, page 31-36](#)

**Note**

The STP commands described in this chapter can be configured on any LAN port, but they are in effect only on LAN ports configured with the **switchport** keyword.

**Caution**

We do not recommend disabling spanning tree, even in a topology that is free of physical loops. Spanning tree serves as a safeguard against misconfigurations and cabling errors. Do not disable spanning tree in a VLAN without ensuring that there are no physical loops present in the VLAN.

Enabling STP

**Note**

STP is enabled by default on VLAN 1 and on all newly created VLANs.

You can enable STP on a per-VLAN basis. The switch maintains a separate instance of STP for each VLAN (except on VLANs on which you disable STP).

To enable STP on a per-VLAN basis, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree vlan <i>vlan_ID</i>	Enables STP on a per-VLAN basis. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see “Default STP Configuration” section on page 31-25).
	Router(config)# default spanning-tree vlan <i>vlan_ID</i>	Reverts all STP parameters to default values for the specified VLAN.
	Router(config)# no spanning-tree vlan <i>vlan_ID</i>	Disables STP on the specified VLAN; see the following Cautions for information regarding this command.
Step 2	Router(config)# end	Exits configuration mode.

**Caution**

Do not disable spanning tree on a VLAN unless all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some switches and bridges in a VLAN and leave it enabled on other switches and bridges in the VLAN. This action can have unexpected results because switches and bridges with spanning tree enabled will have incomplete information regarding the physical topology of the network.

**Caution**

We do not recommend disabling spanning tree, even in a topology that is free of physical loops. Spanning tree serves as a safeguard against misconfigurations and cabling errors. Do not disable spanning tree in a VLAN without ensuring that there are no physical loops present in the VLAN.

This example shows how to enable STP on VLAN 200:

```
Router# configure terminal
Router(config)# spanning-tree vlan 200
Router(config)# end
Router#
```

**Note**

Because STP is enabled by default, entering a **show running** command to view the resulting configuration does not display the command you entered to enable STP.

This example shows how to verify the configuration:

```
Router# show spanning-tree vlan 200

VLAN0200
  Spanning tree enabled protocol ieee
  Root ID    Priority      32768
             Address      00d0.00b8.14c8
             This bridge is the root
             Hello Time    2 sec    Max Age 20 sec    Forward Delay 15 sec

  Bridge ID  Priority      32768
             Address      00d0.00b8.14c8
             Hello Time    2 sec    Max Age 20 sec    Forward Delay 15 sec
             Aging Time    300

Interface                Role Sts Cost          Prio.Nbr Status
-----
Gi1/4                    Desg FWD 200000        128.196 P2p
Gi1/5                    Back BLK 200000        128.197 P2p
```

```
Router#
```

**Note**

You must have at least one interface that is active in VLAN 200 to create a VLAN 200 spanning tree. In this example, two interfaces are active in VLAN 200.

Enabling the Extended System ID

**Note**

- The extended system ID is enabled permanently on chassis that support 64 MAC addresses.
- You must enable the extended system ID to configure extended range VLANs (1006-4094).
- You must enable the extended system ID if it is enabled on any switches in the VTP domain.

You can enable the extended system ID on chassis that support 1024 MAC addresses (see the [“Information about the Bridge ID” section on page 31-3](#)).

To enable the extended system ID, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree extend system-id	Enables the extended system ID. Note You cannot disable the extended system ID on chassis that support 64 MAC addresses or when you have configured extended range VLANs (see the “VLAN Ranges” section on page 26-2).
Step 2	Router(config)# end	Exits configuration mode.

**Note**

When you enable or disable the extended system ID, the bridge IDs of all active STP instances are updated, which might change the spanning tree topology.

This example shows how to enable the extended system ID:

```
Router# configure terminal
Router(config)# spanning-tree extend system-id
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree summary | include Extended
Extended system ID is enabled.
```

Configuring the Root Bridge

The switches supported by Cisco IOS Release 15.4SY maintain a separate instance of STP for each active VLAN. A bridge ID, consisting of the bridge priority and the bridge MAC address, is associated with each instance. For each VLAN, the network device with the lowest bridge ID becomes the root bridge for that VLAN.

To configure a VLAN instance to become the root bridge, enter the **spanning-tree vlan *vlan_ID* root** command to modify the bridge priority from the default value (32768) to a significantly lower value.

When you enter the **spanning-tree vlan *vlan_ID* root** command, the switch checks the bridge priority of the current root bridges for each VLAN. With the extended system ID enabled, the switch sets the bridge priority for the specified VLANs to 24576 if this value will cause the switch to become the root for the specified VLANs.

With the extended system ID enabled, if any root bridge for the specified VLANs has a bridge priority lower than 24576, the switch sets the bridge priority for the specified VLANs to 4096 less than the lowest bridge priority. (4096 is the value of the least significant bit of a 4-bit bridge priority value; see [Table 31-1 on page 31-4](#).)

**Note**

The **spanning-tree vlan *vlan_ID* root** command fails if the value required to be the root bridge is less than 1.

With the extended system ID enabled, if all network devices in, for example, VLAN 20 have the default priority of 32768, entering the **spanning-tree vlan 20 root primary** command on the switch sets the bridge priority to 24576, which causes the switch to become the root bridge for VLAN 20.



Caution

The root bridge for each instance of STP should be a backbone or distribution switch. Do not configure an access switch as the STP primary root.

Use the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of bridge hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically selects an optimal hello time, forward delay time, and maximum age time for a network of that diameter, which can significantly reduce the STP convergence time. You can use the **hello** keyword to override the automatically calculated hello time.



Note

To preserve a stable STP topology, we recommend that you avoid configuring the hello time, forward delay time, and maximum age time manually after configuring the switch as the root bridge.

To configure the switch as the root bridge, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree vlan <i>vlan_ID</i> root primary [diameter <i>hops</i> [hello-time <i>seconds</i>]]	Configures the switch as the root bridge. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see “Default STP Configuration” section on page 31-25).
	Router(config)# no spanning-tree vlan <i>vlan_ID</i> root	Clears the root bridge configuration.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the switch as the root bridge for VLAN 10, with a network diameter of 4:

```
Router# configure terminal
Router(config)# spanning-tree vlan 10 root primary diameter 4
Router(config)# end
Router#
```

Configuring a Secondary Root Bridge

When you configure a switch as the secondary root, the STP bridge priority is modified from the default value (32768) so that the switch is likely to become the root bridge for the specified VLANs if the primary root bridge fails (assuming the other network devices in the network use the default bridge priority of 32768).

With the extended system ID is enabled, STP sets the bridge priority to 28672.

You can run this command on more than one switch to configure multiple backup root bridges. Use the same network diameter and hello time values as you used when configuring the primary root bridge.

To configure the switch as the secondary root bridge, perform this task:

	Command	Purpose
Step 1	Router(config)# [no] spanning-tree vlan <i>vlan_ID</i> root secondary [diameter <i>hops</i> [hello-time <i>seconds</i>]]	Configures the switch as the secondary root bridge. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see “Default STP Configuration” section on page 31-25).
	Router(config)# no spanning-tree vlan <i>vlan_ID</i> root	Clears the root bridge configuring.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the switch as the secondary root bridge for VLAN 10, with a network diameter of 4:

```
Router# configure terminal
Router(config)# spanning-tree vlan 10 root secondary diameter 4
Router(config)# end
Router#
```

Configuring STP Port Priority

If a loop occurs, STP considers port priority when selecting a LAN port to put into the forwarding state. You can assign higher priority values to LAN ports that you want STP to select first and lower priority values to LAN ports that you want STP to select last. If all LAN ports have the same priority value, STP puts the LAN port with the lowest LAN port number in the forwarding state and blocks other LAN ports. The possible priority range is 0 through 240 (default 128), configurable in increments of 16.

Cisco IOS uses the port priority value when the LAN port is configured as an access port and uses VLAN port priority values when the LAN port is configured as a trunk port.

To configure the STP port priority of a Layer 2 LAN interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface { fastethernet gigabitethernet tengigabitethernet } <i>slot/port</i> { port-channel <i>port_channel_number</i> }	Selects an interface to configure.
Step 2	Router(config-if)# spanning-tree port-priority <i>port_priority</i>	Configures the port priority for the LAN interface. The <i>port_priority</i> value can be from 1 to 252 in increments of 4.
Step 3	Router(config-if)# spanning-tree vlan <i>vlan_ID</i> port-priority <i>port_priority</i>	Configures the VLAN port priority for the LAN interface. The <i>port_priority</i> value can be from 1 to 252 in increments of 4. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3).
Step 4	Router(config-if)# end	Exits configuration mode.

This example shows how to configure the STP port priority of Gigabit Ethernet port 1/4:

```
Router# configure terminal
Router(config)# interface gigabitethernet 1/4
Router(config-if)# spanning-tree port-priority 160
Router(config-if)# end
Router#
```

This example shows how to verify the configuration of Gigabit Ethernet port 1/4:

```
Router# show spanning-tree interface gigabitethernet 1/4
Vlan          Role Sts Cost          Prio.Nbr Status
-----
VLAN0001      Back BLK 200000       160.196 P2p
VLAN0006      Back BLK 200000       160.196 P2p
...
VLAN0198      Back BLK 200000       160.196 P2p
VLAN0199      Back BLK 200000       160.196 P2p
VLAN0200      Back BLK 200000       160.196 P2p
Router#
```

Gigabit Ethernet port 1/4 is a trunk. Several VLANs are configured and active as shown in the example. The port priority configuration applies to all VLANs on this interface.



Note

The **show spanning-tree interface** command only displays information if the port is connected and operating. If this condition is not met, enter a **show running-config interface** command to verify the configuration.

This example shows how to configure the VLAN port priority of Gigabit Ethernet port 1/4:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 1/4
Router(config-if)# spanning-tree vlan 200 port-priority 64
Router(config-if)# end
Router#
```

The configuration entered in the example only applies to VLAN 200. All VLANs other than 200 still have a port priority of 160.

This example shows how to verify the configuration:

```
Router# show spanning-tree interface gigabitethernet 1/4
Vlan          Role Sts Cost          Prio.Nbr Status
-----
VLAN0001      Back BLK 200000       160.196 P2p
VLAN0006      Back BLK 200000       160.196 P2p
...
VLAN0199      Back BLK 200000       160.196 P2p
VLAN0200      Desg FWD 200000         64.196 P2p
Router#
```

You also can display spanning tree information for VLAN 200 using the following command:

```
Router# show spanning-tree vlan 200 interface gigabitethernet 1/4
Interface      Role Sts Cost          Prio.Nbr Status
-----
Gi1/4          Desg LRN 200000         64.196 P2p
```

Configuring STP Port Cost

The STP port path cost default value is determined from the media speed of a LAN interface. If a loop occurs, STP considers port cost when selecting a LAN interface to put into the forwarding state. You can assign lower cost values to LAN interfaces that you want STP to select first and higher cost values to LAN interfaces that you want STP to select last. If all LAN interfaces have the same cost value, STP puts the LAN interface with the lowest LAN interface number in the forwarding state and blocks other LAN interfaces. The possible cost range is 0 through 200000000 (the default is media specific).

STP uses the port cost value when the LAN interface is configured as an access port and uses VLAN port cost values when the LAN interface is configured as a trunk port.

To configure the STP port cost of a Layer 2 LAN interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface { fastethernet gigabitethernet tengigabitethernet } <i>slot/port</i> { port-channel <i>port_channel_number</i> }	Selects an interface to configure.
Step 2	Router(config-if)# spanning-tree cost <i>port_cost</i> Router(config-if)# no spanning-tree cost	Configures the port cost for the LAN interface. The <i>port_cost</i> value can be from 1 to 200000000. Reverts to the default port cost.
Step 3	Router(config-if)# spanning-tree vlan <i>vlan_ID</i> cost <i>port_cost</i> Router(config-if)# no spanning-tree vlan <i>vlan_ID</i> cost	Configures the VLAN port cost for the LAN interface. The <i>port_cost</i> value can be from 1 to 200000000. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3). Reverts to the default VLAN port cost.
Step 4	Router(config-if)# end	Exits configuration mode.

This example shows how to change the STP port cost of Gigabit Ethernet port 1/4:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 1/4
Router(config-if)# spanning-tree cost 1000
Router(config-if)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree interface gigabitethernet 1/4
Vlan          Role Sts Cost      Prio.Nbr Status
-----
VLAN0001      Back BLK 1000    160.196 P2p
VLAN0006      Back BLK 1000    160.196 P2p
VLAN0007      Back BLK 1000    160.196 P2p
VLAN0008      Back BLK 1000    160.196 P2p
VLAN0009      Back BLK 1000    160.196 P2p
VLAN0010      Back BLK 1000    160.196 P2p
Router#
```

This example shows how to configure the port priority at an individual port VLAN cost for VLAN 200:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 1/4
Router(config-if)# spanning-tree vlan 200 cost 2000
Router(config-if)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree vlan 200 interface gigabitethernet 1/4
Interface      Role Sts Cost      Prio.Nbr Status
-----
Gi1/4          Desg FWD 2000    64.196 P2p
```


Note

In the following output other VLANs (VLAN 1 for example) have not been affected by this configuration.

```
Router# show spanning-tree vlan 1 interface gigabitethernet 1/4
Interface          Role Sts Cost          Prio.Nbr Status
-----
Gil/4              Back BLK 1000      160.196 P2p
Router#
```


Note

The **show spanning-tree** command only displays information for ports that are in link-up operative state and are appropriately configured for DTP. If these conditions are not met, you can enter a **show running-config** command to confirm the configuration.

Configuring the Bridge Priority of a VLAN


Note

Be careful when using this command. For most situations, we recommend that you enter the **spanning-tree vlan *vlan_ID* root primary** and the **spanning-tree vlan *vlan_ID* root secondary** commands to modify the bridge priority.

To configure the STP bridge priority of a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree vlan <i>vlan_ID</i> priority {0 4096 8192 12288 16384 20480 24576 28672 32768 36864 40960 45056 49152 53248 57344 61440}	Configures the bridge priority of a VLAN when the extended system ID is enabled. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3).
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the bridge priority of VLAN 200 to 33792 when the extended system ID is disabled:

```
Router# configure terminal
Router(config)# spanning-tree vlan 200 priority 32768
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree vlan 200 bridge

Vlan          Bridge ID          Hello Time  Max Age  Fwd Delay  Protocol
-----
VLAN200      32768 0050.3e8d.64c8    2        20        15    ieee
Router#
```


Configuring the Hello Time



Note

Be careful when using this command. For most situations, we recommend that you use the **spanning-tree vlan *vlan_ID* root primary** and **spanning-tree vlan *vlan_ID* root secondary** commands to modify the hello time.

To configure the STP hello time of a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree vlan <i>vlan_ID</i> hello-time <i>hello_time</i>	Configures the hello time of a VLAN. The <i>hello_time</i> value can be from 1 to 10 seconds. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3).
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the hello time for VLAN 200 to 7 seconds:

```
Router# configure terminal
Router(config)# spanning-tree vlan 200 hello-time 7
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree vlan 200 bridge

Vlan                Bridge ID      Hello Time  Max Age  Fwd Delay  Protocol
-----
VLAN200             49152 0050.3e8d.64c8  7         20       15       ieee
Router#
```

Configuring the Forward-Delay Time for a VLAN

To configure the STP forward delay time for a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree vlan <i>vlan_ID</i> forward-time <i>forward_time</i>	Configures the forward time of a VLAN. The <i>forward_time</i> value can be from 4 to 30 seconds. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3).
	Router(config)# no spanning-tree vlan <i>vlan_ID</i> forward-time	Reverts to the default forward time.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the forward delay time for VLAN 200 to 21 seconds:

```
Router# configure terminal
Router(config)# spanning-tree vlan 200 forward-time 21
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree vlan 200 bridge
Vlan                Bridge ID      Hello Time  Max Age  Fwd Delay  Protocol
-----
VLAN200             49152 0050.3e8d.64c8  2         20       21         ieee
Router#
```

Configuring the Maximum Aging Time for a VLAN

To configure the STP maximum aging time for a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree vlan <i>vlan_ID</i> max-age <i>max_age</i>	Configures the maximum aging time of a VLAN. The <i>max_age</i> value can be from 6 to 40 seconds. The <i>vlan_ID</i> value can be 1 through 4094, except reserved VLANs (see Table 26-1 on page 26-3).
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the maximum aging time for VLAN 200 to 36 seconds:

```
Router# configure terminal
Router(config)# spanning-tree vlan 200 max-age 36
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show spanning-tree vlan 200 bridge
Vlan                Bridge ID      Hello Time  Max Age  Fwd Delay  Protocol
-----
VLAN200             49152 0050.3e8d.64c8  2         36       15         ieee
Router#
```

Enabling Rapid-PVST+

- [Rapid-PVST+ Overview, page 31-36](#)
- [Specifying the Link Type, page 31-37](#)
- [Restarting Protocol Migration, page 31-37](#)



Note

Rapid-PVST+ is enabled by default. Use the procedures to reenale Rapid-PVST+.

Rapid-PVST+ Overview

Rapid-PVST+ uses the existing PVST+ framework for configuration and interaction with other features. It also supports some of the PVST+ extensions.

To enable Rapid-PVST+ mode on the switch, enter the **spanning-tree mode rapid-pvst** command in privileged mode. To configure the switch in Rapid-PVST+ mode, see the [“Configuring STP” section on page 31-26](#).

Specifying the Link Type

Rapid connectivity is established only on point-to-point links. Spanning tree views a point-to-point link as a segment connecting only two switches running the spanning tree algorithm. Because the switch assumes that all full-duplex links are point-to-point links and that half-duplex links are shared links, you can avoid explicitly configuring the link type. To configure a specific link type, enter the **spanning-tree linktype** command.

Restarting Protocol Migration

A switch running both MSTP and RSTP supports a built-in protocol migration process that enables the switch to interoperate with legacy 802.1D switches. If this switch receives a legacy 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MSTP switch can also detect that a port is at the boundary of a region when it receives a legacy BPDU, or an MST BPDU (version 3) associated with a different region, or an RST BPDU (version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives 802.1D BPDUs because it cannot determine whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. A switch also might continue to assign a boundary role to a port when the switch to which it is connected has joined the region.

To restart the protocol migration process (force the renegotiation with neighboring switches) on the entire switch, you can use the **clear spanning-tree detected-protocols** privileged EXEC command. To restart the protocol migration process on a specific interface, enter the **clear spanning-tree detected-protocols interface** *interface-id* privileged EXEC command.

Configuring MST

- [Default MST Configuration, page 31-26](#)
- [Specifying the MST Region Configuration and Enabling MST, page 31-38](#) (required)
- [Configuring the Root Bridge, page 31-39](#) (optional)
- [Configuring a Secondary Root Bridge, page 31-30](#) (optional)
- [Configuring STP Port Priority, page 31-31](#) (optional)
- [Configuring Path Cost, page 31-42](#) (optional)
- [Configuring the Switch Priority, page 31-43](#) (optional)
- [Configuring the Hello Time, page 31-44](#) (optional)
- [Configuring the Transmit Hold Count, page 31-45](#) (optional)
- [Configuring the Maximum-Aging Time, page 31-45](#) (optional)
- [Configuring the Maximum-Hop Count, page 31-46](#) (optional)
- [Specifying the Link Type to Ensure Rapid Transitions, page 31-46](#) (optional)
- [Designating the Neighbor Type, page 31-46](#) (optional)
- [Restarting the Protocol Migration Process, page 31-47](#) (optional)
- [Displaying the MST Configuration and Status, page 31-47](#)

Specifying the MST Region Configuration and Enabling MST

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 MST 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.

To specify the MST region configuration and enable MST, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree mst configuration	Enters MST configuration mode.
Step 3	Router(config-mst)# instance <i>instance_id</i> vlan <i>vlan_range</i>	Maps VLANs to an MST instance. - For <i>instance_id</i>, the range is 0 to 4094. - For vlan <i>vlan_range</i>, the range is 1 to 4094. When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped. To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63 maps VLANs 1 through 63 to MST instance 1. To specify a VLAN series, use a comma; for example, instance 1 vlan 10, 20, 30 maps VLANs 10, 20, and 30 to MST instance 1.
Step 4	Router(config-mst)# name <i>instance_name</i>	Specifies the instance name. The <i>name</i> string has a maximum length of 32 characters and is case sensitive.
Step 5	Router(config-mst)# revision <i>version</i>	Specifies the configuration revision number. The range is 0 to 65535.
Step 6	Router(config-mst)# show pending	Verifies your configuration by displaying the pending configuration.
Step 7	Router(config)# exit	Applies all changes, and return to global configuration mode.
Step 8	Router(config)# spanning-tree mode mst	Enables MST and RSTP.  Caution Changing the spanning tree mode can disrupt traffic because all spanning tree instances are stopped for the previous mode and restarted in the new mode. You cannot run both MST and rapid PVST+ at the same time.
Step 9	Router(config)# end	Returns to privileged EXEC mode.

To return to defaults, do the following:

- 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.

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:

```
Router(config)# spanning-tree mst configuration
Router(config-mst)# instance 1 vlan 10-20
Router(config-mst)# name region1
Router(config-mst)# revision 1
Router(config-mst)# show pending
Pending MST configuration
Name      [region1]
Revision  1
Instances configured 2
Instance  Vlans Mapped
-----
0          1-9,21-4094
1          10-20
-----

Router(config-mst)# exit
Router(config)#
```

Configuring the Root Bridge

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 bridge.

To configure a switch to become the root bridge, 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 bridge for the specified spanning tree instance. When you enter this command, the switch checks the switch priorities of the root bridges. Because of 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 bridge for the specified spanning tree instance.

If any root bridge 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 31-1 on page 31-4](#).)

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 bridge. 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 bridge for each spanning tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning tree primary root bridge.

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 Layer 2 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

With the switch configured as the root bridge, do not manually configure the hello time, forward-delay time, and maximum-age time with the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and **spanning-tree mst max-age** global configuration commands.

To configure a switch as the root bridge, perform this task:

	Command	Purpose
Step 1	Router(config)# configure terminal	Enters global configuration mode.
Step 2	Router(config-config)# spanning-tree mst <i>instance_id</i> root primary [diameter <i>net_diameter</i> hello-time <i>seconds</i>]	(Optional) Configures a switch as the root bridge. • 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 Layer 2 hops 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 bridge. The range is 1 to 10 seconds; the default is 2 seconds.
Step 3	Router(config-config)# end	Returns to privileged EXEC mode.

Configuring a Secondary Root Bridge

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 bridge for the specified instance if the primary root bridge fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root bridge.

You can execute this command on more than one switch to configure multiple backup root bridges. Use the same network diameter and hello-time values that you used when you configured the primary root bridge with the **spanning-tree mst instance_id root primary** global configuration command.

To configure a switch as the secondary root bridge, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree mst <i>instance_id</i> root secondary [diameter <i>net_diameter</i> [hello-time <i>seconds</i>]]	(Optional) Configures a switch as the secondary root bridge. - 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 bridge. 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 bridge. See the “Configuring the Root Bridge” section on page 31-39 .
Step 3	Router(config)# end	Returns to privileged EXEC mode.

Configuring Port Priority

If a loop occurs, MST 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, MST puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

To configure the MST port priority of an interface, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface { fastethernet gigabitethernet tengigabitethernet } <i>slot/port</i> { port-channel <i>number</i> }	(Optional) Specifies an interface to configure, and enters interface configuration mode.

	Command	Purpose
Step 3	Router(config-if)# spanning-tree mst <i>instance_id</i> port-priority <i>priority</i>	Configures the port priority. - For <i>instance_id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 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	Router(config-if)# end	Returns to privileged EXEC mode.

**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.

Configuring Path Cost

The MST path cost default value is derived from the media speed of an interface. If a loop occurs, MST 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, MST puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

To configure the MST cost of an interface, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface {{ fastethernet gigabitethernet tengigabitethernet } <i>slot/port</i> { port-channel <i>number</i> }}	(Optional) Specifies an interface to configure, and enters interface configuration mode.

	Command	Purpose
Step 3	Router(config-if)# spanning-tree mst <i>instance_id</i> cost <i>cost</i>	Configures the cost. If a loop occurs, MST uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. - For <i>instance_id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 4094. - For <i>cost</i>, the range is 1 to 2000000000; the default value is derived from the media speed of the interface.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

**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.

Configuring the Switch Priority

You can configure the switch priority so that it is more likely that a switch is chosen as the root bridge.

**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.

To configure the switch priority, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.

	Command	Purpose
Step 2	Router(config)# spanning-tree mst <i>instance_id</i> priority <i>priority</i>	(Optional) Configures the switch priority. - For <i>instance_id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 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 bridge. 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	Router(config)# end	Returns to privileged EXEC mode.

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root bridge by changing the hello time.



Note

Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree mst** *instance_id* **root primary** and the **spanning-tree mst** *instance_id* **root secondary** global configuration commands to modify the hello time.

To configure the hello time for all MST instances, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree mst hello-time <i>seconds</i>	(Optional) Configures the hello time for all MST instances. The hello time is the interval between the generation of configuration messages by the root bridge. These messages mean that the switch is alive. For <i>seconds</i> , the range is 1 to 10; the default is 2.
Step 3	Router(config)# end	Returns to privileged EXEC mode.

Configuring the Forwarding-Delay Time

To configure the forwarding-delay time for all MST instances, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree mst forward-time <i>seconds</i>	(Optional) Configures 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	Router(config)# end	Returns to privileged EXEC mode.

Configuring the Transmit Hold Count

To configure the transmit hold count for all MST instances, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree transmit hold-count <i>hold_count_value</i>	Configures the transmit hold count for all MST instances. For <i>hold_count_value</i> , the range is 1 to 20; the default is 6.
Step 3	Router(config)# end	Returns to privileged EXEC mode.

Configuring the Maximum-Aging Time

To configure the maximum-aging time for all MST instances, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree mst max-age <i>seconds</i>	(Optional) Configures 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	Router(config)# end	Returns to privileged EXEC mode.

Configuring the Maximum-Hop Count

To configure the maximum-hop count for all MST instances, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# spanning-tree mst max-hops <i>hop_count</i>	(Optional) Specifies 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	Router(config)# end	Returns to privileged EXEC mode.

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 31-14](#).

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 MST, you can override the default setting of the link type and enable rapid transitions to the forwarding state.

To override the default link-type setting, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface { fastethernet gigabitethernet tengigabitethernet } <i>slot/port</i> { port-channel <i>number</i> }	(Optional) Specifies an interface to configure, and enters interface configuration mode.
Step 3	Router(config)# spanning-tree link-type point-to-point	Specifies that the link type of a port is point-to-point.
Step 4	Router(config)# end	Returns to privileged EXEC mode.

Designating the Neighbor Type

A topology could contain both prestandard and 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.

To override the default link-type setting, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface { {fastethernet gigabitethernet tengigabitethernet} <i>slot/port</i> {port-channel number} }	(Optional) Specifies an interface to configure, and enters interface configuration mode.
Step 3	Router(config)# spanning-tree mst pre-standard	Specifies that the port can send only prestandard BPDUs.
Step 4	Router(config)# end	Returns to privileged EXEC mode.

Restarting the Protocol Migration Process

A switch running MST supports a built-in protocol migration feature that enables it to interoperate with 802.1D switches. If this switch receives an 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MST switch also can detect that a port is at the boundary of a region when it receives an 802.1D 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 MST mode if it no longer receives 802.1D BPDUs because it cannot detect whether the 802.1D switch has been removed from the link unless the 802.1D 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 that are described in [Table 31-5](#).

Table 31-5 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.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

32

Optional STP Features

- [PortFast](#), page 32-1
- [Bridge Assurance](#), page 32-4
- [BPDU Guard](#), page 32-7
- [PortFast Edge BPDU Filtering](#), page 32-9
- [UplinkFast](#), page 32-11
- [BackboneFast](#), page 32-13
- [EtherChannel Guard](#), page 32-16
- [Root Guard](#), page 32-17
- [Loop Guard](#), page 32-17
- [PVST Simulation](#), page 32-20
- [Verifying the Optional STP Features](#), page 32-21

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- For information on configuring the Spanning Tree Protocol (STP), see [Chapter 31, “Spanning Tree Protocols.”](#)

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

PortFast

- [Information about PortFast](#), page 32-2
- [Enabling PortFast](#), page 32-2

Information about PortFast

STP PortFast causes a Layer 2 LAN port configured as an access port to enter the forwarding state immediately, bypassing the listening and learning states. You can use PortFast on Layer 2 access ports connected to a single workstation or server to allow those devices to connect to the network immediately, instead of waiting for STP to converge. Interfaces connected to a single workstation or server should not receive bridge protocol data units (BPDUs). When configured for PortFast, a port is still running the spanning tree protocol. A PortFast enabled port can immediately transition to the blocking state if necessary (this could happen on receipt of a superior BPDU). PortFast can be enabled on trunk ports. PortFast can have an operational value that is different from the configured value.

You can specifically configure a port as either an edge port, a network port, or a normal port. An edge port, which is connected to a Layer 2 host, can be either an access port or a trunk port. A network port is connected only to a Layer 2 switch or bridge.

Enabling PortFast

- [Configuring the PortFast Default State, page 32-2](#)
- [Enabling PortFast on a Layer 2 Port, page 32-3](#)



Tip

Configure STP BPDU Guard along with STP PortFast to shut down STP PortFast-enabled ports if they receive a BPDU.

Configuring the PortFast Default State

To configure the default PortFast state, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree portfast [edge network normal] default	Configures the default state for all switch access ports to be edge, network, or normal. Bridge Assurance will be enabled on all network access ports by default.
Step 2	Router(config)# end	Exits configuration mode.



Note

- The default spanning tree port type is normal, meaning only that its topology is not specified.
- If you configure a port connected to a Layer 2 switch or bridge as an edge port, you might create a bridging loop.
- If you mistakenly configure a port that is connected to a Layer 2 host as a spanning tree network port, the port will automatically move into the blocking state.

This example shows how to configure the default switch access port state to be edge:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# spanning-tree portfast edge default
```


Enabling PortFast on a Layer 2 Port

- [Enabling PortFast on a Layer 2 Access Port, page 32-3](#)
- [Enabling PortFast on a Layer 2 Network Port, page 32-4](#)

Enabling PortFast on a Layer 2 Access Port



Caution

Enter the **spanning-tree portfast edge [trunk]** command only on ports that are connected to end host devices that terminate VLANs and from which the port should never receive STP BPDUs, such as:

- Workstations.
- Servers.
- Ports on routers that are not configured to support bridging.

To enable PortFast on a Layer 2 access port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port} {port-channel port_channel_number}	Selects a port to configure.
Step 2	Router(config-if)# spanning-tree portfast edge [trunk]	Enables edge behavior on a Layer 2 access port connected to a single workstation or server. Enter the trunk keyword if the link is a trunk.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to enable and verify PortFast on an interface:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/8
Router(config-if)# spanning-tree portfast edge
Router(config-if)# end
Router#
Router# show running-config interface gigabitethernet 5/8
Building configuration...

Current configuration:
!
interface GigabitEthernet5/8
 no ip address
 switchport
 switchport access vlan 200
 switchport mode access
 spanning-tree portfast edge
end
Router#
```

Enabling PortFast on a Layer 2 Network Port

To enable PortFast on a Layer 2 network port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port} {port-channel port_channel_number}	Selects a port to configure.
Step 2	Router(config-if)# spanning-tree portfast network	Configures the port as a network port. Bridge Assurance, if enabled globally, will be enabled on the port.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to enable and verify PortFast on an interface:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/8
Router(config-if)# spanning-tree portfast edge
Router(config-if)# end
Router#
Router# show running-config interface gigabitethernet 5/8
Building configuration...

Current configuration:
!
interface GigabitEthernet5/8
 no ip address
 switchport
 switchport access vlan 200
 switchport mode access
 spanning-tree portfast edge
end
Router#
```

Bridge Assurance

- [Information about Bridge Assurance, page 32-4](#)
- [Enabling Bridge Assurance, page 32-7](#)

Information about Bridge Assurance

You can use Bridge Assurance to protect against certain problems that can cause bridging loops in the network. Specifically, you use Bridge Assurance to protect against a unidirectional link failure or other software failure and a device that continues to forward data traffic when it is no longer running the spanning tree algorithm.

Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance is enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled. If the device on one side of the link has Bridge Assurance enabled and the device on the other side either does not support Bridge Assurance or does not have this feature enabled, the connecting port is blocked.

With Bridge Assurance enabled, BPDUs are sent out on all operational network ports, including alternate and backup ports, for each hello time period. If the port does not receive a BPDU for a specified period, the port moves into an inconsistent state (blocking), and is not used in the root port calculation. Once that port receives a BPDU, it resumes the normal spanning tree transitions.

Figure 32-1 shows a normal STP topology, and Figure 32-2 demonstrates a potential network problem when the device fails and you are not running Bridge Assurance.

Figure 32-1 Network with Normal STP Topology




Figure 32-2 Network Problem without Running Bridge Assurance






Figure 32-3 shows the network with Bridge Assurance enabled, and the STP topology progressing normally with bidirectional BPDUs issuing from every STP network port. Figure 32-4 shows how the potential network problem shown in Figure 32-2 does not happen when you have Bridge Assurance enabled on your network.

Figure 32-3 Network STP Topology Running Bridge Assurance

186526

Figure 32-4 Network Problem Averted with Bridge Assurance Enabled

186411

When using Bridge Assurance, follow these guidelines:

- Bridge Assurance runs only on point-to-point spanning tree network ports. You must configure each side of the link for this feature.
- We recommend that you enable Bridge Assurance throughout your network.

Enabling Bridge Assurance

By default, Bridge Assurance is enabled on all network ports on the switch. Disabling Bridge Assurance causes all configured network ports to behave as normal spanning tree ports. To enable or disable Bridge Assurance globally, perform this task:

Command	Purpose
Router(config)# spanning-tree bridge assurance	Enables Bridge Assurance on all network ports on the switch.

This example shows how to enable PortFast Bridge Assurance on all network ports on the switch, and how to configure a network port:

```
Router(config)# spanning-tree bridge assurance
Router(config)# interface gigabitethernet 5/8
Router(config-if)# spanning-tree portfast network
Router(config-if)# exit
```

BPDU Guard

- [Information about BPDU Guard, page 32-7](#)
- [Enabling BPDU Guard, page 32-7](#)

Information about BPDU Guard

When enabled on a port, BPDU Guard shuts down a port that receives a BPDU. When configured globally, BPDU Guard is only effective on ports in the operational PortFast (edge) state. In a valid configuration, PortFast Layer 2 LAN interfaces (edge ports) do not receive BPDUs. Reception of a BPDU by a PortFast Layer 2 LAN interface signals an invalid configuration, such as connection of an unauthorized device. BPDU Guard provides a secure response to invalid configurations, because the administrator must manually put the Layer 2 LAN interface back in service. BPDU Guard can be configured at the interface level. When configured at the interface level, BPDU Guard shuts the port down as soon as the port receives a BPDU, regardless of the PortFast configuration.

**Note**

When enabled globally, BPDU Guard applies to all interfaces that are in an operational PortFast (edge) state.

Enabling BPDU Guard

- [Enabling BPDU Guard Globally, page 32-8](#)
- [Enabling BPDU Guard on a Port, page 32-8](#)

Enabling BPDU Guard Globally

To enable BPDU Guard globally, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree portfast edge bpduguard default	Enables BPDU Guard globally by default on all edge ports of the switch.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable BPDU Guard:

```
Router# configure terminal
Router(config)# spanning-tree portfast edge bpduguard default
Router(config)# end
```

This example shows how to verify the configuration:

```
Router# show spanning-tree summary totals
Root bridge for: Bridge VLAN0025
EtherChannel misconfiguration guard is enabled
Extended system ID           is enabled
PortFast Edge BPDU Guard Default is enabled
Portfast Edge BPDU Filter Default is disabled
Portfast Default              is edge
Bridge Assurance              is enabled
Loopguard                    is disabled
UplinkFast                   is disabled
BackboneFast                  is disabled
Pathcost method used is long

Name                           Blocking Listening Learning Forwarding STP Active
-----
2 vlans                        0          0          0          3          3
```

Enabling BPDU Guard on a Port

To enable BPDU Guard on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port} {port-channel port_channel_number}	Selects a port to configure.
Step 2	Router(config-if)# spanning-tree bpduguard enable	Enables BPDU Guard on the port.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to enable BPDU Guard:

```
Router# configure terminal
Router(config)# spanning-tree portfast edge bpduguard default
Router(config)# end
```

PortFast Edge BPDU Filtering

- [Information about PortFast Edge BPDU Filtering, page 32-9](#)
- [Enabling PortFast Edge BPDU Filtering, page 32-10](#)

Information about PortFast Edge BPDU Filtering

PortFast edge BPDU filtering allows the administrator to prevent the system from sending or even receiving BPDUs on specified ports.

When configured globally, PortFast edge BPDU filtering applies to all operational PortFast (edge) ports. Ports in an operational PortFast state are supposed to be connected to hosts, which typically drop BPDUs. If an operational PortFast port receives a BPDU, it immediately loses its operational PortFast status and becomes a normal port. In that case, PortFast edge BPDU filtering is disabled on this port and STP resumes sending BPDUs on this port.

PortFast edge BPDU filtering can also be configured on a per-port basis. When PortFast edge BPDU filtering is explicitly configured on a port, it does not send any BPDUs and drops all BPDUs it receives.



Caution

Explicitly configuring PortFast edge BPDU filtering on a port that is not connected to a host can result in bridging loops, as the port will ignore any BPDU it receives and will go to a forwarding state.

When you enable PortFast edge BPDU filtering globally and set the port configuration as the default for PortFast edge BPDU filtering (see the [“Enabling PortFast Edge BPDU Filtering”](#) section on page 32-10), then PortFast enables or disables PortFast edge BPDU filtering.

If the port configuration is not set to default, then the PortFast configuration will not affect PortFast edge BPDU filtering. [Table 32-1](#) lists all the possible PortFast edge BPDU filtering combinations. PortFast edge BPDU filtering allows access ports to move directly to the forwarding state as soon as the end hosts are connected.

Table 32-1 PortFast Edge BPDU Filtering Port Configurations

Per-Port Configuration	Global Configuration	PortFast State	PortFast BPDU Filtering State
Default	Enable	Enable	Enable Note The port transmits at least 10 BPDUs. If this port receives any BPDUs, then PortFast and PortFast edge BPDU filtering are disabled.
Default	Enable	Disable	Disable
Default	Disable	Not applicable	Disable
Disable	Not applicable	Not applicable	Disable
Enable	Not applicable	Not applicable	Enable

Enabling PortFast Edge BPDU Filtering

- [Enabling PortFast Edge BPDU Filtering Globally, page 32-10](#)
- [Enabling PortFast Edge BPDU Filtering on a Nontrunking Port, page 32-11](#)

Enabling PortFast Edge BPDU Filtering Globally

To enable PortFast edge BPDU filtering globally, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree portfast edge bpdufilter default	Enables BPDU filtering globally by default on all edge ports of the switch. Use the no prefix to disable BPDU filtering by default on all edge ports of the switch.
Step 2	Router(config)# end	Exits configuration mode.

BPDU filtering is set to default on each edge port. This example shows how to enable PortFast edge BPDU filtering on the port and verify the configuration in **PVST+** mode:

```
Router(config)# spanning-tree portfast edge bpdufilter default
Router(config)# exit
```

```
Router# show spanning-tree summary totals
Root bridge for: Bridge VLAN0025
EtherChannel misconfiguration guard is enabled
Extended system ID is enabled
PortFast Edge BPDU Guard Default is disabled
Portfast Edge BPDU Filter Default is enabled
Portfast Default is edge
Bridge Assurance is enabled
Loopguard is disabled
UplinkFast is disabled
BackboneFast is disabled
Pathcost method used is long
```

```

Name                  Blocking Listening Learning Forwarding STP Active
-----
2 vlans                0          0          0          3          3
Router#
```


Enabling PortFast Edge BPDU Filtering on a Nontrunking Port

To enable or disable PortFast edge BPDU filtering on a nontrunking port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port}	Selects the interface to configure.
Step 2	Router(config-if)# spanning-tree bpdufilter [enable disable]	Enables or disables BPDU filtering on the port.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to enable PortFast edge BPDU filtering on a nontrunking port:

```
Router(config)# interface gigabitethernet 4/4
Router(config-if)# spanning-tree bpdufilter enable
Router(config-if)# ^Z

Router# show spanning-tree interface gigabitethernet 4/4

Vlan                Role Sts Cost          Prio.Nbr Status
-----
VLAN0010            Desg FWD 1000          160.196  Edge P2p

Router# show spanning-tree interface gigabitethernet 4/4 detail
Port 196 (GigabitEthernet4/4) of VLAN0010 is forwarding
  Port path cost 1000, Port priority 160, Port Identifier 160.196.
  Designated root has priority 32768, address 00d0.00b8.140a
  Designated bridge has priority 32768, address 00d0.00b8.140a
  Designated port id is 160.196, designated path cost 0
  Timers:message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state:1
  The port is in the portfast mode by portfast trunk configuration
  Link type is point-to-point by default
  Bpdu filter is enabled
  BPDU:sent 0, received 0
Router#
```

UplinkFast

- [Information about UplinkFast, page 32-11](#)
- [Enabling UplinkFast, page 32-12](#)

Information about UplinkFast

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 LAN 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.




Note

UplinkFast is most useful in wiring-closet switches. This feature may not be useful for other types of applications.


Figure 32-5 shows an example topology with no link failures. Switch A, the root bridge, is connected directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 LAN interface on Switch C that is connected directly to Switch B is in the blocking state.

Figure 32-5 UplinkFast Example Before Direct Link Failure



If Switch C detects a link failure on the currently active link L2 on the root port (a *direct* link failure), UplinkFast unblocks the blocked port on Switch C and transitions it to the forwarding state without going through the listening and learning states, as shown in Figure 32-6. This switchover takes approximately one to five seconds.

Figure 32-6 UplinkFast Example After Direct Link Failure



Enabling UplinkFast

UplinkFast increases the bridge priority to 49152 and adds 3000 to the STP port cost of all Layer 2 LAN ports on the switch, decreasing the probability that the switch will become the root bridge. UplinkFast cannot be enabled on VLANs that have been configured for bridge priority. To enable UplinkFast on a VLAN with bridge priority configured, restore the bridge priority on the VLAN to the default value by entering a **no spanning-tree vlan *vlan_ID* priority** command in global configuration mode.



Note

When you enable UplinkFast, it affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

To enable UplinkFast, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree uplinkfast	Enables UplinkFast.
	Router(config)# spanning-tree uplinkfast [max-update-rate <i>max_update_rate</i>]	Enables UplinkFast with an update rate in seconds.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable UplinkFast:

```
Router# configure terminal
Router(config)# spanning-tree uplinkfast
Router(config)# exit
```

This example shows how to enable UplinkFast with an update rate of 400 packets per second:

```
Router# configure terminal
Router(config)# spanning-tree uplinkfast
Router(config)# spanning-tree uplinkfast max-update-rate 400
Router(config)# exit
```

This example shows how to verify that UplinkFast is enabled:

```
Router# show spanning-tree uplinkfast
UplinkFast is enabled
```

BackboneFast

- [Information about BackboneFast, page 32-13](#)
- [Enabling BackboneFast, page 32-15](#)

Information about BackboneFast

BackboneFast is initiated when a root port or blocked port on a network device receives inferior BPDUs from its designated bridge. An inferior BPDU identifies one network device as both the root bridge and the designated bridge. When a network device receives an inferior BPDU, it indicates that a link to which the network device is not directly connected (an *indirect* link) has failed (that is, the designated bridge has lost its connection to the root bridge). Under normal STP rules, the network device ignores inferior BPDUs for the configured maximum aging time, as specified by the STP **max-age** command.


The network device tries to determine if it has an alternate path to the root bridge. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the network device become alternate paths to the root bridge. (Self-looped ports are not considered alternate paths to the root bridge.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root bridge. If the inferior BPDU arrives on the root port and there are no blocked ports, the network device assumes that it has lost connectivity to the root bridge, causes the maximum aging time on the root to expire, and becomes the root bridge according to normal STP rules.

If the network device has alternate paths to the root bridge, it uses these alternate paths to transmit a new kind of Protocol Data Unit (PDU) called the Root Link Query PDU. The network device sends the Root Link Query PDU out all alternate paths to the root bridge. If the network device determines that it still has an alternate path to the root, it causes the maximum aging time to expire on the ports on which it received the inferior BPDU. If all the alternate paths to the root bridge indicate that the network device

has lost connectivity to the root bridge, the network device causes the maximum aging times on the ports on which it received an inferior BPDU to expire. If one or more alternate paths can still connect to the root bridge, the network device makes all ports on which it received an inferior BPDU its designated ports and moves them out of the blocking state (if they were in the blocking state), through the listening and learning states, and into the forwarding state.


Figure 32-7 shows an example topology with no link failures. Switch A, the root bridge, connects directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 LAN interface on Switch C that connects directly to Switch B is in the blocking state.

Figure 32-7 BackboneFast Example Before Indirect Link Failure




If link L1 fails, 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 bridge over L1, it detects the failure and elects itself the root and begins sending BPDUs to Switch C indicating itself as the root. When Switch C receives the inferior BPDUs from Switch B, Switch C infers that an indirect failure has occurred. At that point, BackboneFast allows the blocked port on Switch C to move immediately to the listening state without waiting for the maximum aging time for the port to expire. BackboneFast then transitions the Layer 2 LAN interface on Switch C to the forwarding state, providing a path from Switch B to Switch A. This switchover takes approximately 30 seconds, twice the Forward Delay time if the default Forward Delay time of 15 seconds is set. Figure 32-8 shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 32-8 BackboneFast Example After Indirect Link Failure



If a new network device is introduced into a shared-medium topology as shown in Figure 32-9, BackboneFast is not activated because the inferior BPDUs did not come from the recognized designated bridge (Switch B). The new network device begins sending inferior BPDUs that indicate that it is the root bridge. However, the other network devices ignore these inferior BPDUs and the new network device learns that Switch B is the designated bridge to Switch A, the root bridge.

Figure 32-9 Adding a Network Device in a Shared-Medium Topology



11245

Enabling BackboneFast



Note

BackboneFast operates correctly only when enabled on all network devices in the network. BackboneFast is not supported on Token Ring VLANs. This feature is supported for use with third-party network devices.

To enable BackboneFast, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree backbonefast	Enables BackboneFast.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable BackboneFast:

```
Router# configure terminal
Router(config)# spanning-tree backbonefast
Router(config)# end
```

This example shows how to verify that BackboneFast is enabled:

```
Router# show spanning-tree backbonefast
BackboneFast is enabled

BackboneFast statistics
```

```

-----
Number of transition via backboneFast (all VLANs) : 0
Number of inferior BPDUs received (all VLANs)      : 0
Number of RLQ request PDUs received (all VLANs)    : 0
Number of RLQ response PDUs received (all VLANs)   : 0
Number of RLQ request PDUs sent (all VLANs)        : 0
Number of RLQ response PDUs sent (all VLANs)       : 0

```

EtherChannel Guard

- [Information about EtherChannel Guard, page 32-16](#)
- [Enabling EtherChannel Guard, page 32-16](#)

Information about EtherChannel Guard

EtherChannel guard detects a misconfigured EtherChannel when interfaces on the switch are configured as an EtherChannel while interfaces on the other device are not or when not all the interfaces on the other device are in the same EtherChannel.

In response to misconfiguration detected on the other device, EtherChannel guard puts interfaces on the switch into the errdisabled state.

Enabling EtherChannel Guard

To enable EtherChannel guard, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree etherchannel guard misconfig	Enables EtherChannel guard.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable EtherChannel guard:

```

Router# configure terminal
Router(config)# spanning-tree etherchannel guard misconfig
Router(config)# end

```

This example shows how to verify the configuration:

```

Router# show spanning-tree summary | include EtherChannel
EtherChannel misconfiguration guard is enabled

```

To display the interfaces that are in the errdisable state, enter the **show interface status err-disable** command.

After the misconfiguration has been cleared, interfaces in the errdisable state might automatically recover. To manually return a port to service, enter a **shutdown** and then a **no shutdown** command for the interface.

Root Guard

- [Information about Root Guard, page 32-17](#)
- [Enabling Root Guard, page 32-17](#)

Information about Root Guard

The STP root guard feature prevents a port from becoming root port or blocked port. If a port configured for root guard receives a superior BPDU, the port immediately goes to the root-inconsistent (blocked) state.

Enabling Root Guard

To enable root guard, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port} {port-channel port_channel_number}	Selects a port to configure.
Step 2	Router(config-if)# spanning-tree guard root	Enables root guard.
Step 3	Router(config-if)# end	Exits configuration mode.

To display ports that are in the root-inconsistent state, enter the **show spanning-tree inconsistentports** command.

Loop Guard

- [Information about Loop Guard, page 32-17](#)
- [Enabling Loop Guard, page 32-19](#)


Information about Loop Guard

Loop guard helps prevent bridging loops that could occur because of a unidirectional link failure on a point-to-point link. When enabled globally, the loop guard applies to all point-to-point ports on the system. Loop guard detects root ports and blocked ports and ensures that they keep receiving BPDUs from their designated port on the segment. If a loop guard enabled root or blocked port stop a receiving BPDUs from its designated port, it transitions to the loop-inconsistent blocking state, assuming there is a physical link error on this port. The port recovers from this loop-inconsistent state as soon as it receives a BPDU.

You can enable loop guard on a per-port basis. When you enable loop guard, it is automatically applied to all of the active instances or VLANs to which that port belongs. When you disable loop guard, it is disabled for the specified ports. Disabling loop guard moves all loop-inconsistent ports to the listening state.

If you enable loop guard on a channel and the first link becomes unidirectional, loop guard blocks the entire channel until the affected port is removed from the channel. [Figure 32-10](#) shows loop guard in a triangle switch configuration.

Figure 32-10 Triangle Switch Configuration with Loop Guard



[Figure 32-10](#) illustrates the following configuration:

- Switches A and B are distribution switches.
- Switch C is an access switch.
- Loop guard is enabled on ports 3/1 and 3/2 on Switches A, B, and C.

Enabling loop guard on a root switch has no effect but provides protection when a root switch becomes a nonroot switch.

When using loop guard, follow these guidelines:

- You cannot enable loop guard on PortFast-enabled ports.
- You cannot enable loop guard if root guard is enabled.

Loop guard interacts with other features as follows:

- Loop guard does not affect the functionality of UplinkFast or BackboneFast.
- Enabling loop guard on ports that are not connected to a point-to-point link will not work.
- Root guard forces a port to be always designated as the root port. Loop guard is effective only if the port is a root port or an alternate port. You cannot enable loop guard and root guard on a port at the same time.
- Loop guard uses the ports known to spanning tree. Loop guard can take advantage of logical ports provided by the Port Aggregation Protocol (PAgP). However, to form a channel, all the physical ports grouped in the channel must have compatible configurations. PAgP enforces uniform configurations of root guard or loop guard on all the physical ports to form a channel.

These caveats apply to loop guard:

- Spanning tree always chooses the first operational port in the channel to send the BPDUs. If that link becomes unidirectional, loop guard blocks the channel, even if other links in the channel are functioning properly.

- If a set of ports that are already blocked by loop guard are grouped together to form a channel, spanning tree loses all the state information for those ports and the new channel port may obtain the forwarding state with a designated role.
- If a channel is blocked by loop guard and the channel breaks, spanning tree loses all the state information. The individual physical ports may obtain the forwarding state with the designated role, even if one or more of the links that formed the channel are unidirectional.

**Note**

You can enable UniDirectional Link Detection (UDLD) to help isolate the link failure. A loop may occur until UDLD detects the failure, but loop guard will not be able to detect it.

- Loop guard has no effect on a disabled spanning tree instance or a VLAN.

Enabling Loop Guard

To enable loop guard globally on the switch, perform this task:

	Command	Purpose
Step 1	Router(config)# spanning-tree loopguard default	Enables loop guard globally on the switch.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable loop guard globally:

```
Router# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router(config)# spanning-tree loopguard default
```

```
Router(config)# exit
```

```
Router# show spanning-tree interface gigabitethernet 4/4 detail
```

```
Port 196 (GigabitEthernet4/4) of VLAN0010 is forwarding
Port path cost 1000, Port priority 160, Port Identifier 160.196.
Designated root has priority 32768, address 00d0.00b8.140a
Designated bridge has priority 32768, address 00d0.00b8.140a
Designated port id is 160.196, designated path cost 0
Timers:message age 0, forward delay 0, hold 0
Number of transitions to forwarding state:1
The port is in the portfast mode by portfast trunk configuration
Link type is point-to-point by default
Bpdu filter is enabled
Loop guard is enabled by default on the port
BPDU:sent 0, received 0
```

To enable loop guard on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port} { port-channel port_channel_number}	Selects a port to configure.
Step 2	Router(config-if)# spanning-tree guard loop	Configures loop guard.
Step 3	Router(config)# end	Exits configuration mode.

This example shows how to enable loop guard:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 4/4
Router(config-if)# spanning-tree guard loop
Router(config-if)# exit
```

This example shows how to verify the configuration:

```
Router# show spanning-tree interface gigabitethernet 4/4 detail
Port 196 (GigabitEthernet4/4) of VLAN0010 is forwarding
  Port path cost 1000, Port priority 160, Port Identifier 160.196.
  Designated root has priority 32768, address 00d0.00b8.140a
  Designated bridge has priority 32768, address 00d0.00b8.140a
  Designated port id is 160.196, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  The port is in the portfast mode by portfast trunk configuration
  Link type is point-to-point by default
  Bpdu filter is enabled
  Loop guard is enabled on the port
  BPDU: sent 0, received 0
```

PVST Simulation

- [Information about PVST Simulation, page 32-20](#)
- [Configuring PVST Simulation, page 32-21](#)

Information about PVST Simulation

MST interoperates with Rapid PVST+ with no need for user configuration. The PVST simulation feature enables this seamless interoperability.

**Note**

PVST simulation is enabled by default when you enable MST. That is, by default, all interfaces on the device interoperate between MST and Rapid PVST+.

You may want to control the connection between MST and Rapid PVST+ to protect against accidentally connecting an MST-enabled port to a port enabled to run Rapid PVST+. Because Rapid PVST+ is the default STP mode, you may encounter many Rapid PVST+ connections.

Disabling Rapid PVST+ simulation, which can be done per port or globally for the entire device, moves the MST-enabled port to the PVST peer inconsistent (blocking) state once it detects it is connected to a Rapid PVST+-enabled port. This port remains in the inconsistent state until the port stops receiving Shared Spanning Tree Protocol (SSTP) BPDUs, and then the port resumes the normal STP transition process.

The root bridge for all STP instances must all be in either the MST region or the Rapid PVST+ side. If the root bridge for all STP instances are not on one side or the other, the software moves the port into a PVST simulation-inconsistent state.

**Note**

We recommend that you put the root bridge for all STP instances in the MST region.

Configuring PVST Simulation



Note

PVST simulation is enabled by default so that all interfaces on the device interoperate between MST and Rapid PVST+.

To prevent an accidental connection to a device that does not run MST as the default STP mode, you can disable PVST simulation. If you disable PVST simulation, the MST-enabled port moves to the blocking state once it detects it is connected to a Rapid PVST+-enabled port. This port remains in the inconsistent state until the port stops receiving BPDUs, and then the port resumes the normal STP transition process.

To enable or disable PVST simulation globally, enter the command using the **global** keyword, as shown in the following task:

Command	Purpose
Router(config)# spanning-tree mst simulate pvst global	Enables all ports to automatically interoperate with a connected device that is running in Rapid PVST+ mode. The default is enabled; all interfaces will operate seamlessly between Rapid PVST+ and MST.

To override the global PVST simulation setting for a port, enter the command in the interface command mode, as shown in the following task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port}	Selects a port to configure.
Step 2	Router(config-if)# spanning-tree mst simulate pvst	Enables this interface to automatically interoperate with a connected device that is running in Rapid PVST+ mode.

This example shows how to prevent the switch from automatically interoperating with a connecting device that is running Rapid PVST+:

```
Router(config)# no spanning-tree mst simulate pvst global
```

This example shows how to prevent a port from automatically interoperating with a connecting device that is running Rapid PVST+:

```
Router(config)# interface gi3/13
Router(config-if)# spanning-tree mst simulate pvst disable
```

Verifying the Optional STP Features

- [Using the show spanning-tree Commands, page 32-22](#)
- [Examples of the show spanning-tree Commands, page 32-22](#)

Using the show spanning-tree Commands

You can view spanning tree status and configuration information, both global and port-level, using the **show spanning-tree** commands described in this section. To view spanning tree status and configuration information, enter one of the following commands:

Command	Purpose
Router# show spanning-tree	Displays information about the spanning tree, including protocol type and port types.
Router# show spanning-tree summary	Displays a summary of the spanning tree feature settings and the spanning tree states of the VLANs.
Router# show spanning-tree summary totals	Displays a summary of the spanning tree feature settings and totals of the VLAN states.
Router# show spanning-tree interface {type slot/port} detail	Displays the spanning tree status details of an interface.
Router# show spanning-tree interface {type slot/port} portfast edge	Displays the spanning tree portfast edge interface operational state for all the instances.

Examples of the show spanning-tree Commands

This example displays the spanning-tree status with Bridge Assurance enabled but in the inconsistent state:

```
Router# show spanning-tree
VLAN0010
  Spanning tree enabled protocol rstp
  Root ID    Priority    32778
            Address     0002.172c.f400
            This bridge is the root
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778 (priority 32768 sys-id-ext 10)
            Address     0002.172c.f400
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
            Aging Time  300

Interface          Role Sts Cost          Prio.Nbr Type
-----
Gi3/14             Desg BKN*4      128.270 Network, P2p *BA_Inc
Router#
```

The following inconsistency messages can be appended to the Type field:

- ***BA_Inc**—Indicates that Bridge Assurance is in the inconsistent state.
- ***PVST_Peer_Inc**—Indicates that the port is in a peer type Inconsistent state.
- **Dispute**—Indicates that a dispute condition is detected.

This example shows the spanning-tree configuration summary:

```
Router# show spanning-tree summary
```

```
Switch is in rapid-pvst mode
Root bridge for: Bridge VLAN0025
EtherChannel misconfiguration guard is enabled
Extended system ID is enabled
PortFast Edge BPDU Guard Default is disabled
Portfast Edge BPDU Filter Default is disabled
Portfast Default is edge
Bridge Assurance is enabled
Loopguard is disabled
UplinkFast is disabled
BackboneFast is disabled
Pathcost method used is short
PVST Simulation Default is enabled
```

Name	Blocking	Listening	Learning	Forwarding	STP Active
VLAN0025	0	0	0	1	1
VLAN0030	0	0	0	2	2
2 vlans	0	0	0	3	3

Possible states for the Bridge Assurance field are as follows:

- is enabled
- is disabled
- is enabled but not active in the PVST mode

This example shows the spanning tree summary when PVST simulation is disabled in any STP mode:

```
Router# show spanning-tree summary
```

```
Switch is in mst mode (IEEE Standard)
Root bridge for: MST0
EtherChannel misconfig guard is enabled
Extended system ID is enabled
Portfast Default is disabled
PortFast BPDU Guard Default is disabled
Portfast BPDU Filter Default is disabled
Loopguard Default is disabled
UplinkFast is disabled
BackboneFast is disabled
Pathcost method used is long
PVST Simulation Default is disabled
```

Name	Blocking	Listening	Learning	Forwarding	STP Active
MST0	2	0	0	0	2
1 mst	2	0	0	0	2

Possible states for the PVST Simulation Default field are as follows:

- is enabled
- is disabled
- is enabled but not active in rapid-PVST mode

This example shows the spanning tree summary totals:

```
Router# show spanning-tree summary totals
Root bridge for: Bridge VLAN0025
EtherChannel misconfiguration guard is enabled
Extended system ID is enabled
PortFast Edge BPDU Guard Default is enabled
Portfast Edge BPDU Filter Default is disabled
Portfast Default is edge
Bridge Assurance is enabled
Loopguard is disabled
UplinkFast is disabled
BackboneFast is disabled
Pathcost method used is long
```

Name	Blocking	Listening	Learning	Forwarding	STP Active
2 vlans	0	0	0	3	3

Router#

This example shows the spanning-tree configuration details of an edge port:

```
Router# show spanning-tree interface gi3/13 detail
Port 269 (GigabitEthernet3/13) of VLAN0002 is forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.269.
  Designated root has priority 32770, address 0002.172c.f400
  Designated bridge has priority 32770, address 0002.172c.f400
  Designated port id is 128.269, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  Loop guard is enabled by default on the port
  The port is in the portfast edge mode by default
  BPDU: sent 2183, received 0
```

This example shows the spanning-tree configuration details of a trunk port:

```
Router(config-if)# spanning-tree portfast edge trunk
%Warning:portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION
```

```
Router(config-if)# exit
```

```
Router# show spanning-tree interface gi3/13 detail
Port 269 (GigabitEthernet3/13) of VLAN0002 is forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.269.
  Designated root has priority 32770, address 0002.172c.f400
  Designated bridge has priority 32770, address 0002.172c.f400
  Designated port id is 128.269, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  Loop guard is enabled by default on the port
  The port is in the portfast edge trunk mode
  BPDU: sent 2183, received 0
```

This example shows the spanning-tree configuration details of an edge port when a dispute condition has been detected:

```
Router# show spanning-tree interface gi3/13 detail
Port 269 (GigabitEthernet3/13) of VLAN0002 is designated blocking (dispute)
  Port path cost 4, Port priority 128, Port Identifier 128.297.
  Designated root has priority 32769, address 0013.5f20.01c0
```

```
Designated bridge has priority 32769, address 0013.5f20.01c0
Designated port id is 128.297, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDU: sent 132, received 1
```

This example shows the spanning tree portfast edge interface operational state for all the instances:

```
Router# show spanning-tree interface gi3/1 portfast edge
MST0                disabled
MST1                disabled
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

33

IP Unicast Layer 3 Switching

- [Prerequisites for Hardware Layer 3 Switching, page 33-1](#)
- [Restrictions for Hardware Layer 3 Switching, page 33-2](#)
- [Information About Layer 3 Switching, page 33-2](#)
- [Default Settings for Hardware Layer 3 Switching, page 33-4](#)
- [How to Configure Hardware Layer 3 Switching, page 33-4](#)
- [Displaying Hardware Layer 3 Switching Statistics, page 33-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - For information about IP multicast Layer 3 switching, see [Chapter 6, “IPv4 Multicast Layer 3 Features.”](#)
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Hardware Layer 3 Switching

None.

Restrictions for Hardware Layer 3 Switching

- Hardware Layer 3 switching supports the following ingress and egress encapsulations:
 - Ethernet V2.0 (ARPA)
 - 802.3 with 802.2 with 1 byte control (SAP1)

Information About Layer 3 Switching

- [Hardware Layer 3 Switching, page 33-2](#)
- [Layer 3-Switched Packet Rewrite, page 33-2](#)

Hardware Layer 3 Switching

Hardware Layer 3 switching allows the PFC and DFCs, instead of the RP, to forward IP unicast traffic between subnets. Hardware Layer 3 switching provides wire-speed forwarding on the PFC and DFCs, instead of in software on the RP. Hardware Layer 3 switching requires minimal support from the RP. The RP routes any traffic that cannot be hardware Layer 3 switched.

Hardware Layer 3 switching supports the routing protocols configured on the RP. Hardware Layer 3 switching does not replace the routing protocols configured on the RP.

Hardware Layer 3 switching runs equally on the PFC and DFCs to provide IP unicast Layer 3 switching locally on each module. Hardware Layer 3 switching provides the following functions:

- Hardware access control list (ACL) switching for policy-based routing (PBR)
- Hardware flow-based switching for TCP intercept and reflexive ACL forwarding decisions
- Hardware Cisco Express Forwarding (CEF) switching for all other IP unicast traffic

Hardware Layer 3 switching on the PFC supports modules that do not have a DFC. The RP forwards traffic that cannot be Layer 3 switched.

Traffic is hardware Layer 3 switched after being processed by access lists and quality of service (QoS).

Hardware Layer 3 switching makes a forwarding decision locally on the ingress-port module for each packet and sends the rewrite information for each packet to the egress port, where the rewrite occurs when the packet is transmitted from the switch.

Hardware Layer 3 switching generates flow statistics for Layer 3-switched traffic. Hardware Layer 3 flow statistics can be used for NetFlow. (See [Chapter 15, “NetFlow Hardware Support”](#).)

Layer 3-Switched Packet Rewrite

When a packet is Layer 3 switched from a source in one subnet to a destination in another subnet, the switch performs a packet rewrite at the egress port based on information learned from the RP so that the packets appear to have been routed by the RP.

Packet rewrite alters five fields:

- Layer 2 (MAC) destination address
- Layer 2 (MAC) source address
- Layer 3 IP Time to Live (TTL)

- Layer 3 checksum
- Layer 2 (MAC) checksum (also called the frame checksum or FCS)

**Note**

Packets are rewritten with the encapsulation appropriate for the next-hop subnet.

If Source A and Destination B are in different subnets and Source A sends a packet to the RP to be routed to Destination B, the switch recognizes that the packet was sent to the Layer 2 (MAC) address of the RP.

To perform Layer 3 switching, the switch rewrites the Layer 2 frame header, changing the Layer 2 destination address to the Layer 2 address of Destination B and the Layer 2 source address to the Layer 2 address of the RP. The Layer 3 addresses remain the same.

In IP unicast and IP multicast traffic, the switch decrements the Layer 3 TTL value by 1 and recomputes the Layer 3 packet checksum. The switch recomputes the Layer 2 frame checksum and forwards (or, for multicast packets, replicates as necessary) the rewritten packet to Destination B's subnet.

A received IP unicast packet is formatted (conceptually) as follows:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>RP MAC</i>	<i>Source A MAC</i>	<i>Destination B IP</i>	<i>Source A IP</i>	<i>n</i>	<i>calculation1</i>		

After the switch rewrites an IP unicast packet, it is formatted (conceptually) as follows:


Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Destination B MAC</i>	<i>RP MAC</i>	<i>Destination B IP</i>	<i>Source A IP</i>	<i>n-1</i>	<i>calculation2</i>		

Hardware Layer 3 Switching Examples

[Figure 33-1 on page 33-4](#) shows a simple network topology. In this example, Host A is on the Sales VLAN (IP subnet 171.59.1.0), Host B is on the Marketing VLAN (IP subnet 171.59.3.0), and Host C is on the Engineering VLAN (IP subnet 171.59.2.0).

When Host A initiates an HTTP file transfer to Host C, Hardware Layer 3 switching uses the information in the local forwarding information base (FIB) and adjacency table to forward packets from Host A to Host C.

Figure 33-1 Hardware Layer 3 Switching Example Topology



Default Settings for Hardware Layer 3 Switching

Feature	Default Value
Hardware Layer 3 switching enable state	Enabled (cannot be disabled)
Cisco IOS CEF enable state on RP	Enabled (cannot be disabled)
Cisco IOS dCEF enable state on RP	Enabled (cannot be disabled)

How to Configure Hardware Layer 3 Switching



Note

For information on configuring unicast routing on the RP, see [Chapter 35, “Layer 3 Interfaces.”](#)

Hardware Layer 3 switching is permanently enabled. No configuration is required.
To display information about Layer 3-switched traffic, perform this task:

Command	Purpose
Router# show interface {{type slot/port} {port-channel number}} begin L3	Displays a summary of Layer 3-switched traffic.

This example shows how to display information about hardware Layer 3-switched traffic on Gigabit Ethernet port 3/3:

```
Router# show interface gigabitethernet 3/3 | begin L3
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 12 pkt, 778 bytes mcast
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  4046399 packets input, 349370039 bytes, 0 no buffer
  Received 3795255 broadcasts, 2 runts, 0 giants, 0 throttles
<...output truncated...>
Router#
```



Note

The Layer 3 switching packet count is updated approximately every five seconds.

Cisco IOS CEF and dCEF are permanently enabled. No configuration is required to support hardware Layer 3 switching.

With a PFC (and DFCs, if present), hardware Layer 3 switching uses per-flow load balancing based on IP source and destination addresses. Per-flow load balancing avoids the packet reordering that can be necessary with per-packet load balancing. For any given flow, all PFC- and DFC-equipped switches make exactly the same load-balancing decision, which can result in nonrandom load balancing.

The Cisco IOS CEF **ip load-sharing per-packet**, **ip cef accounting per-prefix**, and **ip cef accounting non-recursive** commands on the RP apply only to traffic that is CEF-switched in software on the RP. The commands do not affect traffic that is hardware Layer 3 switched on the PFC or on DFC-equipped switching modules.

Displaying Hardware Layer 3 Switching Statistics

Hardware Layer 3 switching statistics are obtained on a per-VLAN basis.

To display hardware Layer 3 switching statistics, perform this task:

Command	Purpose
Router# show interfaces [{type slot/port} {port-channel number}]	Displays hardware Layer 3 switching statistics.

This example shows how to display hardware Layer 3 switching statistics:

```
Router# show interfaces gigabitethernet 9/5 | include Switched
L2 Switched: ucast: 8199 pkt, 1362060 bytes - mcast: 6980 pkt, 371952 bytes
L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes mcast
L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
```

To display adjacency table information, perform this task:

Command	Purpose
Router# show adjacency [{type slot/port} {port-channel number}] detail internal summary	Displays adjacency table information. The optional detail keyword displays detailed adjacency information, including Layer 2 information.

This example shows how to display adjacency statistics, which are updated approximately every 60 seconds:

```
Router# show adjacency gigabitethernet 9/5 detail
Protocol Interface Address
IP GigabitEthernet9/5 172.20.53.206(11)
504 packets, 6110 bytes
00605C865B82
000164F83FA50800
ARP 03:49:31
```


Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

34

Policy-Based Routing (PBR)

- [Prerequisites for PBR, page 34-1](#)
- [Restrictions for PBR, page 34-1](#)
- [Information About PBR, page 34-2](#)
- [Default Settings for PBR, page 34-3](#)
- [How to Configure PBR, page 34-3](#)
- [Configuration Examples for PBR, page 34-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for PBR

None.

Restrictions for PBR

The PFC and any DFCs provide the hardware support for the following:

- These IPv4 PBR commands:
 - **match ip address**
 - **match length**
 - **set ip next-hop** (2,000 instances)

- **set ip default next-hop**
- **set interface null0**
- **set default interface null0**
- **set ip vrf**
- **set ip default vrf**
- If the RP address falls within the range of a PBR ACL, traffic addressed to the RP is policy routed in hardware instead of being forwarded to the RP. To prevent policy routing of traffic addressed to the RP, configure PBR ACLs to deny traffic addressed to the RP.
- Local PBR.
- IPv4 PBR recursive next-hop with load balancing.
- IPv6 PBR is supported in software.
- IPv6 PBR recursive next-hop is not supported.
- If a PBR route map sequence does not have any set clauses configured, then all the packets matching the entry are punted to the CPU and processed in software.

Information About PBR

- [PBR Overview, page 34-2](#)
- [PBR Recursive Next Hop for IPv4 Traffic, page 34-3](#)

PBR Overview

PBR is an alternative to routing protocols and allows you to configure a policy for unicast traffic flows, which provides more control over routing than a routing protocol does and avoids the need to configure interface-level traffic classification. PBR can route unicast traffic along a different path than a routing protocol would use. PBR can provide:

- Equal access
- Protocol-sensitive routing
- Source-sensitive routing
- Routing based on interactive rather than batch traffic
- Routing based on dedicated links

PBR route maps can be configured to do the following:

- Allow or deny paths based on the identity of a particular end system, an application protocol, or the size of packets or a combination of these values.
- Classify traffic based on extended access list criteria.
- Set IP precedence bits.
- Route packets to specific paths.

PBR applies a route map to all ingress unicast traffic received on a PBR-enabled interface. PBR cannot be applied to egress traffic or to multicast traffic.

If the ingress unicast traffic does not match any route map statements, the route map applies all the configured set clauses. Routing protocols forward traffic that matches a route-map deny statement and traffic that does not match any route-map permit statements.

PBR Recursive Next Hop for IPv4 Traffic

The PBR Recursive Next Hop feature enables configuration of a recursive next-hop address in a PBR route map. The recursive next-hop address is installed in the routing table and can be a subnet that is not directly connected. If the recursive next-hop address is not available, traffic is routed using a default route.

Default Settings for PBR

None.

How to Configure PBR

- [Configuring PBR](#)
- [Configuring Local PBR](#)
- [Configuring PBR Recursive Next Hop](#)

**Note**

For information about Multi-VRF Selection Using Policy Based Routing (PBR VRF), see this document:
http://www.cisco.com/en/US/docs/ios/mpls/configuration/guide/mp_mltvrf_slct_pbr.html

Configuring PBR

To configure PBR on an interface, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# route-map <i>map-tag</i> [permit deny] [<i>sequence-number</i>]	Defines a route map to control where packets are output. This command puts the router into route-map configuration mode.
Step 2	Router(config-route-map)# match length <i>min max</i> Router(config-route-map)# match ip address { <i>access-list-number</i> <i>name</i> } [... <i>access-list-number</i> <i>name</i>]	Specifies the match criteria. Although there are many route-map matching options, here you can specify only length and/or ip address. length matches the Level 3 length of the packet. ip address matches the source or destination IP address that is permitted by one or more standard or extended access lists. If you do not specify a match command, the route map applies to <i>all</i> packets.
Step 3	Router(config-route-map)# set ip precedence [<i>number</i> <i>name</i>] Router(config-route-map)# set ip df Router(config-route-map)# set ip vrf <i>vrf_name</i> Router(config-route-map)# set ip next-hop <i>ip-address</i> [... <i>ip-address</i>] Router(config-route-map)# set ip next-hop recursive <i>ip-address</i> [... <i>ip-address</i>] Router(config-route-map)# set interface <i>interface-type interface-number</i> [... <i>type number</i>] Router(config-route-map)# set ip default next-hop <i>ip-address</i> [... <i>ip-address</i>] Router(config-route-map)# set default interface <i>interface-type interface-number</i> [... <i>type ...number</i>]	Specifies the action(s) to take on the packets that match the criteria. You can specify any or all of the following: precedence: Sets precedence value in the IP header. You can specify either the precedence number or name. df: Sets the ‘Don’t Fragment’ (DF) bit in the ip header. vrf: Sets the VPN Routing and Forwarding (VRF) instance. next-hop: Sets next hop to which to route the packet. next-hop recursive: Sets next hop to which to route the packet if the hop is to a router which is not adjacent. interface: Sets output interface for the packet. default next-hop: Sets next hop to which to route the packet if there is no explicit route for this destination. default interface: Sets output interface for the packet if there is no explicit route for this destination.

	Command	Purpose
Step 4	Router(config-route-map)# interface <i>interface-type</i> <i>interface-number</i>	Specifies the interface, and puts the router into interface configuration mode.
Step 5	Router(config-if)# ip policy route-map <i>map-tag</i>	Identifies the route map to use for PBR. One interface can have only one route map tag; but you can have several route map entries, each with its own sequence number. Entries are evaluated in order of their sequence numbers until the first match occurs. If no match occurs, packets are routed as usual.

The **set** commands can be used in conjunction with each other. They are evaluated in the order shown in Step 3 in the previous task table. A usable next hop implies an interface. Once the local router finds a next hop and a usable interface, it routes the packet.

Configuring Local PBR

To configure PBR for all traffic that originates on the switch, perform this task:

Command	Purpose
Router(config)# ip local policy route-map <i>map-tag</i>	Identifies the route map to use for local PBR.



Note

- Local PBR traffic is processed in software on the RP.
- Use the **show ip local policy** command to display the route map used for local PBR.

Configuring PBR Recursive Next Hop

- [Setting the Recursive Next-Hop IP Address, page 34-5](#)
- [Verifying the Recursive Next-Hop Configuration, page 34-6](#)

Setting the Recursive Next-Hop IP Address



Note

PBR supports only one recursive next-hop IP address per route-map entry.

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	access-list permit source Example: Router(config)# access-list 101 permit 10.60.0.0 0.0.255.255	Configures an access list. The example configuration permits any source IP address that falls within the 10.60.0.0. 0.0.255.255 subnet.
Step 4	route-map map-tag Example: Router(config)# route-map abccomp	Enables policy routing and enters route-map configuration mode.
Step 5	set ip next-hop ip-address Example: Router(config-route-map)# set ip next-hop 10.10.1.1	Sets a next-hop router IP address. Note Set this IP address separately from the next-hop recursive router configuration.
Step 6	set ip next-hop {ip-address [...ip-address] recursive ip-address} Example: Router(config-route-map)# set ip next-hop recursive 10.20.3.3	Sets a recursive next-hop IP address. Note This configuration does not ensure that packets get routed using the recursive IP address if an intermediate IP address is a shorter route to the destination.
Step 7	match ip address access-list-number Example: Router(config-route-map)# match ip address 101	Sets an access list to be matched.
Step 8	end Example: Router(config-route-map)# end	Exits route-map configuration mode and returns to privileged EXEC mode.

Verifying the Recursive Next-Hop Configuration

To verify the recursive next-hop configuration, perform the following steps.

Step 1 show running-config | begin abccomp

Use this command to verify the IP addresses for a next-hop and recursive next-hop IP address, for example:

```
Router# show running-config | begin abccomp
```

```
route-map abccomp permit 10
  match ip address 101 ! Defines the match criteria for an access list.
  set ip next-hop recursive 10.3.3.3 ! If the match criteria are met, the recursive IP
  address is set.
  set ip next-hop 10.1.1.1 10.2.2.2 10.4.4.4
```

Step 2 **show route-map** *map-name*

Use this command to display the route maps, for example:

```
Router# show route-map abccomp
```

```
route-map abccomp, permit, sequence 10
  Match clauses:
    ip address (access-lists): 101
  Set clauses:
    ip next-hop recursive 10.3.3.3
    ip next-hop 10.1.1.1 10.2.2.2 10.4.4.4
  Policy routing matches: 0 packets, 0 bytes
```

Configuration Examples for PBR

- [Equal Access Example](#)
- [Differing Next Hops Example](#)
- [Recursive Next-Hop IP Address: Example](#)



Note

The examples shown below involve the use of the **access-list** command (ACL). The log keyword should not be used with this command in policy-based routing (PBR) because logging is not supported at the interrupt level for ACLs.

Equal Access Example

The following example provides two sources with equal access to two different service providers. Packets arriving on asynchronous interface 1 from the source 209.165.200.225 are sent to the router at 209.165.200.228 if the router has no explicit route for the destination of the packet. Packets arriving from the source 209.165.200.226 are sent to the router at 209.165.200.229 if the router has no explicit route for the destination of the packet. All other packets for which the router has no explicit route to the destination are discarded.

```
access-list 1 permit 209.165.200.225
access-list 2 permit 209.165.200.226
!
interface async 1
  ip policy route-map equal-access
!
route-map equal-access permit 10
  match ip address 1
  set ip default next-hop 209.165.200.228
route-map equal-access permit 20
  match ip address 2
  set ip default next-hop 209.165.200.229
```

```
route-map equal-access permit 30
  set default interface null0
```

Differing Next Hops Example

The following example illustrates how to route traffic from different sources to different places (next hops), and how to set the Precedence bit in the IP header. Packets arriving from source 209.165.200.225 are sent to the next hop at 209.165.200.227 with the Precedence bit set to priority; packets arriving from source 209.165.200.226 are sent to the next hop at 209.165.200.228 with the Precedence bit set to critical.

```
access-list 1 permit 209.165.200.225
access-list 2 permit 209.165.200.226
!
interface ethernet 1
  ip policy route-map Texas
!
route-map Texas permit 10
  match ip address 1
  set ip precedence priority
  set ip next-hop 209.165.200.227
!
route-map Texas permit 20
  match ip address 2
  set ip precedence critical
  set ip next-hop 209.165.200.228
```

Recursive Next-Hop IP Address: Example

The following example shows the configuration of IP address 10.3.3.3 as the recursive next-hop router:

```
route-map abccomp
  set ip next-hop 10.1.1.1
  set ip next-hop 10.2.2.2
  set ip next-hop recursive 10.3.3.3
  set ip next-hop 10.4.4.4
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

35

Layer 3 Interfaces

- [Restrictions for Layer 3 Interfaces, page 35-1](#)
- [How to Configure Subinterfaces on Layer 3 Interfaces, page 35-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Restrictions for Layer 3 Interfaces

When configuring Layer 3 interfaces, follow these guidelines and restrictions:

- We recommend that you configure no more than 2,000 Layer 3 VLAN interfaces.
- The **ip unnumbered** command is supported on Layer 3 VLAN interfaces.
- To support VLAN interfaces, create and configure VLANs and assign VLAN membership to Layer 2 LAN ports. For more information, see [Chapter 26, “Virtual Local Area Networks \(VLANs\)”](#) and [Chapter 25, “VLAN Trunking Protocol \(VTP\).”](#)
- Use bridge groups on VLAN interfaces, sometimes called fall-back bridging, to bridge nonrouted protocols. Bridge groups on VLAN interfaces are supported in software on the route processor (RP).
- Cisco IOS Release 15.4SY does not support the IEEE bridging protocol for bridge groups. Configure bridge groups to use the VLAN-bridge or the DEC spanning-tree protocol.
- The PFC supports these features on LAN port Layer 3 subinterfaces:
 - IPv4 unicast forwarding, including MPLS VPN
 - IPv4 multicast forwarding, including MPLS VPN
 - 6PE
 - EoMPLS

- IPv4 unnumbered
 - Counters for subinterfaces in MIBS and with the **show vlans** command
 - iBGP and eBGP
 - OSPF
 - EIGRP
 - RIPv1/v2
 - RIPv2
 - ISIS
 - Static routing
 - Unidirectional link routing (UDLR)
 - IGMPv1, IGMPv2, IGMPv3
 - PIMv1, PIMv2
 - SSM IGMPv3lite and URD
 - IGMP join
 - IGMP static group
 - Multicast routing monitor (MRM)
 - Multicast source discovery protocol (MSDP)
 - SSM
 - IPv4 Ping
 - IPv6 Ping
- Always use the **native** keyword when the VLAN ID is the ID of the IEEE 802.1Q native VLAN. Do not configure encapsulation on the native VLAN of an IEEE 802.1Q trunk without the **native** keyword.
 - The VLAN IDs used for Layer 2 VLANs and Layer 3 VLAN interfaces are separate from any VLAN IDs configured on Layer 3 subinterfaces. You can configure the same VLAN ID on a Layer 2 VLAN or Layer 3 VLAN interface and on a Layer 3 subinterface.
 - You can configure subinterfaces with any normal range or extended range VLAN ID in VTP transparent mode. Because VLAN IDs 1 to 1005 are global in the VTP domain and can be defined on other network devices in the VTP domain, you can use only extended range VLANs with subinterfaces in VTP client or server mode. In VTP client or server mode, normal range VLANs are excluded from subinterfaces.

**Note**

If you configure normal range VLANs on subinterfaces, you cannot change the VTP mode from transparent.

- Ensure that the Logical Interface Database (LDB) Manager does not pre-allocate the entire range of 4096 LDB Logical Interfaces (LIFs) for ports with Layer 3 subinterfaces, failing which results in a system-wide LDB exhaustion.

How to Configure Subinterfaces on Layer 3 Interfaces

To configure a subinterface, perform this task:

	Command	Purpose
Step 1	Router> enable	Enters privileged EXEC mode.
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# interface {{type slot/port.subinterface} {port-channel port_channel_number.subinterface}}	Selects an interface and enters subinterface configuration mode.
Step 4	Router(config-subif)# encapsulation dot1q vlan_ID [native]	Configures 802.1Q encapsulation for the subinterface.
Step 5	Router(config-if)# exit	Returns to global configuration mode.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

36

Unidirectional Ethernet (UDE) and Unidirectional Link Routing (UDLR)

- [Prerequisites for UDE and UDLR, page 36-1](#)
- [Restrictions for UDE and UDLR, page 36-1](#)
- [Information About UDE and UDLR, page 36-3](#)
- [Default Settings for UDE and UDLR, page 36-4](#)
- [How to Configure UDE and UDLR, page 36-4](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - Cisco IOS Release 15.4SY supports unidirectional Ethernet (UDE) and unidirectional link routing (UDLR) only on the WS-X6704-10GE 4-port 10-Gigabit Ethernet switching module.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for UDE and UDLR

None.

Restrictions for UDE and UDLR

- [UDE Restrictions, page 36-2](#)
- [UDLR Back-Channel Tunnel Restrictions, page 36-2](#)

UDE Restrictions

- STP cannot prevent Layer 2 loops in topologies that include unidirectional links.
- Send-only ports always transition to the STP forwarding state, because send-only ports never receive BPDUs.
- Receive-only ports cannot send BPDUs.
- Unidirectional ports do not support any features or protocols that require negotiation with the port at the other end of the link, including these:
 - Speed and duplex mode autonegotiation
 - Link negotiation
 - IEEE 802.3Z flow control
 - Dynamic trunking protocol (DTP)

You must manually configure the parameters that are typically controlled by Layer 2 protocols.

- A topology that includes unidirectional links only supports the VLAN Trunking Protocol (VTP) when the VTP server can send VTP frames to all switches in the VTP domain.
- Disable VTP pruning on switches that have send-only ports, because VTP pruning depends on a bidirectional exchange of information.
- Unidirectional EtherChannels cannot support PAgP or LACP. To create a unidirectional EtherChannel, you must configure the EtherChannel “on” mode.
- You can configure software-based UDE on the physical ports in an EtherChannel. You cannot configure software-based UDE on any nonphysical interfaces (for example, port-channel interfaces).
- When you implement hardware-based UDE on a port or configure software-based UDE on a port, UDLD is automatically disabled on the port.
- CDP sends CDP frames from send-only ports and receives CDP frames from receive-only ports, which means that the switch on the send-only side of a unidirectional link never receives CDP information.
- SPAN does not restrict configuration of unidirectional ports as sources or destinations.
 - Send-only ports can be SPAN destinations.
 - Receive-only ports can be SPAN sources.
- Unidirectional ports do not support IEEE 802.1X port-based authentication.
- IGMP snooping does not support topologies where there are unidirectional links between the switch and the hosts that are receiving multicast traffic.
- Configure UDLR with UDE to support communication over unidirectional links between IGMP snooping on the switch and a multicast router.
- Unidirectional links do not support ARP.

UDLR Back-Channel Tunnel Restrictions

- The PFC does not provide hardware support for UDLR back-channel tunnels. UDLR back-channel tunnels are supported in software.
- Configure a UDLR back-channel tunnel for each unidirectional link.

- On UDE send-only interfaces, configure the UDLR back-channel tunnel interface to receive.
- On UDE receive-only interfaces, configure the UDLR back-channel tunnel interface to send.
- You must configure IPv4 addresses on UDLR back-channel tunnel interfaces.
- You must configure source and destination IPv4 addresses on UDLR back-channel tunnel interfaces.
- The UDLR back-channel tunnel default mode is GRE.
- UDLR back-channel tunnels do not support IPv6 or MPLS.

Information About UDE and UDLR

- [UDE and UDLR Overview, page 36-3](#)
- [Information about UDE, page 36-3](#)
- [Information about UDLR, page 36-4](#)

UDE and UDLR Overview

Routing protocols support unidirectional links only if the unidirectional links emulate bidirectional links because routing protocols expect to send and receive traffic through the same interface.

Unidirectional links are advantageous because when you transmit mostly unacknowledged unidirectional high-volume traffic (for example, a video broadcast stream) over a high-capacity full-duplex bidirectional link, you use both the link from the source to the receiver and the equally high-capacity reverse-direction link, called the “back channel,” that carries the few acknowledgements from the receiver back to the source.

UDE and UDLR support use of a high-capacity unidirectional link for the high-volume traffic without consuming a similar high-capacity link for the back channel. UDE provides a high-capacity unidirectional link. UDLR provides the back channel through a tunnel that is configured over a regular-capacity link, and also provides bidirectional link emulation by transparently making the back channel appear to be on the same interface as the high-capacity unidirectional link.

Information about UDE

- [UDE Overview, page 36-3](#)
- [Hardware-Based UDE, page 36-4](#)
- [Software-Based UDE, page 36-4](#)

UDE Overview

You can implement UDE with hardware or in software. Hardware-based UDE and software-based UDE both use only one strand of fiber instead of the two strands of fiber required by bidirectional traffic.

The supported unidirectional transceiver (WDM-XENPAK-REC) provides receive-only UDE. You can configure software-based UDE as either transmit-only or receive-only. You do not need to configure software-based UDE on ports where you implement hardware-based UDE.

Hardware-Based UDE

You can create a unidirectional link by using a unidirectional transceiver. Unidirectional transceivers are less expensive than bidirectional transceivers. The supported unidirectional transceiver is WDM-XENPAK-REC.

Software-Based UDE

You can create a unidirectional link by configuring ports equipped with bidirectional transceivers to unidirectionally transmit or receive traffic. You can use software-based UDE when there is no appropriate unidirectional transceiver available. For example, with no supported transmit-only transceivers, you must configure transmit-only links with software-based UDE.

Information about UDLR

UDLR provides a unidirectional tunnel as the back channel of a unidirectional high-capacity link, and transparently emulates a single bidirectional link for unicast and multicast traffic.

UDLR intercepts packets that need to be sent on receive-only interfaces and sends them on UDLR back-channel tunnels. When routers receive these packets over UDLR back-channel tunnels, UDLR makes the packets appear as if received on send-only interfaces.

UDLR back-channel tunnels support these IPv4 features:

- Address Resolution Protocol (ARP)
- Next Hop Resolution Protocol (NHRP)
- Emulation of a bidirectional link for all IPv4 traffic (as opposed to only broadcast and multicast control traffic)
- IPv4 GRE multipoint at a receive-only tunnels



Note

UDLR back-channel tunnels do not support IPv6 or MPLS.

Default Settings for UDE and UDLR

None.

How to Configure UDE and UDLR

- [Configuring UDE, page 36-5](#)
- [Configuring UDLR, page 36-6](#)



Note

This caveat is open in releases that support UDLR: Neighboring ISIS routers are not seen through a UDLR topology. ([CSCee56596](#))

Configuring UDE

- [Configuring Hardware-Based UDE, page 36-5](#)
- [Configuring Software-Based UDE, page 36-5](#)

Configuring Hardware-Based UDE

Install a unidirectional transceiver to implement hardware-based UDE. Hardware-based UDE requires no software configuration procedures.

To verify hardware-based UDE on a port, perform this task:

Command	Purpose
Router# show interfaces [{ gigabitethernet tengigabitethernet } <i>slot/interface</i>] status	Verifies the configuration.

This example shows how to verify the configuration of Gigabit Ethernet port 1/1:

```
Router# show interfaces gigabitethernet 1/1 status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Gil/1		notconnect	1	full	1000	WDM-RXONLY

Configuring Software-Based UDE

To configure software-based UDE on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface [{ gigabitethernet tengigabitethernet } <i>slot/interface</i>]	Selects the interface to configure.
Step 2	Router(config-if)# unidirectional { send-only receive-only }	Configures software-based UDE.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to configure 10-Gigabit Ethernet port 1/1 as a UDE send-only port:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface tengigabitethernet 1/1
Router(config-if)# unidirectional send-only
Router(config-if)# end
```

Warning!

Enable port unidirectional mode will automatically disable port udld. You must manually ensure that the unidirectional link does not create a spanning tree loop in the network.

Enable 13 port unidirectional mode will automatically disable ip routing on the port. You must manually configure static ip route and arp entry in order to route ip traffic.

This example shows how to configure 10-Gigabit Ethernet port 1/2 as a UDE receive-only port:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router(config)# interface tengigabitethernet 1/2
Router(config-if)# unidirectional receive-only
Router(config-if)# end
```

Warning!

Enable port unidirectional mode will automatically disable port udd. You must manually ensure that the unidirectional link does not create a spanning tree loop in the network.

Enable 13 port unidirectional mode will automatically disable ip routing on the port. You must manually configure static ip route and arp entry in order to route ip traffic.

This example shows how to verify the configuration:

```
Router> show interface tengigabitethernet 1/1 unidirectional
Unidirectional configuration mode: send only
CDP neighbour unidirectional configuration mode: receive only
```

This example shows how to disable UDE on 10-Gigabit Ethernet interface 1/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface tengigabitethernet 1/1
Router(config-if)# no unidirectional
Router(config-if)# end
```

This example shows the result of entering the **show interface** command for a port that does not support unidirectional Ethernet:

```
Router# show interface gigabitethernet 6/1 unidirectional
Unidirectional Ethernet is not supported on GigabitEthernet6/1
```

Configuring UDLR

- [Configuring a Receive-Only Tunnel Interface for a UDE Send-Only Port, page 36-6](#)
- [Configuring a Send-Only Tunnel Interface for a UDE Receive-Only Port, page 36-7](#)

Configuring a Receive-Only Tunnel Interface for a UDE Send-Only Port

To configure a receive-only tunnel interface for a UDE send-only port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface tunnel <i>number</i>	Selects the tunnel interface.
Step 2	Router(config-if)# tunnel udlr receive-only <i>ude_send_only_port</i>	Associates the tunnel receive-only interface with the UDE send-only port.
Step 3	Router(config-if)# ip address <i>ipv4_address</i>	Configures the tunnel IPv4 address.
Step 4	Router(config-if)# tunnel source { <i>ipv4_address</i> <i>type number</i> }	Configures the tunnel source.
Step 5	Router(config-if)# tunnel destination { <i>hostname</i> <i>ipv4_address</i> }	Configures the tunnel destination.

Configuring a Send-Only Tunnel Interface for a UDE Receive-Only Port

To configure a send-only tunnel interface for a UDE receive-only port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>tunnel number</i>	Selects the tunnel interface.
Step 2	Router(config-if)# tunnel udlr send-only <i>ude_receive_only_port</i>	Associates the tunnel send-only interface with the UDE receive-only port.
Step 3	Router(config-if)# ip address <i>ipv4_address</i>	Configures the tunnel IPv4 address.
Step 4	Router(config-if)# tunnel source { <i>ipv4_address</i> <i>type number</i> }	Configures the tunnel source.
Step 5	Router(config-if)# tunnel destination { <i>hostname</i> <i>ipv4_address</i> }	Configures the tunnel destination.
Step 6	Router(config-if)# tunnel udlr address-resolution	Enables ARP and NHRP.

In the following UDE and UDLR sample configuration:

- On Router A:
 - Open Shortest Path First (OSPF) and PIM are configured.
 - 10-Gigabit Ethernet port 1/1 is a send-only UDE port.
 - The UDLR back-channel tunnel is configured as receive only and is associated with 10-Gigabit Ethernet port 1/1.
- On Router B:
 - OSPF and PIM are configured.
 - 10-Gigabit Ethernet port 1/2 is a receive-only UDE port.
 - The UDLR back-channel tunnel is configured as send-only and is associated with 10-Gigabit Ethernet port 1/2.
 - ARP and NHRP are enabled.

Router A Configuration

```
ip multicast-routing
!
! tengigabitethernet 1/1 is send-only
!
interface tengigabitethernet 1/1
 unidirectional send-only
 ip address 10.1.0.1 255.255.0.0
 ip pim sparse-dense-mode
!
! Configure tunnel as receive-only UDLR tunnel.
!
interface tunnel 0
 tunnel source 11.0.0.1
 tunnel destination 11.0.0.2
 tunnel udlr receive-only tengigabitethernet 1/1
!
! Configure OSPF.
!
router ospf <pid>
 network 10.0.0.0 0.255.255.255 area 0
```

Router B Configuration

```

ip multicast-routing
!
! tengigabitethernet 1/2 is receive-only
!
interface tengigabitethernet 1/2
 unidirectional receive-only
 ip address 10.1.0.2 255.255.0.0
 ip pim sparse-dense-mode
!
! Configure tunnel as send-only UDLR tunnel.
!
interface tunnel 0
 tunnel source 11.0.0.2
 tunnel destination 11.0.0.1
 tunnel udlr send-only tengigabitethernet 1/2
 tunnel udlr address-resolution
!
! Configure OSPF.
!
router ospf <pid>
 network 10.0.0.0 0.255.255.255 area 0

```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

37

Multiprotocol Label Switching (MPLS)

- [Prerequisites for MPLS, page 37-1](#)
- [Restrictions for MPLS, page 37-1](#)
- [Information About MPLS, page 37-2](#)
- [Default Settings for MPLS, page 37-6](#)
- [How to Configure MPLS Features, page 37-6](#)
- [Configuration Examples for MPLS, page 37-8](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for MPLS

None.

Restrictions for MPLS

- The PFC and DFCs supports up to 16 load-shared paths (Cisco IOS releases for other platforms support only 8 load-shared paths).
- MTU size checking is supported in hardware.
- Fragmentation is supported in software, including traffic that ingresses as IP and egresses as MPLS. To prevent excessive CPU utilization, you can rate-limit the traffic being sent to the RP for fragmentation with the **platform rate-limit all mtu-failure** command.

- MPLS supports these commands:
 - **mpls ip default route**
 - **mpls ip propagate-ttl**
 - **mpls ip ttl-expiration pop**
 - **mpls label protocol**
 - **mpls label range**
 - **mpls ip**
 - **mpls label protocol**
 - **mpls mtu**

For information about these commands, see these publications:

http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html

Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

Information About MPLS

- [MPLS Overview, page 37-2](#)
- [IP to MPLS, page 37-4](#)
- [MPLS to MPLS, page 37-4](#)
- [MPLS to IP, page 37-4](#)
- [MPLS VPN Forwarding, page 37-4](#)
- [Recirculation, page 37-4](#)
- [Hardware Supported Features, page 37-5](#)
- [Supported MPLS Features, page 37-5](#)


MPLS Overview

MPLS uses label switching to forward packets over Ethernet. Labels are assigned to packets based on groupings or forwarding equivalence classes (FECs). The label is added between the Layer 2 and the Layer 3 header.

In an MPLS network, the label edge router (LER) performs a label lookup of the incoming label, swaps the incoming label with an outgoing label, and sends the packet to the next hop at the label switch router (LSR). Labels are imposed (pushed) on packets only at the ingress edge of the MPLS network and are removed (popped) at the egress edge. The core network LSRs (provider, or P routers) read the labels, apply the appropriate services, and forward the packets based on the labels.

Incoming labels are aggregate or nonaggregate. The aggregate label indicates that the arriving MPLS packet must be switched through an IP lookup to find the next hop and the outgoing interface. The nonaggregate label indicates that the packet contains the IP next hop information.


[Figure 37-1](#) shows an MPLS network of a service provider that connects two sites of a customer network.

Figure 37-1 MPLS Network

The route processor (RP) performs Layer 3 control-plane functions, including address resolution and routing protocols. The RP processes information from the Routing and Label Distribution Protocols and builds the IP forwarding (FIB) table and the label forwarding (LFIB) table. The RP distributes the information in both tables to the PFC and DFCs.

The PFC and DFCs receive the information and creates its own copies of the FIB and LFIB tables. Together, these tables comprise the FIB TCAM. The PFC and DFCs look up incoming IP packets and labeled packets against the FIB TCAM table. The lookup result is the pointer to a particular adjacency entry. It is the adjacency entry that contains appropriate information for label pushing (for IP to MPLS path), label swapping (for MPLS to MPLS path), label popping (for MPLS to IP path), and encapsulation.

Figure 37-2 shows the various functional blocks that support MPLS. Routing protocol generates a routing information base (RIB) that is used for forwarding IP and MPLS data packets. For Cisco Express Forwarding (CEF), necessary routing information from the RIB is extracted and built into a forwarding information base (FIB). The label distribution protocol (LDP) obtains routes from the RIB and distributes the label across a label switch path to build a label forwarding information base (LFIB) in each of the LSRs and LERs.

Figure 37-2 MPLS Forwarding, Control and Data Planes

IP to MPLS

At the ingress to the MPLS network, the PFC examines the IP packets and performs a route lookup in the FIB TCAM. The lookup result is the pointer to a particular adjacency entry. The adjacency entry contains the appropriate information for label pushing (for IP to MPLS path) and encapsulation. The PFC generates a result containing the imposition label(s) needed to switch the MPLS packet.

MPLS to MPLS

At the core of an MPLS network, the PFC uses the topmost label to perform a lookup in the FIB TCAM. The successful lookup points to an adjacency that swaps the top label in the packet with a new label as advertised by the downstream label switch router (LSR). If the router is the penultimate hop LSR router (the upstream LSR next to the egress LER), the adjacency instructs the PFCBXL to pop the topmost label, resulting in either an MPLS packet with the remaining label for any VPN or AToM use or a native IP packet.

MPLS to IP

At the egress of the MPLS network there are several possibilities.

For a native IP packet (when the penultimate router has popped the label), the PFC performs a route lookup in the FIB TCAM.

For a MPLS VPN packet, after the Interior Gateway Protocol (IGP) label is popped at penultimate router, the VPN label remains. The operation that the PFC performs depends on the VPN label type. Packets carrying aggregate labels require a second lookup based on the IP header after popping the aggregate label. For a nonaggregate label, the PFC performs a route lookup in the FIB TCAM to obtain the IP next hop information.

For the case of a packet with an IGP label and a VPN label, when there is no penultimate hop popping (PHP), the packet carries the explicit-null label on top of the VPN label. The PFC looks up the top label in the FIB TCAM and recirculates the packet. Then the PFC handles the remaining label as described in the preceding paragraph, depending on whether it is an aggregate or nonaggregate label.

Packets with the explicit-null label for the cases of EoMPLS, MPLS, and MPLS VPN an MPLS are handled the same way.

MPLS VPN Forwarding

There are two types of VPN labels: aggregate labels for directly connected network or aggregate routes, and nonaggregate labels. Packets carrying aggregate labels require a second lookup based on the IP header after popping the aggregate label. The VPN information (VPN-IPv4 address, extended community, and label) is distributed through the Multiprotocol-Border Gateway Protocol (MP-BGP).

Recirculation

In certain cases, the PFC provides the capability to recirculate the packets. Recirculation can be used to perform additional lookups in the ACL or QoS TCAMs, the NetFlow table, or the FIB TCAM table. Recirculation is necessary in these situations:

- To push more than three labels on imposition
- To pop more than two labels on disposition
- To pop an explicit null top label
- When the VPN Routing and Forwarding (VRF) number is more than 511
- For IP ACL on the egress interface (for nonaggregate (per-prefix) labels only)

Packet recirculation occurs only on a particular packet flow; other packet flows are not affected. The rewrite of the packet occurs on the modules; the packets are then forwarded back to the PFC for additional processing.

Hardware Supported Features

The following features are supported in hardware:

- Label operation—Any number of labels can be pushed or popped, although for best results, up to three labels can be pushed, and up to two labels can be popped in the same operation.
- IP to MPLS path—IP packets can be received and sent to the MPLS path.
- MPLS to IP path—Labeled packets can be received and sent to the IP path.
- MPLS to MPLS path—Labeled packets can be received and sent to the label path.
- MPLS Traffic Engineering (MPLS TE)—Enables an MPLS backbone to replicate and expand the traffic engineering capabilities of Layer 2 ATM and Frame Relay networks.
- Time to live (TTL) operation—At the ingress edge of the MPLS network, the TTL value in the MPLS frame header can be received from either the TTL field of the IP packet header or the user-configured value from the adjacency entry. At the egress of the MPLS network, the final TTL equals the minimum (label TTL and IP TTL)-1.



Note

With the Uniform mode, the TTL is taken from the IP TTL; with the Pipe mode, a value of 255, taken from the hardware register, is used for the outgoing label.

- QoS—Information on Differentiated Services (DiffServ) and ToS from IP packets can be mapped to MPLS EXP field.
- MPLS/VPN Support—Up to 1024 VRFs can be supported (over 511 VRFs requires recirculation).
- Ethernet over MPLS—The Ethernet frame can be encapsulated at the ingress to the MPLS domain and the Ethernet frame can be decapsulated at the egress.
- Packet recirculation—The PFC provides the capability to recirculate the packets. See the [“Recirculation” section on page 37-4](#).
- Configuration of MPLS switching is supported on VLAN interfaces with the **mpls ip** command.

Supported MPLS Features

- MPLS features:
 - Basic MPLS
 - MPLS TE
 - MPLS TE DiffServ Aware (DS-TE)

- MPLS TE Forwarding Adjacency
- MPLS TE Interarea Tunnels
- MPLS virtual private networks (VPNs)
- MPLS VPN Carrier Supporting Carrier (CSC)
- MPLS VPN Carrier Supporting Carrier IPv4 BGP Label Distribution
- MPLS VPN Interautonomous System (InterAS) Support
- MPLS VPN Inter-AS IPv4 BGP label distribution

See these publications for more information:

http://www.cisco.com/en/US/docs/ios-xml/ios/mpls/config_library/15-sy/mp-15-sy-library.html

http://www.cisco.com/en/US/tech/tk436/tk428/technologies_configuration_example09186a0080093fcb.shtml

http://www.cisco.com/en/US/tech/tk436/tk428/technologies_configuration_example09186a0080093fd0.shtml

- HSRP Support for MPLS VPNs—See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/ipapp_fhrp/configuration/15-sy/fhp-15-sy-book.html
- OSPF Sham-Link Support for MPLS VPN—See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-sham-link.html
- Multi-VPN Routing and Forwarding (VRF) for CE Routers (VRF Lite)—VRF Lite is supported with the following features:
 - IPv4 forwarding between VRFs interfaces
 - IPv4 ACLs
 - IPv4 HSRP

See this publication:

http://www.cisco.com/en/US/products/hw/routers/ps259/prod_bulletin09186a00800921d7.html

Default Settings for MPLS

None.

How to Configure MPLS Features

- [Configuring MPLS, page 37-6](#)
- [Configuring MUX-UNI Support on LAN Cards, page 37-7](#)

Configuring MPLS

Use these publications to configure MPLS:

http://www.cisco.com/en/US/docs/ios-xml/ios/mpls/config_library/15-sy/mp-15-sy-library.html

Configuring MUX-UNI Support on LAN Cards

A User Network Interface (UNI) is the point where the customer edge (CE) equipment connects to the ingress PE and an attachment VLAN is a VLAN on a UNI port.

The MUX-UNI support on LAN cards feature provides the ability to partition a physical port on an attachment VLAN to provide multiple Layer 2 and Layer 3 services over a single UNI.

To configure MUX-UNI support on LAN cards, perform this task on the provider edge (PE) routers.

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type number</i>	Selects an interface to configure and enters interface configuration mode; valid only for Ethernet ports.
Step 3	Router(config-if)# switchport	Puts an interface that is in Layer 3 mode into Layer 2 mode for Layer 2 configuration.
Step 4	Router(config-if)# switchport trunk encapsulation dot1q	Configures the port to support 802.1Q encapsulation. You must configure each end of the link with the same encapsulation type.
Step 5	Router(config-if)# switchport mode trunk	Configures the port as a VLAN trunk.
Step 6	Router(config-if)# switchport trunk allowed vlan <i>vlan-list</i>	By default, all VLANs are allowed. Use this command to explicitly allow VLANs; valid values for <i>vlan-list</i> are from 1 to 4094. Note Avoid overlapping VLAN assignments between main and subinterfaces. VLAN assignments between the main interface and subinterfaces must be mutually exclusive.
Step 7	Router(config-if)# exit	Exits interface configuration mode.
Step 8	Router(config)# interface <i>type slot/port.subinterface-number</i>	Selects a subinterface to configure and enters interface configuration mode; valid only for Ethernet ports.
Step 9	Router(config-if)# encapsulation dot1q <i>vlan_id</i>	Enables the subinterface to accept 802.1Q VLAN packets. The subinterfaces between the CE and PE routers that are running Ethernet over MPLS must be in the same subnet. All other subinterfaces and backbone routers do not need to be on the same subnet.
Step 10	Router(config-if)# xconnect <i>peer_router_id vcid encapsulation mpls</i>	Binds the attachment circuit to a pseudowire VC. The syntax for this command is the same as for all other Layer 2 transports.

This example shows a physical trunk port used as UNI:

```
Router(config)# interface gigabitethernet 3/1
Router(config-if)# switchport
Router(config-if)# switchport encapsulation dot1q
Router(config-if)# switchport mode trunk
Router(config-if)# switchport trunk allowed vlan 200-250
Router(config-if)# exit
```

```
Router(config)# interface gigabitethernet 3/1.10
Router(config-if)# encaps dot1q 3000
Router(config-if)# xconnect 10.0.0.1 3000 encapsulation mpls
Router(config-if)# exit
```

This example shows a Layer 2 port channel used as UNI:

```
Router(config)# interface port-channel 100
Router(config-if)# switchport
Router(config-if)# switchport trunk encapsulation dot1q
Router(config-if)# switchport trunk allowed vlan 100-200
Router(config-if)# switchport mode trunk
Router(config-if)# no ip address
Router(config-if)# exit

Router(config)# interface port-channel 100.1
Router(config-if)# encapsulation dot1q 3100
Router(config-if)# xconnect 10.0.0.30 100 encapsulation mpls
Router(config-if)# exit
```

This example shows Layer 3 termination and VRF for muxed UNI ports:

```
Router(config)# vlan 200, 300, 400
Router(config)# interface gigabitethernet 3/1
Router(config-if)# switchport
Router(config-if)# switchport encapsulation dot1q
Router(config-if)# switchport mode trunk
Router(config-if)# switchport trunk allowed vlan 200-500
Router(config-if)# exit

Router(config)# interface gigabitethernet 3/1.10
Router(config-if)# encaps dot1q 3000
Router(config-if)# xconnect 10.0.0.1 3000 encapsulation mpls
Router(config-if)# exit

Router(config)# interface vlan 200
Router(config-if)# ip address 1.1.1.3
Router(config-if)# exit

Router(config)# interface vlan 300
Router(config-if)# ip vpn VRF A
Router(config-if)# ip address 3.3.3.1
Router(config-if)# exit

Router(config)# interface vlan 400
Router(config-if)# ip address 4.4.4.1
Router(config-if)# ip ospf network broadcast
Router(config-if)# mpls label protocol ldp
Router(config-if)# mpls ip
Router(config-if)# exit
```

Configuration Examples for MPLS

The following is an example of a basic MPLS configuration:

```
*****
Basic MPLS
*****
```

IP ingress interface:

```
Router# mpls label protocol ldp
```

```
interface GigabitEthernet6/2
 ip address 75.0.77.1 255.255.255.0
 media-type rj45
 speed 1000
end
```

Label egress interface:

```
interface GigabitEthernet7/15
 mtu 9216
 ip address 75.0.67.2 255.255.255.0
 logging event link-status
 mpls ip
```

```
Router# show ip route 188.0.0.0
```

Routing entry for 188.0.0.0/24, 1 known subnets

```
O IA    188.0.0.0 [110/1] via 75.0.77.2, 00:00:10, GigabitEthernet6/2
```

```
Router# show ip routing 88.0.0.0
```

Routing entry for 88.0.0.0/24, 1 known subnets

```
O E2    88.0.0.0 [110/0] via 75.0.67.1, 00:00:24, GigabitEthernet7/15
        [110/0] via 75.0.21.2, 00:00:24, GigabitEthernet7/16
```

```
Router# show mpls forwarding-table 88.0.0.0
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
30	50	88.0.0.0/24	0	Gi7/15	75.0.67.1
	50	88.0.0.0/24	0	Gi7/16	75.0.21.2

```
Router# show platform cef 88.0.0.0 detail
```

Codes: M - mask entry, V - value entry, A - adjacency index, P - priority bit
 D - full don't switch, m - load balancing modnumber, B - BGP Bucket sel
 V0 - Vlan 0, C0 - don't comp bit 0, V1 - Vlan 1, C1 - don't comp bit 1
 RVTEN - RPF Vlan table enable, RVTSEL - RPF Vlan table select

Format: IPV4_DA - (8 | xtag vpn pi cr recirc tos prefix)

Format: IPV4_SA - (9 | xtag vpn pi cr recirc prefix)

```
M(3223 ): E | 1 FFF 0 0 0 0 255.255.255.0
```

```
V(3223 ): 8 | 1 0 0 0 0 0 88.0.0.0 (A:344105 ,P:1,D:0,m:1 ,B:0 )
```

```
M(3223 ): E | 1 FFF 0 0 0 255.255.255.0
```

```
V(3223 ): 9 | 1 0 0 0 0 88.0.0.0 (V0:0 ,C0:0 ,V1:0 ,C1:0 ,RVTEN:0 ,RVTSEL:0 )
```

```
Router# show platform cef adj ent 344105
```

```
Index: 344105 smac: 0005.9a39.a480, dmac: 000a.8ad8.2340
mtu: 9234, vlan: 1031, dindex: 0x0, l3rw_vld: 1
packets: 109478260, bytes: 7006608640
```

```
Router# show platform cef adj ent 344105 detail
```

```
Index: 344105 smac: 0005.9a39.a480, dmac: 000a.8ad8.2340
mtu: 9234, vlan: 1031, dindex: 0x0, l3rw_vld: 1
format: MPLS, flags: 0x1000008418
label0: 0, exp: 0, ovr: 0
label1: 0, exp: 0, ovr: 0
label2: 50, exp: 0, ovr: 0
```

```
op: PUSH_LABEL2  
packets: 112344419, bytes: 7190042816
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

38

MPLS VPN Support

- [Prerequisites for MPLS VPN, page 38-1](#)
- [Restrictions for MPLS VPN, page 38-1](#)
- [Information About MPLS VPN Support, page 38-2](#)
- [How to Configure MPLS VPNs, page 38-3](#)
- [Configuration Example for MPLS VPNs, page 38-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for MPLS VPN

None.

Restrictions for MPLS VPN

- When configuring MPLS VPN, note that VPNs are recirculated when the number of VPNs is over 511.
- MPLS VPN supports these commands:
 - **address-family**
 - **exit-address-family**
 - **import map**
 - **ip route vrf**

- **ip route forwarding**
- **ip vrf**
- **neighbor activate**
- **rd**
- **route-target**

For information about these commands, see these publications:

http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html


Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

Information About MPLS VPN Support

The IP VPN feature for MPLS allows a Cisco IOS network to deploy scalable IP Layer 3 VPN backbone services to multiple sites deployed on a shared infrastructure while also providing the same access or security policies as a private network. VPN based on MPLS technology provides the benefits of routing isolation and security, as well as simplified routing and better scalability. See this publication for more information about MPLS VPNs:

http://www.cisco.com/en/US/docs/ios-xml/ios/mpls/config_library/15-sy/mp-15-sy-library.html

Figure 38-1 *VPNs with MPLS Service Provider Backbone*



119122

At the ingress PE, the PFC makes a forwarding decision based on the packet headers. The PFC contains a table that maps VLANs to VPNs. In the switch architecture, all physical ingress interfaces in the system are associated with a specific VPN. The PFC looks up the IP destination address in the CEF table but only against prefixes that are in the specific VPN. (The table entry points to a specific set of adjacencies and one is chosen as part of the load-balancing decision if multiple parallel paths exist.)

The table entry contains the information on the Layer 2 header that the packet needs, as well as the specific MPLS labels to be pushed onto the frame. The information to rewrite the packet goes back to the ingress module where it is rewritten and forwarded to the egress line interface.

VPN traffic is handled at the egress from the PE based upon the per-prefix labels or aggregate labels. If per-prefix labels are used, then each VPN prefix has a unique label association; this allows the PE to forward the packet to the final destination based upon a label lookup in the FIB.

**Note**

The PFC allocates only one aggregate label per VRF.

If aggregate labels are used for disposition in an egress PE, many prefixes on the multiple interfaces may be associated with the label. In this case, the PFC must perform an IP lookup to determine the final destination. The IP lookup may require recirculation.

How to Configure MPLS VPNs

For information on configuring MPLS VPN, see this publication:

http://www.cisco.com/en/US/docs/ios-xml/ios/mpls/config_library/15-sy/mp-15-sy-library.html

**Note**

If you use a Layer 3 VLAN interface as the MPLS uplink through a Layer 2 port peering with another MPLS device, then you can use another Layer 3 VLAN interface as the VRF interface.

Configuration Example for MPLS VPNs

This sample configuration shows LAN CE-facing interfaces. MPLS switching configuration in Cisco IOS Release 15.4SY is identical to configuration in other releases.

```
!ip vrf blues
  rd 100:10
  route-target export 100:1
  route-target import 100:1
!
mpls label protocol ldp
mpls ldp logging neighbor-changes
!
interface Loopback0
  ip address 10.4.4.4 255.255.255.255
!
interface GigabitEthernet4/2
  description Catalyst link to P2
  no ip address
!
interface GigabitEthernet4/2.42
  encapsulation dot1Q 42
  ip address 10.0.3.2 255.255.255.0
  tag-switching ip
!
interface GigabitEthernet7/3
  description Catalyst link to CE2
  no ip address
!
interface GigabitEthernet7/3.73
```

```

encapsulation dot1Q 73
ip vrf forwarding blues
ip address 10.19.7.1 255.255.255.0
!
router ospf 100
log-adjacency-changes
network 10.4.4.4 0.0.0.0 area 0
network 10.0.0.0 0.0.255.255 area 0
!
router ospf 65000 vrf blues
log-adjacency-changes
redistribute bgp 100 subnets
network 10.19.0.0 0.0.255.255 area 0
!
router bgp 100
no synchronization
bgp log-neighbor-changes
neighbor 10.3.3.3 remote-as 100
neighbor 10.3.3.3 description MP-BGP to PE1
neighbor 10.3.3.3 update-source Loopback0
no auto-summary
!
address-family vpnv4
neighbor 10.3.3.3 activate
neighbor 10.3.3.3 send-community extended
exit-address-family
!
address-family ipv4 vrf blues
redistribute connected
redistribute ospf 65000 match internal external 1 external 2
no auto-summary
no synchronization
exit-address-family
!

```


39

Ethernet over MPLS (EoMPLS)

- [Prerequisites for EoMPLS, page 39-1](#)
- [Restrictions for EoMPLS, page 39-1](#)
- [Information About EoMPLS, page 39-3](#)
- [Default Settings for EoMPLS, page 39-3](#)
- [How to Configure EoMPLS, page 39-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for EoMPLS

Before you configure EoMPLS, ensure that the network is configured as follows:

- Configure IP routing in the core so that the PE routers can reach each other through IP.
- Configure MPLS in the core so that a label switched path (LSP) exists between the PE routers.

Restrictions for EoMPLS

- EoMPLS in Cisco IOS Release 15.4SY does not support load balancing at the tunnel ingress; only one Interior Gateway Protocol (IGP) path is selected even if multiple IGP paths are available, but load balancing is available at the MPLS core.

- Ensure that the maximum transmission unit (MTU) of all intermediate links between endpoints is sufficient to carry the largest Layer 2 packet received.
- EoMPLS supports VLAN packets that conform to the IEEE 802.1Q standard. The 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames.
- When the QoS is enabled on a Layer 2 port, either 802.1q P bits or IP precedence bits can be preserved with the trusted configuration. However, by default the unpreserved bits are overwritten by the value of preserved bits. For instance, if you preserve the P bits, the IP precedence bits are overwritten with the value of the P bits. To preserve the IP precedence bits, use the **no platform qos rewrite ip dscp** command. The **no platform qos rewrite ip dscp** command is not compatible with the MPLS and MPLS VPN features.
- EoMPLS is not supported with private VLANs.
- The following restrictions apply to using trunks with EoMPLS:
 - To support Ethernet spanning tree bridge protocol data units (BPDUs) across an EoMPLS cloud, you must disable spanning tree for the Ethernet-over-MPLS VLAN. This ensures that the EoMPLS VLANs are carried only on the trunk to the customer switch. Otherwise, the BPDUs are not directed to the EoMPLS cloud.
 - The native VLAN of a trunk must not be configured as an EoMPLS VLAN.
- In Cisco IOS Release 15.4SY, all protocols (for example, CDP, VTP, BPDUs) are tunneled across the MPLS cloud without conditions.
- Unique VLANs are required across interfaces. You cannot use the same VLAN ID on different interfaces.
- EoMPLS tunnel destination route in the routing table and the CEF table must be a /32 address (host address where the mask is 255.255.255.255) to ensure that there is a label-switched path (LSP) from PE to PE.
- For a particular EoMPLS connection, both the ingress EoMPLS interface on the ingress PE and the egress EoMPLS interface on the egress PE have to be subinterfaces with dot1Q encapsulation or neither is a subinterface.
- 802.1Q in 802.1Q over EoMPLS is supported if the outgoing interface connecting to MPLS network is a port on an Layer 2 card.
- Shaping EoMPLS traffic is not supported if the egress interface connecting to an MPLS network is a Layer 2 LAN port (a mode known as PFC-based EoMPLS).
- EoMPLS based on a PFC does not perform any Layer 2 lookup to determine if the destination MAC address resides on the local or remote segment and does not perform any Layer 2 address learning (as traditional LAN bridging does).
- The AToM control word is not supported.
- Ethernet packets with hardware-level cyclic redundancy check (CRC) errors, framing errors, and runt packets are discarded on input.
- You must configure VLAN-based EoMPLS on subinterfaces.
- Port-based EoMPLS and VLAN-based EoMPLS are mutually exclusive. If you enable a main interface for port-to-port transport, you also cannot enter commands on a subinterface.
- EoMPLS is not supported on Layer 3 VLAN interfaces.
- Point-to-point EoMPLS works with a physical interface and subinterfaces.

Information About EoMPLS

- [AToM Overview, page 39-3](#)
- [EoMPLS Overview, page 39-3](#)

AToM Overview

Any Transport over MPLS (AToM) transports Layer 2 packets over an MPLS backbone. AToM uses a directed Label Distribution Protocol (LDP) session between edge routers for setting up and maintaining connections. Forwarding occurs through the use of two level labels that provide switching between the edge routers. The external label (tunnel label) routes the packet over the MPLS backbone to the egress PE at the ingress PE. The VC label is a demuxing label that determines the connection at the tunnel endpoint (the particular egress interface on the egress PE as well as the VLAN identifier for an Ethernet frame).

EoMPLS Overview

EoMPLS is one of the AToM transport types. EoMPLS works by encapsulating Ethernet PDUs in MPLS packets and forwarding them across the MPLS network. Each PDU is transported as a single packet. Cisco IOS Release 15.4SY supports two EoMPLS modes:

- VLAN mode—Transports Ethernet traffic from a source 802.1Q VLAN to a destination 802.1Q VLAN through a single VC over an MPLS network. VLAN mode uses VC type 5 as default (no dot1q tag) and VC type 4 (transport dot1 tag) if the remote PE does not support VC type 5 for subinterface (VLAN) based EoMPLS.
- Port mode—Allows all traffic on a port to share a single VC across an MPLS network. Port mode uses VC type 5.

**Note**

For both VLAN mode and port mode, EoMPLS in Cisco IOS Release 15.4SY does not allow local switching of packets between interfaces unless you use loopback interfaces.

LAN ports can receive Layer 2 traffic, impose labels, and switch the frames into the MPLS core.

Default Settings for EoMPLS

None.

How to Configure EoMPLS

- [Configuring VLAN-Based EoMPLS, page 39-4](#)
- [Configuring Port-Based EoMPLS, page 39-6](#)

Configuring VLAN-Based EoMPLS

To configure VLAN-based EoMPLS, perform this task on the provider edge (PE) routers:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface gigabitethernet <i>slot/interface.subinterface</i>	Specifies the Gigabit Ethernet subinterface. Make sure that the subinterface on the adjoining CE router is on the same VLAN as this PE router.
Step 3	Router(config-if)# encapsulation dot1q <i>vlan_id</i>	Enables the subinterface to accept 802.1Q VLAN packets. - The subinterfaces between the CE and PE routers that are running Ethernet over MPLS must be in the same subnet. - All other subinterfaces and backbone routers do not need to be on the same subnet.
Step 4	Router(config-if)# xconnect <i>peer_router_id vcid</i> encapsulation mpls	Binds the attachment circuit to a pseudowire VC. The syntax for this command is the same as for all other Layer 2 transports.

This is a VLAN-based EoMPLS configuration sample:

```
interface GigabitEthernet7/4.2
encapsulation dot1Q 3
xconnect 13.13.13.13 3 encapsulation mpls
no shut
```



Note

The IP address is configured on subinterfaces of the CE devices.

To verify and display the configuration of Layer 2 VLAN transport over MPLS tunnels, perform the following:

- To display a single line for each VLAN, naming the VLAN, status, and ports, enter the **show vlan brief** command.

```
Router# show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	
2	VLAN0002	active	
3	VLAN0003	active	
1002	fddi-default	act/unsup	
1003	token-ring-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trnet-default	act/unsup	

- To verify that the PE router endpoints have discovered each other, enter the **show mpls ldp discovery** command. When an PE router receives an LDP hello message from another PE router, it considers that router and the specified label space to be “discovered.”

```
Router# show mpls ldp discovery
Local LDP Identifier:
 13.13.13.13:0
Discovery Sources:
Interfaces:
  GE-WAN3/3 (ldp): xmit/rcv
    LDP Id: 12.12.12.12:0
Targeted Hellos:
 13.13.13.13 -> 11.11.11.11 (ldp): active/passive, xmit/rcv
    LDP Id: 11.11.11.11:0
```

- To verify that the label distribution session has been established, enter the **show mpls ldp neighbor** command. The third line of the output shows that the state of the LDP session is operational and shows that messages are being sent and received.

```
Router# show mpls ldp neighbor
Peer LDP Ident: 12.12.12.12:0; Local LDP Ident 13.13.13.13:0
TCP connection: 12.12.12.12.646 - 13.13.13.13.11010
State: Oper; Msgs sent/rcvd: 1649/1640; Downstream
Up time: 23:42:45
LDP discovery sources:
  GE-WAN3/3, Src IP addr: 34.0.0.2
Addresses bound to peer LDP Ident:
 23.2.1.14      37.0.0.2      12.12.12.12      34.0.0.2
 99.0.0.1
Peer LDP Ident: 11.11.11.11:0; Local LDP Ident 13.13.13.13:0
TCP connection: 11.11.11.11.646 - 13.13.13.13.11013
State: Oper; Msgs sent/rcvd: 1650/1653; Downstream
Up time: 23:42:29
LDP discovery sources:
  Targeted Hello 13.13.13.13 -> 11.11.11.11, active, passive
Addresses bound to peer LDP Ident:
 11.11.11.11    37.0.0.1      23.2.1.13
```

- To verify that the label forwarding table is built correctly, enter the **show mpls forwarding-table** command to verify that a label has been learned for the remote PE and that the label is going from the correct interface to the correct next-hop.

```
Router# show mpls forwarding-table
Local  Outgoing  Prefix          Bytes tag  Outgoing     Next Hop
tag    tag or VC   or Tunnel Id   switched  interface
16     Untagged   223.255.254.254/32  \
20     Untagged   12ckt (2)       133093    V12          point2point
21     Untagged   12ckt (3)       185497    V13          point2point
24     Pop tag    37.0.0.0/8      0         GE3/3        34.0.0.2
25     17         11.11.11.11/32  0         GE3/3        34.0.0.2
26     Pop tag    12.12.12.12/32  0         GE3/3        34.0.0.2
```

The output shows the following data:

- Local tag—Label assigned by this router.
- Outgoing tag or VC—Label assigned by next hop.
- Prefix or Tunnel Id—Address or tunnel to which packets with this label are going.
- Bytes tag switched— Number of bytes switched out with this incoming label.
- Outgoing interface—Interface through which packets with this label are sent.
- Next Hop—IP address of neighbor that assigned the outgoing label.

- To display the state of the currently routed VCs, enter the **show mpls l2transport vc** command.

```
Router# show mpls l2transport vc
```

Local intf	Local circuit	Dest address	VC ID	Status
Vl2	Eth VLAN 2	11.11.11.11	2	UP
Vl3	Eth VLAN 3	11.11.11.11	3	UP

To display detailed information about each VC, add the keyword **detail**.

```
Router# show mpls l2transport vc detail
```

```
Local interface: Vl2 up, line protocol up, Eth VLAN 2 up
Destination address: 11.11.11.11, VC ID: 2, VC status: up
Tunnel label: 17, next hop 34.0.0.2
Output interface: GE3/3, imposed label stack {17 18}
Create time: 01:24:44, last status change time: 00:10:55
Signaling protocol: LDP, peer 11.11.11.11:0 up
MPLS VC labels: local 20, remote 18
Group ID: local 71, remote 89
MTU: local 1500, remote 1500
Remote interface description:
Sequencing: receive disabled, send disabled
VC statistics:
packet totals: receive 1009, send 1019
byte totals:   receive 133093, send 138089
packet drops:  receive 0, send 0
```

```
Local interface: Vl3 up, line protocol up, Eth VLAN 3 up
Destination address: 11.11.11.11, VC ID: 3, VC status: up
Tunnel label: 17, next hop 34.0.0.2
Output interface: GE3/3, imposed label stack {17 19}
Create time: 01:24:38, last status change time: 00:10:55
Signaling protocol: LDP, peer 11.11.11.11:0 up
MPLS VC labels: local 21, remote 19
Group ID: local 72, remote 90
MTU: local 1500, remote 1500
Remote interface description:
Sequencing: receive disabled, send disabled
VC statistics:
packet totals: receive 1406, send 1414
byte totals:   receive 185497, send 191917
packet drops:  receive 0, send 0
```

Configuring Port-Based EoMPLS

To support 802.1Q-in-802.1Q traffic and Ethernet traffic over EoMPLS in Cisco IOS Release 15.4SY, configure port-based EoMPLS by performing this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface gigabitethernet slot/interface	Specifies the Gigabit Ethernet interface. Make sure that the interface on the adjoining CE router is on the same VLAN as this PE router.
Step 3	Router(config-if)# xconnect peer_router_id vcid encapsulation mpls	Binds the attachment circuit to a pseudowire VC. The syntax for this command is the same as for all other Layer 2 transports.

The following is an example of a port-based configuration:

```
Router# show mpls l2transport vc
```

Local intf	Local circuit	Dest address	VC ID	Status
Gi8/48	Ethernet	75.0.78.1	1	UP
Gi7/11.2000	Eth VLAN 2000	75.0.78.1	2000	UP

```
Router# show run interface gigabitethernet 8/48
```

```
Building configuration...
```

```
Current configuration : 86 bytes
!
interface GigabitEthernet8/48
  no ip address
  xconnect 75.0.78.1 1 encapsulation mpls
end
```

```
Router# show run interface gigabitethernet 7/11
```

```
Building configuration...
```

```
Current configuration : 118 bytes
!
interface GigabitEthernet7/11
  description Traffic-Generator
  no ip address
  logging event link-status
  speed nonegotiate
end
```

```
Router# show run int gigabitethernet 7/11.2000
```

```
Building configuration...
```

```
Current configuration : 112 bytes
!
interface GigabitEthernet7/11.2000
  encapsulation dot1Q 2000
  xconnect 75.0.78.1 2000 encapsulation mpls
end
```

```
Router# show mpls l2transport vc 1 detail
```

```
Local interface: Gi7/47 up, line protocol up, Ethernet up
Destination address: 75.0.80.1, VC ID: 1, VC status: up
Tunnel label: 5704, next hop 75.0.83.1
Output interface: Te8/3, imposed label stack {5704 10038}
Create time: 00:30:33, last status change time: 00:00:43
Signaling protocol: LDP, peer 75.0.80.1:0 up
MPLS VC labels: local 10579, remote 10038
Group ID: local 155, remote 116
MTU: local 1500, remote 1500
Remote interface description:
Sequencing: receive disabled, send disabled
VC statistics:
  packet totals: receive 26, send 0
  byte totals:   receive 13546, send 0
  packet drops:  receive 0, send 0
```

To display the VC type:

```
Router# remote command switch show mpls l2transport vc 1 de
```

```
Local interface: GigabitEthernet7/47, Ethernet
Destination address: 75.0.80.1, VC ID: 1
VC status: receive UP, send DOWN
VC type: receive 5, send 5
Tunnel label: not ready, destination not in LFIB
Output interface: unknown, imposed label stack {}
MPLS VC label: local 10579, remote 10038
Linecard VC statistics:
packet totals: receive: 0 send: 0
byte totals: receive: 0 send: 0
packet drops: receive: 0 send: 0
Control flags:
receive 1, send: 31
```

!

To verify and display the configuration of Layer 2 VLAN transport over MPLS tunnels, perform the following:

- To display a single line for each VLAN, naming the VLAN, status, and ports, enter the **show vlan brief** command.

```
Router# show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	
2	VLAN0002	active	Gi1/4
1002	fddi-default	act/unsup	
1003	token-ring-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trnet-default	act/unsup	

- To verify that the PE router endpoints have discovered each other, enter the **show mpls ldp discovery** command. When an PE router receives an LDP Hello message from another PE router, it considers that router and the specified label space to be “discovered.”

```
Router# show mpls ldp discovery
```

```
Local LDP Identifier:
13.13.13.13:0
Discovery Sources:
Interfaces:
GE-WAN3/3 (ldp): xmit/rcv
LDP Id: 12.12.12.12:0
Targeted Hellos:
13.13.13.13 -> 11.11.11.11 (ldp): active/passive, xmit/rcv
LDP Id: 11.11.11.11:0
```

- To verify that the label distribution session has been established, enter the **show mpls ldp neighbor** command. The third line of the output shows that the state of the LDP session is operational and shows that messages are being sent and received.

```
Router# show mpls ldp neighbor
```

```
Peer LDP Ident: 12.12.12.12:0; Local LDP Ident 13.13.13.13:0
TCP connection: 12.12.12.12.646 - 13.13.13.13.11010
State: Oper; Msgs sent/rcvd: 1715/1706; Downstream
Up time: 1d00h
LDP discovery sources:
GE-WAN3/3, Src IP addr: 34.0.0.2
Addresses bound to peer LDP Ident:
23.2.1.14      37.0.0.2      12.12.12.12      34.0.0.2
99.0.0.1
```



```

Peer LDP Ident: 11.11.11.11:0; Local LDP Ident 13.13.13.13:0
TCP connection: 11.11.11.11.646 - 13.13.13.13.11013
State: Oper; Msgs sent/rcvd: 1724/1730; Downstream
Up time: 1d00h
LDP discovery sources:
  Targeted Hello 13.13.13.13 -> 11.11.11.11, active, passive
Addresses bound to peer LDP Ident:
  11.11.11.11      37.0.0.1      23.2.1.13

```

- To verify that the label forwarding table is built correctly, enter the **show mpls forwarding-table** command.

```

Router# show mpls forwarding-table
Local   Outgoing   Prefix           Bytes tag   Outgoing     Next Hop
tag     tag or VC    or Tunnel Id     switched    interface
16      Untagged     223.255.254.254/32 \
                                0           Gi2/1       23.2.0.1
20      Untagged     12ckt (2)        55146580    V12         point2point
24      Pop tag      37.0.0.0/8       0           GE3/3       34.0.0.2
25      17          11.11.11.11/32   0           GE3/3       34.0.0.2
26      Pop tag      12.12.12.12/32   0           GE3/3       34.0.0.2

```

- The output displays the following data:
 - Local tag—Label assigned by this router.
 - Outgoing tag or VC—Label assigned by next hop.
 - Prefix or Tunnel Id—Address or tunnel to which packets with this label are going.
 - Bytes tag switched— Number of bytes switched out with this incoming label.
 - Outgoing interface—Interface through which packets with this label are sent.
 - Next Hop—IP address of neighbor that assigned the outgoing label.
- To display the state of the currently routed VCs, enter the **show mpls l2transport vc** command:

```

Router# show mpls l2transport vc

Local intf   Local circuit      Dest address       VC ID      Status
-----
V12          Eth VLAN 2        11.11.11.11       2          UP

```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

Participate in the [Technical Documentation Ideas forum](#)

1

Virtual Private LAN Services (VPLS)

- [Prerequisites for VPLS, page 1-2](#)
- [Restrictions for VPLS, page 1-2](#)
- [Information About VPLS, page 1-2](#)
- [Default Settings for VPLS, page 1-7](#)
- [How to Configure VPLS, page 1-7](#)
- [How to Configure VPLS BGP-Based Autodiscovery, page 1-18](#)
- [Configuration Examples for VPLS, page 1-20](#)
- [Configuration Examples for VPLS BGP-Based Autodiscovery, page 1-22](#)
- [Additional References for Virtual Private LAN Services \(VPLS\), page 1-24](#)
- [Feature Information for Virtual Private LAN Services \(VPLS\), page 1-24](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11845/prod_command_reference_list.html
- Cisco IOS Release 15.0SY supports only Ethernet interfaces. Cisco IOS Release 15.0SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for VPLS

Before you configure VPLS, ensure that the network is configured as follows:

- Configure IP routing in the core so that the PE routers can reach each other via IP.
- Configure MPLS in the core so that a label switched path (LSP) exists between the PE routers.
- Configure a loopback interface for originating and terminating Layer 2 traffic. Make sure the PE routers can access the other router's loopback interface. Note that the loopback interface is not needed in all cases. For example, tunnel selection does not need a loopback interface when VPLS is directly mapped to a TE tunnel.

VPLS configuration requires you to identify peer PE routers and to attach Layer 2 circuits to the VPLS at each PE router.

Restrictions for VPLS

- With a Supervisor Engine 2T, Layer 2 protocol tunneling is not supported with VPLS ([CSCue45974](#)).
- Split horizon is the default configuration to avoid broadcast packet looping and to isolate Layer 2 traffic. Split horizon prevents packets received from an emulated VC from being forwarded into another emulated VC. This technique is important for creating loop-free paths in a full-meshed network.
- Supported maximum values:
 - Total number of VFIs: 4,096 (4K)
 - Maximum combined number of edge and the core peer PEs per VFI:
 - VPLS: 250
 - H-VPLS 500
 - Total number of VC: 12,288 (12K)
- No software-based data plane is supported.
- Load sharing and failover on redundant CE-PE links are not supported.
- The addition or removal of MAC addresses with Label Distribution Protocol (LDP) is not supported.
- The virtual forwarding instance (VFI) is supported only with the **interface vlan** command.


Information About VPLS

- [VPLS Overview, page 1-3](#)
- [Full-Mesh Configuration, page 1-3](#)
- [H-VPLS, page 1-4](#)
- [VPLS BGP Based Autodiscovery, page 1-4](#)
- [Supported Features, page 1-5](#)

VPLS Overview

VPLS (Virtual Private LAN Service) enables enterprises to link together their Ethernet-based LANs from multiple sites via the infrastructure provided by their service provider. From the enterprise perspective, the service provider's public network looks like one giant Ethernet LAN. For the service provider, VPLS provides an opportunity to deploy another revenue-generating service on top of their existing network without major capital expenditures. Operators can extend the operational life of equipment in their network.

Virtual Private LAN Services (VPLS) uses the provider core to join multiple attachment circuits together to simulate a virtual bridge that connects the multiple attachment circuits together. From a customer point of view, there is no topology for VPLS. All of the CE devices appear to connect to a logical bridge emulated by the provider core.



Full-Mesh Configuration

The full-mesh configuration requires a full mesh of tunnel label switched paths (LSPs) between all the PEs that participate in the VPLS. With full-mesh, signaling overhead and packet replication requirements for each provisioned VC on a PE can be high.

You set up a VPLS by first creating a virtual forwarding instance (VFI) on each participating PE router. The VFI specifies the VPN ID of a VPLS domain, the addresses of other PE routers in the domain, and the type of tunnel signaling and encapsulation mechanism for each peer PE router.

The set of VFIs formed by the interconnection of the emulated VCs is called a *VPLS instance*; it is the VPLS instance that forms the logic bridge over a packet switched network. The VPLS instance is assigned a unique VPN ID.

The PE routers use the VFI to establish a full-mesh LSP of emulated VCs to all the other PE routers in the VPLS instance. PE routers obtain the membership of a VPLS instance through static configuration using the Cisco IOS CLI.

The full-mesh configuration allows the PE router to maintain a single broadcast domain. Thus, when the PE router receives a broadcast, multicast, or unknown unicast packet on an attachment circuit, it sends the packet out on all other attachment circuits and emulated circuits to all other CE devices participating in that VPLS instance. The CE devices see the VPLS instance as an emulated LAN.

To avoid the problem of a packet looping in the provider core, the PE devices enforce a "split-horizon" principle for the emulated VCs. That means if a packet is received on an emulated VC, it is not forwarded on any other emulated VC.

After the VFI has been defined, it needs to be bound to an attachment circuit to the CE device.

The packet forwarding decision is made by looking up the Layer 2 virtual forwarding instance (VFI) of a particular VPLS domain.

A VPLS instance on a particular PE router receives Ethernet frames that enter on specific physical or logical ports and populates a MAC table similarly to how an Ethernet switch works. The PE router can use the MAC address to switch those frames into the appropriate LSP for delivery to the another PE router at a remote site.

If the MAC address is not in the MAC address table, the PE router replicates the Ethernet frame and floods it to all logical ports associated with that VPLS instance, except the ingress port where it just entered. The PE router updates the MAC table as it receives packets on specific ports and removes addresses not used for specific periods.

H-VPLS

Hierarchical VPLS (H-VPLS) reduces both signaling and replication overhead by using both full-mesh as well as hub and spoke configurations. Hub and spoke configurations operate with split horizon to allow packets to be switched between pseudo-wires (PWs), effectively reducing the number of PWs between PEs.



Note

Split horizon is the default configuration to avoid broadcast packet looping. To avoid looping when using the **no-split-horizon** keyword, be very mindful of your network configuration.

VPLS BGP Based Autodiscovery

VPLS Autodiscovery enables each Virtual Private LAN Service (VPLS) provider edge (PE) device to discover other PE devices that are part of the same VPLS domain. VPLS Autodiscovery also tracks PE devices when they are added to or removed from a VPLS domain. As a result, with VPLS Autodiscovery enabled, you no longer need to manually configure a VPLS domain and maintain the configuration when a PE device is added or deleted. VPLS Autodiscovery uses the Border Gateway Protocol (BGP) to discover VPLS members and set up and tear down pseudowires in a VPLS domain.

BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path. When BGP distributes the endpoint provisioning information in an update message to all its BGP neighbors, this endpoint information is used to configure a pseudowire mesh to support L2VPN-based services.

The BGP autodiscovery mechanism facilitates the configuration of L2VPN services, which are an integral part of the VPLS feature. VPLS enables flexibility in deploying services by connecting geographically dispersed sites as a large LAN over high-speed Ethernet in a robust and scalable IP Multiprotocol Label Switching (MPLS) network.

The VPLS BGP based Autodiscovery feature was introduced on Cisco Catalyst 6500 series Switches starting with the 15.1(1)SY release.

The VPLS BGP based Autodiscovery High Availability feature was introduced on Catalyst 6500 series Switches starting with the Release 15.5(1)SY2. This feature provides Stateful Switchover (SSO) support for VPLS BGP based Autodiscovery feature. SSO minimizes the amount of time a network is unavailable to its users following a RP switchover.

Supported Features

- [Multipoint-to-Multipoint Support, page 1-5](#)
- [Non-Transparent Operation, page 1-5](#)
- [Circuit Multiplexing, page 1-5](#)
- [MAC-Address Learning Forwarding and Aging, page 1-5](#)
- [Jumbo Frame Support, page 1-5](#)
- [Q-in-Q Support and Q-in-Q to EoMPLS Support, page 1-6](#)
- [VPLS Services, page 1-6](#)

Multipoint-to-Multipoint Support

Two or more devices are associated over the core network. No one device is designated as the Root node, but all devices are treated as Root nodes. All frames can be exchanged directly between nodes.

Non-Transparent Operation

A virtual Ethernet connection (VEC) can be transparent or non-transparent with respect to Ethernet PDUs (that is, BPDUs). The purpose of VEC non-transparency is to allow the end user to have a Frame Relay-type service between Layer 3 devices.

Circuit Multiplexing

Circuit Multiplexing allows a node to participate in multiple services over a single Ethernet connection. By participating in multiple services, the Ethernet connection is attached to multiple logical networks. Some examples of possible service offerings are VPN services between sites, Internet services, and third-party connectivity for intercompany communications.

MAC-Address Learning Forwarding and Aging

PEs must learn remote MAC addresses and directly attached MAC addresses on customer facing ports. MAC address learning accomplishes this by deriving topology and forwarding information from packets originating at customer sites. A timer is associated with stored MAC addresses. After the timer expires, the entry is removed from the table.

Jumbo Frame Support

Jumbo frame support provides support for frame sizes between 1548 through 9216 bytes. You use the CLI to establish the jumbo frame size for any value specified in the above range. The default value is 1500 bytes in any Layer 2/VLAN interface. You can configure jumbo frame support on a per-interface basis.

Q-in-Q Support and Q-in-Q to EoMPLS Support

With 802.1Q tunneling (Q-in-Q), the CE issues VLAN-tagged packets and the VPLS forwards the packets to a far-end CE. Q-in-Q refers to the fact that one or more 802.1Q tags may be located in a packet within the interior of the network. As packets are received from a CE device, an additional VLAN tag is added to incoming Ethernet packets to segregate traffic from different CE devices. Untagged packets originating from the CE use a single tag within the interior of the VLAN switched network, while previously tagged packets originating from the CE use two or more tags.

VPLS Services

- [Transparent LAN Service, page 1-6](#)
- [Ethernet Virtual Connection Service, page 1-6](#)

Transparent LAN Service

Transparent LAN Service (TLS) is an extension to the point-to-point port-based EoMPLS, used to provide bridging protocol transparency (for example, bridge protocol data units [BPDUs]) and VLAN values. Bridges see this service as an Ethernet segment. With TLS, the PE router forwards all Ethernet packets received from the customer-facing interface (including tagged, untagged, and BPDUs) as follows:

- To a local Ethernet interface or an emulated VC if the destination MAC address is found in the Layer 2 forwarding table.
- To all other local Ethernet interfaces and emulated VCs belonging to the same VPLS domain if the destination MAC address is a multicast or broadcast address or if the destination MAC address is not found in the Layer 2 forwarding table.

**Note**

With a Supervisor Engine 2T, Layer 2 protocol tunneling is not supported with VPLS, which prevents use of the Cisco Discovery Protocol (CDP), the VLAN Trunking Protocol (VTP), and the Spanning-Tree Protocol (STP) over VPLS ([CSCue45974](#)).

Ethernet Virtual Connection Service

Ethernet Virtual Connection Service (EVCS) is an extension to the point-to-point VLAN-based EoMPLS that allows routers to reach multiple intranet and extranet locations from a single physical port. Routers see subinterfaces through which they access other routers. With EVCS, the PE router forwards all Ethernet packets with a particular VLAN tag received from the customer-facing interface (excluding BPDUs) as follows:

- To a local Ethernet interface or to an emulated VC if the destination MAC address is found in the Layer 2 forwarding table.
- To all other local Ethernet interfaces and emulated VCs belonging to the same VPLS domain if the destination MAC address is a multicast or broadcast address or if the destination MAC address is not found in the Layer 2 forwarding table.

**Note**

Because it has only local significance, the demultiplexing VLAN tag that identifies a VPLS domain is removed before forwarding the packet to the outgoing Ethernet interfaces or emulated VCs.

Default Settings for VPLS

None.

How to Configure VPLS

- [Configuring PE Layer 2 Interfaces to CEs, page 1-7](#)
- [Configuring Layer 2 VLAN Instances on a PE, page 1-11](#)
- [Configuring MPLS in the PE, page 1-11](#)
- [Configuring the VFI in the PE, page 1-12](#)
- [Associating the Attachment Circuit with the VSI at the PE, page 1-13](#)
- [H-VPLS with MPLS Edge, page 1-14](#)
- [VPLS Integrated Routing and Bridging, page 1-17](#)



Note

- Use the procedures in the [QoS chapters](#) to configure QoS for VPLS traffic.
- Provisioning a VPLS link involves provisioning the associated attachment circuit and the VFI on the PE.

Configuring PE Layer 2 Interfaces to CEs

- [Configuring 802.1Q Trunks for Tagged Traffic from a CE, page 1-7](#)
- [Configuring 802.1Q Access Ports for Untagged Traffic from CE, page 1-8](#)
- [Configuring Q-in-Q to Place All VLANs into a Single VPLS Instance, page 1-9](#)



Note

- It is important to define the trunk VLANs; use the **switchport trunk allow vlan** command as shown in the first example below.
- You must configure the Layer 2 interface as a switchport for local bridging. You have the option of selecting tagged or untagged traffic from the CE device.

Configuring 802.1Q Trunks for Tagged Traffic from a CE



Note

When EVCS is configured, the PE router forwards all Ethernet packets with a particular VLAN tag to a local Ethernet interface or emulated VC if the destination MAC address is found in the Layer 2 forwarding table.

Command or Action	Purpose

Step 1	Router(config)# interface <i>type number</i>	Selects an interface to configure.
Step 2	Router(config)# no ip address <i>ip_address mask</i> [secondary]	Disables IP processing and enters interface configuration mode.
Step 3	Router(config-if)# switchport	Modifies the switching characteristics of the Layer 2-switched interface.
Step 4	Router(config-if)# switchport trunk encapsulation dot1q	Sets the switch port encapsulation format to 802.1Q.
Step 5	Router(config-if)# switchport trunk allow vlan <i>vlan_ID</i>	Sets the list of allowed VLANs.
Step 6	Router(config-if)# switchport mode trunk	Sets the interface to a trunking VLAN Layer 2 interface.

This example shows how to configure the tagged traffic.

```
Router(config)# interface GigabitEthernet4/4
Router(config)# no ip address
Router(config-if)# switchport
Router(config-if)# switchport trunk encapsulation dot1q
Router(config-if)# switchport trunk allow vlan 501
Router(config-if)# switchport mode trunk
```

This example shows how to use the **show run interface** command to verify the configuration.

```
Router# show run interface GigabitEthernet4/4
Building configuration...

Current configuration : 212 bytes
!
interface GigabitEthernet4/4
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 500-1999
 switchport mode trunk
end
```

Configuring 802.1Q Access Ports for Untagged Traffic from CE

	Command or Action	Purpose
Step 1	Router(config)# interface <i>type number</i>	Selects an interface to configure.
Step 2	Router(config)# no ip address <i>ip_address mask</i> [secondary]	Disables IP processing and enters interface configuration mode.
Step 3	Router(config-if)# speed [1000 nonegotiate]	Sets the port speed for an Ethernet interface; enables or disables the link negotiation protocol on the Gigabit Ethernet ports.
Step 4	Router(config-if)# switchport	Modifies the switching characteristics of the Layer 2-switched interface.
Step 5	Router(config-if)# switchport mode access	Sets the interface type to nontrunking, nontagged single VLAN Layer 2 interface.
Step 6	Router(config-if)# switchport access vlan <i>vlan_id</i>	Sets the VLAN when the interface is in Access mode.

This example shows how to configure the untagged traffic.

```
Router(config)# interface GigabitEthernet4/4
Router(config)# no ip address
Router(config-if)# speed nonegotiate
Router(config-if)# switchport
Router(config-if)# switchport mode access
Router(config-if)# switchport access vlan 501
```

This example shows how to use the **show run interface** command to verify the configuration.

```
Router# show run interface GigabitEthernet4/4
Building configuration...

Current configuration : 212 bytes
!
interface GigabitEthernet4/4
  speed nonegotiate
  switchport
  switchport mode access
  switchport access vlan 501
end
```

Configuring Q-in-Q to Place All VLANs into a Single VPLS Instance



Note

When TLS is configured, the PE router forwards all Ethernet packets received from the CE device to all local Ethernet interfaces and emulated VCs belonging to the same VPLS domain if the MAC address is not found in the Layer 2 forwarding table.

	Command or Action	Purpose
Step 1	Router(config)# interface <i>type number</i>	Selects an interface to configure.
Step 2	Router(config)# no ip address <i>ip_address mask</i> [secondary]	Disables IP processing and enters interface configuration mode.
Step 3	Router(config-if)# speed [1000 nonegotiate]	Sets the port speed for an Ethernet interface; enables or disables the link negotiation protocol on the Gigabit Ethernet ports.
Step 4	Router(config-if)# switchport	Modifies the switching characteristics of the Layer 2-switched interface.
Step 5	Router(config-if)# switchport access vlan <i>vlan_id</i>	Sets the VLAN when the interface is in Access mode.
Step 6	Router(config-if)# switchport mode dot1q-tunnel	Sets the interface as an 802.1Q tunnel port.
Step 7	Router(config-if)# l2protocol-tunnel [cdp stp vtp]	Enables protocol tunneling on an interface.

This example shows how to configure the tagged traffic.

```
Router(config)# interface GigabitEthernet4/4
Router(config)# no ip address
Router(config-if)# speed nonegotiate
Router(config-if)# switchport
Router(config-if)# switchport access VLAN 501
Router(config-if)# switchport mode dot1q-tunnel
Router(config-if)# l2protocol-tunnel cdp
```

This example shows how to use the **show run interface** command to verify the configuration.

```
Router# show run interface GigabitEthernet4/4
Building configuration...
```

```
Current configuration : 212 bytes
!
interface GigabitEthernet4/4
 no ip address
 speed nonegotiate
 switchport
 switchport access vlan 501
 switchport mode dot1q-tunnel
 l2protocol-tunnel cdp
end
```

Use the **show spanning-tree vlan** command to verify the port is not in a blocked state.

```
Router# show spanning-tree vlan 501
```

```
VLAN0501
Spanning tree enabled protocol ieee
  Root ID    Priority    33269
             Address     0001.6446.2300
             This bridge is the root
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    33269  (priority 32768 sys-id-ext 501)
             Address     0001.6446.2300
             Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  0
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Gi4/4	Desg	FWD	4	128.388	P2p

Use the **show vlan id** command to verify that a specific port is configured to send and receive a specific VLAN's traffic.

```
Router# show vlan id 501
```

VLAN Name	Status	Ports
501 VLAN0501	active	Gi4/4

VLAN Type	SAID	MTU	Parent	RingNo	BridgeNo	Stp	BrdgMode	Trans1	Trans2
501	enet	100501	1500	-	-	-	-	0	0

```
Remote SPAN VLAN
-----
Disabled
```

Primary	Secondary	Type	Ports
-----	-----	-----	-----

Configuring Layer 2 VLAN Instances on a PE

Configuring the Layer 2 VLAN interface on the PE enables the Layer 2 VLAN instance on the PE router to the VLAN database to set up the mapping between the VPLS and VLANs.

	Command or Action	Purpose
Step 1	vlan <i>vlan-id</i> Router(config)# vlan 809	Configures a specific virtual LAN (VLAN).
Step 2	interface vlan <i>vlan-id</i> Router(config)# interface vlan 501	Configures an interface on the VLAN.

This is an example of configuring a Layer 2 VLAN instance.

```
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# vlan 501
Router(config)# interface vlan 501
Router(config-if)#
```

Use the **show interfaces vlan** command to verify the VLAN is in the up state (example not shown).

Configuring MPLS in the PE

To configure MPLS in the PE, you must provide the required MPLS parameters.



Note

Before configuring MPLS, ensure that you have IP connectivity between all PEs by configuring Interior Gateway Protocol (IGP) (Open Shortes Path First [OSPF] or Intermediate System to Intermediate System [IS-IS]) between the PEs.

	Command or Action	Purpose
Step 1	enable Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Router# configure terminal	Enters global configuration mode.
Step 3	mpls label protocol {ldp tdp} Router(config)# mpls label protocol ldp	Specifies the default Label Distribution Protocol for a platform.
Step 4	mpls ldp logging neighbor-changes Router(config)# mpls ldp logging neighbor-changes	(Optional) Determines logging neighbor changes.
Step 5	tag-switching tdp discovery {hello directed hello} {holdtime interval} seconds Router(config)# tag-switching tdp discovery hello holdtime 5	Configures the interval between transmission of LDP (TDP) discovery hello messages, or the hold time for a LDP transport connection
Step 6	tag-switching tdp router-id Loopback0 force Router(config)# tag-switching tdp router-id Loopback0 force	Configures MPLS.

This example shows global MPLS configuration.

```
Router(config)# mpls label protocol ldp
Router(config)# tag-switching tdp discovery directed hello
Router(config)# tag-switching tdp router-id Loopback0 force
```

Use the **show ip cef** command to verify that the LDP label is assigned.

```
Router# show ip cef 192.168.17.7
192.168.17.7/32, version 272, epoch 0, cached adjacency to POS4/1
0 packets, 0 bytes
  tag information set
    local tag: 8149
    fast tag rewrite with P04/1, point2point, tags imposed: {4017}
  via 11.3.1.4, POS4/1, 283 dependencies
    next hop 11.3.1.4, POS4/1
    valid cached adjacency
    tag rewrite with P04/1, point2point, tags imposed: {4017}
```

Configuring the VFI in the PE

The virtual switch instance (VFI) specifies the VPN ID of a VPLS domain, the addresses of other PE routers in this domain, and the type of tunnel signaling and encapsulation mechanism for each peer. (This is where you create the VSI and associated VCs.) Configure a VFI as follows:



Note

Only MPLS encapsulation is supported.

	Command or Action	Purpose
Step 1	<pre>l2 vfi name manual Router(config)# l2 vfi vfi17 manual</pre>	Enables the Layer 2 VFI manual configuration mode.
Step 2	<pre>vpn id vpn-id Router(config-vfi)# vpn id 17</pre>	Configures a VPN ID for a VPLS domain. The emulated VCs bound to this Layer 2 VRF use this VPN ID for signaling.
Step 3	<pre>neighbor remote router id {encapsulation mpls} [no-split-horizon] Router(config-vfi)# neighbor 1.5.1.1 encapsulation mpls</pre>	Specifies the remote peering router ID and the tunnel encapsulation type or the pseudo-wire property to be used to set up the emulated VC. Note Split horizon is the default configuration to avoid broadcast packet looping and to isolate Layer 2 traffic. Use the no-split-horizon keyword to disable split horizon and to configure multiple VCs per spoke into the same VFI.
Step 4	<pre>shutdown Router(config-vfi)# shutdown</pre>	Disconnects all emulated VCs previously established under the Layer 2 VFI and prevents the establishment of new attachment circuits. Note It does not prevent the establishment of new attachment circuits configured with the Layer 2 VFI using CLI.

The following example shows a VFI configuration.

```
Router(config)# l2 vfi VPLSA manual
```

```

Router(config-vfi)# vpn id 100
Router(config-vfi)# neighbor 11.11.11.11 encapsulation mpls
Router(config-vfi)# neighbor 33.33.33.33 encapsulation mpls
Router(config-vfi)# neighbor 44.44.44.44 encapsulation mpls

```

The following example shows a VFI configuration for hub and spoke.

```

Router(config)# 12 vfi VPLSA manual
Router(config-vfi)# vpn id 100
Router(config-vfi)# neighbor 9.9.9.9 encapsulation mpls
Router(config-vfi)# neighbor 12.12.12.12 encapsulation mpls
Router(config-vfi)# neighbor 33.33.33.33 encapsulation mpls no-split-horizon

```

The **show mpls 12transport vc** command displays various information related to PE1.



Note

The **show mpls 12transport vc [detail]** command is also available to show detailed information about the VCs on a PE router as in the following example.

```
VPLS-PE2# show mpls 12transport vc 201
```

Local intf	Local circuit	Dest address	VC ID	Status
VFI test1	VFI	153.1.0.1	201	UP
VFI test1	VFI	153.3.0.1	201	UP
VFI test1	VFI	153.4.0.1	201	UP



Note

The VC ID in the output represents the VPN ID; the VC is identified by the combination of the Dest address and the VC ID as in the example below.

The **show vfi vfi name** command shows VFI status.

```

nPE-3# show vfi VPLS-2
VFI name: VPLS-2, state: up
  Local attachment circuits:
    Vlan2
  Neighbors connected via pseudowires:
Peer Address    VC ID    Split-horizon
1.1.1.1         2        Y
1.1.1.2         2        Y
2.2.2.3         2        N

```

Associating the Attachment Circuit with the VSI at the PE

After defining the VFI, you must bind it to one or more attachment circuits (interfaces, subinterfaces, or virtual circuits).

	Command or Action	Purpose
Step 1	interface vlan <i>vlan-id</i> Router(config-if)# interface vlan 100	Creates or accesses a dynamic switched virtual interface (SVI).
Step 2	no ip address Router(config-if)# no ip address	Disables IP processing. (You configure a Layer 3 interface for the VLAN if you configure an IP address.)
Step 3	xconnect vfi <i>vfi name</i> Router(config-if)# xconnect vfi vfi16	Specifies the Layer 2 VFI that you are binding to the VLAN port.

This example shows an interface VLAN configuration.

```
Router(config-if)# interface vlan 100
Router(config-if)# no ip address
Router(config-if)# xconnect vfi VPLS_501
```

Use the **show vfi** command for VFI status.

```
Router# show vfi VPLS_501
VFI name: VPLS_501, state: up
Local attachment circuits:
  vlan 100
Neighbors connected via pseudowires:
  192.168.11.1 192.168.12.2 192.168.13.3 192.168.16.6
  192.168.17.7
```

H-VPLS with MPLS Edge

- [Overview, page 1-14](#)
- [Configuration on PE1, page 1-15](#)
- [Configuration on PE2, page 1-15](#)
- [Configuration on PE3, page 1-16](#)


Overview

The Hierarchical VPLS model comprises hub and spoke and full-mesh networks. In a full-mesh configuration, each PE router creates a multipoint-to-multipoint forwarding relationship with all other PE routers in the VPLS domain using VFIs.

In the hub and spoke configuration, a PE router can operate in a non-split-horizon mode that allows inter-VC connectivity without the requirement to add a Layer 2 port in the VLAN.

In the example below, the VLANs on CE1, CE2, CE3, and CE4 (in red) connect through a full-mesh network. The VLANs on CE2, CE5, and ISP POP connect through a hub and spoke network where the ISP POP is the hub and CE2 and CE5 are the spokes. [Figure 1-1](#) shows the configuration example.

Figure 1-1 H-VPLS Configuration



Configuration on PE1

- [Configuring VSIs and VCs, page 1-15](#)
- [Configuring the CE Device Interface, page 1-15](#)
- [Associating the Attachment Circuit with the VFI, page 1-15](#)

Configuring VSIs and VCs

This sample configuration shows the creation of the virtual switch instances (VSIs) and associated VCs. Note that the VCs in green require the **no-split-horizon** keyword. The **no-split-horizon** command disables the default Layer 2 split horizon in the data path.

```
12 vfi Internet manual
  vpn id 100
  neighbor 120.0.0.3 encapsulation mpls no-split-horizon
  neighbor 162.0.0.2 encapsulation mpls no-split-horizon

12 vfi PE1-VPLS-A manual
  vpn id 200
  neighbor 120.0.0.3 encapsulation mpls
  neighbor 162.0.0.2 encapsulation mpls

interface Loopback 0
  ip address 20.0.0.1 255.255.255.255
```

Configuring the CE Device Interface

This sample configuration shows the CE device interface (there can be multiple Layer 2 interfaces in a VLAN).

```
interface GigEthernet1/1
  switchport
  switchport mode trunk
  switchport trunk encap dot1q
  switchport trunk allow vlan 1001,1002-1005
```

Associating the Attachment Circuit with the VFI

This sample configuration shows how the attachment circuit (VLAN) is associated with the VFI.

```
interface Vlan 1001
  xconnect vfi Internet

interface FastEthernet2/1
  switchport
  switchport mode trunk
  switchport trunk encap dot1q
  switchport trunk allow vlan 211,1002-1005

interface Vlan 211
  xconnect vfi PE1-VPLS-A
```

Configuration on PE2

- [Configuring VSIs and VCs, page 1-16](#)
- [Configuring the CE Device Interface, page 1-16](#)
- [Associating the Attachment Circuit with the VFI, page 1-16](#)

Configuring VSIs and VCs

This sample configuration shows the creation of the virtual switch instances (VSIs) and associated VCs.

```
12 vfi Internet manual
  vpn id 100
  neighbor 20.0.0.1 encapsulation mpls

12 vfi PE2-VPLS-A manual
  vpn id 200:1
  neighbor 120.0.0.3 encapsulation mpls
  neighbor 20.0.0.1 encapsulation mpls

interface Loopback 0
  ip address 162.0.0.2 255.255.255.255
```

Configuring the CE Device Interface

This sample configuration shows the CE device interface (there can be multiple Layer 2 interfaces in a VLAN).

```
interface GigEthernet2/1
  switchport
  switchport mode trunk
  switchport trunk encap dot1q
  switchport trunk allow vlan 211,1001,1002-1005
```

Associating the Attachment Circuit with the VFI

This sample configuration shows how the attachment circuit (VLAN) is associated with the VFI.

```
interface Vlan 1001
  xconnect vfi Internet

interface Vlan 211
  xconnect vfi PE2-VPLS-A
```

Configuration on PE3

- [Configuring VSIs and VCs, page 1-16](#)
- [Configuring the CE Device Interface, page 1-17](#)
- [Configuring the Attachment Circuits, page 1-17](#)
- [Configuring Port-based EoMPLS on the uPE Device, page 1-17](#)

Configuring VSIs and VCs

This sample configuration shows the creation of the virtual switch instances (VSIs) and associated VCs.

```
12 vfi Internet manual
  vpn id 100
  neighbor 20.0.0.1 encapsulation mpls
  neighbor 162.0.0.2 encapsulation mpls
  neighbor 30.0.0.1 encapsulation mpls no-split horizon

12 vfi PE3-VPLS-A manual
  vpn id 200
  neighbor 162.0.0.2 encapsulation mpls
  neighbor 20.0.0.1 encapsulation mpls
```

```
interface Loopback 0
 ip address 120.0.0.3 255.255.255.255
```

Configuring the CE Device Interface

This sample configuration shows the CE device interface (there can be multiple Layer 2 interfaces in a VLAN).

```
interface GigEthernet6/1
 switchport
 switchport mode trunk
 switchport trunk encap dot1q
 switchport trunk allow vlan 211
```

Configuring the Attachment Circuits

This sample configuration shows the attachment circuits.

```
interface Vlan 1001
 xconnect vfi Internet

interface Vlan 211
 xconnect vfi PE3-VPLS-A
```

Configuring Port-based EoMPLS on the uPE Device

This sample configuration shows port-based EoMPLS on the uPE device.

```
interface GigEthernet 1/1
 xconnect 120.0.0.3 100 encapsulation mpls
```

VPLS Integrated Routing and Bridging

VPLS integrated routing and bridging can route Layer 3 traffic as well as switch Layer 2 frames for pseudowire connections between provider edge (PE) devices using Virtual Private LAN Services (VPLS) multipoint PE. The ability to route frames to and from these interfaces supports termination of a pseudowire into a Layer 3 network (VPN or global) on the same switch, or to tunnel Layer 3 frames over a Layer 2 tunnel (VPLS).



Note

- VPLS integrated routing and bridging is also known as routed pseudowire and routed VPLS.
- VPLS integrated routing and bridging does not support multicast routing.

To configure routing support for the pseudowire, configure an IP address and other Layer 3 features for the Layer 3 domain (VPN or global) in the virtual LAN (VLAN) interface configuration.

- The following example assigns the IP address 10.10.10.1 to the VLAN 100 interface. (Layer 2 forwarding is defined by the VFI VFI100.)

```
interface vlan 100
 xconnect vfi VFI100
 ip address 10.10.10.1 255.255.255.0
```

- The following example assigns an IP address 20.20.20.1 of the VPN domain VFI200. (Layer 2 forwarding is defined by the VFI VFI200.)

```
interface vlan 200
```

```
xconnect vfi VFI200
ip vrf forwarding VFI200
ip address 20.20.20.1 255.255.255.0
```

How to Configure VPLS BGP-Based Autodiscovery

- [Enabling VPLS BGP-based Autodiscovery, page 1-18](#)
- [Configuring BGP to enable VPLS Autodiscovery, page 1-18](#)

Enabling VPLS BGP-based Autodiscovery

Perform this task to enable Virtual Private LAN Service (VPLS) PE devices to discover other PE devices that are part of the same VPLS domain.

	Command or Action	Purpose
Step 1	Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	Device# configure terminal	Enters global configuration mode.
Step 3	Device(config)# l2 vfi vfi-name autodiscovery	Enables VPLS Autodiscovery on a PE device and enters L2 VFI configuration mode.
Step 4	Device(config-vfi)# vpn id vpn-id	Configures VPN ID for the VPLS domain.
Step 5	Device(config-vfi)# end	Returns to privileged EXEC mode.

Configuring BGP to enable VPLS Autodiscovery

The Border Gateway Protocol (BGP) Layer 2 VPN (L2VPN) address family supports a separate L2VPN Routing Information Base (RIB) that contains endpoint provisioning information for Virtual Private LAN Service (VPLS) Autodiscovery. BGP learns the endpoint provisioning information from the L2VPN database, which is updated each time a Layer 2 virtual forwarding instance (VFI) is configured. When BGP distributes the endpoint provisioning information in an update message to all its BGP neighbors, the endpoint information is used to configure a pseudowire mesh to support L2VPN-based services.


	Command or Action	Purpose
Step 1	Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	Device# configure terminal	Enters global configuration mode.
Step 3	Device(config)# router bgp <i>autonomous-system-number</i>	Enters router configuration mode for the specified routing process.

Step 4	Device(config-router)# no bgp default ipv4-unicast	Disables the IPv4 unicast address family for the BGP routing process.  Note Routing information for the IPv4 unicast address family is advertised by default for each BGP routing session configured using the neighbor remote-as router configuration command unless you configure the no bgp default ipv4-unicast router configuration command before configuring the neighbor remote-as command. Existing neighbor configurations are not affected.
Step 5	Device(config-router)# bgp log-neighbor-changes	Enables logging of BGP neighbor resets.
Step 6	Device(config-router)# neighbor {ip-address peer-group-name} remote-as autonomous-system-number	Adds the IP address or peer group name of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local device. • If the autonomous-system-number argument matches the autonomous system number specified in the router bgp command, the neighbor is an internal neighbor. • If the autonomous-system-number argument does not match the autonomous system number specified in the router bgp command, the neighbor is an external neighbor.
Step 7	Device(config-router)# neighbor {ip-address peer-group-name} update-source interface-type interface-number	(Optional) Configures a device to select a specific source or interface to receive routing table updates.
Step 8	Repeat Steps 6 and 7 to configure other BGP neighbors.	
Step 9	Device(config-router)# address-family l2vpn vpls number	Specifies the L2VPN address family and enters address family configuration mode. The optional vpls keyword specifies that the VPLS endpoint provisioning information is to be distributed to BGP peers.
Step 10	Device(config-router-af)# neighbor {ip-address peer-group-name} activate	Enables the exchange of information with a BGP neighbor.
Step 11	Device(config-router-af)# neighbor {ip-address peer-group-name} send-community {both standard extended}	Specifies that a communities attribute should be sent to a BGP neighbor.
Step 12	Repeat Steps 10 and 11 to activate other BGP neighbors under an L2VPN address family.	
Step 13	Device(config-router-af)# exit-address-family	Exits address family configuration mode and returns to router configuration mode.
Step 14	Device(config-router)# end	Exits router configuration mode and returns to privileged EXEC mode.

Configuration Examples for VPLS

In a full-mesh configuration, each PE router creates a multipoint-to-multipoint forwarding relationship with all other PE routers in the VPLS domain using a VFI. An Ethernet or VLAN packet received from the customer network can be forwarded to one or more local interfaces and/or emulated VCs in the VPLS domain. To avoid broadcasted packets looping around in the network, no packet received from an emulated VC can be forwarded to any emulated VC of the VPLS domain on a PE router. That is, the Layer 2 split horizon should always be enabled as the default in a full-mesh network.

Figure 1-2 VPLS Configuration Example



Configuration on PE 1

This shows the creation of the virtual switch instances (VSIs) and associated VCs.

```
12 vfi PE1-VPLS-A manual
   vpn id 100
   neighbor 2.2.2.2 encapsulation mpls
   neighbor 3.3.3.3 encapsulation mpls
!
interface Loopback 0
 ip address 1.1.1.1 255.255.255.255
```

This configures the CE device interface (there can be multiple Layer 2 interfaces in a VLAN).

```
interface FastEthernet0/0
 switchport
 switchport mode dot1qtunnel
 switchport access vlan 100
```

Here the attachment circuit (VLAN) is associated with the VSI.

```
interface vlan 100
 no ip address
 xconnect vfi PE1-VPLS-A
```

This is the enablement of the Layer 2 VLAN instance.

```
vlan 100
 state active
```

Configuration on PE 2

This shows the creation of the virtual switch instances (VSIs) and associated VCs.

```
12 vfi PE2-VPLS-A manual
   vpn id 100
```

```

neighbor 1.1.1.1 encapsulation mpls
neighbor 3.3.3.3 encapsulation mpls
!
interface Loopback 0
ip address 2.2.2.2 255.255.255.255

```

This configures the CE device interface (there can be multiple Layer 2 interfaces in a VLAN).

```

interface FastEthernet0/0
switchport
switchport mode dot1qtunnel
switchport access vlan 100

```

Here the attachment circuit (VLAN) is associated with the VSI.

```

interface vlan 100
no ip address
xconnect vfi PE2-VPLS-A

```

This is the enablement of the Layer 2 VLAN instance.

```

vlan 100
state active

```

Configuration on PE 3

This shows the creation of the virtual switch instances (VSIs) and associated VCs.

```

l2 vfi PE3-VPLS-A manual
vpn id 100
neighbor 1.1.1.1 encapsulation mpls
neighbor 2.2.2.2 encapsulation mpls
!
interface Loopback 0
ip address 3.3.3.3 255.255.255.255

```

This configures the CE device interface (there can be multiple Layer 2 interfaces in a VLAN).

```

interface FastEthernet0/1
switchport
switchport mode dot1qtunnel
switchport access vlan 100
!

```

Here the attachment circuit (VLAN) is associated with the VSI.

```

interface vlan 100
no ip address
xconnect vfi PE3-VPLS-A .
!

```

This is the enablement of the Layer 2 VLAN instance.

```

vlan 100
state active

```

The **show mpls l2 vc** command provides information on the status of the VC.

VPLS1# **show mpls l2 vc**

Local intf	Local circuit	Dest address	VC ID	Status
-----	-----	-----	-----	-----
Vi1	VFI	22.22.22.22	100	DOWN
Vi1	VFI	22.22.22.22	200	UP
Vi1	VFI	33.33.33.33	100	UP
Vi1	VFI	44.44.44.44	100	UP
Vi1	VFI	44.44.44.44	200	UP

The **show vfi** command provides information on the VFI.

```
PE-1# show vfi PE1-VPLS-A
VFI name: VPLSA, state: up
  Local attachment circuits:
    Vlan100
  Neighbors connected via pseudowires:
    2.2.2.2  3.3.3.3
```

The **show mpls 12transport vc** command provides information the virtual circuits.

```
Router# show mpls 12 vc det
Local interface: VFI vfi17 up
  Destination address: 1.3.1.1, VC ID: 17, VC status: up
    Tunnel label: imp-null, next hop point2point
    Output interface: PO3/4, imposed label stack {18}
  Create time: 3d15h, last status change time: 1d03h
  Signaling protocol: LDP, peer 1.3.1.1:0 up
    MPLS VC labels: local 18, remote 18
    Group ID: local 0, remote 0
    MTU: local 1500, remote 1500
    Remote interface description:
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 0, send 0
    byte totals:   receive 0, send 0
    packet drops:  receive 0, send 0
```

Configuration Examples for VPLS BGP-Based Autodiscovery

PE Configuration

This shows the configuration of VPLS BGP-Autodiscovery on PE.

```
router bgp 1000
  bgp log-neighbor-changes
  bgp graceful-restart
  neighbor 44.254.44.44 remote-as 1000 neighbor 44.254.44.44 update-source loopback300
!
  address-family l2vpn vpls
    neighbor 44.254.44.44 activate
    neighbor 44.254.44.44 send-community both
  exit-address-family
!
12 vfi 2128 autodiscovery
  vpn id 2128
interface Vlan2128
  no ip address
  xconnect vfi 2128
!
```

The following is a sample output for the **show bgp l2vpn vpls all** command:

```
BGP table version is 6, local router ID is 222.5.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

Network	Next Hop	Metric	LocPrf	Weight	Path


```

Route Distinguisher: 1000:2128
*> 1000:2128:1.1.1.72/96
                                0.0.0.0                32768 ?
*>i 1000:2128:44.254.44.44/96
                                44.254.44.44           0      100      0 ?

```

This displays the Routing Information Base (RIB) High Availability (HA) checkpoint information:

```

Xconnect RIB Active RP:
  Checkpointing      : Allowed
  Checkpointing epoch: 1
  ISSU Client id: 2102, Session id: 1329, Compatible with peer

Add entries send ok      :      0
Add entries send fail    :      0
Delete entries send ok   :      0
Delete entries send fail :      0

+- Checkpointed to standby (Y/N)?
| +- Origin of entry (i=iBGP/e=eBGP)
| |
v v
C O      VPLS-ID      Target VPLS-ID      Next-Hop      Router-Target
-+-+-----+-----+-----+-----+-----+-----+-----+-----+
N  I      66:66      10.1.1.1      10.1.1.3      66:66
N  I      66:66      10.1.1.2      10.1.1.3      66:66
Y  I      1:1       10.1.1.1      10.1.1.1      2:2
Y  I      1:1       10.1.1.1      10.1.1.3      2:2
Y  I      1:1       10.1.1.2      10.1.1.3      2:2

```

Additional References for Virtual Private LAN Services (VPLS)

MIBs

MIB	MIBs Link
	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Virtual Private LAN Services (VPLS)

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

[Table 1-1](#) lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Table 1-1 Feature Information for Power over Ethernet

Feature Name	Releases	Feature Information
Virtual Private LAN Services (VPLS)	Cisco IOS Release 12.2(50)SY	This feature was introduced.
	Cisco IOS Release 15.1(1)SY	BGP based Autodiscovery was introduced.
	Cisco IOS Release 15.5(1)SY2	High Availability support was added.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

2

Configuring A-VPLS

- [Prerequisites for A-VPLS, page 2-1](#)
- [Restrictions for A-VPLS, page 2-1](#)
- [Information About A-VPLS, page 2-2](#)
- [How to Configure A-VPLS, page 2-2](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for A-VPLS

None.

Restrictions for A-VPLS

- These are the supported configurations:
 - MPLS core with configuration of PE routers through the **neighbor** command under transport vpls mode.
 - MPLS core with configuration of PE routers through MPLS traffic engineering tunnels using explicit paths.
 - IP core with configuration of PE routers through MPLS over GRE tunnels.

Other configuration methods, including use of the **route-via** command, BGP autodiscovery, or explicit VLAN assignment to a PE egress port, are not supported.

- A-VPLS supports the following:
 - Up to 32 EtherChannel port-channel interfaces.
 - Up to 60 VPLS neighbors, minus the number of neighbors configured with the **load-balance flow** command.
- A-VPLS requires nonstop forwarding and stateful switchover.
- A-VPLS works with following:
 - MPLS Traffic Engineering tunnels that are configured with explicit paths.
 - Generic Routing Encapsulation (GRE tunnels) that are configured with static routes to the tunnel destination.

For information about MPLS traffic engineering and GRE tunnels, see the following documents:

- [MPLS Traffic Engineering and Enhancements](#)
- [Implementing Tunnels](#)
- The **ping** and **traceroute** commands that support the Any Transport over MPLS Virtual Circuit Connection Verification (VCCV) feature are not supported over FAT pseudowires.
- The VPLS Autodiscovery feature is not supported with A-VPLS.
- Load-balancing is not supported in the core routers when the core uses IP to transport packets.

Information About A-VPLS

A-VPLS introduces the following enhancements to VPLS:

- Ability to load-balance traffic at the provider edge (PE) among multiple equal-cost core-facing paths and at core interfaces using flow labels.
- Support for redundant PE routers.

A-VPLS uses the Flow Aware Transport (FAT) Pseudowire feature to achieve PE redundancy and load-balancing on both PE and core routers. FAT pseudowires are used to load-balance traffic in the core when equal cost multipaths are used. The PE router adds an additional MPLS Label to the each packet (the flow label). Each flow has a unique flow label. For more information about FAT pseudowires, see PWE3 Internet-Draft [Flow Aware Transport of MPLS Pseudowires](#) (draft-bryant-filsfils-fat-pw).

How to Configure A-VPLS

- [Enabling Load-Balancing with ECMP and FAT Pseudowires, page 2-3](#) (Required)
- [Enabling Port-Channel Load-Balancing, page 2-3](#) (Required)
- [Explicitly Specifying the PE Routers As Part of Virtual Ethernet Interface Configuration, page 2-3](#) (Optional)
- [Configuring an MPLS Traffic Engineering Tunnel, page 2-4](#) (Optional)
- [Configuring a GRE Tunnel, page 2-5](#) (Optional)

Enabling Load-Balancing with ECMP and FAT Pseudowires

The following steps explain how to configure load-balancing on the provider edge (PE) routers, which enables it on the core P routers. No configuration is required on the core P routers.

To enable load-balancing on the edge routers, issue the **load-balance flow** command. The load-balancing rules are configured through the **port-channel load-balance** command parameters (see the “[Enabling Port-Channel Load-Balancing](#)” section on page 2-3).

To enable core load-balancing, issue the **flow-label enable** command on both PE routers. You must issue the **load-balance flow** command with the **flow-label enable** command.

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# pseudowire-class name	Establishes a pseudowire class with a name that you specify and enters pseudowire class configuration mode.
Step 4	Router(config-pw)# encapsulation mpls	Specifies the MPLS tunneling encapsulation type.
Step 5	Router(config-pw)# load-balance flow	Enables load-balancing on ECMPs.
Step 6	Router(config-pw)# flow-label enable	Enables the imposition and disposition of flow labels for the pseudowire.
Step 7	Router(config-pw)# end	Exits pseudowire class configuration mode and enters privileged EXEC mode.

Enabling Port-Channel Load-Balancing

The following task explains how to enable port channel load-balancing, which sets the load-distribution method among the ports in the bundle. If the **port-channel load-balance** command is not configured, load-balancing occurs with default parameters.

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# port-channel load-balance method	Specifies the load distribution method among the ports in a bundle.
Step 4	Router(config)# exit	Exits global configuration mode and enters privileged EXEC mode.

Explicitly Specifying the PE Routers As Part of Virtual Ethernet Interface Configuration

There are several ways to specify the route through which traffic should pass.

- Explicitly specify the PE routers as part of the virtual Ethernet interface configuration
- Configure an MPLS Traffic Engineering tunnel
- Configure a GRE tunnel

The following task explains how to explicitly specify the PE routers as part of the virtual Ethernet interface configuration.

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# interface virtual-ethernet <i>num</i>	Creates a virtual Ethernet interface and enters interface configuration mode.
Step 4	Router(config-if)# transport vpls mesh	Create a full mesh of pseudowires and enters VPLS transport mode.
Step 5	Router(config-if-transport)# neighbor <i>remote-router-id</i> [pw-class <i>pw-class-name</i>]	Specifies the PE routers to be used in the pseudowire.
Step 6	Router(config-if-transport)# exit	Exits VPLS transport configuration mode and enters interface configuration mode.
Step 7	Router(config-if)# switchport	Configures the port for Layer 2 switching.
Step 8	Router(config-if)# switchport mode trunk	Enables permanent trunking mode and negotiates to convert the link into a trunk link.
Step 9	Router(config-if)# switchport trunk allowed vlan { add except none remove } <i>vlan</i> [, <i>vlan</i> [, <i>vlan</i> [,...]]]	Configures the list of VLANs allowed on the trunk.
Step 10	Router(config)# exit	Exits interface configuration mode and enters privileged EXEC mode.

Configuring an MPLS Traffic Engineering Tunnel

There are several ways to specify the route through which traffic should pass.

- Explicitly specify the PE routers as part of the virtual Ethernet interface configuration
- Configure an MPLS Traffic Engineering tunnel
- Configure a GRE tunnel

The following task explains how to configure an MPLS Traffic Engineering tunnel. For more information about MPLS Traffic Engineering tunnels, see [MPLS Traffic Engineering and Enhancements](#).

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode (enter your password if prompted).
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# interface tunnel <i>number</i>	Configures an interface type and enters interface configuration mode.

	Command	Purpose
Step 4	Router(config-if)# ip unnumbered <i>type number</i>	Assigns an IP address to the tunnel interface. An MPLS traffic engineering tunnel interface should be unnumbered because it represents a unidirectional link.
Step 5	Router(config-if)# tunnel destination <i>ip-address</i>	Specifies the destination for a tunnel. The <i>ip-address</i> keyword is the IP address of the host destination expressed in dotted decimal notation.
Step 6	Router(config-if)# tunnel mode mpls traffic-eng	Configures the tunnel encapsulation mode to MPLS traffic engineering.
Step 7	Router(config-if)# tunnel mpls traffic-eng autoroute announce	Configures the IGP to use the tunnel in its enhanced SPF calculation.
Step 8	Router(config-if)# tunnel mpls traffic-eng path-option <i>number</i> { dynamic explicit { <i>name path-name</i> } identifier <i>path-number</i> } [lockdown]	Configures the tunnel to use a named IP explicit path or a path dynamically calculated from the traffic engineering topology database. A dynamic path is used if an explicit path is currently unavailable.
Step 9	Router(config-if)# exit	Exits interface configuration mode and returns to privileged EXEC mode.

Configuring a GRE Tunnel

There are several ways to specify the route through which traffic should pass.

- Explicitly specify the PE routers as part of the virtual Ethernet interface configuration
- Configure an MPLS Traffic Engineering tunnel
- Configure a GRE tunnel

The following task explains how to configure a GRE tunnel. For more information on GRE tunnels, see [Implementing Tunnels](#).

	Command	Purpose
Step 1	Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# interface <i>type number</i>	Specifies the interface type and number and enters interface configuration mode. To configure a tunnel, use tunnel for the type argument.
Step 4	Router(config-if)# tunnel mode { gre ip gre multipoint }	Specifies the encapsulation protocol to be used in the tunnel.
Step 5	Router(config-if)# mpls ip	Enables MPLS on the tunnel.

	Command	Purpose
Step 6	outer(config-if)# tunnel source {ip-address interface-type interface-number}	Configures the tunnel source. - Use the <i>ip-address</i> argument to specify the source IP address. - Use the <i>interface-type</i> and <i>interface-number</i> arguments to specify the interface to use. Note The tunnel source and destination IP addresses must be defined on both PE routers.
Step 7	Router(config-if)# tunnel destination {hostname ip-address}	Configures the tunnel destination. - Use the hostname argument to specify the name of the host destination. - Use the ip-address argument to specify the IP address of the host destination. Note The tunnel source and destination IP addresses must be defined on both PE routers.
Step 8	Router(config-if)# exit	Exits interface configuration mode and returns to privileged EXEC mode.
Step 9	Router(config)# ip route ip-address tunnel num	Creates a static route.

These examples show the three supported methods of configuring A-VPLS.

Explicitly Specifying Peer PE Routers

The following example shows how to create two VPLS domains under VLANs 10 and 20. Each VPLS domain includes two pseudowires to peer PE routers 10.2.2.2 and 10.3.3.3. Load-balancing is enabled through the **load-balance flow** and **flow-label enable** commands.

```
pseudowire-class c11
  encaps mpls
  load-balance flow
  flow-label enable
!
port-channel load-balance src-mac
!
interface virtual-ethernet 1
  transport vpls mesh
  neighbor 10.2.2.2 pw-class c11
  neighbor 10.3.3.3 pw-class c11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10, 20
```

Using MPLS Traffic Engineering Tunnels

The following example shows the creation of two VPLS domains and uses MPLS Traffic Engineering tunnels to specify the explicit path.

```
pseudowire-class c11
  encaps mpls
  load-balance flow
  flow-label enable
!
port-channel load-balance src-mac
!
```

```

interface Tunnel1
    ip unnumbered Loopback0
    tunnel mode mpls traffic-eng
    tunnel destination 192.168.1.1
    tunnel mpls traffic-eng autoroute announce
    tunnel mpls traffic-eng path-option 1 explicit name LSP1
!
ip explicit-path name LSP1 enable
    next-address 192.168.2.2
    next-address loose 192.168.1.1
!
interface Tunnel2
    ip unnumbered Loopback0
    tunnel mode mpls traffic-eng
    tunnel destination 172.16.1.1
    tunnel mpls traffic-eng autoroute announce
    tunnel mpls traffic-eng path-option 1 explicit name LSP2
!
ip explicit-path name LSP2 enable
    next-address 172.16.2.2
    next-address loose 172.16.1.1
!
interface virtual-ethernet 1
    transport vpls mesh
        neighbor 10.2.2.2 pw-class c11
        neighbor 10.3.3.3 pw-class c11
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 10,20

```

Using MPLS over GRE Tunnels

The following example shows the creation of two VPLS domains under VLANs 10 and 20. Each VPLS domain includes two pseudowires to peer PEs 10.2.2.2 and 10.3.3.3. The pseudowires are MPLS over GRE tunnels because the core is IP.

```

pseudowire-class c11
    encaps mpls
    load-balance flow
!
port-channel load-balance src-mac
!
interface tunnel 1
    tunnel mode gre ip
    mpls ip
    tunnel source 10.1.1.1
    tunnel destination 10.2.2.2
!
interface tunnel 2
    tunnel mode gre ip
    mpls ip
    tunnel source 10.1.1.1
    tunnel destination 10.3.3.3
!
interface virtual-ethernet 1
    transport vpls mesh
        neighbor 10.2.2.2 pw-class c11
        neighbor 10.3.3.3 pw-class c11
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 10, 20

ip route 10.2.2.2 255.255.255.255 Tunnel1
ip route 10.3.3.3 255.255.255.255 Tunnel2

```

Routed Pseudo-Wire (RPW) and Routed VPLS

RPW and Routed VPLS can route Layer 3 traffic as well as switch Layer 2 frames for pseudowire connections between provider edge (PE) devices. Both point-to-point PE connections, in the form of Ethernet over MPLS (EoMPLS), and Virtual Private LAN Services (VPLS) multipoint PE connections are supported. The ability to route frames to and from these interfaces supports termination of a pseudowire into a Layer 3 network (VPN or global) on the same switch, or to tunnel Layer 3 frames over a Layer 2 tunnel (EoMPLS or VPLS). The feature supports faster network convergence in the event of a physical interface or device failure through the MPLS Traffic Engineering (MPLS-TE) and Fast Reroute (FRR) features. In particular, the feature enables MPLS TE-FRR protection for Layer 3 multicast over a VPLS domain.



Note

When the RPW is configured in A-VPLS mode, TE/FRR is not supported because A-VPLS runs over ECMP and the ECMP convergence is comparable to TE/FRR.

To configure routing support for the pseudowire, configure an IP address and other Layer 3 features for the Layer 3 domain (VPN or global) in the virtual LAN (VLAN) interface configuration. The following example assigns the IP address 10.10.10.1 to the VLAN 100 interface, and enables Multicast PIM. (Layer 2 forwarding is defined by the VFI VFI100.)

```
interface vlan 100
  xconnect vfi VFI100
  ip address 10.10.10.1 255.255.255.0
  ip pim sparse-mode
```

The following example assigns an IP address 20.20.20.1 of the VPN domain VFI200. (Layer 2 forwarding is defined by the VFI VFI200.)

```
interface vlan 200
  xconnect vfi VFI200
  ip vrf forwarding VFI200
  ip address 20.20.20.1 255.255.255.0
```

3

Ethernet Virtual Connections (EVCs)

- [Prerequisites for EVCs, page 3-1](#)
- [Restrictions for EVCs, page 3-1](#)
- [Information About EVCs, page 3-3](#)
- [Default Settings for EVCs, page 3-9](#)
- [How to Configure EVCs, page 3-9](#)
- [Monitoring EVCs, page 3-13](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for EVCs

None.

Restrictions for EVCs

- LACP EtherChannels and the 802.1ad provider-bridge mode are mutually exclusive. LACP EtherChannels cannot transmit traffic when the 802.1ad provider-bridge mode is enabled.
- Maximum EFPs per switch: 10K.
- Maximum EFPs per bridge domain: 124.
- Maximum EFPs per interface: 4K.
- Maximum bridge domains per switch: 4K.

- Bridge domain configuration is supported only as part of the EVC service instance configuration.
- EVC support requires the following:
 - The spanning tree mode must be MST.
 - The **dot1ad** global configuration mode command must be configured.
- Service instances can be configured only on ports configured to trunk unconditionally with the **switchport nonegotiate** command.
- You can configure PFC QoS to support EVC ports.
- These are the supported EVC features:
 - Service instances—You create, delete, and modify EFP service instances on Ethernet interfaces.
 - Ethernet service protection on EVCs:
 - Ethernet Operations, Administration, and Maintenance (EOAM)
 - Connectivity fault management (CFM)
 - Ethernet Local Management Interface (E-LMI)
 - IPv6 access control lists (ACLs).
 - Encapsulation—You can map traffic to EFPs based on 802.1Q VLANs (a single VLAN or a list or range of VLANs).
 - You can configure EFPs as members of a bridge domain.
 - Bridge domains support push symmetric only: the supported rewrite configuration implies egress pushing (adding a tag)
 - Bridge domains support ingress rewrite
 - EVC forwarding
 - MAC address learning and aging
 - EVCs on EtherChannels
 - EVC MAC address security
 - Bridging between switchports and EFPs
 - MSTP (MST on EVC bridge domain)
 - EFP statistics (packets and bytes)
 - QoS aware EVC/EFP per service instance
- These Layer 2 port-based features can run with EVC configured on a port:
 - PAGP
 - LACP
 - UDLD
 - LLDP
 - CDP
 - MSTP
- These features are not supported on EVCs:
 - Layer 2 multicast frame flooding
 - Layer 2 protocol tunneling

- QinQ tagging
- VLAN Translation
- EoMPLS
- Bridge domain routing
- Split horizon
- Service instance groups; also called Ethernet flow point (EFP) groups
- IPv6 access control lists (ACLs)

Information About EVCs

- [EVC Overview, page 3-3](#)
- [Ethernet Flow Points, page 3-3](#)
- [Service Instances and EFPs, page 3-4](#)
- [Encapsulation \(Flexible Service Mapping\), page 3-5](#)
- [EFPs and MSTP, page 3-7](#)
- [Bridge Domains, page 3-7](#)
- [Rewrite Operations, page 3-8](#)
- [Layer 3 and Layer 4 ACL Support, page 3-9](#)
- [Advanced Frame Manipulation, page 3-9](#)
- [Egress Frame Filtering, page 3-9](#)

EVC Overview

Ethernet virtual circuits (EVCs) define a Layer 2 bridging architecture that supports Ethernet services. An EVC is defined by the Metro-Ethernet Forum (MEF) as an association between two or more user network interfaces that identifies a point-to-point or multipoint-to-multipoint path within the service provider network. An EVC is a conceptual service pipe within the service provider network. A bridge domain is a local broadcast domain that exists separately from VLANs.

Ethernet Flow Points

An Ethernet flow point (EFP) service instance is a logical interface that connects a bridge domain to a physical port or to an EtherChannel. Configuring a service instance on a Layer 2 port creates a pseudoport or EFP on which you configure EVC features. Each service instance has a unique number per interface, but you can use the same number on different interfaces because service instances on different ports are not related.

An EFP classifies frames from the same physical port to one of the multiple service instances associated with that port, based on user-defined criteria. Each EFP can be associated with different forwarding actions and behavior.

The three major characteristics (or parameters) of an EFP are

- Encapsulation

- Rewrite Information
- Forwarding instance or method (bridge-domain or xconnect)

An EVC broadcast domain is determined by a bridge domain and the EFPs that are attached to it. An incoming frame is matched against EFP matching criteria on the interface, learned on the matching EFP, and forwarded to one or more EFPs in the bridge domain. If there are no matching EFPs, the frame is dropped.

You can use EFPs to configure VLAN translation. For example, if there are two EFPs egressing the same interface, each EFP can have a different VLAN rewrite operation, which is more flexible than the traditional switch port VLAN translation model.

When an EFP is created, the initial state is UP. The state changes to DOWN under the following circumstances:

- The EFP is explicitly shut down by a user.
- The main interface to which the EFP is associated is down or removed.
- If the EFP belongs to a bridge domain, the bridge domain is down.
- The EFP is forced down as an error-prevention measure of certain features.

Service Instances and EFPs

Configuring a service instance on a Layer 2 port or EtherChannel creates a pseudoport or Ethernet flow point (EFP) on which you configure EVC features. Each service instance has a unique number per interface, but you can use the same number on different interfaces because service instances on different ports are not related.

If you have defined an EVC by entering the **ethernet evc evc-id** global configuration command, you can associate the EVC with the service instance (optional). There is no default behavior for a service instance. You can configure a service instance only on trunk ports with no allowed VLANs. Any other configuration is not allowed. After you have configured a service instance on an interface, switchport commands are not allowed on the interface. You can also configure a service instance on an EtherChannel group.

Use the **service instance number ethernet [name]** interface configuration command to create an EFP on a Layer 2 interface or EtherChannel and to enter service instance configuration mode. You use service instance configuration mode to configure all management and control data plane attributes and parameters that apply to the service instance on a per-interface basis.

- The **service instance number** is the EFP identifier, an integer from 1 to 4000.
- The optional **ethernet name** is the name of a previously configured EVC. You do not need to enter an EVC *name*, but you must enter **ethernet**. Different EFPs can share the same name when they correspond to the same EVC. EFPs are tied to a global EVC through the common name.

When you enter service instance configuration mode, you can configure these options:

- **default**—Sets a command to its defaults
- **description**—Adds a service instance specific description
- **encapsulation**—Configures Ethernet frame match criteria
- **errdisable**—Configures error disable
- **ethernet**—Configures Ethernet-lmi parameters
- **exit**—Exits from service instance configuration mode

- **l2protocol**—Configures Layer 2 control protocol processing
- **mac**—Commands for MAC address-based features
- **no**—Negates a command or sets its defaults
- **service-policy**—Attaches a policy-map to an EFP
- **shutdown**—Takes the service instance out of service

Enter the **[no] shutdown** service-instance configuration mode to shut down or bring up a service instance.

On a Layer 2 port with no service instance configured, multiple **switchport** commands are available (**access**, **backup**, **block**, **host**, **mode**, and **trunk**). When one or more service instances are configured on a Layer 2 port, no **switchport** commands are accepted on that interface.

Encapsulation (Flexible Service Mapping)

Encapsulation defines the matching criteria that map any of these in any combination to a service instance:

- A VLAN
- A range of VLANs
- The class of service (CoS) bits
- The Ethertype

VLAN tags and CoS can be a single value, a range, or a list. Ethertype can be a single type or a list of types. These are the encapsulation types:

- default
- dot1q
- priority-tagged
- untagged

Priority-tagged frames are always single-tagged. All Ethernet traffic is supported. The encapsulation classification options are:

- inner tag CoS
- inner tag VLAN

When you configure an encapsulation method, you enable flexible service mapping, which allows you to map an incoming packet to an EFP based on the configured encapsulation.

The default behavior for flexible service mapping based on the outer 802.1q VLAN tag value is nonexact, meaning that when the EFP encapsulation configuration does not explicitly specify an inner (second) VLAN tag matching criterion, the software maps both single-tagged and double-tagged frames to the EFP as long as the frames fulfill the criteria of outer VLAN tag values. The command-line interface (CLI) does allow you to specify exact mapping with the **exact** keyword. If this keyword is specified, the EFP is designated as single-tagged-frame-only and double-tagged frames are not classified to that EFP.

Using the CLI **encapsulation** command in service-instance configuration mode, you can set encapsulation criteria. You must configure one encapsulation command per EFP (service instance). After you have configured an encapsulation method, these commands are available in service instance configuration mode:

- **bridge-domain**—Configures a bridge domain.

- **rewrite**—Configures Ethernet rewrite criteria.

Table 3-1 Supported Encapsulation Types

Command	Description
encapsulation dot1q { any <i>vlan-id</i> [, <i>vlan-id</i> [- <i>vlan-id</i>]]}	Defines the matching criteria to be used to map 802.1q frames ingressing on an interface to the appropriate EFP. The options are a single VLAN, a range of VLANs, or lists of VLANs or VLAN ranges. VLAN IDs are 1 to 4094. • Enter the any keyword to match or all VLANs (1-4094) • Enter a single VLAN ID for an exact match of the outermost tag. • Enter a VLAN range for a ranged outermost match.
encapsulation dot1q <i>vlan-id</i> cos <i>cos-value</i>	CoS value encapsulation defines match criteria after including the CoS for the C-Tag. The CoS value is a single digit between 1 and 7. You cannot configure CoS encapsulation with the encapsulation untagged command, but you can configure it with the encapsulation priority-tagged command. The result is an exact outermost VLAN and CoS match. You can also use VLAN ranges.
encapsulation untagged	Matching criteria to be used to map untagged Ethernet frames entering an interface to the appropriate EFP. Only one EFP per port can have untagged encapsulation. However, a port that hosts EFP matching untagged traffic can also host other EFPs that match tagged frames. Note Not supported with the encapsulation priority-tagged command.
encapsulation priority-tagged	Specifies priority-tagged frames. A priority-tagged packet has VLAN ID 0 and a CoS value of 0 to 7. Note Not supported with the encapsulation untagged command.
encapsulation default	Configures the default EFP on a port, which matches all otherwise unmatched packets. If the default EFP is the only one configured on a port, it matches all ingress frames on that port. If you configure the default EFP on a port, you cannot configure any other EFP on the same port with the same bridge domain.

If a packet entering a port does not match any of the encapsulations on that port, the packet is dropped, resulting in filtering of the packet. The encapsulation must match the packet on the wire to determine filtering criteria. On the wire refers to packets ingressing the switch before any rewrites and to packets egressing the switch after all rewrites.

EFPs and MSTP

EFP bridge domains are supported by the Multiple Spanning Tree Protocol (MSTP). These restrictions apply when running STP with bridge domains.

- All incoming VLANs (outer-most or single) mapped to a bridge domain must belong to the same MST instance or loops could occur.
- For all EFPs that are mapped to the same MST instance, you must configure backup EFPs on every redundant path to prevent loss of connectivity due to STP blocking a port.
- When STP mode is PVST+ or PVRST, EFP information is not passed to the protocol. EVC only supports only MSTP.
- Changing STP mode from MST to PVST+ or PVRST for a multicast port is not allowed.

Bridge Domains

- [Bridge Domain Overview, page 3-7](#)
- [Ethernet MAC Address Learning, page 3-7](#)
- [Flooding of Layer 2 Frames for Unknown MAC and Broadcast Addresses, page 3-8](#)
- [Layer 2 Destination MAC Address-Based Forwarding, page 3-8](#)
- [MAC Address Aging, page 3-8](#)
- [MAC Address Table, page 3-8](#)

Bridge Domain Overview

A bridge domain defines a broadcast domain internal to a platform and allows the decoupling of a broadcast domain from a VLAN. This decoupling enables per-port VLAN significance, thus removing the scalability limitations associated with a single per-device VLAN ID space. Frames received from one of the EFPs participating in a bridge domain matches are bridged.

A service instance must be attached to a bridge domain. Flooding and communication behavior of a bridge domain is similar to that of a VLAN domain. Bridge-domain membership is determined by which service instances have joined it (based on encapsulation criteria), while VLAN domain membership is determined by the VLAN tag in the packet.

**Note**

You must configure encapsulation before you can configure the bridge domain.

IGMP snooping is enabled by default on the switch and on all VLANs but is automatically disabled on a VLAN when you configure a bridge domain under 4094. The switches support up to 124 bridge domains.

Ethernet MAC Address Learning

MAC address learning is always enabled and cannot be disabled.

Flooding of Layer 2 Frames for Unknown MAC and Broadcast Addresses

A Layer 2 frame with an unknown unicast or broadcast destination MAC address is flooded to all the EFPs in the bridge domain except to the originating EFP.

Replication of frames involves recirculating the frame several times. Recirculation negatively affect forwarding performance and reduce the packet forwarding rate for all features.

Layer 2 Destination MAC Address-Based Forwarding

When bridging is configured, a unicast frame received from an EFP is forwarded based on the destination Layer 2 MAC address. If the destination address is known, the frame is forwarded only to the EFP/NNI associated with the destination address.

Because the bridge and EFP configurations are interrelated, bridging is supported only on EFPs. To support multiple bridge domains, MAC address entries are associated with the bridge domain of the EFP. Only unicast MAC addresses need to be dynamically learned.

The EVC infrastructure does not modify frame contents.

MAC Address Aging

The dynamically learned MAC address entries in the MAC table are periodically aged out and entries that are inactive for longer than the configured time period are removed from the table. The supported range of aging-time values, in seconds, is 5 through 1000000, with a granularity of 1. The default is 8 minutes. The **aging-time** parameter can be configured per bridge domain and is a relative value. The value is the aging time relative to the time a frame was received with that MAC address.

MAC Address Table

The MAC address table is used to forward frames based on Layer 2 destination MAC addresses. The table consists of static MAC addresses downloaded from the route processor (RP) and the MAC addresses dynamically learned by the data path.

While the MAC Learning feature is enabled, an entry is added to the MAC table when a new unique MAC address is learned on the data path and an entry is deleted from the table when it is aged out.

Rewrite Operations

The **rewrite** command pushes the 802.1ad tag onto ingress packets to forward the packet on the 802.1ad cloud.

Enter the **rewrite ingress tag push dot1ad vlan-id symmetric** service-instance configuration mode command to specify the encapsulation of additional dot1ad tag on the frame ingress to the EFP.



Note

The **symmetric** keyword is required to complete rewrite to configuration.

When you enter the **symmetric** keyword, the egress counterpart performs the inverse action and pushes (adds) the encapsulation VLAN.

Layer 3 and Layer 4 ACL Support

Configuring an ACL on an EFP is the same as configuring an ACL on other types of interfaces.



Note

ACLs are not supported for packets prefixed with a Multiprotocol Label Switching (MPLS) header, including when an MPLS packet contains either Layer 3 or Layer 4 headers of supported protocols.

Advanced Frame Manipulation

The Advanced Frame Manipulation feature supports a PUSH operation that adds one VLAN tag to both the incoming and outgoing frames of an EFP.

When a VLAN tag exists and a new one is added, the CoS field of the new tag is set to the same value as the CoS field of the existing VLAN tag; otherwise, the CoS field is set to a default of 0. Using QoS marking configuration commands, you can change the CoS marking.

Egress Frame Filtering

Egress frame filtering is performed to ensure that frames exiting an EFP contain a Layer 2 header that matches the encapsulation characteristics associated with the EFP. This filtering is done primarily to prevent unintended frame leaks and is always enabled on EFPs.

Default Settings for EVCs

None.

How to Configure EVCs

Configuring a service instance on a Layer 2 port creates an EFP on which you can configure EVC features. Perform this task to configure an EFP.

	Command or Action	Purpose
Step 1	Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Router# configure terminal	Enters global configuration mode.
Step 3	Router(config)# dot1ad	Enables 802.1ad provider-bridge mode. Note LACP EtherChannels and the 802.1ad provider-bridge mode are mutually exclusive. LACP EtherChannels cannot transmit traffic when the 802.1ad provider-bridge mode is enabled.
Step 4	Router(config)# interface <i>type number</i>	Enters interface configuration mode.

	Command or Action	Purpose
Step 5	Router(config-if)# switchport	Configures the port for Layer 2 switching. Note You must enter the switchport command once without any keywords to configure the LAN port as a Layer 2 port before you can enter additional switchport commands with keywords.
Step 6	Router(config-if)# switchport mode trunk	Configures the port to trunk unconditionally.
Step 7	Router(config-if)# switchport nonegotiate	Configures the trunk not to use DTP.
Step 8	Router(config-if)# switchport trunk encapsulation dot1q	Configures the trunk encapsulation as 802.1Q.
Step 9	Router(config-if)# switchport trunk allowed vlan <i>vlan</i> [, <i>vlan</i> [, <i>vlan</i> [, ...]]	Configures the list of VLANs allowed on the trunk. Note If VLAN locking is enabled, enter VLAN names instead of VLAN numbers. For more information, see the “ VLAN Locking ” section on page 26-4.
Step 10	Router(config-if)# dot1ad uni	Configures the port as an 802.1ad provider-bridge user-to-network interface (UNI) port. Note The dot1ad uni interface mode command imposes some SPAN restrictions (see “ Feature Incompatibilities ” section on page 18-2).
Step 11	Router(config-if)# no cdp enable	Disables CPD on the port.
Step 12	Router(config-if)# no lldp transmit	(Required on PE ports) Disables LLDP.
Step 13	Router(config-if)# spanning-tree bpdufilter enable	Enables BPDU filtering on the port.
Step 14	Router(config-if)# service instance <i>number</i> ethernet [<i>name</i>]	Configures an Ethernet service instance (EFP) and enters service instance configuration mode. - The <i>number</i> is the EFP identifier, an integer from 1 through 4000. (Optional) ethernet <i>name</i> is the name of a previously configured EVC. You do not need to use an EVC name in a service instance.
Step 15	Router(config-if)# ip access-group <i>access-list-number</i> <i>access-list-name</i> { in out }	(Optional) Applies an IP access list or object group access control list (OGACL) to an interface.
Step 16	Router(config-if-srv)# encapsulation <i>encapsulation-type</i> <i>vlan-id</i> [cos <i>cos_value</i>]	Configures the encapsulation type for the service instance. default—Configures matching for all otherwise unmatched packets. dot1q—Configures 802.1Q encapsulation. See Table 3-1 for more information. priority-tagged—Specify priority-tagged frames, VLAN-ID 0 and CoS value of 0 to 7. untagged—Map to untagged VLANs. Only one EFP per port can have untagged encapsulation. - The CoS value defines match criterion, an integer from 1 through 7.

	Command or Action	Purpose
Step 17	Router(config-if-srv)# rewrite ingress tag push dot1ad <i>vlan-id</i> [symmetric]	(Optional) Specifies the encapsulation adjustment to be performed on a frame ingressing a service instance.
Step 18	Router(config-if-srv)# bridge-domain <i>bridge-id</i>	Configures the bridge domain.
Step 19	Router(config-if-srv)# end	Returns to privileged EXEC mode.

Configuring Multiple Service Instances

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# switchport mode trunk
Router(config-if)# switchport nonegotiate
Router(config-if)# service instance 1 ethernet evc1
Router(config-if-srv)# encapsulation dot1q 201 cos 1
Router(config-if-srv)# rewrite ingress tag push dot1ad 300 symmetric
Router(config-if-srv)# bridge-domain 300
Router(config-if-srv)# end
Router(config-if)# service instance 2 ethernet evc2
Router(config-if-srv)# encapsulation default
Router(config-if-srv)# rewrite ingress tag push dot1ad 301 symmetric
Router(config-if-srv)# bridge-domain 301
Router(config-if-srv)# end
Router(config-if)# service instance 3 ethernet evc3
Router(config-if-srv)# encapsulation priority-tagged cos 1
Router(config-if-srv)# rewrite ingress tag push dot1ad 302 symmetric
Router(config-if-srv)# bridge-domain 302
```

Configuring a Service Instance

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# switchport mode trunk
Router(config-if)# switchport nonegotiate
Router(config-if)# switchport trunk allowed vlan none
Router(config-if)# service instance 22 ethernet evc1
Router(config-if-srv)# encapsulation dot1q 100
Router(config-if-srv)# rewrite ingress tag push dot1ad 10 symmetric
Router(config-if-srv)# bridge-domain 10
```

Encapsulation Using a VLAN Range

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# service instance 1 ethernet
Router(config-if-srv)# encapsulation dot1q 22-44 cos 1
Router(config-if-srv)# rewrite ingress tag push dot1ad 10 symmetric
Router(config-if-srv)# bridge-domain 10
```

Two Service Instances Joining the Same Bridge Domain

In this example, service instance 1 on interfaces Gigabit Ethernet 1/1 and 1/2 can bridge between each other.

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# service instance 1 Ethernet
Router(config-if-srv)# encapsulation dot1q 10
Router(config-if-srv)# rewrite ingress tag push dot1ad 10 symmetric
Router(config-if-srv)# bridge-domain 10
```

```
Router(config)# interface gigabitethernet1/2
Router(config-if)# service instance 1 Ethernet
Router(config-if-srv)# encapsulation dot1q 10
Router(config-if-srv)# rewrite ingress tag push dot1ad 10 symmetric
```

```
Router(config-if-srv)# bridge-domain 10
```

Bridge Domains and VLAN Encapsulation

Use the VLAN ID configured with the **rewrite ingress tag push dot1ad** command as the bridge-domain number, rather than the VLAN ID configured with the **encapsulation dot1q** command, which can be the same or a different value.

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# service instance 1 ethernet evc1
Router(config-if-srv)# encapsulation dot1q 10
Router(config-if-srv)# rewrite ingress tag push dot1ad 4000 symmetric
Router(config-if-srv)# bridge-domain 4000
```

```
Router(config)# interface gigabitethernet1/2
Router(config-if)# service instance 1 ethernet evc1
Router(config-if-srv)# encapsulation dot1q 10
Router(config-if-srv)# rewrite ingress tag push dot1ad 4000 symmetric
Router(config-if-srv)# bridge-domain 4000
```

Traffic cannot be forwarded if the the VLAN ID configured with the **encapsulation dot1q** commands do not match in a bridge domain. In this example, the service instances on Gigabit Ethernet 1/1 and 1/2 cannot forward between each other, because the encapsulation VLAN IDs do not match (filtering criteria). You can use the **rewrite** command to allow communication between these two.

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# service instance 1 ethernet evc1
Router(config-if-srv)# encapsulation dot1q 10
Router(config-if-srv)# rewrite ingress tag push dot1ad 4000 symmetric
Router(config-if-srv)# bridge-domain 4000

Router(config)# interface gigabitethernet1/2
Router(config-if)# service instance 1 ethernet evc1
Router(config-if-srv)# encapsulation dot1q 99
Router(config-if-srv)# rewrite ingress tag push dot1ad 4000 symmetric
Router(config-if-srv)# bridge-domain 4000
```

Rewrite

In this example, the VLAN ID configured in the **rewrite ingress tag push dot1ad** command (4000 in the example) is pushed onto packets that match the VLAN ID configured in the **encapsulation dot1q** command (10 in the example). The **symmetric** keyword enables the inverse action on packets in the reverse direction: packets that egress from this service instance with VLAN ID 4000 are deencapsulated, resulting in VLAN ID 10 with cos 1.

```
Router(config)# interface gigabitethernet1/1
Router(config-if)# service instance 1 ethernet
Router(config-if-srv)# encapsulation dot1q 10 cos 1
Router(config-if-srv)# rewrite ingress tag push dot1ad 4000 symmetric
Router(config-if-srv)# bridge-domain 4000
```


Monitoring EVCs

Table 3-2 Supported show Commands

Command	Description
<code>show ethernet service evc [id evc-id interface interface-id] [detail]</code>	Displays information about all EVCs, or a specific EVC when you enter an EVC ID, or all EVCs on an interface when you enter an interface ID. The detail option provides additional information about the EVC.
<code>show ethernet service instance [id instance-id interface interface-id interface interface-id] {[detail] [stats]}</code>	Displays information about one or more service instance (EFPs). If you specify an EFP ID and interface, only data pertaining to that particular EFP is displayed. If you specify only an interface ID, data is displayed for all EFPs on the interface.
<code>show bridge-domain [n]</code>	When you enter <i>n</i> , this command displays all the members of the specified bridge-domain, if a bridge-domain with the specified number exists. If you do not enter <i>n</i> , the command displays all the members of all bridge-domains in the system.
<code>show ethernet service instance detail</code>	This command displays detailed service instance information, including Layer 2 protocol information. This is an example of the output: Router# show ethernet service instance detail Service Instance ID: 2 Associated Interface: GigabitEthernet7/2 Associated EVC: evc2 L2protocol drop CE-Vlans: Encapsulation: dot1q 2 vlan protocol type 0x8100 Rewrite: ingress tag push dot1ad 2 vlan-type 0x88A8 symmetric Interface Dot1q Tunnel Ethertype: 0x8100 State: Up EFP Statistics: Pkts In Bytes In Pkts Out Bytes Out 0 0 0 0
<code>show mac address-table</code>	This command displays dynamically learned or statically configured MAC security addresses.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

4

Layer 2 over Multipoint GRE (L2omGRE)

- [Prerequisites for L2omGRE, page 4-1](#)
- [Restrictions for L2omGRE, page 4-1](#)
- [Information About L2omGRE, page 4-2](#)
- [Default Settings for L2omGRE, page 4-2](#)
- [How to Configure L2omGRE, page 4-2](#)
- [Verifying the L2omGRE Configuration, page 4-4](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for L2omGRE

None.


Restrictions for L2omGRE

- The L2omGRE feature is supported in VSS mode.
- The VLAN interface used to support L2omGRE cannot also be configured to support any Layer 3 features.
- You can configure QoS to police (rate limit) Layer 2 traffic that is flooded in a VLAN because it is addressed to a currently unlearned MAC-Layer destination address (see **match l2 miss** in the [“Traffic Classification” section on page 25-2](#)).

Information About L2omGRE

The L2omGRE feature provides Layer 2 connectivity between multiple separate network sites by extending the topology of a Layer 2 broadcast domain through an mGRE tunnel (see [Figure 4-1](#)).

Figure 4-1 L2omGRE Topology



330528

The L2omGRE feature associates a VLAN interface with an mGRE tunnel interface. mGRE tunnel interfaces configured to support the L2omGRE feature function like Layer 2 LAN ports to provide Layer 2 switching (bridging) for traffic of all types in the VLAN, so that the tunnel only carries traffic addressed to devices accessible through the mGRE tunnel (see the [“Information about Layer 2 Ethernet Switching”](#) section on page 21-2). The switch will learn MAC addresses that are accessible through mGRE tunnels and the **show mac address-table** command displays the learned addresses.

The PFC and any DFCs support bridging and mGRE tunnel encapsulation and decapsulation in hardware.

Each tunnel can carry multiple L2omGRE-connected VLANs. The tunnel encapsulation includes the VLAN ID. When the tunneled traffic is decapsulated, the VLAN ID in the tunneled traffic is used to select the correct VLAN for the traffic.

Default Settings for L2omGRE

None.

How to Configure L2omGRE

- [Configuring a Loopback Interface, page 4-3](#)
- [Configuring an mGRE Tunnel Interface, page 4-3](#)
- [Configuring a VLAN Interface, page 4-4](#)
- [L2omGRE Configuration Examples, page 4-4](#)

Configuring a Loopback Interface

To configure a loopback interface to support L2omGRE, perform this task:

	Command or Action	Purpose
Step 1	Router> <code>enable</code>	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Router# <code>configure terminal</code>	Enters global configuration mode.
Step 3	Router(config)# <code>interface loopback number</code>	Creates the loopback interface and enters interface configuration mode.
Step 4	Router(config-if)# <code>ip address address mask</code>	Configures an IP address on the interface. Note Tunnel interfaces on other switches refer to this IP address on this switch.
Step 5	Router(config-if)# <code>end</code>	Returns to global configuration mode.

Configuring an mGRE Tunnel Interface

To configure an mGRE tunnel interface, perform this task:

	Command or Action	Purpose
Step 1	Router(config)# <code>interface tunnel number</code>	Creates a tunnel interface and enters interface configuration mode. Note VLAN interfaces on this switch refer to this tunnel number.
Step 2	Router(config-if)# <code>ip address address mask</code>	Configures an IP address on the interface.
Step 3	Router(config-if)# <code>ip nhrp map tunnel_address loopback_address</code>	Configures the IP-to-nonbroadcast multiaccess (NBMA) address mapping to the loopback interface IP addresses configured on the other switches with L2omGRE tunnel interfaces. - Enter the <i>tunnel_address loopback_address</i> values configured on the other L2omGRE switches. - Repeat this command for each of the other L2omGRE switches.
Step 4	Router(config-if)# <code>ip nhrp network-id ID</code>	Enables the Next Hop Resolution Protocol (NHRP) for the mGRE tunnels. Use the same ID value on all of the L2omGRE tunnel interfaces.
Step 5	Router(config-if)# <code>tunnel source loopback number</code>	Associates this tunnel interface with the loopback interface .
Step 6	Router(config-if)# <code>tunnel mode gre multipoint</code>	Configures multipoint GRE as the tunnel mode.
Step 7	Router(config-if)# <code>end</code>	Returns to global configuration mode.

Configuring a VLAN Interface

To configure a VLAN interface to support L2omGRE, perform this task:

	Command or Action	Purpose
Step 1	Router(config)# interface <i>vlan number</i>	Creates the VLAN interface and enters interface configuration mode.
Step 2	Router(config-if)# no ip address	Ensures that no IP address is configured on the interface.
Step 3	Router(config-if)# platform xconnect l2gre tunnel <i>number</i>	Associates the VLAN interface with an L2omGRE tunnel interface.
Step 4	Router(config-if)# end	Returns to global configuration mode.

L2omGRE Configuration Examples

The L2omGRE configuration commands must be coordinated as follows:

- The **ip nhrp map** commands on each switch point to the tunnel and loopback IP addresses on each of the other switches.
- On each switch, the **tunnel source loopback** command points to the loopback interface configured on that switch.
- On each switch, the **platform xconnect l2gre tunnel** command points to the L2omGRE tunnel interface on that switch.

Enter the **show platform l2transport gre summary** command to display information about the L2omGRE configuration on a switch.

Switch One	Switch Two	Switch Three
<pre>interface loopback 1 ip address 10.1.1.1 255.255.255.255 interface tunnel 10 ip address 10.10.10.1 255.255.255.0 no ip redirects ip nhrp map 10.20.20.2 10.2.2.2 ip nhrp map 10.30.30.3 10.3.3.3 ip nhrp network-id 10 tunnel source loopback 1 tunnel mode gre multipoint interface vlan 10 no ip address platform xconnect l2gre tunnel 10</pre>	<pre>interface loopback 1 ip address 10.2.2.2 255.255.255.255 interface tunnel 10 ip address 10.20.20.2 255.255.255.0 no ip redirects ip nhrp map 10.10.10.1 10.1.1.1 ip nhrp map 10.30.30.3 10.3.3.3 ip nhrp network-id 10 tunnel source loopback 1 tunnel mode gre multipoint interface vlan 10 no ip address platform xconnect l2gre tunnel 10</pre>	<pre>interface loopback 1 ip address 10.3.3.3 255.255.255.255 interface tunnel 10 ip address 10.30.30.3 255.255.255.0 no ip redirects ip nhrp map 10.10.10.1 10.1.1.1 ip nhrp map 10.20.20.2 10.2.2.2 ip nhrp network-id 10 tunnel source loopback1 tunnel mode gre multipoint interface vlan 10 no ip address platform xconnect l2gre tunnel 10</pre>

Verifying the L2omGRE Configuration

Enter the **show platform l2transport gre** command to display information about L2omGRE, including traffic statistics.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

5

Campus Fabric

Campus Fabric provides the basic infrastructure for building virtual networks based on policy-based segmentation constructs. Fabric Overlay provides services such as host mobility and enhanced security, which are additional to normal switching and routing capabilities.

Campus Fabric Overlay provisioning consists of three main components:

- Control-Plane
- Data-Plane
- Policy-Plane

Understanding Fabric Domain Elements

The following figure displays the elements that make up the fabric domain.

- **Fabric Edge Devices** — Provide connectivity to users and devices that connect to the fabric domain. Fabric edge devices identify and authenticate endpoints, and register endpoint ID information in the fabric host-tracking database. They encapsulate at ingress and decapsulate at egress, to forward traffic to and from endpoints connected to the fabric domain.
- **Fabric Control-Plane Devices** — Provide overlay reachability information and endpoints-to-routing-locator mapping, in the host-tracking database. The control-plane device receives registrations from fabric edge devices with local endpoints, and resolves requests from edge devices to locate remote endpoints.
- **Fabric Border Devices** — Connect traditional Layer 3 networks or different fabric domains to the local domain, and translate reachability and policy information, such as VRF and SGT information, from one domain to another.
- **Virtual Contexts** — Provide virtualization at the device level, using virtual routing and forwarding (VRF) to create multiple instances of Layer 3 routing tables. Contexts or VRFs provide segmentation across IP addresses, allowing for overlapped address space and traffic separation.
- **Host-Pools** — Group endpoints in the fabric domain into IP pools, and identify them with a VLAN ID and an IP subnet.

Campus Fabric Configuration Guidelines

Consider the following guidelines and limitations when configuring campus fabric elements:

- Catalyst 6500 switch can act as a border router or control plane device.
- Fabric edge functionality will not be supported on Catalyst 6500 switch.
- If RBACL enforcement is not done on the border router, then the user needs to configure SGT caching on the border and transport the SGT via SXP to a further upstream router.

- Configure no more than 250 virtual contexts in each fabric border node.

How to Configure Fabric Overlay

Configuring Fabric Edge Devices

For detailed instructions, see this publication:

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/16-3/configuration_guide/b_163_consolidated_3850_cg/b_163_consolidated_3850_cg_chapter_010100010.html

Configuring Control-Plane Devices

To configure control-plane devices, use the following LISP commands:

Before You Begin

Configure a loopback IP address for the control-plane device to ensure that the device is reachable.

	Command	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters the global configuration mode.
Step 3	router lisp Example: Device(config)# router lisp	Enters Locator ID/Separation Protocol (LISP) configuration mode.
Step 4	site site-name Example: Device(config-router-lisp)# site FD_Default	Configures a LISP site on a control-plane device and enters LISP site configuration mode.
Step 5	authentication-key key Example: Device(config-router-lisp-site)# authentication-key examplekey	Configures the password used to create the Hash-based Message Authentication Code (HMAC) Secure Hash Algorithm (SHA-1) hash for authenticating the map-register messages sent by edge devices when registering with the control-plane device.
Step 6	eid-prefix [instance-id instance-id] eid-prefix accept-more-specifics Example: Device(config-router-lisp-site)# eid-prefix 10.1.0.0/16 accept-more-specifics Device(config-router-lisp-site)# eid-prefix instance-id 10 10.1.0.0/16 accept-more-specifics	Configures a host-pool or a list of endpoint identifier (EID) prefixes that are allowed in a map-register message sent by the edge device when registering with the control-plane device. Specifies that an EID prefix that is more specific than the EID prefix configured is accepted and tracked. The instance-id keyword includes the specified instance ID (the instance ID used by the context you want to include in the host-pool) with the host-pool when it is registered with the control-plane device.

	Command	Purpose
Step 7	exit Example: Device(config-router-lisp-site)# exit	Exits LISP site configuration mode and returns to LISP configuration mode.
Step 8	Repeat Step 4 to Step 7 to configure another LISP site.	
Step 9	ipv4 map-server Example: Device(config-router-lisp)# ipv4 map-server	Configures a device to act as an IPv4 control-plane device.
Step 10	Step 10ipv4 map-resolver Example: Device(config-router-lisp)# ipv4 map-resolver	In the fabric domain, the control-plane device acts as the map-server and the map-resolver. Enables the control-plane device with IPv4 LISP map resolver capabilities.
Step 11	Step 11end Example: Device(config-router-lisp)# end	Exits LISP configuration mode and returns to privileged EXEC mode.

Configuring Border Devices

To configure a border device, use the following LISP commands.

Before You Begin

Configure a loopback IP address for the border device to ensure that the device is reachable.

	Command	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters the global configuration mode.
Step 3	router lisp Example: Device(config)# router lisp	Enters LISP configuration mode.
Step 4	encapsulation vxlan Example: Device(config-router-lisp)# encapsulation vxlan	Specifies VXLAN-based encapsulation.
Step 5	eid-table default instance-id instance-id Example: Device(config-router-lisp)# eid-table default instance-id 0	Associates the default EID table with the specified instance ID. Control-plane device messages include this instance ID along with the associated EID prefixes.

	Command	Purpose
Step 6	map-cache eid-prefix ipv4 address/subnet mask map-request Example: Device(config-router-lisp)# map-cache 10.1.1.0/24 map-request	Configures a static IPv4 EID-to-RLOC mapping relationship by adding a map-cache with action send-map-request for the specified dynamic EID or host pool.
Step 7	ipv4 sgt Example: Device(config-router-lisp)# ipv4 sgt	Enables the transport of Security Group Tags (SGT) in the fabric . For more information on SGTs, see Cisco TrustSec Configuration Guide.
Step 8	ipv4 proxy-etr Example: Device(config-router-lisp)# ipv4 proxy-etr	Enables the border device service in the fabric domain.
Step 9	ipv4 proxy-itr ipv4 address Example: Device(config-router-lisp)# ipv4 proxy-itr 10.1.1.1	Configures the device to operate as an IPv4 proxy ingress tunnel router (PITR), and configures the interface IP address used as a source address for encapsulation of data packets. The IPv4 locator address is used as the source address for data packets or a map-request messages.
Step 10	ipv4 itr map-resolver ipv4 address Example: Device(config-router-lisp)# ipv4 itr map-resolver 10.1.1.2	Configures the map-resolver IP from where it needs to query the RLOC corresponding to destination EID IP
Step 11	exit Example: Device(config-router-lisp)# exit	Exits LISP configuration mode and enters global configuration mode.
Step 12	ip route ipv4-prefix next-hop Example: Device(config)# ip route 0.0.0.0 0.0.0.0 10.10.10.1	Configures an IPv4 static route.
Step 13	exit Example: Device(config)# exit	Exits global configuration mode and returns to privileged EXEC mode.

Example: Configuring Fabric Border and Control-Plane Devices

Fabric Border Router

```

vrf definition abcd
rd 1:1
!
address-family ipv4
exit-address-family
!
address-family ipv6
exit-address-family
!
interface Loopback10
 ip address 1.3.3.3 255.255.255.255
!
interface TenGigabitEthernet2/8 < ----- North upstream facing
 ip address 119.1.1.1 255.255.255.0

```

```

end
!
interface TenGigabitEthernet2/9 < ----- North upstream facing
vrf forwarding abcd
ip address 19.1.1.1 255.255.255.0
end
!
interface TenGigabitEthernet 4/12 < ----- South RLOC facing
ip address 141.1.1.2 255.255.255.0
ip lisp source-locator Loopback10
!
router lisp
encapsulation vxlan
locator-set set1
    1.3.3.3 priority 1 weight 100
exit
!
eid-table default instance-id 0
map-cache 120.1.1.0/24 map-request <----- EID subnet at fabric edge
exit
!
eid-table vrf abcd instance-id 1
map-cache 20.1.1.0/24 map-request <----- EID subnet at fabric edge
exit
!
ipv4 sgt
ipv4 proxy-etr
ipv4 proxy-itr 1.3.3.3
ipv4 itr map-resolver 1.4.4.4
exit

```

Control Plane

```

interface Loopback10
ip address 1.4.4.4 255.255.255.255
!
router lisp
encapsulation vxlan
locator-set set1
    1.4.4.4 priority 1 weight 100
exit
!
site site1
    authentication-key key1
    eid-prefix 120.1.1.0/24 accept-more-specifics
    eid-prefix instance-id 1 20.1.1.0/24 accept-more-specifics
exit
!
ipv4 map-server
ipv4 map-resolver
exit

```


6

IPv4 Multicast Layer 3 Features

- [Prerequisites for IPv4 Multicast Layer 3 Features, page 6-1](#)
- [Restrictions for IPv4 Multicast Layer 3 Features, page 6-1](#)
- [Information About IPv4 Multicast Layer 3 Features, page 6-2](#)
- [Default Settings for IPv4 Multicast Layer 3 Features, page 6-15](#)
- [How to Configure IPv4 Multicast Layer 3 Features, page 6-15](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IPv4 Multicast Layer 3 Features

None.

Restrictions for IPv4 Multicast Layer 3 Features

IP multicast Layer 3 switching is not provided for an IP multicast flow in the following situations:

- For IP multicast groups that fall into the range 224.0.0.* (where * is in the range 0 to 255), which is used by routing protocols. Layer 3 switching is supported for groups 225.0.0.* through 239.0.0.* and 224.128.0.* through 239.128.0.*.

**Note**

Groups in the 224.0.0.* range are reserved for routing control packets and must be flooded to all forwarding ports of the VLAN. These addresses map to the multicast MAC address range 01-00-5E-00-00-xx, where xx is in the range 0–0xFF.

- For PIM auto-RP multicast groups (IP multicast group addresses 224.0.1.39 and 224.0.1.40).
- For packets with IP options. However, packets in the flow that do not specify IP options are hardware switched.
- If a (S,G) entry for sparse mode does not have the SPT-bit, RPT-bit, or Pruned flag set.
- A (*,G) entry is not hardware switched if at least one (S,G) entry has an RPF different from the (*,G) entry's RPF and the (S,G) is not hardware switched.
- If the ingress interface of a (S,G) or (*,G) entry is null, except if the (*,G) entry is a IPv4 bidirectional PIM entry and the switch is the RP for the group.
- If you enable IP multicast Layer 3 switching, IP accounting for Layer 3 interfaces does not report accurate values. The **show ip accounting** command is not supported.

Information About IPv4 Multicast Layer 3 Features

- [IPv4 Multicast Layer 3 Features Overview, page 6-2](#)
- [Distributed MRIB and MFIB Infrastructure, page 6-3](#)
- [Multicast Layer 3 Hardware Features Entries, page 6-4](#)
- [Layer 3-Switched Multicast Statistics, page 6-4](#)
- [Layer 3-Switched Multicast Packet Rewrite, page 6-5](#)
- [Replication Modes, page 6-5](#)
- [Local Egress Replication Mode, page 6-6](#)
- [PIM-SM hardware register support, page 6-6](#)
- [PIM-SM hardware SPT-switchover support, page 6-6](#)
- [Control Plane Policing \(CoPP\), page 6-7](#)
- [Non-RPF Traffic Processing, page 6-7](#)
- [Multicast Boundary, page 6-7](#)
- [IPv4 Bidirectional PIM, page 6-8](#)
- [Supported Multicast Features, page 6-8](#)

IPv4 Multicast Layer 3 Features Overview

Multicast Layer 3 switching on a Supervisor Engine 6T uses application-specific integrated circuits (ASICs) to provide hardware-supported forwarding for IP multicast data packet flows between IP subnets, which offloads processor-intensive multicast forwarding and replication from the route processor.

The Policy Feature Card (PFC) and Distributed Forwarding Cards (DFCs) use the forwarding information base (FIB) and adjacency table to switch IP Multicast flows in hardware. The FIB table provides support for different entry and mask values; for example, (S/32, G/32) and (*/0, G/32). The RPF RAM is used to identify packets arriving on a directly connected subnet.

The result of a FIB lookup is an adjacency, which provides the replication list for the entry. Different from earlier supervisor engines, the Supervisor Engine 6T performs packet rewrite after replication, which provides enhanced sharing capabilities for outgoing interfaces.

Also different from earlier supervisor engines, the Supervisor Engine 6T performs Layer 2 and Layer 3 forwarding decisions separately. For a routed interface, the final LTL for a packet is determined based only on Layer 3 information. For VLAN interfaces, the LTL for the packet is determined based only on the Layer 2 lookup.

Distributed MRIB and MFIB Infrastructure

The Supervisor Engine 6T uses an MRIB- and MFIB-based software model for IPv4 multicast traffic. The MRIB/MFIB infrastructure supports multiple Layer 3 multicast protocols (for example, PIM Sparse-mode, SSM, and bidirectional PIM) for both IPv4 and IPv6 flows and provides consistent CLI for both IPv4 and IPv6 addresses.

The Multicast Routing Information Base (MRIB) is a collection of multicast entries keyed on source, group, and group mask. The multicast control plane programs the entries. An entry can have one or more associated interfaces with different flags indicating the role of the interface in the forwarding plane. On the Supervisor Engine 6T, the PFC and each DFC has an instance of the Multicast Forwarding Information Base (MFIB), which registers as a client of the MRIB. The MFIB registers its interest with the MRIB for the entries and associated flags, and maintains a local database keyed on source, group, and group mask based on MRIB updates.

Without hardware support, the MFIB would forward the multicast traffic. On the Supervisor Engine 6T, in addition to any required software forwarding, the MFIB also sends multicast route updates for the Layer 3 switching hardware tables. The communication between the MRIB, MFIB, and the Layer 3 switching hardware tables is based on the following:

- Multicast entries
- Flags set on multicast entries
- Interfaces associated with multicast entries



Note

With other supervisor engines, some flows are switched partially in hardware and partially in software. With a Supervisor Engine 6T, multicast flows are switched either in software or in hardware.

Examples of mroute, MRIB, and MFIB information:

```
Router# show ip mroute
(100.1.1.3, 239.2.1.1), 00:00:53/00:02:08, flags: sTI
  Incoming interface: TenGigabitEthernet3/0, RPF nbr 0.0.0.0
  Outgoing interface list:
    TenGigabitEthernet3/1, Forward/Sparse-Dense, 00:00:51/00:02:08
    TenGigabitEthernet3/2, Forward/Sparse-Dense, 00:00:51/00:02:08
    TenGigabitEthernet3/3, Forward/Sparse-Dense, 00:00:51/00:02:08

Router# show ip mrrib route
(100.1.1.3,239.2.1.1)
  TenGigabitEthernet3/1 Flags: A
  TenGigabitEthernet3/2 Flags: F
```

```

TenGigabitEthernet3/3 Flags: F
TenGigabitEthernet3/4 Flags: F

Router# show ip mfib
(100.1.1.3,239.2.1.1) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  TenGigabitEthernet3/1 Flags: A
    Pkts: 0/0
  TenGigabitEthernet3/2 Flags: F
    Pkts: 0/0
  TenGigabitEthernet3/3 Flags: F
    Pkts: 0/0
  TenGigabitEthernet3/4 Flags: F
    Pkts: 0/0

```

Multicast Layer 3 Hardware Features Entries

This section describes how the PFC and the DFCs maintain Layer 3 switching information in hardware tables.

The PFC and DFC populate the (S,G) or (*,G) flows in the hardware FIB table with the appropriate masks; for example, (S/32, G/32) and (*,G/32). The RPF interface and the adjacency pointer information is also stored in each entry. The adjacency table contains the rewrite and a pointer to the replication entries. If a flow matches a FIB entry, the RPF check compares the incoming interface/VLAN with the entry. A mismatch is an RPF failure, which can be rate limited if this feature is enabled. In the event of a forwarding information database (FIB) fatal error, the default error action is for the system to reset and the FIB to reload.

The PFC and any DFCs run an MFIB client, which registers as a client to the MRIB on the RP. The hardware tables are programmed based on MFIB updates to install or delete entries for a traffic flow, or the addition and deletion of outgoing interfaces to an existing hardware entry.

These commands affect the Layer 3 switching entries:

- When you clear the multicast routing table using the **clear ip mroute** command, all multicast Layer 3 switching cache entries are cleared.
- When you disable IP multicast routing on the RP using the **no ip multicast-routing** command, all multicast Layer 3 switching cache entries on the PFC are purged.
- When you disable hardware-supported multicast Layer 3 switching on an interface with the **no platform multicast forwarding ip** interface mode command, flows that use the interface as the RPF interface are routed only by the RP in software.
- When you disable MFIB forwarding on an individual interface basis using the **no ip mfib forwarding ip** command, flows that use this interface as the RPF interface will not be forwarded in hardware or by MFIB.

Layer 3-Switched Multicast Statistics

While a flow is being switched in hardware, the entry is kept alive in software on the basis of the packet forwarding statistics for the flow. The MFIB entry on the PFC periodically pulls and aggregates the entry statistics from the PFC and every DFC. The statistics increment the hardware forwarding counters for an entry. A software or hardware forwarding counter increment resets the expiration timer for that multicast route.

**Note**

A (*,G) state is created on the PIM-RP or for PIM-dense mode but is not used for forwarding the flows, and Layer 3 switching entries are not created for these flows.

Layer 3-Switched Multicast Packet Rewrite

When a multicast packet is Layer 3 switched from a multicast source to a destination multicast group, the PFC and the DFCs perform a packet rewrite that is based on information learned from the RP and stored in the adjacency table.

For example, Server A sends a multicast packet addressed to IP multicast group G1. If there are members of group G1 on VLANs other than the source VLAN, the PFC must perform a packet rewrite when it replicates the traffic to the other VLANs (the switch also bridges the packet in the source VLAN).

When the PFC receives the multicast packet, it is (conceptually) formatted as follows:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Group G1 MAC</i>	<i>Source A MAC</i>	<i>Group G1 IP</i>	<i>Source A IP</i>	<i>n</i>	<i>calculation1</i>		
Note In this example, Destination B is a member of Group G1.							

The PFC rewrites the packet as follows:

- Changes the source MAC address in the Layer 2 frame header from the MAC address of the host to the MAC address of the RP (This is the burned-in MAC address of the system. This MAC address will be the same for all outgoing interfaces and cannot be modified. This MAC address can be displayed using the **show platform multicast statistics** command.)
- Decrements the IP header Time to Live (TTL) by one and recalculates the IP header checksum

The result is a rewritten IP multicast packet that appears to have been routed. The PFC replicates the rewritten packet onto the appropriate destination VLANs, where it is forwarded to members of IP multicast group G1.

After the PFC performs the packet rewrite, the packet is (conceptually) formatted as follows:

Frame Header		IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Group G1 MAC</i>	<i>RP MAC</i>	<i>Group G1 IP</i>	<i>Source A IP</i>	<i>n-1</i>	<i>calculation2</i>		

Replication Modes

The Supervisor Engine 6T supports these replication modes:

- Ingress mode—The ingress module performs the replication for the egress interfaces on all modules.
- Egress mode (default)—Ingress multicast traffic is distributed over the fabric to the egress modules. The egress modules perform the replication for the egress interfaces.

The Supervisor Engine 6T performs egress-mode replication for ingress-only legacy switching modules, so that the presence of an ingress-only legacy switching module does not force a replication mode change.

**Note**

Intranet and extranet MVPN are supported in egress and ingress replication mode.

Local Egress Replication Mode

DFC-equipped modules with dual switch-fabric connections host two packet replication engines, one per fabric connection. Each replication engine is responsible for forwarding packets to and from the interfaces associated with the switch-fabric connections. The interfaces that are associated with a switch-fabric connection are considered to be “local” from the perspective of the packet replication engine. Without local egress replication mode, both replication engines have the complete outgoing interface list for all modules, and the replication engines process and then drop traffic for nonlocal interfaces. The Supervisor Engine 6T provides local egress replication for CFC-equipped switching modules.

Local egress replication mode limits the outgoing interface list to only the local interfaces that each replication engine supports, which prevents unnecessary processing of multicast traffic.

In Cisco IOS Release 15.4SY, local egress replication mode is enabled by default.

PIM-SM hardware register support

The PIM-SM protocol requires the first-hop router for a multicast source to send a register packet to the Rendezvous Point (RP) when the source starts sending packets. On the Supervisor Engine 6T, the PIM control-plane uses the MFIB infrastructure to represent the PIM register tunnel as an interface and adds it to the outgoing interface list when PIM register packets are sent. When a new PIM RP is configured or learned, a PIM register tunnel interface is created and used for PIM register packets. The interface is deleted when a register-stop packet is received. On the RP, an additional interface is created and used to decapsulate the PIM register packet when it is received.

The Supervisor Engine 6T supports PIM register packet transmission and reception in hardware. PIM register packets are sent to the RP for the PIM registration process. CoPP can rate limit the PIM register packets sent to the RP for the PIM registration process.

PIM-SM hardware SPT-switchover support

SPT-switchover occurs in a PIM-SM network when the shortest path to the source diverges from the shortest path to the RP for a multicast group. During the transition from the shortest path tree to a source based tree, packets are received on both the shared and source tree and the multicast entry must be programmed to accept packets from two interfaces.

Hardware support is implemented with a multicast FIB dual-RPF mode. When a multicast FIB entry is programmed with a dual-RPF, packets received on the RPF towards the RP are forwarded and packets received on the RPF towards the source are sent to be processed in software for the PIM protocol to complete the SPT-switchover. After switching to the source tree, the multicast entry transitions back to a single RPF interface. To avoid excessive CPU utilization during the switchover process, you can configure CoPP to rate limit multicast packets sent to the RP during SPT-switchover.

Control Plane Policing (CoPP)

You can configure CoPP to protect the CPU from unnecessary or DoS traffic. With Cisco IOS Release 15.4SY, CoPP is configured and enabled by default. CoPP provides hardware rate limiting for the packets sent to the CPU for processing in software. Although multicast rate-limiters are still supported, CoPP is more efficient.




Note

For any particular type of traffic, configure CoPP or a rate limiter, not both.

Non-RPF Traffic Processing

In a redundant configuration where multiple routers connect to the same LAN segment, only one router forwards the multicast traffic from the source to the receivers on the outgoing interfaces (see [Figure 6-1](#)). In this kind of topology, only the PIM designated router (PIM DR) forwards the data in the common VLAN, but the non-PIM DR receives the forwarded multicast traffic. The redundant router (non-PIM DR) must drop this traffic because it has arrived on the wrong interface and fails the RPF check. Traffic that fails the RPF check is called non-RPF traffic.

Figure 6-1 Redundant Multicast Router Configuration in a Stub Network



Non-RPF protection, which is enabled by default, keeps the CPU from being overwhelmed with non-RPF traffic while still allowing some packets to reach the CPU to support correct multicast routing protocol operation. The non-RPF protection feature uses a leak and drop mechanism: a non-RPF packet is leaked to the CPU, and subsequent non-RPF packets are dropped. The sequence repeats periodically.

Multicast RPF-fail packets that conform to the configured CoPP policy reach the CPU. Nonconforming packets are dropped. Rate limiting is configured in the CoPP policy.

Multicast Boundary

The multicast boundary feature allows you to configure an administrative boundary for multicast group addresses. By restricting the flow of multicast data packets, you can reuse the same multicast group address in different administrative domains.

You configure the multicast boundary on an interface. A multicast data packet is blocked from flowing across the interface if the packet's multicast group address matches the access control list (ACL) associated with the multicast boundary feature.

Multicast boundary ACLs can be processed in hardware by the Policy Feature Card (PFC), a Distributed Forwarding Card (DFC), or in software by the RP. The multicast boundary ACLs are programmed to match the destination address of the packet. These ACLs are applied to traffic on the interface in both directions (input and output).

To support multicast boundary ACLs in hardware, the switch creates new ACL TCAM entries or modifies existing ACL TCAM entries (if other ACL-based features are active on the interface). To verify TCAM resource utilization, enter the **show tcam counts ip** command.

If you configure the **filter-autorp** keyword, the administrative boundary also examines auto-RP discovery and announcement messages and removes any auto-RP group range announcements from the auto-RP packets that are denied by the boundary ACL.

IPv4 Bidirectional PIM

The PFC and DFCs support hardware forwarding of IPv4 bidirectional PIM groups. To support IPv4 bidirectional PIM groups, the PFC and DFCs support the designated forwarder (DF) mode. The designated forwarder is the router elected to forward packets to and from a segment for a IPv4 bidirectional PIM group. In DF mode, the switch accepts packets from the RPF and from the DF interfaces.

When the switch is forwarding IPv4 bidirectional PIM groups, the RPF interface is always included in the outgoing interface list of (*,G) entry, and the DF interfaces are included depending on IGMP/PIM joins.

If the route to the RP becomes unavailable, the group is changed to dense mode. Should the RPF link to the RP become unavailable, the IPv4 bidirectional PIM flow is removed from the hardware FIB.

For information on configuring IPv4 bidirectional PIM, see the [“Configuring IPv4 Bidirectional PIM” section on page 6-26](#).

Supported Multicast Features

- [Hardware-Supported IPv4 Layer 3 Features, page 6-9](#)
- [Unsupported IPv4 Layer 3 Features, page 6-10](#)
- [Hardware-Supported IPv6 Layer 3 Multicast Features, page 6-11](#)
- [Partially Hardware-Supported IPv6 Layer 3 Multicast Features, page 6-11](#)
- [Software-Supported IPv6 Layer 3 Multicast Features, page 6-11](#)
- [Unsupported IPv6 Layer 3 Multicast Features, page 6-12](#)
- [Hardware-Supported Layer 2 Common Multicast Features, page 6-12](#)
- [Unsupported Layer 2 Common Multicast Features, page 6-12](#)
- [Hardware-Supported Layer 2 Enterprise Multicast Features, page 6-13](#)
- [Unsupported Layer 2 Enterprise Multicast Features, page 6-13](#)
- [Hardware-Supported Layer 2 Metro Multicast Features, page 6-13](#)
- [Unsupported Layer 2 Metro Multicast Features, page 6-13](#)

- [Unsupported MPLS Multicast Features, page 6-14](#)
- [Hardware-Supported Security Multicast Features, page 6-14](#)
- [Software-Supported Security Multicast Features, page 6-14](#)
- [Unsupported Security Multicast Features, page 6-14](#)

Hardware-Supported IPv4 Layer 3 Features

- Control plane policing (CoPP)
- Egress forced replication mode
- Egress replication local
- Egress replication mode
- HW assisted SPT switchover
- Input ACL logging
- Input and output ACL filtering
- IPv4 multicast over multipoint IPv4 GRE tunnel
- IPv4 multicast over P2P IPv4 GRE tunnel
- IPv4 multicast over P2P IPv4 GRE tunnel with tunnel endpoints in VRF
- IPv4 multicast over P2P IPv4 VRF GRE tunnel
- Load-Balancing of multicast packets on port-channels
- Multicast boundary
- Multicast Layer 3 forwarding on routed ports
- Multicast Layer 3 forwarding on subinterfaces
- Multicast Layer 3 forwarding on SVI
- Multicast load-splitting across parallel links
- Multicast VPN for IPv4 extranet support
- Multicast VPN for IPv4 intranet support
- Multicast VRF-lite
- MVPN over P2P IPv4 GRE tunnel
- Netflow accounting
- Non-RPF protection
- PIM Register decapsulation over IPv4
- PIM Register encapsulation over IPv4
- PIM-SM (S,G) and (*,G) forwarding
- PIM-SSM
- QoS policing for ingress mode
- Rate limiters
- Statistics
- UDLR - unidirectional link routing

- URD - URL rendezvous directory
- Partially Hardware-Supported IPv4 Layer 3 Features
- Egress replication mode and QoS marking
- QoS marking for ingress mode
- Software-Supported IPv4 Layer 3 Features
- IGMPv3/v2/v1
- MET sharing
- MRM/mrinfo/mmon
- MSDP/MBGP
- Mtrace/Mping
- PGM router assist
- PGM router assist in VRF
- Platform Dependent MIB support
- Platform Independent MIB support
- SSM Mapping
- SSO/NSF

Unsupported IPv4 Layer 3 Features

- 6PE IPv6 over IPv4 infrastructure (using MDT tunnels)
- Destination IP NAT multicast
- MTR multicast ToS-based lookup
- Multicast stub (Supported in 12.2SX)
- Partial shortcut (Supported in 12.2SX)
- Service reflect
- Source IP NAT multicast
- Egress replication mode and QoS policing
- Output ACL logging
- QoS ingress or egress shaping
- QoS marking for multicast bridged frames undergoing routing
- DVMRP interoperability (supported in Release 12.2SX)
- ISSU/MDR
- MFIB consistency checker
- MFIB to HW consistency checker
- MTR multicast: separate RPF table for group range
- MTR multicast: separate URIB RPF table
- Multicast helper map (supported in Release 12.2SX)
- RPF change tracking

- PIM-DM (S,G) forwarding
- The **ip multicast rate-limit** command (supported in Release 12.2SX)
- The **ip multicast ttl-threshold** command (supported in Release 12.2SX)

Hardware-Supported IPv6 Layer 3 Multicast Features

- Control plane policing (CoPP)
- Egress forced replication mode
- Egress replication local
- Egress replication mode
- HW assisted SPT switchover
- Input ACL logging
- Input and output ACL filtering
- IPv6 Multicast over P2P IPv4 GRE/IP-in-IP tunnel (6over4)
- Load-Balancing of multicast packets on port-channels
- Multicast Layer 3 forwarding on routed ports
- Multicast Layer 3 forwarding on subinterfaces
- Multicast Layer 3 forwarding on SVI
- Multicast load-splitting across parallel links
- Netflow accounting
- Non-RPF protection
- PIM register decapsulation over IPv6
- PIM register encapsulation over IPv6
- PIM-SM (S,G) and (*,G) forwarding
- PIM-SSM
- QoS ingress mode marking
- QoS ingress mode policing
- Rate limiters
- Scope checking
- Statistics

Partially Hardware-Supported IPv6 Layer 3 Multicast Features

- Egress replication mode and QoS marking

Software-Supported IPv6 Layer 3 Multicast Features

- SSM mapping
- MET sharing
- MLDv1/v2

Unsupported IPv6 Layer 3 Multicast Features

- BIDIR PIM over P2P GRE tunnel
- Destination IP NAT multicast
- IPv4 multicast over P2P IPv6 GRE tunnel (4over6)
- IPv6 multicast over multipoint IPv4 GRE tunnel (6over4 mGRE)
- IPv6 multicast over multipoint IPv6 GRE tunnel
- IPv6 multicast over P2P IPv6 GRE tunnel
- IPv6 multicast over P2P IPv6 GRE tunnel with tunnel endpoints in VRF
- IPv6 multicast over P2P IPv6 VRF GRE tunnel
- MTR multicast: TOS based lookup
- Multicast VPN for IPv6 extranet support
- Multicast VPN for IPv6 intranet support
- Multicast VRF-lite
- MVPN over P2P IPv6 GRE tunnel
- PIM-DM (S,G) forwarding
- Source IP NAT multicast
- Egress replication mode and QoS policing
- QoS ingress and egress: shaping support
- MIB support
- Multicast boundary
- Multicast helper map
- Output ACL logging
- PGM router assist
- PGM router assist in VRF
- PIM-BIDIR
- QoS Marking for multicast bridged frames undergoing routing

Hardware-Supported Layer 2 Common Multicast Features

- NSF/SSO support for IGMP/PIM/MLD snooping

Unsupported Layer 2 Common Multicast Features

- ISSU/MDR support for IGMP/PIM/MLD snooping
- MIB support

Hardware-Supported Layer 2 Enterprise Multicast Features

- IGMPv2/v1 snooping - IP based constrain
- IGMPv2/v1 snooping - MAC based constrain
- IGMPv3 snooping with S,G constrain
- MLD v1 snooping - IP based constrain
- MLD v1 snooping - MAC based constrain
- MLD v2 snooping - IP based constrain
- MLD v2 snooping - MAC based constrain
- Multicast support for PVLAN over BD
- Optimized flooding for unknown IP multicast frames
- PIM snooping - IP based constrain
- PIM snooping - MAC based constrain
- IGMP snooping querier
- IGMPv3 snooping Explicit tracking
- MLD snooping querier
- PIM snooping - DR flooding

Unsupported Layer 2 Enterprise Multicast Features

- RGMP
- Source-Only detection
- Multicast Flood Protection
- CGMP compatibility mode for IGMP snooping
- CGMP redirect suppression

Hardware-Supported Layer 2 Metro Multicast Features

- Multicast VLAN Registration (MVR)

Unsupported Layer 2 Metro Multicast Features

- MLD v1/v2 snooping over Q-in-Q tagged frames
- Optimized VPLS multicast constrains on LAN ports
- PIM snooping over Q-in-Q
- IGMPv3/v2/v1 snooping over Q-in-Q tagged frames
- MTP multicast optimization
- Multicast routing for BD with EOM

Unsupported MPLS Multicast Features

- Inter-AS IPv4 Multicast VPN using mLDP
- Inter-AS IPv6 Multicast VPN using mLDP
- IPv4 Multicast Traffic at the Edge (via Global routing table) using mLDP
- IPv4 Multicast Traffic at the Edge (via Global routing table) over a P2MP RSVP TE LSP
- IPv4 Multicast Traffic in a VRF over a P2MP RSVP TE LSP
- IPv4 Multicast Traffic in a VRF using mLDP
- IPv6 Multicast Traffic at the Edge (via Global routing table) using mLDP
- IPv6 Multicast Traffic at the edge (via Global routing table) over a P2MP RSVP TE LSP
- IPv6 Multicast Traffic in a VRF over a P2MP RSVP TE LSP
- IPv6 Multicast Traffic in a VRF using mLDP
- ISSU Support
- Link protection for P2MP TE LSPs (500 msec)
- Node protection for P2MP TE LSPs (500 msec)
- SSO/NSF Support
- Carrier Supporting Carrier (CSC) for IPv4 Multicast VPN using mLDP
- Carrier Supporting Carrier (CSC) for IPv6 Multicast VPN using mLDP
- Extranet support for IPv4 Multicast VPN for Label Switched Multicast
- Extranet support for IPv6 Multicast VPN for Label Switched Multicast
- Link protection for mLDP trees (500 msec)
- Node protection for mLDP trees (500 msec)
- MIB support

Hardware-Supported Security Multicast Features

- Multicast and service blade interaction
- P2P GRE tunnel and VPNSM/IPSEC SPA module interaction

Software-Supported Security Multicast Features

- IGMP Snooping filtering

Unsupported Security Multicast Features

- CTS LinkSecurity for multicast
- RBACL and IPv4/IPv6 multicast data packets
- CTS implicit tunnel for multicast
- GDOI key distribution

Default Settings for IPv4 Multicast Layer 3 Features

- Multicast routing: disabled globally.
- PIM routing: disabled on all interfaces.
- IP multicast Layer 3 switching: enabled when multicast routing is enabled and PIM is enabled on the interface.
- IP MFIB forwarding: enabled when PIM is enabled on the interface.

Internet Group Management Protocol (IGMP) snooping is enabled by default on all VLAN interfaces. If you disable IGMP snooping on an interface, multicast Layer 3 flows are still switched by the hardware. Bridging of the flow on an interface with IGMP snooping disabled causes flooding to all forwarding interfaces of the VLAN. For details on configuring IGMP snooping, see [Chapter 7, “IGMP Snooping for IPv4 Multicast Traffic.”](#)

How to Configure IPv4 Multicast Layer 3 Features

- [Enabling IPv4 Multicast Routing Globally, page 6-16](#)
- [Enabling IPv4 PIM on Layer 3 Interfaces, page 6-16](#)
- [Enabling IP Multicast Layer 3 Switching on Layer 3 Interfaces, page 6-16](#)
- [Enabling IP MFIB forwarding on Layer 3 Interfaces, page 6-17](#)
- [Configuring the Replication Mode, page 6-18](#)
- [Configuring Multicast Boundary, page 6-18](#)
- [Verifying Local Egress Replication, page 6-19](#)
- [Displaying IPv4 Multicast PIM-SM register tunnel information, page 6-20](#)
- [Displaying the IPv4 Multicast Routing Table, page 6-20](#)
- [Displaying IPv4 MRIB Information, page 6-21](#)
- [Displaying IPv4 MFIB Information, page 6-22](#)
- [Viewing Directly Connected Entries, page 6-23](#)
- [Displaying IPv4 Hardware Switching Information, page 6-24](#)
- [Displaying IPv4 CoPP Information, page 6-25](#)
- [Source-Specific Multicast with IGMPv3, IGMP v3lite, and URD, page 6-26](#)
- [Configuring IPv4 Bidirectional PIM, page 6-26](#)
- [Enabling IPv4 Bidirectional PIM Globally, page 6-27](#)
- [Configuring the Rendezvous Point for IPv4 Bidirectional PIM Groups, page 6-27](#)
- [Displaying IPv4 Bidirectional PIM Information, page 6-27](#)
- [Using IPv4 Debug Commands, page 6-31](#)
- [Redundancy for Multicast Traffic, page 6-31](#)

Enabling IPv4 Multicast Routing Globally

You must enable IP multicast routing globally before you can enable IP multicast Layer 3 switching on Layer 3 interfaces.

To enable IP multicast routing globally, perform this task:

Command	Purpose
Router(config)# ip multicast-routing	Enables IP multicast routing globally.

This example shows how to enable multicast routing globally:

```
Router(config)# ip multicast-routing
Router(config)#
```

Enabling IPv4 PIM on Layer 3 Interfaces

You must enable PIM on the Layer 3 interfaces before IP multicast Layer 3 switching functions on those interfaces.

To enable IP PIM on a Layer 3 interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port}}	Selects an interface to configure.
Step 2	Router(config-if)# ip pim {dense-mode sparse-mode sparse-dense-mode}	Enables IP PIM on a Layer 3 interface.

This example shows how to enable PIM on an interface using the default mode (**sparse-dense-mode**):

```
Router(config-if)# ip pim
```

This example shows how to enable PIM sparse mode on an interface:

```
Router(config-if)# ip pim sparse-mode
```



Note

- You must enable PIM on all participating Layer 3 interfaces before IP multicast Layer 3 switching will function. For information on configuring PIM on Layer 3 interfaces, see the [“Enabling IPv4 PIM on Layer 3 Interfaces”](#) section on page 6-16.
- PIM can be enabled on any Layer 3 interface, including VLAN interfaces.

Enabling IP Multicast Layer 3 Switching on Layer 3 Interfaces

IP multicast Layer 3 switching is enabled by default on the Layer 3 interface when you enable PIM on the interface. Perform this task only if you disabled IP multicast Layer 3 switching on the interface and you want to reenabling it.

To enable IP multicast Layer 3 switching on a Layer 3 interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port}}	Selects an interface to configure.
Step 2	Router(config-if)# platform multicast forwarding ip	Enables IP multicast Layer 3 switching on a Layer 3 interface.
Step 3	Router(config-if)# exit	Returns you to global configuration mode.
Step 4	Router # platform ip multicast syslog	(Optional) Enables display of multicast related syslog messages on console.

This example shows how to enable IP multicast Layer 3 switching on a Layer 3 interface:

```
Router(config-if)# platform multicast forwarding ip
Router(config-if)#
```

Enabling IP MFIB forwarding on Layer 3 Interfaces

Disabling MFIB forwarding on the interface is another way to disable IP multicast Layer 3 switching for an interface. By default MFIB forwarding in and out are enabled by when PIM is enabled on the interface. To enable MFIB forwarding on a Layer 3 interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port}}	Selects an interface to configure.
Step 2	Router(config-if)# ip mfib forwarding in	Enables MFIB forwarding on a Layer 3 interface.
Step 3	Router(config-if)# exit	Returns you to global configuration mode.
Step 4	Router # [no] platform ip multicast syslog	(Optional) Enables display of multicast related syslog messages on console.

This example shows how to enable MFIB forwarding on a Layer 3 interface:

```
Router(config-if)# ip mfib forwarding in
Router(config-if)#
```

Configuring the Replication Mode

The default for Cisco IOS Release 15.4SY is egress replication mode. Egress replication mode is always supported because the Supervisor Engine 6T provides egress replication for ingress-only legacy switching modules. Modules installed or added do not constrain the replication mode. You can change the configured replication mode.

To configure the replication mode, perform this task:

Command	Purpose
Router(config)# [no] platform ip multicast routing replication egress	Configures the replication mode. • The platform multicast routing replication egress command configures egress replication mode. • The no platform multicast routing replication egress command configures ingress replication mode. • Changing the replication mode might interrupt traffic.

This example shows how to configure ingress replication mode:

```
Router(config)# no platform multicast routing replication egress
```

This example shows how to display the replication mode:

```
Router# show platform multicast routing replication
Current mode of replication is Ingress
Configured mode of replication is Ingress

Slot                Multicast replication capability
2                   Egress
5                   Egress
```

Configuring Multicast Boundary

To configure a multicast boundary, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{ vlan <i>vlan_ID</i> } { <i>type slot/port</i> } { port-channel <i>number</i> }}	Selects an interface to configure.
Step 2	Router(config-if)# ip multicast boundary <i>access_list</i> [filter-autorp]	Enables an administratively scoped boundary on an interface. • For <i>access_list</i>, specify the access list that you have configured to filter the traffic at this boundary. • (Optional) Specify filter-autorp to filter auto-RP messages at this boundary.

**Note**

If you configure the **filter-autorp** keyword, the administrative boundary examines auto-RP discovery and announcement messages and removes any auto-RP group range announcements from the auto-RP packets that are denied by the boundary ACL. An auto-RP group range announcement is permitted and passed by the boundary only if all addresses in the auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the auto-RP message before the auto-RP message is forwarded.

The following example sets up a multicast boundary for all administratively scoped addresses:

```
Router (config)# access-list 1 deny 239.0.0.0 0.255.255.255
Router (config)# access-list 1 permit 224.0.0.0 15.255.255.255
Router (config)# interface gigabitethernet 5/2
Router (config-if)# ip multicast boundary 1
```

Verifying Local Egress Replication

DFC-equipped modules with dual switch-fabric connections host two packet replication engines, one per fabric connection. Each replication engine is responsible for forwarding packets to and from the interfaces associated with the switch-fabric connections. The interfaces that are associated with a switch-fabric connection are considered to be “local” from the perspective of the packet replication engine. When local egress replication mode is not enabled, both replication engines have the complete outgoing interface list for all modules, and the replication engines process and then drop traffic for nonlocal interfaces.

Local egress replication mode limits the outgoing interface list to only the local interfaces that each replication engine supports, which prevents unnecessary processing of multicast traffic.

Local egress replication is supported with the following software configuration and hardware:

- Egress replication mode.
- Dual fabric-connection DFC-equipped modules.
- CFC-equipped modules (functionality provided by the Supervisor Engine 6T).
- Members of Layer 3 EtherChannels and VLAN interfaces.

This example shows how to verify the replication engine selected for local egress replication:

```
Router# show platform multicast routing replication
Current mode of replication is Ingress
Configured mode of replication is Ingress

Slot                Multicast replication capability
2                   Egress
5                   Egress
```

Displaying IPv4 Multicast PIM-SM register tunnel information

To view the register tunnel interfaces associated with a PIM RP, enter the **show ip pim tunnel** command.

Command	Purpose
Router# show ip pim rp mapping	Displays PIM RP information.
Router# show ip pim tunnel	Displays PIM register encapsulation tunnel information on all routers. Displays additional PIM register decapsulation tunnel on the RP.

This example shows how to display the PIM register tunnel information:

```
Router# show ip pim rp mapping
PIM Group-to-RP Mappings
```

```
Group(s): 224.0.0.0/4, Static
RP: 11.1.1.1 (?)
```

```
Router# show ip pim tunnel
```

```
Tunnel0
  Type   : PIM Encap
  RP     : 11.1.1.1*
  Source : 11.1.1.1
Tunnel1*
  Type   : PIM Decap
  RP     : 11.1.1.1*
  Source : -
```

Displaying the IPv4 Multicast Routing Table

To view a mroute entry, enter the **show ip mroute** command.

Command	Purpose
Router# show ip mroute [<i>hostname</i> <i>group_number</i>]	Displays the IP multicast routing table.

This example shows how to display the IP multicast routing table:

```
Router# show ip mroute 225.1.1.1
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner
```

```
Timers: Uptime/Expires
```

```
Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(*, 225.1.1.1), 00:25:35/00:02:52, RP 11.1.1.1, flags: SJC
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
GigabitEthernet1/4, Forward/Sparse-Dense, 00:21:43/00:02:52
Vlan546, Forward/Sparse-Dense, 00:25:36/00:02:51

(22.2.2.1, 225.1.1.1), 00:25:36/00:01:49, flags: T
Incoming interface: Vlan546, RPF nbr 0.0.0.0
Outgoing interface list:
GigabitEthernet1/4, Forward/Sparse-Dense, 00:21:46/00:03:24
```

Displaying IPv4 MRIB Information

The **show ip mrib** command displays detailed information about IP MRIB. To display detailed MRIB information, perform one of these tasks:

Command	Purpose
Router# show ip mrib client	Displays the clients registered with MRIB.
Router# show ip mrib route [<i>hostname</i> <i>group_number</i> summary reserved]	Displays the MRIB route information.

To view the MRIB clients, enter the **show ip mrib client** command. MFIB running on the PFC and each DFC must present as clients of the MRIB. This example shows how to display the MRIB clients:

```
Router# show ip mrib client
IP MRIB client-connections
MRIB Trans for MVRP #0      table:434      (connection id 1)
IPv4_mfib(0x57D354C8):6.642  (connection id 2)
IPv4_mfib(0x555FC7D8):1.362  (connection id 3)
```

This example shows how to display the MRIB table:

```
Router# show ip mrib route 225.1.1.1
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
             C - Directly-Connected Check, S - Signal, IA - Inherit Accept, D - Drop
             ET - Data Rate Exceeds Threshold, K - Keepalive, DDE - Data Driven Event
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
                NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
                II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
                LD - Local Disinterest, MD - mCAC Denied

(*,225.1.1.1) RPF nbr: 0.0.0.0 Flags: C
Tunnell Flags: A
GigabitEthernet1/4 Flags: F NS
Vlan546 Flags: F NS

(22.2.2.1,225.1.1.1) RPF nbr: 0.0.0.0 Flags:
Vlan546 Flags: A
GigabitEthernet1/4 Flags: F NS
```

Displaying IPv4 MFIB Information

The **show ip mfib** command displays detailed information about IP MFIB. To display detailed MFIB information, perform one of these tasks:

Command	Purpose
Router# show ip mfib interface [<i>type name</i>]	Displays MFIB interface information
Router# show ip mfib [<i>hostname</i> <i>group_number</i> global verbose]	Displays the MFIB route information.
Router# show ip mfib summary	Displays a summary of MFIB state
Router# show ip mfib status	Displays MFIB status
Router# show ip mfib count	Displays route and packet count data

This example shows how to display the MFIB interface information:

```
Router# show ip mfib interface
IPv4 Multicast Forwarding (MFIB) status:
  Configuration Status: enabled
  Operational Status: running
  Initialization State: Running
  Total signalling packets queued: 0
  Process Status: may enable - 3 - pid 642
  Tables 1/1/0 (active/mrib/io)

MFIB interface          status    CEF-based output
                        [configured,available]
GigabitEthernet1/1      up        [yes      ,no      ]
GigabitEthernet1/2      up        [yes      ,no      ]
GigabitEthernet1/4      up        [yes      ,yes     ]
Loopback1               up        [yes      ,yes     ]
Tunnel0                 up        [yes      ,yes     ]
Tunnel1                 up        [yes      ,no      ]
Vlan546                 up        [yes      ,yes     ]
```

This example shows how to display the MFIB table:

```
Router# show ip mfib 225.1.1.1
Entry Flags:      C - Directly Connected, S - Signal, IA - Inherit A flag,
                  ET - Data Rate Exceeds Threshold, K - Keepalive
                  DDE - Data Driven Event, HW - Hardware Installed
I/O Item Flags:  IC - Internal Copy, NP - Not platform switched,
                  NS - Negate Signalling, SP - Signal Present,
                  A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                  MA - MFIB Accept
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   FS Pkt Count/PS Pkt Count
Default
(*,225.1.1.1) Flags: C HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding: 0/0/0/0, Other: 0/0/0
  Tunnel1 Flags: A
  GigabitEthernet1/4 Flags: F NS
    Pkts: 0/0
  Vlan546 Flags: F NS
    Pkts: 0/0
(22.2.2.1,225.1.1.1) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 37/0/37
```

```

HW Forwarding: 536649628/234964/512/939858, Other: 0/0/0
Vlan546 Flags: A
GigabitEthernet1/4 Flags: F NS
Pkts: 0/0

```

This example shows how to display the MFIB summary:

```

Router# show ip mfib summary
Default
211 prefixes (211/0/0 fwd/non-fwd/deleted)
343 ioitems (343/0/0 fwd/non-fwd/deleted)
Forwarding prefixes: [101 (S,G), 108 (*,G), 2 (*,G/m)]
Table id 0x0, instance 0x57D354C8
Database: epoch 2

```

This example shows how to display the MFIB status:

```

Router# show ip mfib status
IPv4 Multicast Forwarding (MFIB) status:
Configuration Status: enabled
Operational Status: running
Initialization State: Running
Total signalling packets queued: 0
Process Status: may enable - 3 - pid 642
Tables 1/1/0 (active/mrib/io)

```

This example shows how to display the MFIB count:

```

Router# show ip mfib count
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts:      Total/RPF failed/Other drops(OIF-null, rate-limit etc)
Default
211 routes, 108 (*,G)s, 2 (*,G/m)s
Group: 224.0.0.0/4
RP-tree,
SW Forwarding: 0/0/0/0, Other: 0/0/0
HW Forwarding: 0/0/0/0, Other: 0/0/0
Group: 224.0.1.40
RP-tree,
SW Forwarding: 0/0/0/0, Other: 0/0/0
Group: 225.1.1.1
RP-tree,
SW Forwarding: 0/0/0/0, Other: 0/0/0
HW Forwarding: 0/0/0/0, Other: 0/0/0
Source: 22.2.2.1,
SW Forwarding: 0/0/0/0, Other: 37/0/37
HW Forwarding: 38970288942/234961/512/939847, Other: 0/0/0
Totals - Source count: 1, Packet count: 38970288942

```

Viewing Directly Connected Entries

This example shows how to display the directly connected subnet entries:

```

Router# show platform hardware cef | include receive
224 0.0.0.0/32 receive
225 255.255.255.255/32 receive
226 22.2.2.2/32 receive
227 22.2.2.0/32 receive
228 22.2.2.255/32 receive
229 10.10.10.10/32 receive
230 11.1.1.1/32 receive
231 11.1.1.0/32 receive
232 11.1.1.255/32 receive

```

```

3200    224.0.0.0/24          receive
198433  224.0.0.0/4          receive

```

Displaying IPv4 Hardware Switching Information

The **show platform hardware multicast routing ip** command displays detailed information about IP multicast Layer 3 switching. To display detailed IP multicast Layer 3 switching information, perform one of these tasks:

Command	Purpose
Router# show platform hardware multicast routing ip control	Displays Layer 3 IP multicast control entries.
Router# show platform hardware multicast routing ip [source <i>ip_address</i>] [group <i>ip_address</i>] [detail verbose]	Displays IP multicast Layer 3 switching details for all interfaces.
Router# show platform ip multicast summary	Displays a summary of IP multicast Layer 3 switching information.

This example shows how to display the Layer 3 IP multicast control entries:

```

Router# show platform hardware multicast routing ip control
IPv4 Multicast CEF Entries for VPN#0
Flags: C - Control, B - Bidir, c - CoPP ELIF, Q - QoS ELIF, n - Non-primary Input
Source      Destination/mask      RPF/DF      Flags  #packets  #bytes
Output LIFs/Info
+-----+-----+-----+-----+-----+-----+
--++-----+-----+-----+-----+-----+
*           224.0.0.0/24      -           C      -        -
*           224.0.1.39/32    -           C      -        -
*           224.0.1.40/32    -           C      -        -
Found 3 entries.

```

This example shows how to display the Layer 3 IP multicast switching information:

```

Router# show platform hardware multicast routing ip source 22.2.2.1 group 225.1.1.1 de
IPv4 Multicast CEF Entries for VPN#0

(22.2.2.1, 225.1.1.1/32)
FIBAddr: 0x20 IOSVPN: 0 RpfType: SglRpfChk SrcRpf: V1546
CPx: 0 s_star_pri: 1 non-rpf drop: 0

PIAdjPtr: 0x1C001 Format: IP rdt: off elif: 0xC5409
fltr_en: off idx_sel/bndl_en: 0 dec_ttl: on mtu_idx: 2(1518)
PV: 1 rwtype: MCAST_L3_REWRITE
met3: 0x8000 met2: 0x0
Packets: 0          Bytes: 0

NPIAdjPtr: 0x1C002 Format: IP rdt: off elif: 0xC5409
fltr_en: off idx_sel/bndl_en: 0 dec_ttl: off
PV: 0 rwtype: NO_REWRITE
smac_rwt: 0 ip_to_mac: 1
Packets: 0          Bytes: 0
MET offset: 0x8000
      OIF      AdjPtr      Elif      CR
+-----+-----+-----+-----+
      EDT-2C001      0x2C001      0x8400A      6T1/T2
Found 1 entries.

```

This example shows how to display the Layer 3 IP multicast switching summary:

```
Router# show platform hardware multicast routing ip summary
IPv4 Multicast CEF Entries Summary for VPN#0
  Slot    #shcut    # (S,G)    # (*,G)    # (*,G/m)    #Ctrl
+-----+-----+-----+-----+-----+
  6        203        101        101         1           3
```

Displaying IPv4 CoPP Information

The following commands can be used to check the multicast CoPP information:

Command	Purpose
Router# show platform hardware multicast routing ip control	Displays Layer 3 IP multicast control entries.
Router# show policy-map control-plane input class <i>class_name</i>	Displays class-map information for the control-plane policy.

This example shows how to display CoPP information:

```
Router# show platform hardware multicast routing ip control detail
IPv4 Multicast CEF Entries for VPN#0

(*, 224.0.0.0/24)
  FIBAddr: 0x2A0 IOSVPN: 0 RpfType: SkipRpf SrcRpf:
  CPx: 0 s_star_pri: 1 non-rpf drop: 0

  PIAAdjPtr: 0x13110(IPv4 Control) Format: IP rdt: off elif: 0x9FC01
  fltr_en: off idx_sel/bndl_en: 0 dec_ttl: off mtu_idx: 0(9234)
  PV: 1 rwtype: NO_REWRITE
  smac_rwt: 0 ip_to_mac: 0

(*, 224.0.1.39/32)
  FIBAddr: 0x1E0 IOSVPN: 0 RpfType: SkipRpf SrcRpf:
  CPx: 0 s_star_pri: 1 non-rpf drop: 0

  PIAAdjPtr: 0x13110(IPv4 Control) Format: IP rdt: off elif: 0x9FC01
  fltr_en: off idx_sel/bndl_en: 0 dec_ttl: off mtu_idx: 0(9234)
  PV: 1 rwtype: NO_REWRITE
  smac_rwt: 0 ip_to_mac: 0

(*, 224.0.1.40/32)
  FIBAddr: 0x1E2 IOSVPN: 0 RpfType: SkipRpf SrcRpf:
  CPx: 0 s_star_pri: 1 non-rpf drop: 0

  PIAAdjPtr: 0x13110(IPv4 Control) Format: IP rdt: off elif: 0x9FC01
  fltr_en: off idx_sel/bndl_en: 0 dec_ttl: off mtu_idx: 0(9234)
  PV: 1 rwtype: NO_REWRITE
  smac_rwt: 0 ip_to_mac: 0
Found 3 entries.
```

```

Router# show policy-map control-plane input class class-copp-match-igmp

Control Plane Interface

Service-policy input: policy-default-autocopp

Hardware Counters:

class-map: class-copp-match-igmp (match-any)
  Match: access-group name acl-copp-match-igmp
  police :
    10000 pps 10000 limit 10000 extended limit
  Earl in slot 6 :
    0 packets
    5 minute offered rate 0 pps
    aggregate-forwarded 0 packets
    action: set-discard-class-transmit
    exceeded 0 packets action: transmit
    aggregate-forward 0 pps exceed 0 pps

Software Counters:

Class-map: class-copp-match-igmp (match-any)
  7138 packets, 267084 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group name acl-copp-match-igmp
    7138 packets, 267084 bytes
    5 minute rate 0 bps
  police:
    rate 10000 pps, burst 10000 packets
    conformed 7138 packets, 7138 bytes; action:
      set-discard-class-transmit 48
    exceeded 0 packets, 0 bytes; action:
      transmit
    conformed 0 pps, exceeded 0 pps

```

Source-Specific Multicast with IGMPv3, IGMP v3lite, and URD

For complete information and procedures about source-specific multicast (SSM) with IGMPv3, IGMP v3lite, and URL Rendezvous Directory (URD), see this URL:

http://www.cisco.com/en/US/docs/ios-xml/ios/ipmulti_igmp/configuration/15-sy/imc-igmp-15-sy-book.html

Configuring IPv4 Bidirectional PIM

- [Enabling IPv4 Bidirectional PIM Globally, page 6-27](#)
- [Configuring the Rendezvous Point for IPv4 Bidirectional PIM Groups, page 6-27](#)
- [Displaying IPv4 Bidirectional PIM Information, page 6-27](#)

Enabling IPv4 Bidirectional PIM Globally

To enable IPv4 bidirectional PIM, perform this task:

Command	Purpose
Router(config)# ip pim bidir-enable	Enables IPv4 bidirectional PIM globally on the switch.

This example shows how to enable IPv4 bidirectional PIM on the switch:

```
Router(config)# ip pim bidir-enable
Router(config)#
```

Configuring the Rendezvous Point for IPv4 Bidirectional PIM Groups

To statically configure the rendezvous point for an IPv4 bidirectional PIM group, perform this task:

	Command	Purpose
Step 1	Router(config)# ip pim rp-address <i>ip_address</i> <i>access_list</i> [override]	Statically configures the IP address of the rendezvous point for the group. When you specify the override option, the static rendezvous point is used.
Step 2	Router(config)# access-list <i>access-list</i> [permit deny] <i>ip_address</i>	Configures an access list.
Step 3	Router(config)# ip pim send-rp-announce <i>type</i> <i>number</i> scope <i>ttl_value</i> [group-list <i>access-list</i>] [interval <i>seconds</i>] [bidir]	Configures the system to use auto-RP to configure groups for which the router will act as a rendezvous point (RP).
Step 4	Router(config)# ip access-list standard <i>access-list-name</i> permit deny <i>ip_address</i>	Configures a standard IP access list.
Step 5	Router(config)# platform ip multicast	Enables hardware-supported IP multicast.

This example shows how to configure a static rendezvous point for an IPv4 bidirectional PIM group:

```
Router(config)# ip pim rp-address 10.0.0.1 10 bidir override
Router(config)# access-list 10 permit 224.1.0.0 0.0.255.255
Router(config)# ip pim send-rp-announce Loopback0 scope 16 group-list c21-rp-list-0 bidir
Router(config)# ip access-list standard c21-rp-list-0 permit 230.31.31.1 0.0.255.255
```

Displaying IPv4 Bidirectional PIM Information

To display IPv4 bidirectional PIM information, perform one of these tasks:

Command	Purpose
Router# show ip pim rp mapping [<i>in-use</i>]	Displays mappings between PIM groups and rendezvous points and shows learned rendezvous points in use.
Router# show ip mfib	Displays MFIB information for bidirectional PIM.
Router# show platform hardware multicast routing ip	Displays IP multicast Layer 3 switching details.

Command	Purpose
Router# show platform software multicast routing cmrp info	Displays information about the DF indices and DF masks allocated.
Router# show platform ip multicast bidir	Displays IPv4 bidirectional PIM information.

This example shows how to display information about the PIM group and rendezvous point mappings:

```
Router# show ip pim rp mapping
PIM Group-to-RP Mappings
This system is an RP (Auto-RP)
This system is an RP-mapping agent
Group(s) 230.31.0.0/16
  RP 60.0.0.60 (?), v2v1, bidir
    Info source:60.0.0.60 (?), elected via Auto-RP
    Uptime:00:03:47, expires:00:02:11
  RP 50.0.0.50 (?), v2v1, bidir
    Info source:50.0.0.50 (?), via Auto-RP
    Uptime:00:03:04, expires:00:02:55
  RP 40.0.0.40 (?), v2v1, bidir
    Info source:40.0.0.40 (?), via Auto-RP
    Uptime:00:04:19, expires:00:02:38
```

This example shows how to display information in the IP multicast routing table that is related to IPv4 bidirectional PIM:

```
Router# show ip mroute bidirectional
(*,224.0.0.0/4), 00:17:18/-, RP 20.2.2.2, flags: B
  Bidir-Upstream: Loopback2, RPF nbr: 20.2.2.2
  Incoming interface list:
    TenGigabitEthernet3/8, Accepting/Sparse-Dense
    TenGigabitEthernet3/7, Accepting/Sparse-Dense
    GigabitEthernet1/12, Accepting/Dense
    GigabitEthernet1/1, Accepting/Sparse-Dense
    Port-channel2, Accepting/Sparse-Dense
    Loopback100, Accepting/Sparse-Dense
    Loopback10, Accepting/Sparse-Dense
    Loopback2, Accepting/Sparse-Dense

(*, 224.1.1.1), 00:17:18/00:02:26, RP 20.2.2.2, flags: BC
  Bidir-Upstream: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    TenGigabitEthernet3/8, Forward/Sparse-Dense, 00:17:17/00:02:26
```

This example shows how to display the MFIB table related to IPv4 bidirectional PIM:

```
Router# show ip mfib
Entry Flags:      C - Directly Connected, S - Signal, IA - Inherit A flag,
                  ET - Data Rate Exceeds Threshold, K - Keepalive
                  DDE - Data Driven Event, HW - Hardware Installed
I/O Item Flags:  IC - Internal Copy, NP - Not platform switched,
                  NS - Negate Signalling, SP - Signal Present,
                  A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                  MA - MFIB Accept
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   FS Pkt Count/PS Pkt Count
Default
(*,224.0.0.0/4) Flags: HW  ?- Indicates that Entry is installed in hardware
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding: 0/0/0/0, Other: 0/0/0
  GigabitEthernet1/1 Flags: A
```

```

GigabitEthernet1/12 Flags: A
TenGigabitEthernet3/7 Flags: A
TenGigabitEthernet3/8 Flags: A
Port-channel2 Flags: A
Loopback100 Flags: A
Loopback10 Flags: A
Loopback2 Flags: A F
  Pkts: 0/0
Null0 Flags: A
(*,224.1.1.1) Flags: IA HW
SW Forwarding: 0/0/0/0, Other: 3/3/0
HW Forwarding: 3357168/3000/100/2343, Other: 0/0/0 ?--- Hw Forwarding Statistics
TenGigabitEthernet3/7 Flags: F
  Pkts: 0/0
TenGigabitEthernet3/8 Flags: F
  Pkts: 0/0

```

This example shows how to display the hardware-switching table related to IPv4 bidirectional PIM:

```

Router# show platform hardware multicast routing ip group 224.0.0.0/4
IPv4 Multicast CEF Entries for VPN#0
Flags: C - Control, B - Bidir, c - CoPP ELIF, Q - QoS ELIF, n - Non-primary Input
Source      Destination/mask      RPF/DF      Flags #packets      #bytes
Output LIFs/Info
+-----+-----+-----+-----+-----+-----+
--+-----+-----+-----+-----+-----+-----+
*          224.0.0.0/4          0      B      0          0
*          224.1.1.1/32        0      B      6539001     653900100
Te3/7, Te3/8 [2 oif(s)]
Found 2 entries.

```

This example shows how to display the DF-index and DF mask information related to IPv4 bidirectional PIM:

```

Router# show platform software multicast routing cmrp info

Replication Mode Information:
=====
  repl_mode: Ingress; pending_repl_mode: Ingress;
  notify_repl_mode: Ingress; repl_mode_chng_in_prog: 0;
  repl_mode_notif_pending: 0

Resource Information:
=====
  fib_full_mask: 0x0; adj_full_mask: 0x0;
  met_full_mask: 0x0; met_unavail_mask: 0x0

Global HA Information:
=====
  reconstruct_in_prog: 0; reconstruct_lc_mask: 0x0

Vrf Name: R; Vrf ID: 3; Df Idx Allocated: 8; mfib_sweep_mask: 0x0
=====
  DF Idx Info for df_idx: 0
  =====
  coll_obj: 0x530493CC; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x19D1D1A8; vrf_id: 3; df_idx: 0; cleanup: 0
  DF Set
  =====
  Tu100 (0x320004072), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),

```

```

DF Idx Info for df_idx: 1
=====
coll_obj: 0x530493EC; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x19C3C740; vrf_id: 3; df_idx: 1; cleanup: 0
  DF Set
  =====
  Tu101 (0x320004073), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),
DF Idx Info for df_idx: 2
=====
coll_obj: 0x5304936C; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x19D4D938; vrf_id: 3; df_idx: 2; cleanup: 0
  DF Set
  =====
  Tu102 (0x320004074), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),
DF Idx Info for df_idx: 3
=====
coll_obj: 0x530492EC; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x19CE6338; vrf_id: 3; df_idx: 3; cleanup: 0
  DF Set
  =====
  Tu103 (0x320004075), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),
DF Idx Info for df_idx: 4
=====
coll_obj: 0x5304930C; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x5C092894; vrf_id: 3; df_idx: 4; cleanup: 0
  DF Set
  =====
  Tu104 (0x320004076), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),
DF Idx Info for df_idx: 5
=====
coll_obj: 0x530493AC; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x19D4D7FC; vrf_id: 3; df_idx: 5; cleanup: 0
  DF Set
  =====
  Tu105 (0x320004077), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),
DF Idx Info for df_idx: 6
=====
coll_obj: 0x5304934C; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x19D4DC58; vrf_id: 3; df_idx: 6; cleanup: 0
  DF Set
  =====
  Tu106 (0x320004078), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),
DF Idx Info for df_idx: 7
=====
coll_obj: 0x5304938C; af:ipv4; df_state: USED
  DF Collection Obj Info
  =====
  ref_count: 1; id_count: 4; app_data: 0x5C09AD64; vrf_id: 3; df_idx: 7; cleanup: 0
  DF Set
  =====
  Tu107 (0x320004079), Tu108 (0x32000407A), Nu0 (0x32000407B), Tu3 (0x32000407F),

```

Using IPv4 Debug Commands

Table 6-1 describes IPv4 multicast Layer 3 switching debug commands that you can use to troubleshoot IP multicast Layer 3 switching problems.

Table 6-1 *IP Multicast Layer 3 Switching Debug Commands*

Command	Description
<code>[no] debug platform software multicast routing cmrp {event error}</code> <code>[verbose]</code>	Debug CMRP related events and errors.
<code>[no] debug platform software multicast routing edc server</code> <code>[event error]</code>	Debug errors and events for the egress distribution server component.
<code>[no] debug platform software multicast routing edc client</code> <code>[event error]</code>	Debug errors and events for the egress distribution client component.
<code>[no] debug platform software multicast routing hal [event error]</code>	Debug errors and events for the multicast hardware abstraction layer.
<code>[no] debug platform software multicast routing cmfib [event error]</code>	Debug errors and events for the Constellation MFIB component.
<code>[no] debug platform software met [event error detail all]</code>	Debug errors and events for MET manager.
<code>[no] debug platform software filter filter_id</code> <code>{ip {destination source} ip_address [mask]</code> <code> string {exclude include} text_string}</code>	Turns on filtering for debug messages based on IPv4 destination address or source address or an input string.

Redundancy for Multicast Traffic

Redundancy for multicast traffic requires the following conditions:

- Unicast routing protocol such as OSPF or EIGRP:
PIM uses RPF checks on the unicast routing table to determine the proper paths for multicast data to traverse. If a unicast routing path changes, PIM relies upon the unicast routing protocol (OSPF) to properly converge, so that the RPF checks used by PIM continue to work and show valid unicast paths to and from the source IP address of the server sourcing the multicast stream.
- PIM configured on all related Layer 3 interfaces:
The unicast routing table is used to do path selection for PIM. PIM uses RPF checks to ultimately determine the shortest path tree (SPT) between the client (receiver VLAN) and the source (multicast VLAN). Therefore, the objective of PIM is to find the shortest unicast path between the receiver subnet and the source subnet. You do not need to configure anything else for multicast when the unicast routing protocol is working as expected and PIM is configured on all the Layer 3 links associated with the unicast routing protocol.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

7

IGMP Snooping for IPv4 Multicast Traffic

- [Prerequisites for IGMP Snooping, page 7-1](#)
- [Restrictions for IGMP Snooping, page 7-1](#)
- [Information About IGMP Snooping, page 7-2](#)
- [Default Settings for IGMP Snooping, page 7-8](#)
- [How to Configure IGMP Snooping, page 7-8](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - To constrain IPv6 Multicast traffic, see [Chapter 14, “IPv6 MLD Snooping.”](#)
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IGMP Snooping

None.

Restrictions for IGMP Snooping

- [General IGMP Snooping Restrictions, page 7-2](#)
- [IGMP Snooping Querier Restrictions, page 7-2](#)

General IGMP Snooping Restrictions

- Multicast packets are not bridged in a VLAN to local receivers that send IGMP joins when PIM snooping is enabled in the VLAN and IGMP snooping is disabled in the VLAN. ([CSCta03980](#))
- For more information on IP multicast and IGMP, see RFC 1112 and RFC 2236.
- IGMP snooping supports private VLANs. Private VLANs do not impose any restrictions on IGMP snooping.
- IGMP snooping constrains traffic in MAC multicast groups 0100.5e00.0001 to 0100.5eff.ffff.
- IGMP snooping does not constrain Layer 2 multicasts generated by routing protocols.

IGMP Snooping Querier Restrictions

- The IGMP snooping querier does not support querier elections. Enable the IGMP snooping querier on only one switch in the VLAN. ([CSCsk48795](#))
- Configure the VLAN in global configuration mode (see [Chapter 26, “Virtual Local Area Networks \(VLANs\)”](#)).
- Configure an IP address on the VLAN interface (see [Chapter 35, “Layer 3 Interfaces”](#)). 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 does not start. The IGMP snooping querier disables itself if the IP address is cleared. When enabled, the IGMP snooping querier restarts if you configure an IP address.
- The IGMP snooping querier sends IGMPv3 querier messages. Although the IGMP version of the querier messages is not configurable, the querier is compatible with IGMPv2 hosts.
- When enabled, the IGMP snooping querier starts after 60 seconds with no IGMP traffic detected from a multicast router. If IGMP traffic from a multicast router, or from another IGMP snooping querier in the VLAN, is detected after the IGMP snooping querier has started, the querier will disable itself.
- QoS does not support IGMP packets when IGMP snooping is enabled.
- With Release 15.1(1)SY1 and later releases, IGMP snooping and PIM snooping constrain VPLS multicast traffic.

Information About IGMP Snooping

- [IGMP Snooping Overview, page 7-3](#)
- [Joining a Multicast Group, page 7-3](#)
- [Leaving a Multicast Group, page 7-5](#)
- [Information about the IGMP Snooping Querier, page 7-6](#)
- [Information about IGMP Version 3 Support, page 7-6](#)

IGMP Snooping Overview

IGMP snooping allows switches to examine IGMP packets and make forwarding decisions based on their content. You can configure the switch to use IGMP snooping in subnets that receive IGMP queries from either IGMP or the IGMP snooping querier. IGMP snooping constrains IPv4 multicast traffic at Layer 2 by configuring Layer 2 LAN ports dynamically to forward IPv4 multicast traffic only to those ports that want to receive it.

Some applications use a single unicast cluster IP address and multicast cluster MAC address. Multicast traffic addressed to a unicast cluster IP address is forwarded to the last-hop router that is configured with the shared multicast MAC address. To support cluster-addressed multicast traffic, assign a static multicast MAC address for the destination IP address of the end host or cluster.

You can configure the IGMP snooping lookup method for each VLAN. Layer 3 IGMP snooping lookup uses destination IP addresses in the Layer 2 multicast table (this is the default). Layer 2 IGMP snooping lookup uses destination MAC addresses in the Layer 2 multicast table.

**Note**

Changing the lookup mode is disruptive. Multicast forwarding is not optimal until all multicast entries are programmed with the new lookup mode. Also, if 32 IP addresses are mapped to a single MAC address, forwarding on the device might be suboptimal.

IGMP, which runs at Layer 3 on a multicast router, generates Layer 3 IGMP queries in subnets where the multicast traffic needs to be routed.

You can configure the IGMP snooping querier on the switch to support IGMP snooping in subnets that do not have any multicast router interfaces. For more information about the IGMP snooping querier, see the [“Enabling the IGMP Snooping Querier” section on page 7-9](#).

IGMP (on a multicast router) or, locally, the IGMP snooping querier, sends out periodic general IGMP queries that the switch forwards through all ports in the VLAN and to which hosts respond. IGMP snooping monitors the Layer 3 IGMP traffic.

**Note**

If a multicast group has only sources and no receivers in a VLAN, IGMP snooping constrains the multicast traffic to only the multicast router ports.

Joining a Multicast Group

Hosts join multicast groups either by sending an unsolicited IGMP join message or by sending an IGMP join message in response to a general query from a multicast router (the switch forwards general queries from multicast routers to all ports in a VLAN).


In response to an IGMP join request, the switch creates an entry in its Layer 2 forwarding table for the VLAN on which the join request was received. When other hosts that are interested in this multicast traffic send IGMP join requests, the switch adds them to the existing Layer 2 forwarding table entry. The switch creates only one entry per VLAN in the Layer 2 forwarding table for each multicast group for which it receives an IGMP join request.

IGMP snooping suppresses all but one of the host join messages per multicast group and forwards this one join message to the multicast router.

The switch forwards multicast traffic for the multicast group specified in the join message to the interfaces where join messages were received (see [Figure 7-1](#)).

Layer 2 multicast groups learned through IGMP snooping are dynamic. However, you can statically configure Layer 2 multicast groups using the **mac address-table static** command. When you specify group membership for a multicast group address statically, the static setting supersedes any IGMP snooping learning. Multicast group membership lists can consist of both static and IGMP snooping-learned settings.

Figure 7-1 Initial IGMP Join Message



Multicast router A sends a general query to the switch, which forwards the query to ports 2 through 5 (all members of the same VLAN). Host 1 wants to join multicast group 224.1.2.3 and multicasts an IGMP membership report (IGMP join message) to the group with the equivalent MAC destination address of 0x0100.5E01.0203. When the CPU receives the IGMP report multicast by Host 1, the CPU uses the information in the IGMP report to set up a forwarding-table entry, that includes the port numbers of Host 1, the multicast router, and the switch internal CPU.

CPU installs snooping forwarding entry based on lookup type, either IP or MAC (by default, it is IP-based). Using IP-based forwarding can avoid group address aliasing problem and optimize per group or per group and source forwarding.

If IP-based is configured, IGMP snooping forwarding table has the following entry. The switch engine matches on the destination IP address of multicast data packets. If they are 224.1.2.3, send them to the host that has joined the group and multicast routers.


vlan	mac/ip address	LTL	ports
200	(*,224.1.2.3)	0x924	Router, Gi3/11

If MAC-based is configured, the entry is as follows. In this case, the switch engine matches on the destination MAC address of the packets. The packets with 0100.5e01.0203 are sent to the host that has joined the group and multicast routers.

vlan	mac/ip address	LTL	ports
200	0100.5e01.0203	0x92C	Router, Gi3/11

If another host (for example, Host 4) sends an unsolicited IGMP join message for the same group (Figure 7-2), the CPU receives that message and adds the port number of Host 4 to the forwarding table. Because the forwarding table directs IGMP messages only to the CPU, the message is not flooded to other ports. Any known multicast traffic is forwarded to the group and not to the CPU.

Figure 7-2 Second Host Joining a Multicast Group



Leaving a Multicast Group

- [Normal Leave Processing, page 7-5](#)
- [Immediate-Leave Processing, page 7-6](#)

Normal Leave Processing

Interested hosts must continue to respond to the periodic general IGMP queries. As long as at least one host in the VLAN responds to the periodic general IGMP queries, the multicast router continues forwarding the multicast traffic to the VLAN. When hosts want to leave a multicast group, they can either ignore the periodic general IGMP queries (called a “silent leave”), or they can send a group-specific IGMPv2 leave message.

When IGMP snooping receives a group-specific IGMPv2 leave message from a host, it sends out a MAC-based general query to determine if any other devices connected to that interface are interested in traffic for the specific multicast group. If IGMP snooping does not receive an IGMP Join message in response to the general query, it assumes that no other devices connected to the interface are interested in receiving traffic for this multicast group, and it removes the interface from its Layer 2 forwarding table entry for that multicast group. If the leave message was from the only remaining interface with hosts interested in the group and IGMP snooping does not receive an IGMP Join in response to the general

query, it removes the group entry and relays the IGMP leave to the multicast router. If the multicast router receives no reports from a VLAN, the multicast router removes the group for the VLAN from its IGMP cache.

The interval for which the switch waits before updating the table entry is called the “last member query interval.” To configure the interval, enter the **ip igmp snooping last-member-query-interval** *interval* command.

Immediate-Leave Processing

IGMP snooping immediate-leave processing allows IGMP snooping to remove a Layer 2 LAN interface from the forwarding-table entry without first sending out IGMP group-specific queries to the interface. Upon receiving a group-specific IGMPv2 leave message, IGMP snooping immediately removes the interface from the Layer 2 forwarding table entry for that multicast group, unless a multicast router was learned on the port. Immediate-leave processing improves bandwidth management for all hosts on a switched network.



Note

Use immediate-leave processing only on VLANs where only one host is connected to each Layer 2 LAN port. If immediate-leave is enabled in VLANs where more than one host is connected to a Layer 2 LAN port, some hosts might be dropped inadvertently. Immediate-leave processing is supported only with IGMP version 2 and 3 hosts.

Information about the IGMP Snooping Querier

Use the IGMP snooping querier to support IGMP snooping in a VLAN where PIM and IGMP are not configured because the multicast traffic does not need to be routed.

In a network where IP multicast routing is configured, the IP multicast router acts as the IGMP querier. If the IP-multicast traffic in a VLAN only needs to be Layer 2 switched, an IP-multicast router is not required, but without an IP-multicast router on the VLAN, you must configure another switch as the IGMP querier so that it can send queries.

When enabled, the IGMP snooping querier sends out periodic IGMPv3 queries that trigger IGMP report messages from the switch that wants to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

Configure one switch as the IGMP snooping querier in each VLAN that is supported on switches that use IGMP to report interest in IP multicast traffic.

You can configure a switch to generate IGMP queries on a VLAN regardless of whether or not IP multicast routing is enabled.

Information about IGMP Version 3 Support

- [IGMP Version 3 Support Overview, page 7-7](#)
- [IGMPv3 Immediate-Leave Processing, page 7-7](#)
- [Proxy Reporting, page 7-7](#)
- [Explicit Host Tracking, page 7-8](#)

IGMP Version 3 Support Overview

IGMP snooping supports IGMP version 3 (IGMPv3). IGMPv3 uses source-based filtering, which enables hosts and routers to specify which source addresses should be allowed or blocked for a specific multicast group. When you enable IGMPv3 snooping, the switch maintains IGMPv3 states based on messages it receives for a particular group in a particular VLAN and either allows or blocks traffic based on the following information in these messages:

- Source lists
- Allow (include) or block (exclude) filtering options

When a host wants to receive multicast traffic only from specific sources, it can send IGMPv3 joins with source filtering. For example, when host1 on port 3/11 sends join to group 224.1.2.3 from sources 10.1.1.1 and host2 on port 3/12 sends join to the same group but from a different source 20.1.1.1, the following entries are installed in forwarding table when IP-based lookup is configured:

vlan	mac/ip address	LTL	ports
200	(*, 224.1.2.3)	0x920	
200	(10.1.1.1, 224.1.2.3)	0x93E	Gi3/11
200	(20.1.1.1, 224.1.2.3)	0x940	Gi3/12

The second entry constrain group traffic from source 10.1.1.1 to host1 only and the third entry constrain traffic from source 20.1.1.1 to Host2. The first entry drops group traffic from any other sources since there is no receiver interesting in other sources.

IGMPv3 Immediate-Leave Processing

IGMPv3 immediate-leave processing is active if explicit-host tracking is enabled. The **ip igmp snooping immediate-leave** command that enables IGMP version 2 immediate-leave processing does not affect IGMPv3 immediate-leave processing.

Immediate-leave processing with IGMPv3 is implemented by maintaining source-group based membership information in software while also allocating LTL indexes on a MAC GDA basis.

When immediate-leave processing is active, hosts send **BLOCK_OLD_SOURCES{src-list}** messages for a specific group when they no longer want to receive traffic from that source. When the switch receives such a message from a host, it parses the list of sources for that host for the given group. If this source list is exactly the same as the source list received in the leave message, the switch removes the host from the LTL index and stops forwarding this multicast group traffic to this host.

If the source lists do not match, the switch does not remove the host from the LTL index until the host is no longer interested in receiving traffic from any source.

Proxy Reporting

IGMP supports proxy reporting for IGMPv1 and IGMPv2 messages to handle group-specific queries. These queries are not sent downstream, but the switch does respond to them directly. When the switch receives a group-specific query, the switch terminates the query and sends an IGMP proxy report if there is a receiver for the group. There is no proxy reporting for IGMPv3 messages. For IGMPv3, a group-specific query or a group source-specific query is flooded to all VLAN member ports. The database for the IGMPv3 membership report is built based on the reports received.

Host reports responding to a specific query can be suppressed by the report suppression feature. Report suppression is supported for IGMPv1, IGMPv2, and IGMPv3 messages. With report suppression enabled (by default), when the switch receives a general query, the switch starts a suppression cycle for

reports from all hosts to each group or channel (S,G). Only the first report to the discovered multicast routers is forwarded; the rest of the reports are suppressed. For IGMPv1 and IGMPv2, the time of suppression is the report response time indicated in the general query message. For IGMPv3, suppression occurs for the entire general query interval.

**Note**

The states are maintained only in software and used for explicit host tracking and statistics collection.

Explicit Host Tracking

IGMPv3 supports explicit tracking of membership information on any port. The explicit-tracking database is used for immediate-leave processing for IGMPv3 hosts, proxy reporting, and statistics collection. When explicit tracking is enabled on a VLAN, the IGMP snooping software processes the IGMPv3 report it receives from a host and builds an explicit-tracking database that contains the following information:

- The port connected to the host
- The channels reported by the host
- The filter mode for each group reported by the host
- The list of sources for each group reported by the hosts
- The router filter mode of each group
- For each group, the list of hosts requesting the source

**Note**

- When explicit tracking is enabled and the switch is working in proxy-reporting mode, the router may not be able to track all the hosts behind a VLAN interface.

Default Settings for IGMP Snooping

None.

How to Configure IGMP Snooping

- [Enabling the IGMP Snooping Querier, page 7-9](#)
- [Enabling IGMP Snooping, page 7-9](#)
- [Configuring the IGMP Snooping Lookup Method, page 7-11](#)
- [Configuring a Static Connection to a Multicast Receiver, page 7-11](#)
- [Configuring a Multicast Router Port Statically, page 7-12](#)
- [Configuring the IGMP Snooping Query Interval, page 7-12](#)
- [Enabling IGMP Snooping Immediate-Leave Processing, page 7-13](#)
- [Configuring IGMPv3 Snooping Explicit Host Tracking, page 7-13](#)
- [Displaying IGMP Snooping Information, page 7-14](#)

**Note**

To use IGMP snooping, configure a Layer 3 interface in the subnet for multicast routing (see [Chapter 6, “IPv4 Multicast Layer 3 Features”](#)) or enable the IGMP snooping querier in the subnet (see the [“Enabling the IGMP Snooping Querier”](#) section on page 7-9).

Enabling the IGMP Snooping Querier

Use the IGMP snooping querier to support IGMP snooping in a VLAN where PIM and IGMP are not configured because the multicast traffic does not need to be routed. To enable the IGMP snooping querier in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# ip igmp snooping querier	Enables the IGMP snooping querier globally.
Step 2	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 3	Router(config-vlan-config)# ip igmp snooping querier address <i>ip_address</i>	Assigns the IP address.
Step 4	Router(config-vlan-config)# ip igmp snooping querier	Enables the IGMP snooping querier on the VLAN.
Step 5	Router(config-vlan-config)# end	Exits configuration mode.

This example shows how to enable the IGMP snooping querier on VLAN 200 and verify the configuration:

```
Router(config)# ip igmp snooping querier
Router(config)# vlan configuration 200
Router(config-vlan-config)# ip igmp snooping querier address 10.1.1.1
Router(config-vlan-config)# igmp snooping querier
Router(config-vlan-config)# end
```

Enabling IGMP Snooping

- [Enabling IGMP Snooping Globally, page 7-9](#)
- [Enabling IGMP Snooping in a VLAN, page 7-10](#)

Enabling IGMP Snooping Globally

To enable IGMP snooping globally, perform this task:

	Command	Purpose
Step 1	Router(config)# ip igmp snooping	Enables IGMP snooping.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable IGMP snooping globally and verify the configuration:

```
Router(config)# ip igmp snooping
Router(config)# end
Router# show ip igmp interface vlan 200 | include globally
```

```

    IGMP snooping is globally enabled
Router#

```

Enabling IGMP Snooping in a VLAN

To enable IGMP snooping in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ip igmp snooping	Enables IGMP snooping.
Step 3	Router(config-vlan-config)# end	Exits configuration mode.

This example shows how to enable IGMP snooping on VLAN 25 and verify the configuration:

```

Router# vlan configuration 25
Router(config-vlan-config)# ip igmp snooping
Router(config-vlan-config)# end
Router# show ip igmp snooping vlan 25
Global IGMP Snooping configuration:
-----
IGMP snooping Oper State      : Enabled
IGMPv3 snooping               : Enabled
Report suppression            : Disabled
EHT DB limit/count            : 100000/2
TCN solicit query             : Disabled
Robustness variable           : 2
Last member query count        : 3
Last member query interval     : 1000
Check TTL=1                    : No
Check Router-Alert-Option      : No

Vlan 25:
-----
IGMP snooping Admin State      : Enabled
IGMP snooping Oper State      : Enabled
IGMPv2 immediate leave         : Disabled
Explicit host tracking          : Enabled
Report suppression             : Enabled
Robustness variable            : 2
Last member query count         : 2
Last member query interval      : 1000
EHT DB limit/count             : 100000/2
Check TTL=1                    : Yes
Check Router-Alert-Option       : Yes
Query Interval                  : 100
Max Response Time               : 10000
Router#

```


Configuring the IGMP Snooping Lookup Method

To configure the IGMP snooping lookup method for a VLAN, perform this task:

Command	Purpose
Router(config-vlan-config)# multicast snooping lookup { ip mac }	Configures the IGMP snooping lookup method for the VLAN. - Enter the ip keyword to use IP addresses to forward multicast traffic. - Enter the mac keyword to use destination MAC addresses to forward multicast traffic.



Note

Changing the lookup mode is disruptive. Multicast forwarding is not optimal until all multicast entries are programmed with the new lookup mode. Also, if 32 IP addresses are mapped to a single MAC address, forwarding on the device might be suboptimal.

Configuring a Static Connection to a Multicast Receiver

To configure a static connection to a multicast receiver, perform this task:

Command	Purpose
Router(config)# mac address-table static <i>mac_addr</i> vlan <i>vlan_id</i> interface <i>type slot/port</i> [disable-snooping]	Configures a static connection to a multicast receiver.

When you configure a static connection, enter the **disable-snooping** keyword to prevent multicast traffic addressed to the statically configured multicast MAC address from also being sent to other ports in the same VLAN.

This example shows how to configure a static connection to a multicast receiver:

```
Router(config)# mac address-table static 0050.3e8d.6400 vlan 12 interface gigabitethernet 5/7
```

The above static mac command can be used when the lookup type in the VLAN is MAC-base.

Irrespective of lookup type, the following commands can be used to configure static connection to a multicast receiver for a group or a group and from a specific source.

```
Router(config)# vlan configuration 200
Router(config-vlan-config)# ip igmp snooping static 224.1.2.3 interface g3/11
Router(config-vlan-config)# ip igmp snooping static 224.1.2.3 source 20.1.1.1 interface Gi3/12
```

Configuring a Multicast Router Port Statically

To configure a static connection to a multicast router, perform this task:

	Command	Purpose
Step 1	Router(config-vlan-config)# ip igmp snooping mrouter interface <i>type slot/port</i>	Configures a static connection to a multicast router.
Step 2	Router(config-vlan-config)# end	Exits configuration mode.

The interface to the router must be in the VLAN where you are entering the command, the interface must be administratively up, and the line protocol must be up.

This example shows how to configure a static connection to a multicast router:

```
Router(config-if)# ip igmp snooping mrouter interface gigabitethernet 5/6
```

Configuring the IGMP Snooping Query Interval

You can configure the interval for which the switch waits after sending a group-specific query to determine if hosts are still interested in a specific multicast group.



Note

When both IGMP immediate-leave processing and the IGMP query interval are configured, immediate-leave processing takes precedence.

To configure the interval for the IGMP snooping queries sent by the switch, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ip igmp snooping last-member-query-interval <i>interval</i>	Configures the interval for the IGMP snooping queries sent by the switch. Default is 1 second. Valid range is 100 to 999 milliseconds.

This example shows how to configure the IGMP snooping query interval:

```
Router(config-vlan-config)# ip igmp snooping last-member-query-interval 200
Router(config-vlan-config)# exit
Router# show ip igmp interface vlan 200 | include last
IGMP snooping last member query interval on this interface is 200 ms
```

Enabling IGMP Snooping Immediate-Leave Processing

Fast-leave configuration applies to IGMP version 2 hosts only. To enable IGMP snooping fast-leave processing in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ip igmp snooping	Enables IGMP snooping. This step is only necessary if IGMP snooping is not already enabled on this VLAN.
Step 3	Router(config-vlan-config)# ip igmp snooping immediate-leave	Enables IGMP immediate-leave processing in the VLAN.

This example shows how to enable IGMP snooping immediate-leave processing for IGMP version 2 hosts on the VLAN 200 interface, and how to verify the configuration:

```
Router# interface vlan 200
Router(config-vlan-config)# ip igmp snooping
Router(config-vlan-config)# ip igmp snooping immediate-leave
Configuring immediate leave on vlan 200
Router(config-vlan-config)# end
Router# show ip igmp interface vlan 200 | include immediate-leave
IGMP snooping immediate-leave is enabled on this interface
```

Configuring IGMPv3 Snooping Explicit Host Tracking

To enable IGMPv3 snooping explicit host tracking on a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ip igmp snooping explicit-tracking limit <i>limit</i>	Enable IGMPv3 snooping explicit host tracking in a VLAN.

This example shows how to enable IGMPv3 snooping explicit host tracking:

```
Router(config-vlan-config)# ip igmp snooping explicit-tracking limit 400
Router# show ip igmp snooping vlan 200
Global IGMP Snooping configuration:
-----
IGMP snooping Oper State      : Enabled
IGMPv3 snooping               : Enabled
Report suppression            : Disabled
EHT DB limit/count            : 100000/2
TCN solicit query             : Disabled
Robustness variable           : 2
Last member query count        : 3
Last member query interval     : 1000
Check TTL=1                   : No
Check Router-Alert-Option      : No

Vlan 200:
-----
IGMP snooping Admin State     : Enabled
IGMP snooping Oper State      : Enabled
```

```

IGMPv2 immediate leave      : Disabled
Explicit host tracking      : Enabled
Report suppression         : Enabled
Robustness variable        : 2
Last member query count    : 2
Last member query interval : 1000
EHT DB limit/count         : 100000/2
Check TTL=1                : Yes
Check Router-Alert-Option  : Yes
Query Interval             : 100
Max Response Time          : 10000
Router(config-vlan-config)# ip igmp snooping static 224.1.2.3 source 10.1.1.1 interface Gi3/11
Router# show ip igmp snooping groups vlan 200
Flags: I -- IGMP snooping, S -- Static, P -- PIM snooping, A -- ASM mode
Vlan      Group/source      Type      Version      Port List
-----
200        224.1.2.3          S          v3           Gi3/11
           /10.1.1.1
Router#

```

Displaying IGMP Snooping Information

- [Displaying Multicast Router Interfaces, page 7-14](#)
- [Displaying MAC Address Multicast Entries, page 7-15](#)
- [Displaying IGMP Snooping Information for a VLAN Interface, page 7-15](#)

Displaying Multicast Router Interfaces

When you enable IGMP snooping, the switch automatically learns to which interface the multicast routers are connected. To display multicast router interfaces, perform this task:

Command	Purpose
Router# show ip igmp snooping vlan <i>vlan_ID</i>	Displays multicast router interfaces.

This example shows how to display the multicast router interfaces in VLAN 1:

```

Router# show ip igmp snooping vlan 200
Global IGMP Snooping configuration:
-----
IGMP snooping Oper State      : Enabled
IGMPv3 snooping              : Enabled
Report suppression            : Disabled
EHT DB limit/count           : 100000/2
TCN solicit query            : Disabled
Robustness variable           : 2
Last member query count       : 3
Last member query interval    : 1000
Check TTL=1                   : No
Check Router-Alert-Option     : No

Vlan 200:
-----
IGMP snooping Admin State     : Enabled
IGMP snooping Oper State      : Enabled
IGMPv2 immediate leave        : Disabled
Explicit host tracking         : Enabled

```

```

Report suppression      : Enabled
Robustness variable     : 2
Last member query count : 2
Last member query interval : 1000
EHT DB limit/count     : 100000/2
Check TTL=1            : Yes
Check Router-Alert-Option : Yes
Query Interval         : 100
Max Response Time      : 10000
Router#

```

Displaying MAC Address Multicast Entries

To display MAC address multicast entries for a VLAN, perform this task:

Command	Purpose
Router# show mac address-table multicast <i>vlan_ID</i> [<i>count</i>]	Displays MAC address multicast entries for a VLAN.

This example shows how to display MAC address multicast entries for VLAN 1:

```

Router# show mac address-table multicast vlan 1
vlan  mac address      type    qos    ports
-----+-----+-----+-----+-----
  1  0100.5e02.0203  static  --    Gi1/1,Gi2/1,Gi3/48,Router
  1  0100.5e00.0127  static  --    Gi1/1,Gi2/1,Gi3/48,Router
  1  0100.5e00.0128  static  --    Gi1/1,Gi2/1,Gi3/48,Router
  1  0100.5e00.0001  static  --    Gi1/1,Gi2/1,Gi3/48,Router,Switch
Router#

```

This example shows how to display a total count of MAC address entries for a VLAN:

```

Router# show mac address-table multicast 1 count

Multicast MAC Entries for vlan 1:    4
Router#

```

Displaying IGMP Snooping Information for a VLAN Interface

To display IGMP snooping information for a VLAN interface, perform this task:

Command	Purpose
Router# show ip igmp interface <i>vlan_ID</i>	Displays IGMP snooping information on a VLAN interface.

This example shows how to display IGMP snooping information on the VLAN 200 interface:

```

Router# show ip igmp interface vlan 43
Vlan43 is up, line protocol is up
Internet address is 43.0.0.1/24
IGMP is enabled on interface
Current IGMP host version is 2
Current IGMP router version is 2
IGMP query interval is 60 seconds
IGMP querier timeout is 120 seconds

```

```
IGMP max query response time is 10 seconds
Last member query count is 2
Last member query response interval is 1000 ms
Inbound IGMP access group is not set
IGMP activity:1 joins, 0 leaves
Multicast routing is enabled on interface
Multicast TTL threshold is 0
Multicast designated router (DR) is 43.0.0.1 (this system)
IGMP querying router is 43.0.0.1 (this system)
Multicast groups joined by this system (number of users):
  224.0.1.40(1)
IGMP snooping is globally enabled
IGMP snooping is enabled on this interface
IGMP snooping immediate-leave is disabled and querier is disabled
IGMP snooping explicit-tracking is enabled on this interface
IGMP snooping last member query interval on this interface is 1000 ms
Router#
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

8

PIM Snooping

- [Prerequisites for PIM Snooping, page 8-1](#)
- [Restrictions for PIM Snooping, page 8-1](#)
- [Information About PIM Snooping, page 8-2](#)
- [Default Settings for PIM Snooping, page 8-4](#)
- [How to Configure PIM Snooping, page 8-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for PIM Snooping

None.

Restrictions for PIM Snooping

- Multicast packets are not bridged in a VLAN to local receivers that send IGMP joins when PIM snooping is enabled in the VLAN and IGMP snooping is disabled in the VLAN. (CSCta03980)
- When you use the PIM-sparse mode (PIM-SM) feature, downstream routers only see traffic if they previously indicated interest through a PIM join or prune message. An upstream router only sees traffic if it was used as an upstream router during the PIM join or prune process.
- Join or prune messages are not flooded on all router ports but are sent only to the port corresponding to the upstream router mentioned in the payload of the join or prune message.

- Directly connected sources are supported for bidirectional PIM groups. Traffic from directly connected sources is forwarded to the designated router and designated forwarder for a VLAN. In some cases, a nondesignated router (NDR) can receive a downstream (S, G) join. For source-only networks, the initial unknown traffic is flooded only to the designated routers and designated forwarders.
- Dense group mode traffic is seen as unknown traffic and is dropped.
- The AUTO-RP groups (224.0.1.39 and 224.0.1.40) are always flooded.
- The switch snoops on designated forwarder election and maintains a list of all designated forwarder routers for various RPs for the VLAN. All traffic is sent to all designated forwarders which ensures that bidirectional functionality works properly.
- PIM snooping and IGMP snooping can be enabled at the same time in a VLAN. Either RGMP or PIM snooping can be enabled in a VLAN but not both.
- Any non-PIMv2 multicast router will receive all traffic.
- You can enable or disable PIM snooping on a per-VLAN basis.
- All mroute and router information is timed out based on the hold-time indicated in the PIM hello and join/prune control packets. All mroute state and neighbor information is maintained per VLAN.
- With Release 15.1(1)SY1 and later releases, IGMP snooping and PIM snooping constrain VPLS multicast traffic.

Information About PIM Snooping

In networks where a Layer 2 switch interconnects several routers, such as an Internet exchange point (IXP), the switch floods IP multicast packets on all multicast router ports by default, even if there are no multicast receivers downstream. With PIM snooping enabled, the switch restricts multicast packets for each IP multicast group to only those multicast router ports that have downstream receivers joined to that group. When you enable PIM snooping, the switch learns which multicast router ports need to receive the multicast traffic within a specific VLAN by listening to the PIM hello messages, PIM join and prune messages, and bidirectional PIM designated forwarder-election messages.



Note

To use PIM snooping, you must enable IGMP snooping on the switch. IGMP snooping restricts multicast traffic that exits through the LAN ports to which hosts are connected. IGMP snooping does not restrict traffic that exits through the LAN ports to which one or more multicast routers are connected.

The following illustrations show the flow of traffic and flooding that results in networks without PIM snooping enabled and the flow of traffic and traffic restriction when PIM snooping is enabled.

[Figure 8-1](#) shows the flow of a PIM join message without PIM snooping enabled. In the figure, the switches flood the PIM join message intended for Router B to all connected routers.


Figure 8-1 PIM Join Message Flow without PIM Snooping

Figure 8-2 shows the flow of a PIM join message with PIM snooping enabled. In the figure, the switches restrict the PIM join message and forward it only to the router that needs to receive it (Router B).


Figure 8-2 PIM Join Message Flow with PIM Snooping

Figure 8-3 shows the flow of data traffic without PIM snooping enabled. In the figure, the switches flood the data traffic intended for Router A to all connected routers.



Figure 8-3 Data Traffic Flow without PIM Snooping

Figure 8-4 shows the flow of data traffic with PIM snooping enabled. In the figure, the switches forward the data traffic only to the router that needs to receive it (Router A).

Figure 8-4 Data Traffic Flow with PIM Snooping

Default Settings for PIM Snooping

PIM snooping is disabled by default.

How to Configure PIM Snooping

- [Enabling PIM Snooping Globally, page 8-5](#)
- [Enabling PIM Snooping in a VLAN, page 8-5](#)
- [Disabling PIM Snooping Designated-Router Flooding, page 8-6](#)

Enabling PIM Snooping Globally

To enable PIM snooping globally, perform this task:

	Command	Purpose
Step 1	Router(config)# ip pim snooping	Enables PIM snooping.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable PIM snooping globally and verify the configuration:

```
Router(config)# ip pim snooping
Router(config)# end
Router# show ip pim snooping
Global runtime mode: Enabled
Global admin mode   : Enabled
Number of user enabled VLANs: 1
User enabled VLANs: 10
Router#
```

**Note**

You do not need to configure an IP address or IP PIM in order to run PIM snooping.

Enabling PIM Snooping in a VLAN

To enable PIM snooping in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ip pim snooping	Enables PIM snooping.
Step 3	Router(config-vlan-config)# end	Exits configuration mode.

This example shows how to enable PIM snooping on VLAN 10 and verify the configuration:

```
Router# vlan configuration 10
Router(config-vlan-config)# ip pim snooping
Router(config-vlan-config)# end
Router# show ip pim snooping vlan 10
3 neighbors (0 DR priority incapable, 0 Bi-dir incapable)
6 mroutes, 3 mac entries
DR is 10.10.10.4
RP DF Set
Router#
```

Disabling PIM Snooping Designated-Router Flooding

**Note**

Do not disable designated-router flooding on switches in a Layer 2 broadcast domain that supports multicast sources.

By default, switches that have PIM snooping enabled will flood multicast traffic to the designated router (DR). This method of operation can send unnecessary multicast packets to the designated router. The network must carry the unnecessary traffic, and the designated router must process and drop the unnecessary traffic.

To reduce the traffic sent over the network to the designated router, disable designated-router flooding. With designated-router flooding disabled, PIM snooping only passes to the designated-router traffic that is in multicast groups for which PIM snooping receives an explicit join from the link towards the designated router.

To disable PIM snooping designated-router flooding, perform this task:

	Command	Purpose
Step 1	Router(config)# no ip pim snooping dr-flood	Disables PIM snooping designated-router flooding.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to disable PIM snooping designated-router flooding:

```
Router(config)# no ip pim snooping dr-flood
Router(config)# end
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

9

Multicast VLAN Registration (MVR)

- [Restrictions for MVR, page 9-1](#)
- [Restrictions for MVR, page 9-1](#)
- [Information About MVR, page 9-2](#)
- [Default MVR Configuration, page 9-5](#)
- [How to Configure MVR, page 9-5](#)
- [Displaying MVR Information, page 9-8](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Restrictions for MVR

None.

Restrictions for MVR

- Only one MVR VLAN can be present in a switch, and you should configure the same VLAN as the MVR VLAN for all the switches in the same network.
- Source ports must be in the MVR VLAN.
- Receiver ports on a switch can be in different VLANs, but must not be in the MVR VLAN.
- Receiver ports can only be access ports; they cannot be trunk ports.
- When using private VLANs, you cannot configure a secondary VLAN as the MVR VLAN.

- Do not connect a multicast router to a receiver port.
- The MVR VLAN must not be a reverse path forwarding (RPF) interface for any multicast route.
- MVR data received on an MVR receiver port is not forwarded to MVR source ports.
- 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 8000.
- MVR is available only on native systems.
- VTP pruning should be disabled if the MVR VLAN number is between 1 and 1000.
- MVR can coexist with IGMP snooping on a switch.
- MVR supports IGMPv3 messages.

Information About MVR

- [MVR Overview, page 9-2](#)
- [Using MVR in a Multicast Television Application, page 9-3](#)

MVR Overview

MVR is designed for applications that use 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 feature can be enabled or disabled without affecting the operation 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.

MVR does the following:

- Identifies the MVR IP multicast streams and their associated IP multicast group in the Layer 2 forwarding table.
- Intercepts the IGMP messages.
- Modifies the Layer 2 forwarding table to include or remove the subscriber as a receiver of the multicast stream, even though the receivers might be in a different VLAN from the source.


This forwarding behavior selectively allows traffic to cross between different VLANs.

The switch will forward multicast data for MVR IP multicast streams only to MVR ports on which hosts have joined, either by IGMP reports or by MVR static configuration. The switch will forward IGMP reports received from MVR hosts only to the source (uplink) port. This eliminates using unnecessary bandwidth on MVR data port links.

Only Layer 2 ports participate in MVR. You must configure ports as MVR receiver ports. Only one MVR multicast VLAN per switch.

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 9-1](#) 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 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 9-1 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 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.

Unless the Immediate Leave feature is enabled, 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 the Immediate Leave feature enabled, 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.

Default MVR Configuration

- MVR: Disabled globally and per interface
- Multicast addresses: None configured
- Query response time: 1 second
- Multicast VLAN: VLAN 1
- Interface (per port) default: Neither a receiver nor a source port
- Immediate Leave: Disabled on all ports

How to Configure MVR

- [Configuring MVR Global Parameters, page 9-5](#)
- [Configuring MVR Interfaces, page 9-6](#)
- [Displaying MVR Information, page 9-8](#)
- [Clearing MVR Counters, page 9-8](#)

Configuring MVR Global Parameters

To configure the MVR global parameters, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# mvr	Enables MVR on the switch.
Step 3	Router(config)# mvr max-groups <i>max-groups</i>	Specifies the maximum number of MVR groups. The range is 1 to 8000. The default is 1000.
Step 4	Router(config)# mvr group <i>ip-address</i> [<i>count</i>]	Configures an IP multicast address on the switch or uses 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.

	Command	Purpose
Step 5	Router(config)# mvr querytime <i>value</i>	(Optional) Defines 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 10 tenths or one second.
Step 6	Router(config)# mvr vlan <i>vlan-id</i>	(Optional) Specifies 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.
Step 7	Router(config)# end	Returns to privileged EXEC mode.

You do not need to set the optional MVR parameters if you choose to use the default settings. Before changing the default parameters (except for the MVR VLAN), you must first enable MVR.

To return the switch to its default settings, use the **no mvr [group ip-address | querytime | vlan]** global configuration command.

This example shows how to enable MVR, configure the group address, set the query time to 1 second (10 tenths), and specify the MVR multicast VLAN as VLAN 22:

```
Router(config)# mvr
Router(config)# mvr group 228.1.23.4
Router(config)# mvr querytime 10
Router(config)# mvr vlan 22
Router(config)# end
```

You can use the **show mvr groups** privileged EXEC command to verify the MVR multicast group addresses on the switch.

Configuring MVR Interfaces

To configure Layer 2 MVR interfaces, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# mvr	Enables MVR on the switch.
Step 3	Router(config)# interface <i>interface-id</i>	Specifies the Layer 2 port to configure, and enters interface configuration mode.

	Command	Purpose
Step 4	Router(config-if)# mvr type {source receiver}	Configures an MVR port as one of these types of ports: • source—Configures 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—Configures 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. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails. The default configuration is as a non-MVR port.
Step 5	Router(config-if)# mvr immediate	(Optional) Enables the Immediate Leave feature of MVR on the port. The Immediate Leave feature is disabled by default. 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 6	Router(config-if)# end	Returns to privileged EXEC mode.

To return the interface to its default settings, use the **no mvr [type | immediate]** interface configuration commands.

This example shows how to configure a source port and a receiver port and to configure Immediate Leave on the receiver port:

```
Router(config)# mvr
Router(config)# interface gigabitethernet 3/48
Router(config-if)# switchport
Router(config-if)# switchport access vlan 22
Router(config-if)# mvr type source
Router(config-if)# exit
Router(config)# interface gigabitethernet 3/47
Router(config-if)# switchport
Router(config-if)# switchport access vlan 30
Router(config-if)# mvr type receiver
Router(config-if)# mvr immediate
Router(config-if)# exit
```

Clearing MVR Counters

You can clear MVR join counters for the switch, for source or receiver ports, or for a specified interface. To clear MVR counters, perform this task:

Command	Purpose
Router# clear mvr counters [[receiver-ports source-ports] [<i>type module/port</i>]]	Clears the join counters of all the MVR ports, or source or receiver ports, or of a specified MVR interface port.

This example clears the join counters for the receiver port on GigabitEthernet port 1/7:

```
Router# clear mvr receiver-ports GigabitEthernet 1/7
Router# show mvr receiver-ports GigabitEthernet 1/7
Joins: v1,v2,v3 counter shows total IGMP joins
       v3 counter shows IGMP joins received with both MVR and non-MVR groups
Port   VLAN Status      Immediate      Joins
      (v1,v2,v3) (v3)
-----
Gi1/7  202 INACTIVE/UP  ENABLED              0          0
```

Displaying MVR Information

You can display MVR information for the switch or for a specified interface. To display MVR configurations, perform one or more of these tasks:

Command	Purpose
Router# show mvr	Displays MVR status and these values for the switch: whether MVR is enabled or disabled, the multicast VLAN, the configured maximum and current number of multicast groups, and the query response time.
Router# show mvr groups	Displays the MVR group configuration.
Router# show mvr interface [<i>type module/port</i>]	Displays all MVR interfaces and their MVR configurations. When a specific interface is entered, displays this information: - Type—Receiver or Source - Status—One of these: - Active—At least one IGMP join has been received for an MVR group on the port. - Inactive—The port is not participating in any MVR groups. - Up/Down—The port is forwarding (Up) or nonforwarding (Down). - Immediate Leave—Enabled or Disabled
Router# show mvr members [[vlan <i>vlan-id</i>] [<i>type module/port</i>]]	Displays details of all MVR members or MVR members on a specified VLAN or port.

Command	Purpose
Router# show mvr	Displays MVR status and these values for the switch: whether MVR is enabled or disabled, the multicast VLAN, the configured maximum and current number of multicast groups, and the query response time.
Router# show mvr groups	Displays the MVR group configuration.
Router# show mvr members [[vlan <i>vlan-id</i>] [type <i>module/port</i>]] count	Displays number of MVR members in all active MVR groups, or on a specified VLAN or port.
Router# show mvr { receiver-ports source-ports } [type <i>module/port</i>]	Displays all receiver or source ports that are members of any IP multicast group or those on the specified interface port.

This example displays MVR status and values for the switch:

```
Router# show mvr
MVR Running: TRUE
MVR multicast vlan: 22
MVR Max Multicast Groups: 1000
MVR Current multicast groups: 256
MVR Global query response time: 10 (tenths of sec)
```

This example displays the MVR group configuration:

```
Router# show mvr groups
MVR max Multicast Groups allowed: 8000
MVR current multicast groups: 8000
MVR groups:
      Group start      Group end      Type  Count/Mask
      -----
      225.0.7.226      225.0.7.226      count  1
      225.0.7.227      225.0.7.227      count  1
      225.0.7.228      225.0.7.228      count  1
      225.0.7.229      225.0.7.229      count  1
      225.0.7.230      225.0.7.230      count  1
      225.0.7.231      225.0.7.231      count  1
      236.8.7.0         236.8.7.255      mask   255.255.255.0
      237.8.7.0         237.8.7.255      mask   255.255.255.0
      237.8.8.0         237.8.8.255      mask   255.255.255.0
```

This example displays all MVR interfaces and their MVR configurations:

```
Router# show mvr interface
Port      VLAN  Type      Status      Immediate Leave
----      -
Gi1/20    2    RECEIVER  ACTIVE/UP    DISABLED
Gi1/21    2    SOURCE    ACTIVE/UP    DISABLED
```

This example displays all MVR members on VLAN 2:

```
Router# show mvr members vlan 2
MVR Group IP      Status  Members
-----
224.0.0.001.001    ACTIVE  Gi1/20(u),Gi1/21(u)
224.0.0.001.002    ACTIVE  Gi3/2(d),Gi1/12(u)
```

This example displays the number of MVR members on all MVR VLANs:

```
Router# show mvr members count
```

```
Count of active MVR groups:
Vlan 490: 400
```

```
Vlan 600: 400
Vlan 700: 0
Vlan 950: 0
```

This example displays all receiver ports that are members of any IP multicast group:

```
Router# show mvr receiver-ports
```

```
Joins: v1,v2,v3 counter shows total IGMP joins
```

```
      v3 counter shows IGMP joins received with both MVR and non-MVR groups
```

Port	VLAN	Status	Immediate Leave	Joins (v1,v2,v3)	(v3)
Gi1/7	202	INACTIVE/UP	ENABLED	305336	0
Gi1/8	202	ACTIVE/UP	DISABLED	4005	0
Gi1/9	203	INACTIVE/DOWN	DISABLED	53007	0
Gi1/10	203	ACTIVE/UP	DISABLED	6204	0
Gi1/11	204	ACTIVE/UP	DISABLED	0	940
Gi1/12	205	INACTIVE/UP	ENABLED	8623	0



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

10

IPv4 IGMP Filtering

- [Prerequisites for IGMP Filtering, page 10-1](#)
- [Restrictions for IGMP Filtering, page 10-1](#)
- [Information About IGMP Filtering, page 10-2](#)
- [Default Settings for IGMP Filtering, page 10-4](#)
- [How to Configure IGMP Filters, page 10-4](#)
- [Verifying the IGMP Filtering Configuration, page 10-6](#)
- [Configuration Examples for IGMP Filtering, page 10-8](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IGMP Filtering

None.

Restrictions for IGMP Filtering

None.

Information About IGMP Filtering

- [IGMP Filtering Overview, page 10-3](#)
- [IGMP Filter Precedence, page 10-4](#)

IGMP Filtering Overview

**Note**

IGMP, which runs at Layer 3 on a multicast router, generates Layer 3 IGMP queries in subnets where the multicast traffic needs to be routed. For information about IGMP, see [Chapter 6, “Information About IPv4 Multicast Layer 3 Features.”](#)

IGMP snooping is a protocol that learns and maintains multicast group membership at the Layer 2 level. IGMP snooping looks at IGMP traffic to decide which ports should be allowed to receive multicast traffic from certain sources and for certain groups. This information is used to forward multicast traffic to only interested ports. The main benefit of IGMP snooping is to reduce flooding of packets. For information about IGMP snooping, see [“Information About IGMP Filtering” section on page 10-2.](#)

IGMP filtering allows users to configure filters on a switch virtual interface (SVI), a per-port, or a per-port per-VLAN basis to control the propagation of IGMP traffic through the network. By managing the IGMP traffic, IGMP filtering provides the capability to manage IGMP snooping, which in turn controls the forwarding of multicast traffic.

When an IGMP packet is received, IGMP filtering uses the filters configured by the user to determine whether the IGMP packet should be discarded or allowed to be processed by the existing IGMP snooping code. With a IGMP version 1 or version 2 packet, the entire packet is discarded. With a IGMPv3 packet, the packet is rewritten to remove message elements that were denied by the filters.

The IGMP filtering feature is SSO compliant.

IGMP traffic filters control the access of a port to multicast traffic. Access can be restricted based on the following:

- Which multicast groups or channels can be joined on a port. Channels are joined by IGMPv3 hosts that specify both the group and the source of the multicast traffic.
- Maximum number of groups or channels allowed on a specific port or interface (regardless of the number of hosts requesting service).
- IGMP protocol versions (for example, disallow all IGMPv1 messages).

When you enter an IGMP filtering command, a user policy is applied to a Layer 3 SVI interface, a Layer 2 port, or a particular VLAN on a Layer 2 trunk port. The Layer 2 port may be an access port or a trunk port. The IGMP filtering features will work only if IGMP snooping is enabled (either on the interface or globally).

IGMP filtering is typically used in access switches connected to end-user devices.

There are three different types of IGMP filters: IGMP group and channel access control, several IGMP groups and channels limit, and an IGMP minimum version. These filters are configurable and operate differently on different types of ports:

- Per SVI
- Per port
- Per VLAN basis on a trunk port

You can configure filters separately for each VLAN passing through a trunk port.

IGMP Filter Precedence

- [Access Mode, page 10-4](#)
- [Trunk Mode, page 10-4](#)

Access Mode

In access mode, filters can be configured on both the port and the SVI. When an IGMP packet is received on a port in access mode, the port filter is checked first. If the port filter exists, it is applied and the SVI filter is ignored. If no per-port filter exists, the SVI filter is used.

This hierarchy is applied separately for each type of filter. For example, a limit filter configured on the port overrides the default limit filter on the SVI, but has no affect on any of the other filters.

Trunk Mode

With ports in trunk mode, a filter can be configured for an SVI corresponding to one of the VLANs on the trunk port, another filter configured on the trunk port itself, and a third filter configured on one of the Layer 2 VLANs passing through the trunk. When an IGMP packet is received, the trunk-per-VLAN specific filter will be checked first. If this filter exists, it is applied. The main trunk port filter and SVI filter will be ignored. If no trunk-per-VLAN filter exists, the main trunk port filter will be used. If neither of these filters exist, the SVI filter for the VLAN will be used as a final default for ports in trunk mode.

Default Settings for IGMP Filtering

None.

How to Configure IGMP Filters

- [Configuring IGMP Group and Channel Access Control, page 10-4](#)
- [Configuring IGMP Group and Channel Limits, page 10-5](#)
- [Configuring IGMP Version Filtering, page 10-5](#)
- [Clearing IGMP Filtering Statistics, page 10-6](#)

Configuring IGMP Group and Channel Access Control

Filtering on the IGMP group or channel allows the user to control which IGMP groups or channels can be joined on a port or on a per VLAN basis on a trunk port.

To configure filtering on the IGMP group or channel use the following CLI command:

```
ip igmp snooping access-group acl [vlan vlan_id]
```

To allow or deny several groups or channels, you must configure multiple access control entries in the access control list. Depending on whether the ACL is configured as permit or deny, the corresponding group or channel is allowed or denied. The ACL specified may be either a simple or extended ACL.

Filtering by IGMP group or channel is configurable on a Layer 3 SVI as a default filter for all ports in access mode under that SVI and for the corresponding VLAN on all trunk ports carrying that VLAN. This filter is also configurable on a Layer 2 port. If the port is in access mode, this filter will override any default SVI filter. If the port is in trunk mode, this filter will act as a default for all VLANs on that trunk and will override the SVI filter for each corresponding VLAN.

The **vlan** keyword can apply the filter only to IGMP packets arriving on the specified Layer 2 VLAN if the port is a trunk port. This per-VLAN filter (configured using the **vlan** keyword) will override any interface level filter and any SVI filter for the same VLAN.

Configuring IGMP Group and Channel Limits

Limiting the number of IGMP groups or channels allows you to control how many IGMP groups or channels can be joined on a port or on a per-VLAN basis on a trunk port.

To limit the number of IGMP groups or channels, use the following interface command CLI:

```
ip igmp snooping limit n [except acl] [vlan vlan_id]
```

A maximum of *n* groups or channels are allowed on the port or interface. The **except** keyword allows you to specify groups or channels that are exempt from the configured limit. The ACL used with the **except** keyword may be either a simple or extended ACL.

If joins are received for (*,G1) and (S1,G1) on the same interface, these are counted as two separate joins. If the limit on an interface has been set to 2, and joins are received for (*,G1) and (S1,G1), all other joins (for groups or channels different from these two) will then be discarded.

This filter is configurable on a Layer 3 SVI as a default filter for all ports in access mode under that SVI and for the corresponding VLAN on all trunk ports carrying that VLAN. This filter is also configurable on a Layer 2 port. If the Layer 2 port is in access mode, this filter will override any default SVI filter. If the Layer 2 switch port is in trunk mode, this filter will act as a default for all VLANs on that trunk and will override the SVI filter for each corresponding VLAN. The **vlan** keyword allows the user to apply the filter only to IGMP packets arriving on the specified Layer 2 VLAN if the Layer 2 switch port is a trunk port. This per-VLAN filter, configured using the **vlan** keyword, will override any interface level filter and any SVI filter for the same VLAN.

Configuring IGMP Version Filtering

Filtering on the IGMP protocol allows you to configure the minimum version of IGMP hosts allowed on the SVI. For example, you may want to disallow all IGMPv1 hosts (such as, allow a minimum IGMP version of 2) or all IGMPv1 and IGMPv2 hosts (such as, allow a minimum IGMP version of 3). This filtering applies only to membership reports.

To configure filtering on the IGMP protocol, use the following CLI command:

```
ip igmp snooping minimum-version 2 | 3
```

This filter is only configurable on a Layer 3 SVI as a default filter for all ports in access mode under that SVI and for the corresponding VLAN on all trunk ports.

Clearing IGMP Filtering Statistics

To clear IGMP filtering statistics, perform one of these tasks:

Command	Purpose
Router# clear ip igmp snooping filter statistics	Clears IGMP filtering statistics for all access ports and for all VLANs on all trunk ports.
Router# clear ip igmp snooping filter statistics interface <i>interface_name</i>	Clears statistics for one particular access port or for all VLANs on one particular trunk port.
Router# clear ip igmp snooping filter statistics interface <i>interface_name</i> vlan <i>vlan_ID</i>	Clears statistics for one particular VLAN on a trunk port.

Verifying the IGMP Filtering Configuration

- [Displaying IGMP Filtering Configuration, page 10-6](#)
- [Displaying IGMP Filtering Statistics, page 10-7](#)

Displaying IGMP Filtering Configuration

To display IGMP filtering rules, perform this task:

Command	Purpose
Router(config-if)# show ip igmp snooping filter interface <i>interface-name</i> [details]	Displays the filters configured for the specified interface.

This example shows how to display the default filters configured on the SVI:

```
Router# show ip igmp snooping filter interface vlan 20
Access-Group: Channel1-Acl
Groups/Channels Limit:100 (Exception List: Channel6-Acl)
IGMP Minimum-Version:Not Configured
```

This example shows how to display the filters configured for all ports in access mode under this SVI and for all trunk ports carrying the corresponding VLAN:

```
Router# show ip igmp snooping filter interface g3/48
Access-Group: Channel4-Acl
Groups/Channels Limit:10 (Exception List: Channel3-Acl)
```

This example shows how to display the filters configured for all ports in access mode under this SVI:

```
Router# show ip igmp snooping filter interface vlan 20 detail
GigabitEthernet3/47 :
Access-Group: Not Configured
Groups/Channels Limit: Not Configured
GigabitEthernet3/48 :
Access-Group: Channel4-ACL
Groups/Channels Limit: 10 (Exception-list: Channel3-Acl)
```

This example shows how to display the default trunk port filters:

```
Router# show ip igmp snooping filter interface g3/46
```

```
Access-Group: Channel1-Acl
Groups/Channels Limit: 10 (Exception List: Channel3-Acl)
```

This example shows how to display the per-VLAN filters for all VLANs on this trunk:

```
Router# show ip igmp snooping filter interface g3/46 detail
Vlan 10 :
  Access-Group: Not Configured
  Groups/Channels Limit: Not Configured
Vlan 20 :
  Access-Group: Not Configured
  Groups/Channels Limit: 8 (Exception List: Channel4-Acl)
```

This example shows how to display the per-VLAN filters for a specific VLAN on this trunk:

```
Router# show ip igmp snooping filter interface g3/46 vlan 20
  Access-Group: Not Configured
  Groups/Channels Limit: 8 (Exception List: Channel4-Acl)
```



Note

If the port is in the shutdown state, filter status will not be displayed because it cannot be determined whether the port is in trunk mode or access mode. In this situation, you can use the **show running-config interface xxxx** command to view the configuration.

Displaying IGMP Filtering Statistics

Statistics are maintained on an interface basis for ports in access mode and on a per-VLAN basis for ports in trunk mode.

To display IGMP filtering statistics, perform this task:

Command	Purpose
Switch(config-if)# show ip igmp snooping filter interface interface-name [statistics]	Displays the filtering statistic collected for the specified interface.

This example shows how to display statistics for each port in access mode under the SVI:

```
Router# show ip igmp snooping filter interface vlan 20 statistics
GigabitEthernet3/47 :
  IGMP Filters are not configured

GigabitEthernet3/48 :
  Access-group denied : 0
  Limit denied : 2
  Limit status : 0 active out of 2 max
  Minimum-version denied : 0
```

This example shows how to display statistics for a specific port in access mode:

```
Router# show ip igmp snooping filter interface g3/48 statistics
  Access-group denied : 0
  Limit denied : 2
  Limit status : 0 active out of 2 max
  Minimum-version denied : 0
```

This example shows how to display statistics for Gigabit Ethernet port 3/47 in access mode with no default SVI filter and no port filter:

```
Router# show ip igmp snooping filter interface g3/47 statistics
```

```
IGMP Filters are not configured
```

This example shows how to display statistics for all VLANs under a trunk:

```
Router# show ip igmp snooping filter interface g3/46 statistics
Vlan 10   :
IGMP Filters are not configured
```

```
Vlan 20   :
  Access-group denied : 0
  Limit denied : 0
  Minimum-version denied : 0
```

This example shows how to display statistics for a specific VLAN under a trunk:

```
Router# show ip igmp snooping filter interface g3/46 vlan 20 statistics
  Access-group denied : 0
  Limit denied : 0
  Minimum-version denied : 0
```

This example shows how to display statistics for a specific VLAN under a trunk port with no trunk and no VLAN filter:

```
Router# show ip igmp snooping filter interface g3/46 vlan 10 statistics
IGMP Filters are not configured
```


Note

If the port is in the shutdown state, filter statistics will not be displayed because it cannot be determined whether the port is in trunk mode or access mode.

Configuration Examples for IGMP Filtering

This example shows the filter hierarchy. The following configuration of SVI VLAN 100 contains three access ports g1/1, g1/2, and g1/3:

VLAN 100:

```
Router(config-if)# ip igmp snooping limit 20
```

Port g1/1:

```
Router(config-if)# ip igmp snooping limit 35
```

Port g1/2:

```
Router(config-if)# no limit filter
```

Port g1/3:

```
Router(config-if)# no limit filter
```

In this example, the limit value for g1/1 is 35, the limit value for g1/2 is 20, and the limit value for g1/3 is also 20.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

11

IPv4 Router Guard

- [Prerequisites for Router Guard, page 11-1](#)
- [Restrictions for Router Guard, page 11-1](#)
- [Information About Router Guard, page 11-1](#)
- [Default Settings for Router Guard, page 11-2](#)
- [How to Configure Router Guard, page 11-2](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Router Guard

None.

Restrictions for Router Guard

None.

Information About Router Guard

The Router Guard feature allows you to designate a specified port only as a multicast host port and not as a multicast router port. Multicast router control packets received on this port are dropped.

Any port can become a multicast router port if the switch receives one of the multicast router control packets, such as IGMP general query, PIM hello, or CGMP hello. When a port becomes a multicast router port, all multicast traffic (both known and unknown source traffic) is sent to all multicast router ports. This cannot be prevented without Router Guard.

When configured, the Router Guard feature makes the specified port a host port only. The port is prevented from becoming a router port, even if a multicast router control packets are received.

In addition, any control packets normally received from multicast routers, such as IGMP queries and PIM joins, will also be discarded by this filter.

A Router Guard command applies a user policy to a Layer 3 SVI interface, a Layer 2 port, or a particular VLAN on a Layer 2 trunk port. The Layer 2 port may be an access port or a trunk port.

The Router Guard feature does not require IGMP snooping to be enabled.

Router Guard is implemented only for IPv4.

Router Guard is typically used in access switches connected to end-user boxes in Ethernet-to-home deployment scenarios.

The IPv4 multicast Router Guard feature is SSO-compliant.

The following packet types are discarded if they are received on a port that has Router Guard enabled:

- IGMP query messages
- IPv4 PIMv2 messages
- IGMP PIM messages (PIMv1)
- IGMP DVMRP messages
- RGMP messages
- CGMP messages

When these packets are discarded, statistics are updated indicating that packets are being dropped due to Router Guard.

Router Guard can be configured globally and per-interface. The global configuration initiates Router Guard for all Layer 2 ports, which can be modified with the interface configuration commands, for example, on ports where multicast routers are connected.

Default Settings for Router Guard

None.

How to Configure Router Guard

- [Enabling Router Guard Globally, page 11-3](#)
- [Disabling Router Guard on Ports, page 11-3](#)
- [Clearing Router Guard Statistics, page 11-3](#)
- [Displaying Router Guard Configuration, page 11-4](#)
- [Displaying Router Guard Interfaces, page 11-5](#)

Enabling Router Guard Globally

To enable Router Guard globally, perform this task:

Command	Purpose
Router# router-guard ip multicast switchports	Enables Router Guard globally.

Disabling Router Guard on Ports

To disable Router Guard on a Layer 2 port to which a multicast router is connected, perform this task:

Command	Purpose
Router(config-if)# no router-guard ip multicast [vlan <i>vlan_id</i>]	Disables Router Guard on a Layer 2 port. Note The vlan keyword is effective only if the port is in trunk mode. You can use this keyword to override Router Guard only for specific VLANs on the trunk port.

This example shows how to allow multicast router messages on trunk port Gigabit Ethernet 3/46, VLAN 20:

```
Router# configure terminal
Router(config)# interface gigabitethernet 3/46
Router(config-if)# no router-guard ip multicast vlan 20
```

Clearing Router Guard Statistics

To clear Router Guard statistics, perform one of these tasks:

Command	Purpose
Router(config)# clear router-guard ip multicast statistics	Clears statistics for all access ports and for all VLANs on all trunk ports.
Router(config)# clear router-guard ip multicast statistics interface <i>interface_name</i>	Clears statistics for an access port and for all VLANs on a trunk port.
Router(config)# clear router-guard ip multicast statistics interface <i>interface_name</i> vlan <i>v</i>	Clears statistics for one particular VLAN on a trunk port.

This example shows how to clear statistics for one particular VLAN on a trunk port:

```
Router# clear router-guard ip multicast statistics interface interface_name vlan v
```

Verifying the Router Guard Configuration

- [Displaying Router Guard Configuration, page 11-4](#)
- [Displaying Router Guard Interfaces, page 11-5](#)

Displaying Router Guard Configuration

To display the global Router Guard configuration and the Router Guard configuration for a specific interface, perform these tasks:

Command	Purpose
Router# show router-guard	Displays the global Router Guard configuration.
Router# show router-guard interface <i>interface_name</i>	Displays the Router Guard configuration for a specific interface.

This example shows how to display the interface command output for a port in access mode with Router Guard not active:

```
Router# show router-guard interface g3/48
  Router Guard for IP Multicast:
  Globally enabled for all switch ports
  Enabled on this interface
  Packets denied:
    IGMP Queries:
    PIMv2 Messages:
    PIMv1 Messages:
    DVMRP Messages:
    RGMP Messages:
    CGMP Messages:
```

This example shows how to display the interface command output for a port in trunk mode:

```
Router# show router-guard interface g3/48
  Router Guard for IP Multicast:
  Globally enabled for all switch ports
  Disabled on this interface
```

This example shows how to verify that a trunk port is carrying VLANs 10 and 20:

```
Router# show router-guard interface g3/46
  Router Guard for IP Multicast:
  Globally enabled for all switch ports
  Default: Enabled for all VLANs on this interface
  VLAN 10:
  Enabled on this VLAN
  Packets denied:
    IGMP Queries:
    PIMv2 Messages:
    PIMv1 Messages:
    DVMRP Messages:
    RGMP Messages:
    CGMP Messages:
  VLAN 20 :
  Disabled on this VLAN
```

**Note**

If the port is in the shutdown state, the status will not be displayed because it cannot be determined whether the port is in trunk mode or access mode. You can use the **show running-config interface xxxx** command to display the Router Guard configuration.

Displaying Router Guard Interfaces

To display a list of all interfaces for which Router Guard is disabled, perform this task:

Command	Purpose
Router# show router-guard interface Router Guard for IP Multicast: Globally enabled for all switchports Interfaces: Gi3/46: Disabled on this port for VLANs: ALL	Displays a list of all interfaces for which Router Guard is disabled.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

12

IPv4 Multicast VPN Support

- [Prerequisites for mVPNs, page 12-1](#)
- [Restrictions for mVPNs, page 12-1](#)
- [Information About mVPN, page 12-3](#)
- [Default Settings for mVPNs, page 12-10](#)
- [How to Configure mVPNs, page 12-11](#)
- [Configuration Examples for mVPNs, page 12-27](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for mVPNs

None.

Restrictions for mVPNs

- [General Restrictions, page 12-2](#)
- [mVPN with L3VPN over mGRE Restrictions, page 12-3](#)

General Restrictions

- All PE routers in the multicast domain need to be running a Cisco IOS software image that supports the mVPN feature. There is no requirement for mVPN support on the P and CE routers.
- Support for IPv4 multicast traffic must be enabled on all backbone routers.
- The Border Gateway Protocol (BGP) routing protocol must be configured and operational on all routers supporting multicast traffic. In addition, BGP extended communities must be enabled (using the **neighbor send-community both** or **neighbor send-community extended** command) to support the use of MDTs in the network.
- When the switch is acting as a PE, and receives a multicast packet from a customer router with a time-to-live (TTL) value of 2, it drops the packet instead of encapsulating it and forwarding it across the mVPN link. Because such packets would normally be dropped by the PE at the other end of the mVPN link, this does not affect traffic flow.
- If the core multicast routing uses SSM, then the data and default multicast distribution tree (MDT) groups must be configured within the SSM range of IPv4 addresses.
- The update source interface for the BGP peerings must be the same for all BGP peerings configured on the router in order for the default MDT to be configured properly. If you use a loopback address for BGP peering, then PIM sparse mode must be enabled on the loopback address.
- The **ip mroute-cache** command must be enabled on the loopback interface used as the BGP peering interface in order for distributed multicast switching to function on the platforms that support it. The **no ip mroute-cache** command must *not* be present on these interfaces.
- Data MDTs are not created for VRF PIM dense mode multicast streams because of the flood and prune nature of dense mode multicast flows and the resulting periodic bring-up and tear-down of such data MDTs.
- Data MDTs are not created for VRF PIM bidirectional mode because source information is not available.
- mVPN does not support multiple BGP peering update sources, and configuring them can break mVPN RPF checking. The source IPv4 address of the mVPN tunnels is determined by the highest IPv4 address used for the BGP peering update source. If this IPv4 address is not the IPv4 address used as the BGP peering address with the remote PE router, mVPN will not function properly.
- MDT tunnels do not carry unicast traffic.
- If mVPN uses the infrastructure of an MPLS VPN network, you cannot apply MPLS tags or labels to multicast traffic over the VPNs.
- Each mVRF that is configured with a default MDT uses three hidden VLANs (one each for encapsulation, decapsulation, and interface), in addition to external, user-visible VLANs. This means that an absolute maximum of 1,000 mVRFs are supported on each router. (mVRFs without a configured MDT still use one internal VLAN, so unused mVRFs should be deleted to conserve VLAN allocation.)
- If your MPLS VPN network already contains a network of VRFs, you do not need to delete them or recreate them to be able to support mVRF traffic. Instead, configure the **mdt default** and **mdt data** commands, as listed in the following procedure, to enable multicast traffic over the VRF.
- The same mVRF must be configured on each PE router that is to support a particular VPN connection.
- Each PE router that supports a particular mVRF must be configured with the same **mdt default** command.

mVPN with L3VPN over mGRE Restrictions

- In releases earlier than Release 15.1(1)SY, With mVPN with L3VPN over mGRE configured, do not configure IPv4 routing on the supervisor engine ports or on ports on switching modules that have a CFC. (CSCtr05033)
- Ensure that the unicast path to the RP does not use any supervisor engine ports; additionally, in VSS mode, ensure that the unicast path to the RP does not use any ports on switching modules that have a CFC. (CSCts43614)
- When a GRE tunnel has the same destination address and source address as the mGRE tunnel, the GRE tunnel is route-cache switched.
- Packets that require fragmentation are route cache-switched.
- When an L3VPN profile is removed and added back, you should clear the Border Gateway Protocol (BGP) using the **clear ip bgp neighbor_ip_address soft** command.
- When an mGRE tunnel is created, a dummy tunnel is also created.
- The loopback or IP address used in the update source of the BGP configuration should be the same as that of the transport source of the L3VPN profile.
- mGRE is not stateful switchover (SSO) compliant. However, both mGRE and SSO coexist.
- Not all GRE options are supported in hardware (for example, the GRE extended header and GRE key).
- Checking identical VLANs (Internet Control Message Protocol [ICMP] redirect) is not supported on the tunnels.
- Features such as unicast reverse path forwarding (uRPF) and BGP policy accounting are not supported on the tunnels.

Information About mVPN

- [mVPN Overview, page 12-3](#)
- [Multicast Routing and Forwarding and Multicast Domains, page 12-4](#)
- [Multicast Distribution Trees, page 12-4](#)
- [Multicast Tunnel Interfaces, page 12-7](#)
- [PE Router Routing Table Support for mVPN, page 12-8](#)
- [Multicast Distributed Switching Support, page 12-8](#)
- [Hardware-Assisted IPv4 Multicast, page 12-8](#)
- [Information About mVPN with L3VPN over mGRE, page 12-9](#)

mVPN Overview

mVPN is a standards-based feature that transmits IPv4 multicast traffic across a virtualized provider network (for example, MPLS or mGRE tunnels). mVPN uses the IPv4 multicast traffic PFC hardware support to forward multicast traffic over VPNs at wire speeds. mVPN adds support for IPv4 multicast traffic over Layer 3 IPv4 VPNs to the existing IPv4 unicast support.

mVPN routes and forwards multicast packets for each individual VPN routing and forwarding (VRF) instance, as well as transmitting the multicast packets through VPN tunnels across the service provider backbone.

mVPN is an alternative to full-mesh point-to-point GRE tunnels, which is not a readily scalable solution and are limited in the granularity they provide to customers.

Multicast Routing and Forwarding and Multicast Domains

mVPN adds multicast routing information to the VPN routing and forwarding table. When a provider-edge (PE) router receives multicast data or control packets from a customer-edge (CE) router, forwarding is performed according to the information in the multicast VRF (mVRF).

Each mVRF maintains the routing and forwarding information that is needed for its particular VRF instance. An mVRF is created and configured in the same way as existing VRFs, except multicast routing is also enabled on each mVRF.

A multicast domain constitutes the set of hosts that can send multicast traffic to each other within the MPLS network. For example, the multicast domain for a customer that wanted to send certain types of multicast traffic to all global employees would consist of all CE routers associated with that enterprise.

Multicast Distribution Trees

The mVPN feature establishes at least one multicast distribution tree (MDT) for each multicast domain. The MDT provides the information needed to interconnect the same mVRFs that exist on the different PE routers.

mVPN supports two MDT types:

- **Default MDT**—The default MDT is a permanent channel for PIM control messages and low-bandwidth streams between all PE routers in a particular multicast domain. All multicast traffic in the default MDT is replicated to every other PE router in the domain. Each PE router is logically seen as a PIM neighbor (one hop away) from every other PE router in the domain.
- **Data MDT**—Data MDTs are optional. If enabled, they are dynamically created to provide optimal paths for high-bandwidth transmissions, such as full-motion video, that do not need to be sent to every PE router. This allows for on-demand forwarding of high-bandwidth traffic between PE routers, so as to avoid flooding every PE router with every high-bandwidth stream that might be created.

To create data MDTs, each PE router that is forwarding multicast streams to the backbone periodically examines the traffic being sent in each default MDT as follows:

1. Each PE router periodically samples the multicast traffic (approximately every 10 seconds for software switching, and 90 seconds for hardware switching) to determine whether a multicast stream has exceeded the configured threshold. (Depending on when the stream is sampled, this means that in a worst-case scenario, it could take up to 180 seconds before a high-bandwidth stream is detected.)

**Note**

Data MDTs are created only for (S, G) multicast route entries within the VRF multicast routing table. They are not created for (*, G) entries.

2. If a particular multicast stream exceeds the defined threshold, the sending PE router dynamically creates a data MDT for that particular multicast traffic.

3. The sending PE router then transmits a DATA-MDT JOIN request (which is a User Datagram Protocol (UDP) message to port 3232) to the other PE routers, informing them of the new data MDT.
4. Receiving PE routers examine their VRF routing tables to determine if they have any customers interested in receiving this data stream. If so, they use the PIM protocol to transmit a PIM JOIN message for this particular data MDT group (in the global table PIM instance) to accept the stream. Routers that do not currently have any customers for this stream still cache the information, in case any customers request it later on.
5. Three seconds after sending the DATA-MDT JOIN message, the sending PE router removes the high-bandwidth multicast stream from the default MDT and begins transmitting it over the new data MDT.
6. The sending PE router continues to send a DATA-MDT JOIN message every 60 seconds, as long as the multicast stream continues to exceed the defined threshold. If the stream falls below the threshold for more than 60 seconds, the sending PE router stops sending the DATA-MDT JOIN messages, and moves the stream back to the default MDT.
7. Receiving routers age out the cache information for the default MDT when they do not receive a DATA-MDT JOIN message for more than three minutes.


Data MDTs allow for high-bandwidth sources inside the VPN while still ensuring optimal traffic forwarding in the MPLS VPN core.

In the following example, a service provider has a multicast customer with offices in San Jose, New York, and Dallas. The San Jose site is transmitting a one-way multicast presentation. The service provider network supports all three sites associated with this customer, in addition to the Houston site of a different enterprise customer.

The default MDT for the enterprise customer consists of provider routers P1, P2, and P3 and their associated PE routers. Although PE4 is interconnected to these other routers in the MPLS core, PE4 is associated with a different customer and is therefore not part of the default MDT.


Figure 12-1 shows the situation in this network when no one outside of San Jose has joined the multicast broadcast, which means that no data is flowing along the default MDT. Each PE router maintains a PIM relationship with the other PE routers over the default MDT, as well as a PIM relationship with its directly attached PE routers.

Figure 12-1 Default Multicast Distribution Tree Overview



If an employee in New York joins the multicast session, the PE router associated for the New York site sends a join request that flows across the default MDT for the multicast domain. The PE router associated with the multicast session source (PE1) receives the request. Figure 12-2 shows how the PE router forwards the request to the CE router associated with the multicast source (CE1a).

Figure 12-2 *Initializing the Data MDT*



The CE router (CE1a) starts sending the multicast data to the associated PE router (PE1), which recognizes that the multicast data exceeds the bandwidth threshold at which a data MDT should be created. PE1 then creates a data MDT and sends a message to all routers using the default MDT that contains information about the data MDT.

Approximately three seconds later, PE1 begins sending the multicast data for that particular stream using the data MDT. Because only PE2 has receivers who are interested in this source, only PE2 joins the data MDT and receives traffic on it.

Multicast Tunnel Interfaces

The PE router creates a multicast tunnel interface (MTI) for each multicast VRF (mVRF) in the multicast domain. The mVRF uses the tunnel interface to access the multicast domain to provide a conduit that connects an mVRF and the global mVRF.

On the router, the MTI is a tunnel interface (created with the **interface tunnel** command) with a class D multicast address. All PE routers that are configured with a default MDT for this mVRF create a logical network in which each PE router appears as a PIM neighbor (one hop away) to every other PE router in the multicast domain, regardless of the actual physical distance between them.

The MTI is automatically created when an mVRF is configured. The BGP peering address is assigned as the MTI interface source address, and the PIM protocol is automatically enabled on each MTI.

When the router receives a multicast packet from the customer side of the network, it uses the incoming interface's VRF to determine which mVRFs should receive it. The router then encapsulates the packet using GRE encapsulation. When the router encapsulates the packet, it sets the source address to that of the BGP peering interface and sets the destination address to the multicast address of the default MDT, or to the source address of the data MDT if configured. The router then replicates the packet as needed for forwarding on the appropriate number of MTI interfaces.

When the router receives a packet on the MTI interface, it uses the destination address to identify the appropriate default MDT or data MDT, which in turn identifies the appropriate mVRF. It then decapsulates the packet and forwards it out the appropriate interfaces, replicating it as many times as are necessary.

**Note**

- Unlike other tunnel interfaces that are commonly used on Cisco routers, the mVPN MTI is classified as a LAN interface, not a point-to-point interface. The MTI interface is not configurable, but you can use the **show interface tunnel** command to display its status.
- The MTI interface is used exclusively for multicast traffic over the VPN tunnel.
- The tunnel does not carry unicast routed traffic.

PE Router Routing Table Support for mVPN

Each PE router that supports the mVPN feature uses the following routing tables to ensure that the VPN and mVPN traffic is routed correctly:

- Default routing table—Standard routing table used in all Cisco routers. This table contains the routes that are needed for backbone traffic and for non-VPN unicast and multicast traffic (including Generic Routing Encapsulation (GRE) multicast traffic).
- VPN routing/forwarding (VRF) table—Routing table created for each VRF instance. Responsible for routing the unicast traffic between VPNs in the provider network.
- Multicast VRF (mVRF) table—Multicast routing table and multicast routing protocol instance created for each VRF instance. Responsible for routing the multicast traffic in the multicast domain of the network. This table also includes the multicast tunnel interfaces that are used to access the multicast domain.

Multicast Distributed Switching Support

mVPN supports multicast distributed switching (MDS) for multicast support on a per-interface and a per-VRF basis. When configuring MDS, you must make sure that no interface (including loopback interfaces) has the **no ip mroute-cache** command configured.

Hardware-Assisted IPv4 Multicast

Cisco IOS Release 15.4SY supports hardware acceleration for IPv4 multicast over VPN traffic, which forwards multicast traffic to the appropriate VPNs at wire speed without increased RP CPU utilization.

In a customer VRF, PFC hardware acceleration supports multicast traffic in PIM dense, PIM sparse, PIM bidirectional, and PIM Source Specific Multicast (SSM) modes.

In the service provider core, PFC hardware acceleration supports multicast traffic in PIM sparse, PIM bidirectional, and PIM SSM modes. In the service provider core, PFC hardware acceleration does not support multicast traffic in PIM dense mode.

Information About mVPN with L3VPN over mGRE

- [Overview, page 12-9](#)
- [Route Maps, page 12-9](#)
- [Tunnel Endpoint Discovery and Forwarding, page 12-10](#)
- [Tunnel Decapsulation, page 12-10](#)
- [Tunnel Source, page 12-10](#)

**Note**

See the [“Configuring mVPN with L3VPN over mGRE” section on page 12-22](#) for more information.

Overview

Release 15.0(1)SY1 and later releases support multicast virtual private network with Layer 3 virtual private network over multipoint generic routing encapsulation (mVPN with L3VPN over mGRE). mVPN with L3VPN over mGRE provides VPN connectivity between networks that are connected by standards-based IP-only networks. mGRE tunnels overlay the IP network and connect PE devices to transport VPNs that support deployment of L3 PE-based VPN services over the IP core.

**Note**

- Because mGRE is a point-to-multipoint model, fully meshed GRE tunnels are not required to interconnect PE devices.
- Multicast and unicast traffic use separate tunnels, MDT for multicast and mGRE for unicast.

Route Maps

By default, VPN unicast traffic is sent using an LSP. The mVPN with L3VPN over mGRE feature uses user-defined route maps to determine which VPN prefixes are reachable over an mGRE tunnel and which VPN prefixes are reachable using an LSP. The route map is applied to advertisements for VPNv4 and VPNv6 address families. The route map uses a next hop tunnel table to determine the encapsulation method for the VPN traffic.

Traffic routed over an mGRE tunnel uses an alternative address space; all next hops are reached by encapsulating the traffic in an mGRE tunnel. To configure a specific route to use an mGRE tunnel, requires configuration of an entry for that route to the route map. The new entry remaps the Network Layer Reachability Information (NLRI) of the route to the alternative address space. If there is no remap entry in the route map for a route, then traffic on that route is forwarded over an LSP.

The mVPN with L3VPN over mGRE feature automatically provisions the alternative address space, normally held in the tunnel-encapsulated virtual routing and forwarding (VRF) instance. To ensure that all traffic reachable through the address space is encapsulated in an mGRE tunnel, the feature installs a single default route out of a tunnel. The feature also creates a default tunnel on the route map. The default route map can be attached to the appropriate BGP updates.

Tunnel Endpoint Discovery and Forwarding

The mVPN with L3VPN over mGRE feature must be able to discover the remote PEs in the network and construct tunnel forwarding information for the remote PEs. The feature must be able to detect when a remote PE is no longer valid and remove the tunnel forwarding information for that PE.

If an ingress PE receives a VPN advertisement over BGP, it uses the route target attributes (which it inserts into the VRF) and the MPLS VPN label from the advertisement, to associate the prefixes with the appropriate customer. The next hop of the inserted route is set to the NLRI of the advertisement.

The advertised prefixes contain information about remote PEs in the system (in the form of NLRIs), and the PE uses this information to notify the system when an NLRI becomes active or inactive. The system uses this notification to update the PE forwarding information.

When the feature receives notification of a new remote PE, it adds the information to the tunnel endpoint database and creates an adjacency associated with the tunnel interface. The adjacency description includes information on the encapsulation and other processing required to send encapsulated packets to the new remote PE.

The feature places the adjacency information into the tunnel encapsulated VRF. When a VPN NLRI is remapped to a route in the VRF (using the route map), the feature links the NLRI to the adjacency, which links the VPN to the tunnel.

Tunnel Decapsulation

When the egress PE receives a packet from a tunnel interface that uses the mVPN with L3VPN over mGRE feature, the PE decapsulates the packet to create a VPN label tagged packet, and forwards the packet.

Tunnel Source

The mVPN with L3VPN over mGRE feature uses a single tunnel configured as an mGRE tunnel to configure a system with a large number of endpoints (remote PEs). To identify the origin of tunnel-encapsulated packets, the system uses the tunnel source information.

At the transmitting (ingress) PE, when a VPN packet is sent to a tunnel, the tunnel destination is the NLRI. At a receiving (egress) PE, the tunnel source is the address that the packets encapsulated in the mGRE tunnel are received on. Therefore, at the egress PE the packet destination must match the NLRI from the local PE.

Default Settings for mVPNs

None.

How to Configure mVPNs

- [Configuring a Multicast VPN Routing and Forwarding Instance, page 12-11](#)
- [Configuring Multicast VRF Routing, page 12-16](#)
- [Configuring Interfaces for Multicast Routing to Support mVPN, page 12-20](#)
- [Configuring mVPN with L3VPN over mGRE, page 12-22](#)



Note

These configuration tasks assume that BGP is already configured and operational on all routers that are sending or receiving the multicast traffic. In addition, BGP extended communities must be enabled (using the **neighbor send-community both** or **neighbor send-community extended** command) to support the use of MDTs in the network.

Configuring a Multicast VPN Routing and Forwarding Instance

- [Configuring a VRF Entry, page 12-11](#)
- [Configuring the Route Distinguisher, page 12-12](#)
- [Configuring the Route-Target Extended Community, page 12-12](#)
- [Configuring the Default MDT, page 12-13](#)
- [Configuring Data MDTs \(Optional\), page 12-13](#)
- [Enabling Data MDT Logging, page 12-14](#)
- [Sample Configuration, page 12-14](#)
- [Displaying VRF Information, page 12-15](#)

Configuring a VRF Entry

To configure a VRF entry, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip vrf vrf_name	Configures a VRF routing table entry and a Cisco Express Forwarding (CEF) table entry and enters VRF configuration mode.
Step 3	Router(config-vrf)# do show ip vrf vrf_name	Verifies the configuration.

This example show how to configure a VRF named blue and verify the configuration:

```
Router# configure terminal
Router(config)# ip vrf blue
Router(config-vrf)# do show ip vrf blue
Name                Default RD          Interfaces
blue                <not set>
```

Configuring the Route Distinguisher

To configure the route distinguisher, perform this task:

	Command	Purpose
Step 1	Router(config-vrf)# rd <i>route_distinguisher</i>	Specifies the route distinguisher for a VPN IPv4 prefix.
Step 2	Router(config-vrf)# do show ip vrf <i>vrf_name</i>	Verifies the configuration.

When configuring the route distinguisher, enter the route distinguisher in one of the following formats:

- 16-bit AS number:your 32-bit number (101:3)
- 32-bit IPv4 address:your 16-bit number (192.168.122.15:1)

This example show how to configure 55:1111 as the route distinguisher and verify the configuration:

```
Router(config-vrf)# rd 55:1111
Router(config-vrf)# do show ip vrf blue
Name                               Default RD           Interfaces
blue                               55:1111
```

Configuring the Route-Target Extended Community

To configure the route-target extended community, perform this task:

	Command	Purpose
Step 1	Router(config-vrf)# route-target [import export both] <i>route_target_ext_community</i>	Configures a route-target extended community for the VRF.
Step 2	Router(config-vrf)# do show ip vrf detail	Verifies the configuration.

When configuring the route-target extended community, note the following information:

- **import**—Imports routing information from the target VPN extended community.
- **export**—Exports routing information to the target VPN extended community.
- **both**—Imports and exports.
- *route_target_ext_community*—Adds the 48-bit route-target extended community to the VRF. Enter the number in one of the following formats:
 - 16-bit AS number:your 32-bit number (101:3)
 - 32-bit IPv4 address:your 16-bit number (192.168.122.15:1)

This example shows how to configure 55:1111 as the import and export route-target extended community and verify the configuration:

```
Router(config-vrf)# route-target both 55:1111
Router(config-vrf)# do show ip vrf detail
VRF blue; default RD 55:1111; default VPNID <not set>
VRF Table ID = 1
  No interfaces
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:55:1111
  Import VPN route-target communities
    RT:55:1111
  No import route-map
  No export route-map
  CSC is not configured.
```

Configuring the Default MDT

To configure the default MDT, perform this task:

Command	Purpose
Router(config-vrf)# mdt default <i>group_address</i>	Configures the default MDT.

When configuring the default MDT, note the following information:

- The *group_address* is the multicast IPv4 address of the default MDT group. This address serves as an identifier for the mVRF community, because all provider-edge (PE) routers configured with this same group address become members of the group, which allows them to receive the PIM control messages and multicast traffic that are sent by other members of the group.
- This same default MDT must be configured on each PE router to enable the PE routers to receive multicast traffic for this particular mVRF.

This example shows how to configure 239.1.1.1 as the default MDT:

```
Router(config-vrf)# mdt default 239.1.1.1
```

Configuring Data MDTs (Optional)

To configure optional data MDTs, perform this task:

Command	Purpose
Router(config-vrf)# mdt data <i>group_address</i> <i>wildcard_bits</i> [threshold <i>threshold_value</i>] [list <i>access_list</i>]	(Optional) Configures a data MDTs for the specified range of multicast addresses.

When configuring optional data MDTs, note the following information:

- *group_address1*—Multicast group address. The address can range from 224.0.0.1 to 239.255.255.255, but cannot overlap the address that has been assigned to the default MDT.
- *wildcard_bits*—Wildcard bit mask to be applied to the multicast group address to create a range of possible addresses. This allows you to limit the maximum number of data MDTs that each mVRF can support.

- **threshold** *threshold_value*—(Optional) Defines the threshold value in kilobits, at which multicast traffic should be switched from the default MDT to the data MDT. The *threshold_value* parameter can range from 1 through 4294967 kilobits.
- **list** *access_list*—(Optional) Specifies an access list name or number to be applied to this traffic.

This example shows how to configure a data MDT:

```
Router(config-vrf)# mdt data 239.1.2.0 0.0.0.3 threshold 10
```

Enabling Data MDT Logging

To enable data MDT logging, perform this task:

Command	Purpose
Router(config-vrf)# mdt log-reuse	(Optional) Enables the recording of data MDT reuse information, by generating a SYSLOG message whenever a data MDT is reused. Frequent reuse of a data MDT might indicate a need to increase the number of allowable data MDTs by increasing the size of the wildcard bitmask that is used in the mdt data command.

This example shows how to enable data MDT logging:

```
Router(config-vrf)# mdt log-reuse
```

Sample Configuration

The following excerpt from a configuration file shows typical VRF configurations for a range of VRFs. To simplify the display, only the starting and ending VRFs are shown.

```
!
ip vrf mvpn-cus1
 rd 200:1
  route-target export 200:1
  route-target import 200:1
  mdt default 239.1.1.1
!
ip vrf mvpn-cus2
 rd 200:2
  route-target export 200:2
  route-target import 200:2
  mdt default 239.1.1.2
!
ip vrf mvpn-cus3
 rd 200:3
  route-target export 200:3
  route-target import 200:3
  mdt default 239.1.1.3
!
...

ip vrf mvpn-cus249
 rd 200:249
  route-target export 200:249
  route-target import 200:249
  mdt default 239.1.1.249
  mdt data 239.1.1.128 0.0.0.7
```

Displaying VRF Information

To display all of the VRFs that are configured on the switch, use the **show ip vrf** command:

```
Router# show ip vrf
```

Name	Default RD	Interfaces
green	1:52	GigabitEthernet6/1
red	200:1	GigabitEthernet1/1
		GigabitEthernet3/16
		Loopback2

```
Router#
```

To display information about the MDTs that are currently configured for all mVRFs, use the **show ip pim mdt** command. The following example shows typical output for this command:

```
Router# show ip pim mdt
```

MDT Group	Interface	Source	VRF
* 227.1.0.1	Tunnel1	Loopback0	BIDIR01
* 227.2.0.1	Tunnel2	Loopback0	BIDIR02
* 228.1.0.1	Tunnel3	Loopback0	SPARSE01
* 228.2.0.1	Tunnel4	Loopback0	SPARSE02



Note

To display information about a specific tunnel interface, use the **show interface tunnel** command. The IPv4 address for the tunnel interface is the multicast group address for the default MDT of the mVRF.

To display routing information for a particular VRF, use the **show ip route vrf** command:

```
Router# show ip route vrf red
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```

      2.0.0.0/32 is subnetted, 1 subnets
C      2.2.2.2 is directly connected, Loopback2
      3.0.0.0/32 is subnetted, 1 subnets
B      3.3.3.3 [200/0] via 3.1.1.3, 00:20:09
C      21.0.0.0/8 is directly connected, GigabitEthernet3/16
B      22.0.0.0/8 [200/0] via 3.1.1.3, 00:20:09
```

```
Router#
```

To display information about the multicast routing table and tunnel interface for a particular mVRF, use the **show ip mroute vrf** command. The following example shows typical output for a mVRF named BIDIR01:

```
Router# show ip mroute vrf BIDIR01
```

```
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
```

```

X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report, Z - Multicast Tunnel
Y - Joined MDT-data group, y - Sending to MDT-data group
Outgoing interface flags: H - Hardware switched
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 228.1.0.1), 00:16:25/stopped, RP 10.10.10.12, flags: SJCF
Incoming interface: Tunnel1, RPF nbr 10.10.10.12, Partial-SC
Outgoing interface list:
GigabitEthernet3/1.3001, Forward/Sparse-Dense, 00:16:25/00:02:49, H
(6.9.0.100, 228.1.0.1), 00:14:13/00:03:29, flags: FT
Incoming interface: GigabitEthernet3/1.3001, RPF nbr 0.0.0.0, RPF-MFD
Outgoing interface list:
Tunnel1, Forward/Sparse-Dense, 00:14:13/00:02:46, H

```

Router#



Note

In this example, the **show ip mroute vrf** command shows that Tunnel1 is the MDT tunnel interface (MTI) being used by this VRF.

Configuring Multicast VRF Routing

- [Enabling IPv4 Multicast Routing Globally, page 12-16](#)
- [Enabling IPv4 Multicast VRF Routing, page 12-17](#)
- [Specifying the PIM VRF RP Address, page 12-17](#)
- [Configuring a PIM VRF Register Message Source Address \(Optional\), page 12-18](#)
- [Configuring an MSDP Peer \(Optional\), page 12-18](#)
- [Configuring the Maximum Number of Multicast Routes \(Optional\), page 12-18](#)
- [Sample Configuration, page 12-19](#)
- [Displaying IPv4 Multicast VRF Routing Information, page 12-20](#)



Note

BGP should be already configured and operational on all routers that are sending or receiving multicast traffic. In addition, BGP extended communities must be enabled (using the **neighbor send-community both** or **neighbor send-community extended** command) to support the use of MDTs in the network.

Enabling IPv4 Multicast Routing Globally

To enable IPv4 multicast routing globally, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip multicast-routing	Enables IPv4 multicast routing globally.

This example show how to enable IPv4 multicast routing globally:

```

Router# configure terminal
Router(config)# ip multicast-routing

```

Enabling IPv4 Multicast VRF Routing

To enable IPv4 multicast VRF routing, perform this task:

Command	Purpose
Router(config)# ip multicast-routing vrf <i>vrf_name</i> [distributed]	Enables IPv4 multicast VRF routing.

When enabling IPv4 multicast VRF routing, note the following information:

- *vrf_name*—Specifies a particular VRF for multicast routing. The *vrf_name* should see a VRF that has been previously created, as specified in the [“Configuring a Multicast VPN Routing and Forwarding Instance”](#) section on page 12-11.
- **distributed**—(Optional) Enables Multicast Distributed Switching (MDS).

This example show how to enable IPv4 multicast VRF routing:

```
Router# configure terminal  
Router(config)# ip multicast-routing vrf blue
```

Specifying the PIM VRF RP Address

To specify the PIM VRF rendezvous point (RP) address, perform this task:

Command	Purpose
Router(config)# ip pim vrf <i>vrf_name</i> rp-address <i>rp_address</i> [<i>access_list</i>] [override] [bidir]	Specifies the PIM RP IPv4 address for a (required for sparse PIM networks):

When specifying the PIM VRF RP address, note the following information:

- **vrf** *vrf_name*—(Optional) Specifies a particular VRF instance to be used.
- *rp_address*—Unicast IP address for the PIM RP router.
- *access_list*—(Optional) Number or name of an access list that defines the multicast groups for the RP.
- **override**—(Optional) In the event of conflicting RP addresses, this particular RP overrides any RP that is learned through Auto-RP.
- **bidir**—(Optional) Specifies that the multicast groups specified by the *access_list* argument are to operate in bidirectional mode. If this option is not specified, the groups operate in PIM sparse mode.
- Use bidirectional mode whenever possible, because it offers better scalability.

This example show how to specify the PIM VRF RP address:

```
Router(config)# ip pim vrf blue rp-address 198.196.100.33
```

Configuring a PIM VRF Register Message Source Address (Optional)

To configure a PIM VRF register message source address, perform this task:

Command	Purpose
Router(config)# ip pim vrf <i>vrf_name</i> register-source <i>interface_type interface_number</i>	(Optional) Configures a PIM VRF register message source address. You can configure a loopback interface as the source of the register messages.

This example show how to configure a PIM VRF register message source address:

```
Router(config)# ip pim vrf blue register-source loopback 3
```

Configuring an MSDP Peer (Optional)

To configure a multicast source discovery protocol (MSDP) peer, perform this task:

Command	Purpose
Router(config)# ip msdp vrf <i>vrf_name</i> peer { <i>peer_name</i> <i>peer_address</i> } [connect-source <i>interface_type interface_number</i>] [remote-as <i>ASN</i>]	(Optional) Configures an MSDP peer.

When configuring an MSDP peer, note the following information:

- **vrf** *vrf_name*—Specifies a particular VRF instance to be used.
- {*peer_name* | *peer_address*}—Domain Name System (DNS) name or IP address of the MSDP peer router.
- **connect-source** *interface_type interface_number*—Interface name and number for the interface whose primary address is used as the source IP address for the TCP connection.
- **remote-as** *ASN*—(Optional) Autonomous system number of the MSDP peer. This is for display-only purposes.

This example show how to configure an MSDP peer:

```
Router(config)# ip msdp peer router.cisco.com connect-source gigabitethernet 1/1 remote-as 109
```

Configuring the Maximum Number of Multicast Routes (Optional)

To configure the maximum number of multicast routes, perform this task:

Command	Purpose
Router(config)# ip multicast vrf <i>vrf_name</i> route-limit <i>limit</i> [<i>threshold</i>]	(Optional) Configures the maximum number of multicast routes that can be added for multicast traffic.

When configuring the maximum number of routes, note the following information:

- **vrf vrf_name**— Enables route limiting for the specified VRF.
- **limit**—The number of multicast routes that can be added. The range is from 1 to 2147483647, with a default of 2147483647.
- **threshold**—(Optional) Number of multicast routes that can be added before a warning message occurs. The valid range is from 1 to the value of the *limit* parameter.

This example show how to configure the maximum number of multicast routes:

```
Router(config)# ip multicast vrf blue route-limit 200000 20000
```

Configuring IPv4 Multicast Route Filtering (Optional)

To configure IPV4 multicast route filtering, perform this task:

Command	Purpose
Router(config)# ip multicast mrimfo-filter access_list	(Optional) Configures IPV4 multicast route filtering with an access list. The <i>access_list</i> parameter can be the name or number of a access list.

This example show how to configure IPV4 multicast route filtering:

```
Router(config)# ip multicast mrimfo-filter 101
```

Sample Configuration

The following excerpt from a configuration file shows the minimum configuration that is needed to support multicast routing for a range of VRFs. To simplify the display, only the starting and ending VRFs are shown.

```
!
ip multicast-routing
ip multicast-routing vrf lite
ip multicast-routing vrf vpn201
ip multicast-routing vrf vpn202

...

ip multicast-routing vrf vpn249
ip multicast-routing vrf vpn250

...

ip pim rp-address 192.0.1.1
ip pim vrf lite rp-address 104.1.1.2
ip pim vrf vpn201 rp-address 192.200.1.1
ip pim vrf vpn202 rp-address 192.200.2.1

...

ip pim vrf vpn249 rp-address 192.200.49.6
ip pim vrf vpn250 rp-address 192.200.50.6
...
```

Displaying IPv4 Multicast VRF Routing Information

To display the known PIM neighbors for a particular mVRF, use the **show ip pim vrf neighbor** command:

```
Router# show ip pim vrf 98 neighbor
```

```
PIM Neighbor Table
Neighbor      Interface      Uptime/Expires    Ver    DR
Address
40.60.0.11    Tunnel96        00:00:31/00:01:13 v2      1 / S
40.50.0.11    Tunnel96        00:00:54/00:00:50 v2      1 / S
```

```
Router#
```

Configuring Interfaces for Multicast Routing to Support mVPN

- [Multicast Routing Configuration Overview, page 12-20](#)
- [Configuring PIM on an Interface, page 12-21](#)
- [Configuring an Interface for IPv4 VRF Forwarding, page 12-21](#)
- [Sample Configuration, page 12-22](#)

Multicast Routing Configuration Overview

Protocol Independent Multicast (PIM) must be configured on all interfaces that are being used for IPv4 multicast traffic. In a VPN multicast environment, you should enable PIM on at least all of the following interfaces:

- Physical interface on a provider edge (PE) router that is connected to the backbone.
- Loopback interface that is used for BGP peering.
- Loopback interface that is used as the source for the sparse PIM rendezvous point (RP) router address.

In addition, you must also associate mVRFs with those interfaces over which they are going to forward multicast traffic.

BGP should be already configured and operational on all routers that are sending or receiving multicast traffic. In addition, BGP extended communities must be enabled (using the **neighbor send-community both** or **neighbor send-community extended** command) to support the use of MDTs in the network.

Configuring PIM on an Interface

To configure PIM on an interface, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type {slot/port number}</i>	Enters interface configuration mode for the specified interface.
Step 3	Router(config-if)# ip pim {dense-mode sparse-mode sparse-dense-mode}	Enables PIM on the interface.

When configuring PIM on an interface, note the following information:

- You can use one of these interface types:
 - A physical interface on a provider edge (PE) router that is connected to the backbone.
 - A loopback interface that is used for BGP peering.
 - A loopback interface that is used as the source for the sparse PIM network rendezvous point (RP) address.
- These are the PIM modes:
 - dense-mode**—Enables dense mode of operation.
 - sparse-mode**—Enables sparse mode of operation.
 - sparse-dense-mode**—Enables sparse mode if the multicast group has an RP router defined, or enables dense mode if an RP router is not defined.
- Use **sparse-mode** for the physical interfaces of all PE routers that are connected to the backbone, and on all loopback interfaces that are used for BGP peering or as the source for RP addressing.

This example shows how to configure PIM sparse mode on a physical interface:

```
Router# configure terminal
Router(config)# interface gigabitethernet 10/1
Router(config-if)# ip pim sparse-mode
```

This example shows how to configure PIM sparse mode on a loopback interface:

```
Router# configure terminal
Router(config)# interface loopback 2
Router(config-if)# ip pim sparse-mode
```

Configuring an Interface for IPv4 VRF Forwarding

To configure an interface for IPv4 VRF forwarding, perform this task:

Command	Purpose
Router(config-if)# ip vrf forwarding <i>vrf_name</i>	(Optional) Associates the specified VRF routing and forwarding tables with the interface. If this is not specified, the interface defaults to using the global routing table. Note Entering this command on an interface removes the IP address, so reconfigure the IP address.

This example shows how to configure the interface for VRF blue forwarding:

```
Router(config-if)# ip vrf forwarding blue
```

Sample Configuration

The following excerpt from a configuration file shows the interface configuration, along with the associated mVRF configuration, to enable multicast traffic over a single mVRF:

```
ip multicast-routing vrf blue
ip multicast-routing

ip vrf blue
  rd 100:27
  route-target export 100:27
  route-target import 100:27
  mdt default 239.192.10.2

interface GigabitEthernet1/1
  description blue connection
  ip vrf forwarding blue
  ip address 192.168.2.26 255.255.255.0
  ip pim sparse-mode

interface GigabitEthernet1/15
  description Backbone connection
  ip address 10.8.4.2 255.255.255.0
  ip pim sparse-mode

ip pim vrf blue rp-address 192.7.25.1
ip pim rp-address 10.1.1.1
```

Configuring mVPN with L3VPN over mGRE

- [Configuring an L3VPN Encapsulation Profile, page 12-22](#) (required)
- [Configuring BGP and Route Maps, page 12-23](#) (required)



Note

See the [“Information About mVPN with L3VPN over mGRE”](#) section on [page 12-9](#) for more information.

Configuring an L3VPN Encapsulation Profile



Note

Transport protocols such as IPv6, MPLS, IP, and Layer 2 Tunneling Protocol version 3 (L2TPv3) can also be used in this configuration.

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	l3vpn encapsulation ip <i>profile-name</i> Example: Router(config)# l3vpn encapsulation ip tunnel encap	Enters L3 VPN encapsulation configuration mode to create the tunnel.
Step 4	transport ipv4 [source <i>interface-type interface-number</i>] Example: Router(config-l3vpn-encap-ip)# transport ipv4 source loopback 0	(Optional) Specifies IPv4 transport source mode and defines the transport source interface. - If you use the transport ipv4 source <i>interface-type interface-number</i> command, make sure that the specified source address is used as the next hop in BGP updates advertised by the PE. - If you do not use this command, the bgp update source or bgp next-hop command is automatically used as the tunnel source.
Step 5	protocol gre [key <i>gre-key</i>] Example: Router(config-l3vpn-encap-ip)# protocol gre key 1234	Specifies GRE as the tunnel mode and sets the GRE key.
Step 6	end Example: Router(config-l3vpn-encap-ip)# end	Exits L3 VPN encapsulation configuration mode and returns to privileged EXEC mode.
Step 7	show l3vpn encapsulation ip <i>profile-name</i> Example: Router# show l3vpn encapsulation ip tunnel encap	(Optional) Displays the profile health and the underlying tunnel interface.

Configuring BGP and Route Maps

Perform this task to configure BGP and route maps. The following steps also enable you to link the route map to the application template and set up the BGP VPNv4 and VPNv6 exchange so that the updates are filtered through the route map.

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>as-number</i> Example: Router(config)# router bgp 100	Specifies the number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along, and enters router configuration mode.
Step 4	bgp log-neighbor-changes Example: Router(config-router)# bgp log-neighbor-changes	Enables logging of BGP neighbor resets.
Step 5	neighbor <i>ip-address</i> remote-as <i>as-number</i> Example: Router(config-router)# neighbor 209.165.200.225 remote-as 100	Adds an entry to the BGP or multiprotocol BGP neighbor table.
Step 6	neighbor <i>ip-address</i> update-source <i>interface name</i> Example: Router(config-router)# neighbor 209.165.200.225 update-source loopback 0	Allows BGP sessions to use any operational interface for TCP connections.
Step 7	address-family ipv4 Example: Router(config-router)# address-family ipv4	Enters address family configuration mode to configure routing sessions that use IPv4 address prefixes.
Step 8	no synchronization Example: Router(config-router-af)# no synchronization	Enables the Cisco IOS software to advertise a network route without waiting for an IGP.
Step 9	redistribute connected Example: Router(config-router-af)# redistribute connected	Redistributes routes from one routing domain into another routing domain and allows the target protocol to redistribute routes learned by the source protocol and connected prefixes on those interfaces over which the source protocol is running.
Step 10	neighbor <i>ip-address</i> activate Example: Router(config-router-af)# neighbor 209.165.200.225 activate	Enables the exchange of information with a BGP neighbor.

	Command or Action	Purpose
Step 11	no auto-summary Example: Router(config-router-af)# no auto-summary	Disables automatic summarization and sends subprefix routing information across classful network boundaries.
Step 12	exit Example: Router(config-router-af)# exit	Exits address family configuration mode.
Step 13	address-family vpnv4 Example: Router(config-router)# address-family vpnv4	Enters address family configuration mode to configure routing sessions, such as BGP, that use standard VPNv4 address prefixes.
Step 14	neighbor ip-address activate Example: Router(config-router-af)# neighbor 209.165.200.225 activate	Enables the exchange of information with a BGP neighbor.
Step 15	neighbor ip-address send-community both Example: Router(config-router-af)# neighbor 209.165.200.225 send-community both	Specifies that a communities attribute, for both standard and extended communities, should be sent to a BGP neighbor.
Step 16	neighbor ip-address route-map map-name in Example: Router(config-router-af)# neighbor 209.165.200.225 route-map SELECT_UPDATE_FOR_L3VPN in	Applies the named route map to the incoming route.
Step 17	exit Example: Router(config-router-af)# exit	Exits address family configuration mode.
Step 18	address-family vpnv6 Example: Router(config-router)# address-family vpnv6	Enters address family configuration mode to configure routing sessions, such as BGP, that use VPNv6 address prefixes.
Step 19	neighbor ip-address activate Example: Router(config-router-af)# neighbor 209.165.200.252 activate	Enables the exchange of information with a BGP neighbor.
Step 20	neighbor ip-address send-community both Example: Router(config-router-af)# neighbor 209.165.200.252 send-community both	Specifies that a communities attribute, for both standard and extended communities, should be sent to a BGP neighbor.

	Command or Action	Purpose
Step 21	neighbor <i>ip-address</i> route-map <i>map-name</i> in Example: Router(config-router-af)# neighbor 209.165.200.252 route-map SELECT_UPDATE_FOR_L3VPN in	Applies the named route map to the incoming route.
Step 22	exit Example: Router(config-router-af)# exit	Exits address family configuration mode.
Step 23	route-map <i>map-tag</i> permit <i>position</i> Example: Router(config-router)# route-map SELECT_UPDATE_FOR_L3VPN permit 10	Enters route-map configuration mode and defines the conditions for redistributing routes from one routing protocol into another. - The redistribute router configuration command uses the specified map tag to reference this route map. Multiple route maps may share the same map tag name. - If the match criteria are met for this route map, the route is redistributed as controlled by the set actions. - If the match criteria are not met, the next route map with the same map tag is tested. If a route passes none of the match criteria for the set of route maps sharing the same name, it is not redistributed by that set. - The <i>position</i> argument indicates the position a new route map will have in the list of route maps already configured with the same name.
Step 24	set ip next-hop encapsulate l3vpn <i>profile-name</i> Example: Router(config-route-map)# set ip next-hop encapsulate l3vpn my profile	Indicates that output IPv4 packets that pass a match clause of the route map are sent to the VRF for tunnel encapsulation.
Step 25	set ipv6 next-hop encapsulate l3vpn <i>profile-name</i> Example: Router(config-route-map)# set ip next-hop encapsulate l3vpn tunnel encap	Indicates that output IPv6 packets that pass a match clause of the route map are sent to the VRF for tunnel encapsulation.
Step 26	exit Example: Router(config-route-map)# exit	Exits route-map configuration mode and enters global configuration mode.
Step 27	exit Example: Router(config)# exit	Exits global configuration mode.

Configuration Examples for mVPNs

- [mVPN Configuration with Default MDTs Only, page 12-27](#)
- [mVPN Configuration with Default and Data MDTs, page 12-29](#)
- [Verifying the mVPN with L3VPN over mGRE Configuration, page 12-32](#)
- [Configuration Sequence for mVPN with L3VPN over mGRE, page 12-33](#)

mVPN Configuration with Default MDTs Only

The following excerpt from a configuration file shows the lines that are related to the mVPN configuration for three mVRFs. (The required BGP configuration is not shown.)

```
!  
service timestamps debug datetime msec  
service timestamps log datetime msec  
service password-encryption  
service compress-config  
!  
hostname MVPN Router  
!  
boot system flash slot0:  
logging snmp-authfail  
!  
ip subnet-zero  
!  
no ip domain-lookup  
ip host tftp 223.255.254.238  
!  
ip vrf mvpn-cus1  
  rd 200:1  
  route-target export 200:1  
  route-target import 200:1  
  mdt default 239.1.1.1  
!  
ip vrf mvpn-cus2  
  rd 200:2  
  route-target export 200:2  
  route-target import 200:2  
  mdt default 239.1.1.2  
!  
ip vrf mvpn-cus3  
  rd 200:3  
  route-target export 200:3  
  route-target import 200:3  
  mdt default 239.1.1.3  
!  
ip multicast-routing  
ip multicast-routing vrf mvpn-cus1  
ip multicast-routing vrf mvpn-cus2  
ip multicast-routing vrf mvpn-cus3  
ip multicast multipath  
frame-relay switching  
mpls label range 4112 262143  
mpls label protocol ldp  
mpls ldp logging neighbor-changes  
mpls ldp explicit-null  
mpls traffic-eng tunnels  
mpls tdp discovery directed-hello accept from 1
```

```

mpls tdp router-id Loopback0 force
platform flow ip destination
no platform flow ipv6
platform rate-limit unicast cef glean 10 10
platform qos
platform cef error action freeze

...

vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
vlan 2001-2101,3501-3700,4001,4051-4080,4093
!
!
!
interface Loopback0
 ip address 201.252.1.14 255.255.255.255
 ip pim sparse-dense-mode
!
interface Loopback1
 ip address 209.255.255.14 255.255.255.255
!
interface Loopback10
 ip vrf forwarding mvpn-cus1
 ip address 210.101.255.14 255.255.255.255
!
interface Loopback11
 ip vrf forwarding mvpn-cus1
 ip address 210.111.255.14 255.255.255.255
 ip pim sparse-dense-mode
!
interface Loopback12
 ip vrf forwarding mvpn-cus1
 ip address 210.112.255.14 255.255.255.255

...

!
interface GigabitEthernet3/3
 mtu 9216
 ip vrf forwarding mvpn-cus3
 ip address 172.10.14.1 255.255.255.0
 ip pim sparse-dense-mode
!

...

!
interface GigabitEthernet3/19
 ip vrf forwarding mvpn-cus2
 ip address 192.16.4.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp static-group 229.1.1.1
 ip igmp static-group 229.1.1.2
 ip igmp static-group 229.1.1.4
!
interface GigabitEthernet3/20
 ip vrf forwarding mvpn-cus1
 ip address 192.16.1.1 255.255.255.0
 ip pim sparse-dense-mode
!
...

```

mVPN Configuration with Default and Data MDTs

The following sample configuration includes three mVRFs that have been configured for both default and data MDTs. Only the configuration that is relevant to the mVPN configuration is shown.

```
...
!
ip vrf v1
  rd 1:1
  route-target export 1:1
  route-target import 1:1
  mdt default 226.1.1.1
  mdt data 226.1.1.128 0.0.0.7 threshold 1
!
ip vrf v2
  rd 2:2
  route-target export 2:2
  route-target import 2:2
  mdt default 226.2.2.1
  mdt data 226.2.2.128 0.0.0.7
!
ip vrf v3
  rd 3:3
  route-target export 3:3
  route-target import 3:3
  mdt default 226.3.3.1
  mdt data 226.3.3.128 0.0.0.7
!
ip vrf v4
  rd 155.255.255.1:4
  route-target export 155.255.255.1:4
  route-target import 155.255.255.1:4
  mdt default 226.4.4.1
  mdt data 226.4.4.128 0.0.0.7
!
ip multicast-routing
ip multicast-routing vrf v1
ip multicast-routing vrf v2
ip multicast-routing vrf v3
ip multicast-routing vrf v4
mpls label protocol ldp
mpls ldp logging neighbor-changes
mpls tdp router-id Loopback1
platform ip multicast replication-mode ingress
platform ip multicast bidir gm-scan-interval 10
no platform flow ip
no platform flow ipv6
platform cef error action freeze
!
...

vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
!
interface Loopback1
  ip address 155.255.255.1 255.255.255.255
  ip pim sparse-mode
!
interface Loopback11
  ip vrf forwarding v1
  ip address 155.255.255.11 255.255.255.255
```

```

    ip pim sparse-dense-mode
    !
interface Loopback22
    ip vrf forwarding v2
    ip address 155.255.255.22 255.255.255.255
    ip pim sparse-mode
    !
interface Loopback33
    ip vrf forwarding v3
    ip address 155.255.255.33 255.255.255.255
    ip pim sparse-mode
    !
interface Loopback44
    ip vrf forwarding v4
    ip address 155.255.4.4 255.255.255.255
    ip pim sparse-mode
    !
interface Loopback111
    ip vrf forwarding v1
    ip address 1.1.1.1 255.255.255.252
    ip pim sparse-dense-mode
    ip ospf network point-to-point
    !
interface GigabitEthernet1/1
    description Gi1/1 - 155.50.1.155 255.255.255.0 - peer dut50 - mpls
    mtu 9216
    ip address 155.50.1.155 255.255.255.0
    ip pim sparse-mode
    mpls ip
    !
interface GigabitEthernet1/2
    ip vrf forwarding v1
    ip address 155.1.2.254 255.255.255.0
    ip pim sparse-mode
    !
interface GigabitEthernet1/3
    description Gi1/3 - 185.155.1.155/24 - vrf v1 stub peer 185.Gi1/3
    ip vrf forwarding v1
    ip address 185.155.1.155 255.255.255.0
    ip pim sparse-mode
    !
...

    !
interface GigabitEthernet1/48
    ip vrf forwarding v1
    ip address 157.155.1.155 255.255.255.0
    ip pim bsr-border
    ip pim sparse-dense-mode
    !
interface GigabitEthernet6/1
    no ip address
    shutdown
    !
interface GigabitEthernet6/2
    ip address 9.1.10.155 255.255.255.0
    media-type rj45
    !
interface Vlan1
    no ip address
    shutdown
    !
router ospf 11 vrf v1

```

```
router-id 155.255.255.11
log-adjacency-changes
redistribute connected subnets tag 155
redistribute bgp 1 subnets tag 155
network 1.1.1.0 0.0.0.3 area 155
network 155.255.255.11 0.0.0.0 area 155
network 155.0.0.0 0.255.255.255 area 155
network 157.155.1.0 0.0.0.255 area 0
!
router ospf 22 vrf v2
router-id 155.255.255.22
log-adjacency-changes
network 155.255.255.22 0.0.0.0 area 155
network 155.0.0.0 0.255.255.255 area 155
network 157.155.1.0 0.0.0.255 area 0
!
router ospf 33 vrf v3
router-id 155.255.255.33
log-adjacency-changes
network 155.255.255.33 0.0.0.0 area 155
!
router ospf 1
log-adjacency-changes
network 155.50.1.0 0.0.0.255 area 0
network 155.255.255.1 0.0.0.0 area 155
!
router bgp 1
bgp router-id 155.255.255.1
no bgp default ipv4-unicast
bgp log-neighbor-changes
neighbor 175.255.255.1 remote-as 1
neighbor 175.255.255.1 update-source Loopback1
neighbor 185.255.255.1 remote-as 1
neighbor 185.255.255.1 update-source Loopback1
!
address-family vpnv4
neighbor 175.255.255.1 activate
neighbor 175.255.255.1 send-community extended
neighbor 185.255.255.1 activate
neighbor 185.255.255.1 send-community extended
exit-address-family
!
address-family ipv4 vrf v4
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 vrf v3
redistribute ospf 33
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 vrf v2
redistribute ospf 22
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 vrf v1
redistribute ospf 11
no auto-summary
no synchronization
exit-address-family
```

```

!
ip classless
ip route 9.255.254.1 255.255.255.255 9.1.10.254
no ip http server
ip pim bidir-enable
ip pim rp-address 50.255.2.2 MCAST.MVPN.MDT.v2 override bidir
ip pim rp-address 50.255.3.3 MCAST.MVPN.MDT.v3 override bidir
ip pim rp-address 50.255.1.1 MCAST.MVPN.MDT.v1 override bidir
ip pim vrf v1 spt-threshold infinity
ip pim vrf v1 send-rp-announce Loopback11 scope 16 group-list MCAST.GROUP.BIDIR bidir
ip pim vrf v1 send-rp-discovery Loopback11 scope 16
ip pim vrf v1 bsr-candidate Loopback111 0
ip msdp vrf v1 peer 185.255.255.11 connect-source Loopback11
ip msdp vrf v1 cache-sa-state
!
!
ip access-list standard MCAST.ANYCAST.CE
 permit 2.2.2.2
ip access-list standard MCAST.ANYCAST.PE
 permit 1.1.1.1
ip access-list standard MCAST.BOUNDARY.VRF.v1
 deny 226.192.1.1
 permit any
ip access-list standard MCAST.GROUP.BIDIR
 permit 226.192.0.0 0.0.255.255
ip access-list standard MCAST.GROUP.SPARSE
 permit 226.193.0.0 0.0.255.255
ip access-list standard MCAST.MVPN.BOUNDARY.DATA.MDT
 deny 226.1.1.128
 permit any
ip access-list standard MCAST.MVPN.MDT.v1
 permit 226.1.0.0 0.0.255.255
ip access-list standard MCAST.MVPN.MDT.v2
 permit 226.2.0.0 0.0.255.255
ip access-list standard MCAST.MVPN.MDT.v3
 permit 226.3.0.0 0.0.255.255
ip access-list standard MCAST.MVPN.RP.v4
 permit 227.0.0.0 0.255.255.255
!
access-list 1 permit 226.1.1.1
access-list 2 deny 226.1.1.1
access-list 2 permit any
...

```

Verifying the mVPN with L3VPN over mGRE Configuration

Use the following examples to verify that the configuration is working properly:

Endpoint Creation

You can verify the tunnel endpoint that has been created:

```
Router# show tunnel endpoints tunnel 0
```

```
Tunnel0 running in multi-GRE/IP mode
```

```
Endpoint transport 209.165.200.251 Refcount 3 Base 0x2AE93F0 Create Time 00:00:42
overlay 209.165.200.254 Refcount 2 Parent 0x2AE93F0 Create Time 00:00:42
```

Adjacency

You can verify that the corresponding adjacency has been created:

```
Router# show adjacency tunnel 0
```

Protocol	Interface	Address
IP	Tunnel0	209.165.200.251(4)
TAG	Tunnel0	209.165.200.251(3)

Profile Health

You can use **show l3vpn encapsulation *profile-name*** command to get information on the basic state of the application. The output of this command provides you details on the references to the underlying tunnel.

```
Router# show l3vpn encapsulation ip tunnel encap
```

```
Profile: tunnel encap
transport ipv4 source Auto: Loopback0
protocol gre
Tunnel Tunnel0 Created [OK]
Tunnel Linestate [OK]
Tunnel Transport Source (Auto) Loopback0 [OK]
```

Configuration Sequence for mVPN with L3VPN over mGRE

This example shows the configuration sequence for mVPN with L3VPN over mGRE:

```
vrf definition Customer A
  rd 100:110
  route-target export 100:1000
  route-target import 100:1000
  !
  address-family ipv4
  exit-address-family
  !
  address-family ipv6
  exit-address-family
  !
  !
  ipv6 unicast-routing
  !
  l3vpn encapsulation ip sample_profile_name
  transport ipv4 source loopback 0
  !
  !
  interface Loopback0
    ip address 209.165.200.252 255.255.255.224
    ip router isis
  !
  interface gigabitethernet 1/1
    vrf forwarding Customer A
    ip address 209.165.200.253 255.255.255.224
    ipv6 address 3FFE:1001::/64 eui-64
  !
  router bgp 100
    bgp log-neighbor-changes
    neighbor 209.165.200.254 remote-as 100
    neighbor 209.165.200.254 update-source Loopback0
  !
  address-family ipv4
    no synchronization
    redistribute connected
    neighbor 209.165.200.254 activate
```

```

no auto-summary
exit-address-family
!
address-family vpnv4
neighbor 209.165.200.254 activate
neighbor 209.165.200.254 send-community both
neighbor 209.165.200.254 route-map SELECT_UPDATE_FOR_L3VPN in
exit-address-family
!
address-family vpnv6
neighbor 209.165.200.254 activate
neighbor 209.165.200.254 send-community both
neighbor 209.165.200.254 route-map SELECT_UPDATE_FOR_L3VPN in
exit-address-family
!
address-family ipv4 vrf Customer A
no synchronization
redistribute connected
exit-address-family
!
address-family ipv6 vrf Customer A
redistribute connected
no synchronization
exit-address-family
!
!
route-map SELECT_UPDATE_FOR_L3VPN permit 10
set ip next-hop encapsulate sample_profile_name
set ipv6 next-hop encapsulate sample_profile_name

```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

13

IPv6 Multicast Support

- [Prerequisites for IPv6 Multicast, page 13-1](#)
- [Restrictions for IPv6 Multicast, page 13-1](#)
- [Information About IPv6 Multicast Support, page 13-2](#)
- [How to Configure IPv6 Multicast Support, page 13-4](#)
- [Verifying the IPv6 Multicast Layer 3 Configuration, page 13-4](#)

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IPv6 Multicast

None.

Restrictions for IPv6 Multicast

- The PFC and DFCs provide hardware support for the following:
 - Completely switched IPv6 multicast flows
 - IPv6 PIM-Sparse Mode (PIM-SM) (S,G) and (*,G) forwarding
 - Multicast RPF check for IPv6 PIM-SM (S,G) traffic using the NetFlow table
 - Rate limiting of IPv6 PIM-SM (S,G) traffic that fails the multicast RPF check
 - Static IPv6 multicast routes
 - SSM Mapping for IPv6 (PIM-SSM)
 - IPv6 multicast forwarding information base (MFIB) using the NetFlow table
 - IPv6 distributed MFIB (dMFIB) using the NetFlow table
 - Link-local and link-global IPv6 multicast scopes
 - Egress multicast replication with the **ipv6 mfib hardware-switching** command
 - Ingress interface statistics for multicast routes (egress interface statistics not available)
 - RPR and RPR+ redundancy mode (see [Chapter 9, “Route Processor Redundancy \(RPR\)”](#))

- Ingress and egress PFC QoS (see [Chapter 24, “PFC QoS Overview”](#))
 - Input and output Cisco access-control lists (ACLs)
- The PFC and DFCs do not provide hardware support for the following:
 - Partially switched IPv6 multicast flows
 - Multicast RPF check for PIM-SM (*,G) traffic
 - Multicast helper maps
 - Site-local multicast scopes
 - Manually configured IPv6 over IPv4 tunnels
 - IPv6 multicast 6to4 tunnels
 - IPv6 multicast automatic tunnels
 - IPv6 over GRE tunnels
 - IPv6-in-IPv6 PIM register tunnels
 - IPv6 multicast basic ISATAP tunnels
 - ISATAP tunnels with embedded 6to4 tunnels

Information About IPv6 Multicast Support

- [Hardware-Supported IPv6 Layer 3 Multicast Features, page 13-2](#)
- [Partially Hardware-Supported IPv6 Layer 3 Multicast Features, page 13-3](#)
- [Software-Supported IPv6 Layer 3 Multicast Features, page 13-3](#)
- [Unsupported IPv6 Layer 3 Multicast Features, page 13-3](#)

Hardware-Supported IPv6 Layer 3 Multicast Features

- Control plane policing (CoPP)
- Egress forced replication mode
- Egress replication local
- Egress replication mode
- HW assisted SPT switchover
- Input ACL logging
- Input and output ACL filtering
- IPv6 Multicast over P2P IPv4 GRE/IP-in-IP tunnel (6over4)
- Load-Balancing of multicast packets on port-channels
- Multicast Layer 3 forwarding on routed ports
- Multicast Layer 3 forwarding on subinterfaces
- Multicast Layer 3 forwarding on SVI
- Multicast load-splitting across parallel links
- Netflow accounting

- Non-RPF protection
- PIM register decapsulation over IPv6
- PIM register encapsulation over IPv6
- PIM-SM (S,G) and (*,G) forwarding
- PIM-SSM
- QoS ingress mode marking
- QoS ingress mode policing
- Rate limiters
- Scope checking
- Statistics

Partially Hardware-Supported IPv6 Layer 3 Multicast Features

- Egress replication mode and QoS marking

Software-Supported IPv6 Layer 3 Multicast Features

- SSM mapping
- MET sharing
- MLDv1/v2

Unsupported IPv6 Layer 3 Multicast Features

- BIDIR PIM over P2P GRE tunnel
- Destination IP NAT multicast
- IPv4 multicast over P2P IPv6 GRE tunnel (4over6)
- IPv6 multicast over multipoint IPv4 GRE tunnel (6over4 mGRE)
- IPv6 multicast over multipoint IPv6 GRE tunnel
- IPv6 multicast over P2P IPv6 GRE tunnel
- IPv6 multicast over P2P IPv6 GRE tunnel with tunnel endpoints in VRF
- IPv6 multicast over P2P IPv6 VRF GRE tunnel
- MTR multicast: TOS based lookup
- Multicast VPN for IPv6 extranet support
- Multicast VPN for IPv6 intranet support
- Multicast VRF-lite
- MVPN over P2P IPv6 GRE tunnel
- PIM-BIDIR
- PIM-DM (S,G) forwarding

- Source IP NAT multicast
- Egress replication mode and QoS policing
- QoS ingress and egress: shaping support
- MIB support
- Multicast boundary
- Multicast helper map
- Output ACL logging
- PGM router assist
- PGM router assist in VRF
- QoS Marking for multicast bridged frames undergoing routing

How to Configure IPv6 Multicast Support

The PFC and DFCs provide hardware support for IPv6 multicast traffic. Use these publications to configure IPv6 multicast in Cisco IOS Release 15.4SY:

- The *Cisco IOS IPv6 Configuration Library*, “Implementing IPv6 Multicast”:
<http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-15-2mt-book.html>
- The *Cisco IOS IPv6 Command Reference*:
http://www.cisco.com/en/US/docs/ios/ipv6/command/reference/ipv6_book.html

Verifying the IPv6 Multicast Layer 3 Configuration

- [Verifying MFIB Clients, page 13-5](#)
- [Displaying the Switching Capability, page 13-5](#)
- [Verifying the \(S,G\) Forwarding Capability, page 13-5](#)
- [Verifying the \(*,G\) Forwarding Capability, page 13-5](#)
- [Verifying the Subnet Entry Support Status, page 13-5](#)
- [Verifying the Current Replication Mode, page 13-5](#)
- [Displaying the Replication Mode Auto-Detection Status, page 13-5](#)
- [Displaying the Replication Mode Capabilities, page 13-6](#)
- [Displaying Subnet Entries, page 13-6](#)
- [Displaying the IPv6 Multicast Summary, page 13-6](#)
- [Displaying the NetFlow Hardware Forwarding Count, page 13-7](#)
- [Displaying the FIB Hardware Bridging and Drop Counts, page 13-7](#)
- [Displaying the Shared and Well-Known Hardware Adjacency Counters, page 13-8](#)

Verifying MFIB Clients

This example shows the complete output of the **show ipv6 mrib client** command:

```
Router# show ipv6 mrib client
```

Displaying the Switching Capability

This example displays the complete output of the **show platform software ipv6-multicast capability** command:

```
Router# show platform software ipv6-multicast capability
```

Verifying the (S,G) Forwarding Capability

This example shows how to verify the (S,G) forwarding:

```
Router# show platform software ipv6-multicast capability | include (S,G)
(S,G) forwarding for IPv6 supported using Netflow
```

Verifying the (*,G) Forwarding Capability

This example shows how to verify the (*,G) forwarding:

```
Router# show platform software ipv6-multicast capability | include (\*,G)
(*,G) bridging for IPv6 is supported using FIB
```

Verifying the Subnet Entry Support Status

This example shows how to verify the subnet entry support status:

```
Router# show platform software ipv6-multicast capability | include entries
Directly-connected entries for IPv6 is supported using ACL-TCAM.
```

Verifying the Current Replication Mode

This example shows how to verify the current replication mode:

```
Router# show platform software ipv6-multicast capability | include Current
Current System HW Replication Mode : Ingress
```

**Note**

Enter the **no ipv6 mrib hardware-switching replication-mode ingress** command to enable replication mode auto-detection.

Displaying the Replication Mode Auto-Detection Status

This example shows how to display the replication mode auto-detection status:

```
Router# show platform software ipv6-multicast capability | include detection
Auto-detection of Replication Mode : ON
```

Displaying the Replication Mode Capabilities

This example shows how to display the replication mode capabilities of the installed modules:

```
Router# show platform software ipv6-multicast capability | begin ^Slot
Slot Replication-Capability Replication-Mode
  1 Ingress Ingress
  2 Egress Ingress
  6 Egress Ingress
  8 Ingress Ingress
```

Displaying Subnet Entries

This example shows how to display subnet entries:

```
Router# show platform software ipv6-multicast connected
IPv6 Multicast Subnet entries
Flags : H - Installed in ACL-TCAM
        X - Not installed in ACL-TCAM due to
            label-full exception
Interface: Vlan20 [ H ]
           S:20::1 G:FF00::
Interface: Vlan10 [ H ]
           S:10::1 G:FF00::
```



Note In this example, there are subnet entries for VLAN 10 and VLAN 20.

Displaying the IPv6 Multicast Summary

This example shows how to display the IPv6 multicast summary:

```
Router# show platform software ipv6-multicast summary
IPv6 Multicast Netflow SC summary on Slot[1]:
Shortcut Type          Shortcut count
-----+-----
(S, G)                 100
(*, G)                 0
IPv6 Multicast FIB SC summary on Slot[1]:
Shortcut Type          Shortcut count
-----+-----
(*, G/128)             10
(*, G/m)               47

IPv6 Multicast Netflow SC summary on Slot[6]:
Shortcut Type          Shortcut count
-----+-----
(S, G)                 100
(*, G)                 0
IPv6 Multicast FIB SC summary on Slot[6]:
Shortcut Type          Shortcut count
-----+-----
(*, G/128)             10
(*, G/m)               47
```

Displaying the NetFlow Hardware Forwarding Count

This example shows how to display the NetFlow hardware forwarding count:

```
Router# show platform software ipv6-multicast summary
IPv6 Multicast Netflow SC summary on Slot[1]:
Shortcut Type          Shortcut count
-----+-----
(S, G)                100
(*, G)                 0

<...Output deleted...>

IPv6 Multicast Netflow SC summary on Slot[6]:
Shortcut Type          Shortcut count
-----+-----
(S, G)                100
(*, G)                 0

<...Output truncated...>
```



Note

The NetFlow (*, G) count is always zero because PIM-SM (*,G) forwarding is supported in software on the RP.

Displaying the FIB Hardware Bridging and Drop Counts

This example shows how to display the FIB hardware bridging and drop hardware counts:

```
Router# show platform software ipv6-multicast summary | begin FIB
IPv6 Multicast FIB SC summary on Slot[1]:
Shortcut Type          Shortcut count
-----+-----
(*, G/128)            10
(*, G/m)               47

<...Output deleted...>

IPv6 Multicast FIB SC summary on Slot[6]:
Shortcut Type          Shortcut count
-----+-----
(*, G/128)            10
(*, G/m)               47
```



Note

- The (*, G/128) value is a hardware bridge entry count.
- The (*, G/m) value is a hardware bridge/drop entry count.

Displaying the Shared and Well-Known Hardware Adjacency Counters

The **show platform software ipv6-multicast shared-adjacencies** command displays the shared and well-known hardware adjacency counters used for IPv6 multicast by entries in FIB and ACL-TCAM.

Router# **show platform software ipv6-multicast shared-adjacencies**

---- SLOT [1] ----

Shared IPv6 Mcast Adjacencies	Index	Packets	Bytes
Subnet bridge adjacency	0x7F802	0	0
Control bridge adjacency	0x7	0	0
StarG_M bridge adjacency	0x8	0	0
S_G bridge adjacency	0x9	0	0
Default drop adjacency	0xA	0	0
StarG (spt == INF) adjacency	0xB	0	0
StarG (spt != INF) adjacency	0xC	0	0

---- SLOT [6] ----

Shared IPv6 Mcast Adjacencies	Index	Packets	Bytes
Subnet bridge adjacency	0x7F802	0	0
Control bridge adjacency	0x7	0	0
StarG_M bridge adjacency	0x8	0	0
S_G bridge adjacency	0x9	0	0
Default drop adjacency	0xA	28237	3146058
StarG (spt == INF) adjacency	0xB	0	0
StarG (spt != INF) adjacency	0xC	0	0



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

14

IPv6 MLD Snooping

- [Prerequisites for MLD Snooping, page 14-1](#)
- [Restrictions for MLD Snooping, page 14-2](#)
- [Information About MLD Snooping, page 14-3](#)
- [Default MLD Snooping Configuration, page 14-9](#)
- [How to Configure MLD Snooping, page 14-9](#)
- [Verifying the MLD Snooping Configuration, page 14-14](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- To constrain IPv4 multicast traffic, see [Chapter 7, “IGMP Snooping for IPv4 Multicast Traffic.”](#)
- All PFC modes support Multicast Listener Discovery (MLD) version 1 (MLDv1) and MLD version 2 (MLDv2).

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for MLD Snooping

None.

Restrictions for MLD Snooping

- [General MLD Snooping Restrictions, page 14-2](#)
- [MLD Snooping Querier Restrictions, page 14-2](#)

General MLD Snooping Restrictions

- All PFC modes support MLD version 1 (MLDv1) and MLD version 2 (MLDv2).
- MLD is derived from Internet Group Management Protocol version 3 (IGMPv3). MLD protocol operations and state transitions, host and router behavior, query and report message processing, message forwarding rules, and timer operations are exactly same as IGMPv3. See draft-vida-ml-d-.02.txt for detailed information on MLD protocol.
- MLD protocol messages are Internet Control Message Protocol version 6 (ICMPv6) messages.
- MLD message formats are almost identical to IGMPv3 messages.
- IPv6 multicast for Cisco IOS software uses MLD version 2. This version of MLD is fully backward-compatible with MLD version 1 (described in RFC 2710). Hosts that support only MLD version 1 interoperate with a router running MLD version 2. Mixed LANs with both MLD version 1 and MLD version 2 hosts are supported.
- MLD snooping supports private VLANs. Private VLANs do not impose any restrictions on MLD snooping.
- MLD snooping constrains traffic in MAC multicast groups 0100.5e00.0001 to 0100.5eff.ffff.
- MLD snooping does not constrain Layer 2 multicasts generated by routing protocols.

MLD Snooping Querier Restrictions

- Configure an IPv6 address on the VLAN interface (see [Chapter 35, “Layer 3 Interfaces”](#)). When enabled, the MLD snooping querier uses the IPv6 address as the query source address.
- If there is no IPv6 address configured on the VLAN interface, the MLD snooping querier does not start. The MLD snooping querier disables itself if the IPv6 address is cleared. When enabled, the MLD snooping querier restarts if you configure an IPv6 address.
- When enabled, the MLD snooping querier does not start if it detects MLD traffic from an IPv6 multicast router.
- When enabled, the MLD snooping querier starts after 60 seconds with no MLD traffic detected from an IPv6 multicast router.
- When enabled, the MLD snooping querier disables itself if it detects MLD traffic from an IPv6 multicast router.
- QoS does not support MLD packets when MLD snooping is enabled.
- You can enable the MLD snooping querier on all the switches in the VLAN that support it. One switch is elected as the querier.

- To configure redundant MLD snooping queriers, complete the tasks in the “[Enabling the MLD Snooping Querier](#)” section on page 14-10 on more than one switch in the VLAN.

When multiple MLD snooping queriers are enabled in a VLAN, the querier with the lowest IP address in the VLAN is elected as the active MLD snooping querier.

An MLD snooping querier election occurs if the active MLD snooping querier goes down or if there is an IP address change on any of the queriers.

**Note**

To avoid unnecessary active querier time outs, configure the **ipv6 mld snooping last-member-query-interval** command with the same value on all queriers in a VLAN.

Information About MLD Snooping

- [MLD Snooping Overview](#), page 14-3
- [MLD Messages](#), page 14-4
- [Source-Based Filtering](#), page 14-4
- [Explicit Host Tracking](#), page 14-4
- [MLD Snooping Proxy Reporting](#), page 14-5
- [Joining an IPv6 Multicast Group](#), page 14-5
- [Leaving a Multicast Group](#), page 14-7
- [Information about the MLD Snooping Querier](#), page 14-8

MLD Snooping Overview

MLD snooping allows the switch to examine MLD packets and make forwarding decisions based on their content.

You can configure the switch to use MLD snooping in subnets that receive MLD queries from either MLD or the MLD snooping querier. MLD snooping constrains IPv6 multicast traffic at Layer 2 by configuring Layer 2 LAN ports dynamically to forward IPv6 multicast traffic only to those ports that want to receive it.

MLD, which runs at Layer 3 on a multicast router, generates Layer 3 MLD queries in subnets where the multicast traffic needs to be routed. For information about MLD, see this publication:

<http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-multicast.html>

You can configure the MLD snooping querier on the switch to support MLD snooping in subnets that do not have any multicast router interfaces. For more information about the MLD snooping querier, see the “[Enabling the MLD Snooping Querier](#)” section on page 14-10.

MLD (on a multicast router) or, locally, the MLD snooping querier, sends out periodic general MLD queries that the switch forwards through all ports in the VLAN, and to which hosts respond. MLD snooping monitors the Layer 3 MLD traffic.

**Note**

If a multicast group has only sources and no receivers in a VLAN, MLD snooping constrains the multicast traffic to only the multicast router ports.

MLD Messages

- Multicast listener queries
 - General query—Sent by a multicast router to learn which multicast addresses have listeners.
 - Multicast address specific query—Sent by a multicast router to learn if a particular multicast address has any listeners.
 - Multicast address and source specific query—Sent by a multicast router to learn if any of the sources from the specified list for the particular multicast address has any listeners.
- Multicast listener reports
 - Current state record (solicited)—Sent by a host in response to a query to specify the INCLUDE or EXCLUDE mode for every multicast group in which the host is interested.
 - Filter mode change record (unsolicited)—Sent by a host to change the INCLUDE or EXCLUDE mode of one or more multicast groups.
 - Source list change record (unsolicited)—Sent by a host to change information about multicast sources.

Source-Based Filtering

MLD uses source-based filtering, which enables hosts and routers to specify which multicast sources should be allowed or blocked for a specific multicast group. Source-based filtering either allows or blocks traffic based on the following information in MLD messages:

- Source lists
- INCLUDE or EXCLUDE mode

Because the Layer 2 table is (MAC-group, VLAN) based, with MLD hosts it is preferable to have only a single multicast source per MAC-group.

**Note**

Source-based filtering is not supported in hardware. The states are maintained only in software and used for explicit host tracking and statistics collection.

Explicit Host Tracking

MLD supports explicit tracking of membership information on any port. The explicit-tracking database is used for fast-leave processing, proxy reporting, and statistics collection. When explicit tracking is enabled on a VLAN, the MLD snooping software processes the MLD report it receives from a host and builds an explicit-tracking database that contains the following information:

- The port connected to the host
- The channels reported by the host
- The filter mode for each group reported by the host
- The list of sources for each group reported by the hosts
- The router filter mode of each group
- For each group, the list of hosts requesting the source

**Note**

- Disabling explicit host tracking disables fast-leave processing and proxy reporting.
- When explicit tracking is enabled and the switch is in report-suppression mode, the multicast router might not be able to track all the hosts accessed through a VLAN interface.

MLD Snooping Proxy Reporting

Because MLD does not have report suppression, all the hosts send their complete multicast group membership information to the multicast router in response to queries. The switch snoops these responses, updates the database and forwards the reports to the multicast router. To prevent the multicast router from becoming overloaded with reports, MLD snooping does proxy reporting.

Proxy reporting forwards only the first report for a multicast group to the router and suppresses all other reports for the same multicast group.

Proxy reporting processes solicited and unsolicited reports. Proxy reporting is enabled and cannot be disabled.

**Note**

Disabling explicit host tracking disables fast-leave processing and proxy reporting.

Joining an IPv6 Multicast Group


Hosts join IPv6 multicast groups either by sending an unsolicited MLD report or by sending an MLD report in response to a general query from an IPv6 multicast router (the switch forwards general queries from IPv6 multicast routers to all ports in a VLAN). The switch snoops these reports.

In response to a snooped MLD report, the switch creates an entry in its Layer 2 forwarding table for the VLAN on which the report was received. When other hosts that are interested in this multicast traffic send MLD reports, the switch snoops their reports and adds them to the existing Layer 2 forwarding table entry. The switch creates only one entry per VLAN in the Layer 2 forwarding table for each multicast group for which it snoops an MLD report.

MLD snooping suppresses all but one of the host reports per multicast group and forwards this one report to the IPv6 multicast router.

The switch forwards multicast traffic for the multicast group specified in the report to the interfaces where reports were received (see [Figure 14-1](#)).

Layer 2 multicast groups learned through MLD snooping are dynamic. However, you can statically configure Layer 2 multicast groups using the **mac address-table static** command. When you specify group membership for a multicast group address statically, the static setting supersedes any MLD snooping learning. Multicast group membership lists can consist of both static and MLD snooping-learned settings.

Figure 14-1 Initial MLD Listener Report


Multicast router A sends an MLD general query to the switch, which forwards the query to ports 2 through 5 (all members of the same VLAN). Host 1 wants to join an IPv6 multicast group and multicasts an MLD report to the group with the equivalent MAC destination address of 0x0100.5E01.0203. When the switch snoops the MLD report multicast by Host 1, the switch uses the information in the MLD report to create a forwarding-table entry.

Table 14-1 MLD Snooping Forwarding Table

Destination MAC Address	Type of Packet	Ports
0100.5exx.xxxx	MLD	0
0100.5e01.0203	!MLD	1, 2

The switch hardware can distinguish MLD information packets from other packets for the multicast group. The first entry in the table indicates that only MLD packets should be sent to the CPU, which prevents the switch from becoming overloaded with multicast frames. The second entry indicates that frames addressed to the 0x0100.5E01.0203 multicast MAC address that are not MLD packets (!MLD) should be sent to the multicast router and to the host that has joined the group.

If another host (for example, Host 4) sends an unsolicited MLD report for the same group (Figure 14-2), the switch snoops that message and adds the port number of Host 4 to the forwarding table as shown in Table 14-2. Because the forwarding table directs MLD messages only to the switch, the message is not flooded to other ports. Any known multicast traffic is forwarded to the group and not to the switch.

Figure 14-2 Second Host Joining a Multicast Group**Table 14-2** Updated MLD Snooping Forwarding Table

Destination MAC Address	Type of Packet	Ports
0100.5exx.xxxx	MLD	0
0100.5e01.0203	!MLD	1, 2, 5

Leaving a Multicast Group

- [Normal Leave Processing, page 14-7](#)
- [Fast-Leave Processing, page 14-8](#)

Normal Leave Processing

Interested hosts must continue to respond to the periodic MLD general queries. As long as at least one host in the VLAN responds to the periodic MLD general queries, the multicast router continues forwarding the multicast traffic to the VLAN. When hosts want to leave a multicast group, they can either ignore the periodic MLD general queries (called a “silent leave”), or they can send an MLD filter mode change record.

When MLD snooping receives a filter mode change record from a host that configures the EXCLUDE mode for a group, MLD snooping sends out a MAC-addressed general query to determine if any other hosts connected to that interface are interested in traffic for the specified multicast group.

If MLD snooping does not receive an MLD report in response to the general query, MLD snooping assumes that no other hosts connected to the interface are interested in receiving traffic for the specified multicast group, and MLD snooping removes the interface from its Layer 2 forwarding table entry for the specified multicast group.

If the filter mode change record was from the only remaining interface with hosts interested in the group, and MLD snooping does not receive an MLD report in response to the general query, MLD snooping removes the group entry and relays the MLD filter mode change record to the multicast router. If the multicast router receives no reports from a VLAN, the multicast router removes the group for the VLAN from its MLD cache.

The interval for which the switch waits before updating the table entry is called the “last member query interval.” To configure the interval, enter the **ipv6 mld snooping last-member-query-interval** *interval* command.

Fast-Leave Processing

Fast-leave processing is enabled by default. To disable fast-leave processing, turn off explicit-host tracking.

Fast-leave processing is implemented by maintaining source-group based membership information in software while also allocating LTL indexes on a MAC GDA basis.

When fast-leave processing is enabled, hosts send BLOCK_OLD_SOURCES{src-list} messages for a specific group when they no longer want to receive traffic from that source. When the switch receives such a message from a host, it parses the list of sources for that host for the given group. If this source list is exactly the same as the source list received in the leave message, the switch removes the host from the LTL index and stops forwarding this multicast group traffic to this host.

If the source lists do not match, the switch does not remove the host from the LTL index until the host is no longer interested in receiving traffic from any source.



Note

Disabling explicit host tracking disables fast-leave processing and proxy reporting.

Information about the MLD Snooping Querier

Use the MLD snooping querier to support MLD snooping in a VLAN where PIM and MLD are not configured because the multicast traffic does not need to be routed.

In a network where IP multicast routing is configured, the IP multicast router acts as the MLD querier. If the IP-multicast traffic in a VLAN only needs to be Layer 2 switched, an IP-multicast router is not required, but without an IP-multicast router on the VLAN, you must configure another switch as the MLD querier so that it can send queries.

When enabled, the MLD snooping querier sends out periodic MLD queries that trigger MLD report messages from the switch that wants to receive IP multicast traffic. MLD snooping listens to these MLD reports to establish appropriate forwarding.

You can enable the MLD snooping querier on all the switches in the VLAN, but for each VLAN that is connected to switches that use MLD to report interest in IP multicast traffic, you must configure at least one switch as the MLD snooping querier.

You can configure a switch to generate MLD queries on a VLAN regardless of whether or not IP multicast routing is enabled.

Default MLD Snooping Configuration

- MLD snooping querier: disabled
- MLD snooping: enabled
- Multicast routers: none configured
- MLD report suppression: enabled
- MLD snooping router learning method: learned automatically through PIM or MLD packets
- Fast-Leave Processing: enabled
- MLD Explicit Host Tracking: enabled

How to Configure MLD Snooping

- [Enabling the MLD Snooping Querier, page 14-10](#)
- [Configuring the MLD Snooping Query Interval, page 14-10](#)
- [Enabling MLD Snooping, page 14-11](#)
- [Configuring a Static Connection to a Multicast Receiver, page 14-12](#)
- [Configuring a Multicast Router Port Statically, page 14-12](#)
- [Enabling Fast-Leave Processing, page 14-12](#)
- [Enabling SSM Safe Reporting, page 14-13](#)
- [Configuring Explicit Host Tracking, page 14-13](#)
- [Configuring Report Suppression, page 14-14](#)

**Note**

- To use MLD snooping, configure a Layer 3 interface in the subnet for IPv6 multicast routing or enable the MLD snooping querier in the subnet (see the [“Enabling the MLD Snooping Querier” section on page 14-10](#)).
- Except for the global enable command, all MLD snooping commands are supported only on VLAN interfaces.

Enabling the MLD Snooping Querier

Use the MLD snooping querier to support MLD snooping in a VLAN where PIM and MLD are not configured because the multicast traffic does not need to be routed. To enable the MLD snooping querier in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 address <i>prefix/prefix_length</i>	Configures the IPv6 address and subnet.
Step 3	Router(config-vlan-config)# ipv6 mld snooping querier	Enables the MLD snooping querier.
Step 4	Router(config-vlan-config)# end	Exits configuration mode.

This example shows how to enable the MLD snooping querier on VLAN 200 and verify the configuration:

```
Router# vlan configuration 200
Router(config-vlan-config)# ipv6 address 2001:0DB8:0:1::/64 eui-64
Router(config-vlan-config)# ipv6 mld snooping querier
Router(config-vlan-config)# end
Router# show ipv6 mld interface vlan 200 | include querier
      MLD snooping fast-leave is enabled and querier is enabled
```

Configuring the MLD Snooping Query Interval

You can configure the interval for which the switch waits after sending a group-specific query to determine if hosts are still interested in a specific multicast group.



Note

When both MLD snooping fast-leave processing and the MLD snooping query interval are configured, fast-leave processing takes precedence.

To configure the interval for the MLD snooping queries sent by the switch, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping last-member-query-interval <i>interval</i>	Configures the interval for the IGMP queries sent by the switch. Default is 1 second. Valid range is 1000 to 9990 milliseconds.

This example shows how to configure the MLD snooping query interval:

```
Router(config-vlan-config)# ipv6 mld snooping last-member-query-interval 1000
Router(config-vlan-config)# exit
Router# show ipv6 mld interface vlan 200 | include last
      MLD snooping last member query response interval is 1000 ms
```

Enabling MLD Snooping

- [Enabling MLD Snooping Globally, page 14-11](#)
- [Enabling MLD Snooping in a VLAN, page 14-11](#)

Enabling MLD Snooping Globally

To enable MLD snooping globally, perform this task:

	Command	Purpose
Step 1	Router(config)# ipv6 mld snooping	Enables MLD snooping.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable MLD snooping globally and verify the configuration:

```
Router(config)# ipv6 mld snooping
Router(config)# end
Router# show ipv6 mld interface vlan 200 | include globally
    MLD snooping is globally enabled
Router#
```

Enabling MLD Snooping in a VLAN

To enable MLD snooping in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping	Enables MLD snooping.
Step 3	Router(config-vlan-config)# end	Exits configuration mode.

This example shows how to enable MLD snooping on VLAN 25 and verify the configuration:

```
Router# vlan configuration 25
Router(config-vlan-config)# ipv6 mld snooping
Router(config-vlan-config)# end
Router# show ipv6 mld interface vlan 25 | include snooping
    MLD snooping is globally enabled
    MLD snooping is enabled on this interface
    MLD snooping fast-leave is enabled and querier is enabled
    MLD snooping explicit-tracking is enabled
    MLD snooping last member query response interval is 1000 ms
    MLD snooping report-suppression is disabled
Router#
```

Configuring a Static Connection to a Multicast Receiver

To configure a static connection to a multicast receiver, perform this task:

	Command	Purpose
Step 1	Router(config)# mac address-table static <i>mac_addr</i> vlan <i>vlan_id</i> interface <i>type slot/port</i> [disable-snooping]	Configures a static connection to a multicast receiver.
Step 2	Router(config)# end	Exits configuration mode.

When you configure a static connection, enter the **disable-snooping** keyword to prevent multicast traffic addressed to the statically configured multicast MAC address from also being sent to other ports in the same VLAN.

This example shows how to configure a static connection to a multicast receiver:

```
Router(config)# mac address-table static 0050.3e8d.6400 vlan 12 interface gigabitethernet 5/7
```

Configuring a Multicast Router Port Statically

To configure a static connection to a multicast router, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping mrouter interface <i>type slot/port</i>	Configures a static connection to a multicast router.
Step 3	Router(config-vlan-config)# end	Exits configuration mode.

The interface to the router must be in the VLAN where you are entering the command, the interface must be administratively up, and the line protocol must be up.

This example shows how to configure a static connection to a multicast router:

```
Router(config-vlan-config)# ipv6 mld snooping mrouter interface gigabitethernet 5/6
Router(config-vlan-config)#
```

Enabling Fast-Leave Processing

To enable fast-leave processing in a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping fast-leave	Enables fast-leave processing in the VLAN.

This example shows how to enable fast-leave processing on the VLAN 200 interface and verify the configuration:

```
Router# vlan configuration 200
Router(config-vlan-config)# ipv6 mld snooping fast-leave
Configuring fast leave on vlan 200
Router(config-vlan-config)# end
Router# show ipv6 mld interface vlan 200 | include fast-leave
      MLD snooping fast-leave is enabled and querier is enabled
Router#
```

Enabling SSM Safe Reporting

To enable source-specific multicast (SSM) safe reporting, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping ssm-safe-reporting	Enables SSM safe reporting.

This example shows how to SSM safe reporting:

```
Router(config)# vlan configuration 10
Router(config-vlan-config)# ipv6 mld snooping ssm-safe-reporting
```

Configuring Explicit Host Tracking



Note

Disabling explicit host tracking disables fast-leave processing and proxy reporting.

To enable explicit host tracking on a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping explicit-tracking	Enables explicit host tracking.

This example shows how to enable explicit host tracking:

```
Router(config)# vlan configuration 25
Router(config-vlan-config)# ipv6 mld snooping explicit-tracking
Router(config-vlan-config)# end
Router# show ipv6 mld snooping explicit-tracking vlan 25
Source/Group          Interface    Reporter    Filter_mode
-----
10.1.1.1/226.2.2.2    V125:1/2    16.27.2.3    INCLUDE
10.2.2.2/226.2.2.2    V125:1/2    16.27.2.3    INCLUDE
```

Configuring Report Suppression

To enable report suppression on a VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# vlan configuration <i>vlan_ID</i>	Selects a VLAN.
Step 2	Router(config-vlan-config)# ipv6 mld snooping report-suppression	Enables report suppression.

This example shows how to enable explicit host tracking:

```
Router(config)# vlan configuration 25
Router(config-vlan-config)# ipv6 mld snooping report-suppression
Router(config-vlan-config)# end
Router# Router# show ipv6 mld interface vlan 25 | include report-suppression
MLD snooping report-suppression is enabled
```

Verifying the MLD Snooping Configuration

- [Displaying Multicast Router Interfaces, page 14-14](#)
- [Displaying MAC Address Multicast Entries, page 14-14](#)
- [Displaying MLD Snooping Information for a VLAN Interface, page 14-15](#)

Displaying Multicast Router Interfaces

When you enable IGMP snooping, the switch automatically learns to which interface the multicast routers are connected.

To display multicast router interfaces, perform this task:

Command	Purpose
Router# show ipv6 mld snooping mrouter <i>vlan_ID</i>	Displays multicast router interfaces.

This example shows how to display the multicast router interfaces in VLAN 1:

```
Router# show ipv6 mld snooping mrouter vlan 1
vlan          ports
-----+-----
1             Gi1/1,Gi2/1,Gi3/48,Router
Router#
```

Displaying MAC Address Multicast Entries

To display MAC address multicast entries for a VLAN, perform this task:

Command	Purpose
Router# show mac address-table multicast <i>vlan_ID</i> [<i>count</i>]	Displays MAC address multicast entries for a VLAN.

This example shows how to display MAC address multicast entries for VLAN 1:

```
Router# show mac address-table multicast vlan 1
vlan    mac address      type    qos    ports
-----+-----+-----+-----+-----
  1    0100.5e02.0203    static  --    Gi1/1,Gi2/1,Gi3/48,Router
  1    0100.5e00.0127    static  --    Gi1/1,Gi2/1,Gi3/48,Router
  1    0100.5e00.0128    static  --    Gi1/1,Gi2/1,Gi3/48,Router
  1    0100.5e00.0001    static  --    Gi1/1,Gi2/1,Gi3/48,Router,Switch
Router#
```

This example shows how to display a total count of MAC address entries for a VLAN:

```
Router# show mac address-table multicast 1 count

Multicast MAC Entries for vlan 1:    4
Router#
```

Displaying MLD Snooping Information for a VLAN Interface

To display MLD snooping information for a VLAN interface, perform this task:

Command	Purpose
Router# show ipv6 mld snooping [{ explicit-tracking <i>vlan_ID</i> } { mrouter [<i>vlan</i> <i>vlan_ID</i>]} { report-suppression <i>vlan</i> <i>vlan_ID</i> } { statistics <i>vlan</i> <i>vlan_ID</i> }]	Displays MLD snooping information on a VLAN interface.

This example shows how to display explicit tracking information on VLAN 25:

```
Router# show ipv6 mld snooping explicit-tracking vlan 25
Source/Group      Interface    Reporter    Filter_mode
-----+-----+-----+-----
10.1.1.1/226.2.2.2    Vl25:1/2    16.27.2.3    INCLUDE
10.2.2.2/226.2.2.2    Vl25:1/2    16.27.2.3    INCLUDE
```

This example shows how to display the multicast router interfaces in VLAN 1:

```
Router# show ipv6 mld snooping mrouter vlan 1
vlan    ports
-----+-----
  1      Gi1/1,Gi2/1,Gi3/48,Router
```

This example shows IGMP snooping statistics information for VLAN 25:

```
Router# show ipv6 mld snooping statistics interface vlan 25

Snooping statictics for Vlan25
#channels:2
#hosts    :1

Source/Group      Interface    Reporter    Uptime      Last-Join    Last-Leave
-----+-----+-----+-----+-----+-----
10.1.1.1/226.2.2.2    Gi1/2:Vl25    16.27.2.3    00:01:47    00:00:50    -
10.2.2.2/226.2.2.2    Gi1/2:Vl25    16.27.2.3    00:01:47    00:00:50    -
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

15

NetFlow Hardware Support

- [Prerequisites for NetFlow Hardware Support, page 15-1](#)
- [Restrictions for NetFlow Hardware Support, page 15-1](#)
- [Information About NetFlow Hardware Support, page 15-2](#)
- [Default Settings for NetFlow Hardware Support, page 15-2](#)
- [How to Configure NetFlow Hardware Support, page 15-2](#)
- [Verifying the NetFlow Table Aging Configuration, page 15-4](#)

**Note**

In Cisco IOS Release 15.4SY, the Flexible NetFlow feature provides statistics collection and data export. See these publications:

<http://www.cisco.com/en/US/docs/ios-xml/ios/fnetflow/command/fnf-cr-book.html>

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for NetFlow Hardware Support

None.

Restrictions for NetFlow Hardware Support

- Cisco IOS Release 15.4SY and later releases do not support NetFlow version 7 or NetFlow version 8. Flexible NetFlow has limited support for NetFlow version 5.
- No statistics are available for flows that are forwarded when the NetFlow table is full.
- If the NetFlow table utilization exceeds the recommended utilization levels, there is an increased probability that there will be insufficient room to store statistics. [Table 15-1](#) lists the recommended maximum utilization levels.

Table 15-1 NetFlow Table Utilization

PFC Mode	Effective NetFlow Table Utilization		Total NetFlow Table Capacity	
PFC4XL	506,184 ingress entries	506,184 egress entries	524,288 (512k) ingress entries	524,288 (512k) egress entries
PFC4	515,032 ingress+egress entries		524,288 (512k) ingress+egress entries	

- If a flow is destined to an address in the PBR range or is sourced from an address in the PBR range, the input and output interface will be the default route (if configured) or be null.

Information About NetFlow Hardware Support

The NetFlow table on the PFC and any DFCs captures data for flows forwarded in hardware. These are some of the features that use the NetFlow table:

- Flexible NetFlow
- Network address translation (NAT)
- QoS microflow policing
- Reflexive ACLS
- WCCP

To limit NetFlow CPU usage, you can configure aging timers to identify stale flows that can be deleted from the table. NetFlow deletes the stale entries to clear table space for new entries.

Default Settings for NetFlow Hardware Support

- Inactive Flow Aging: enabled (300 seconds)
- Fast Aging: disabled
- Active Flow Aging: enabled (1920 seconds)

How to Configure NetFlow Hardware Support

- [Configuring Inactive Flow Aging, page 15-3](#)
- [Configuring Fast Aging, page 15-3](#)
- [Configuring Active Flow Aging, page 15-4](#)

**Note**

- NetFlow table aging keeps the NetFlow table size below the recommended utilization. If the number of NetFlow table entries exceeds the recommended utilization (see the [“Restrictions for NetFlow Hardware Support” section on page 15-1](#)), only adjacency statistics might be available for some flows.
- Network events (for example, routing changes or a link state change) can also purge NetFlow table entries.

Configuring Inactive Flow Aging

To configure inactive flow aging, perform this task:

Command	Purpose
Router(config)# flow platform cache timeout inactive <i>seconds</i>	Configures the aging time for NetFlow table entries that have been inactive longer than the configured time value. • Default: enabled; value: 300 seconds. • Range for the <i>seconds</i> value: 32–512.

This example displays how to configure the aging time for NetFlow table entries that have been inactive longer than the configured time value:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# flow platform cache timeout inactive 300
```

Configuring Fast Aging

To configure fast aging, perform this task:

Command	Purpose
Router(config)# flow platform cache timeout fast [[time <i>seconds</i>] [threshold <i>packets</i>]]	Configures an aging time for NetFlow table entries that have been inactive longer than the configured time value and that have forwarded fewer packets than the configured threshold value. • Default: disabled. • Default if time <i>seconds</i> not entered: 32 seconds; range for the <i>seconds</i> value: 60–4092. • Default if threshold <i>packets</i> not entered: 100 packets range for the <i>packets</i> value: 1–4000.

**Note**

If you enable fast aging, initially set the value to 128 seconds. If the size of the NetFlow table continues to grow over the recommended utilization, decrease the setting until the table size stays below the recommended utilization. If the table continues to grow over the recommended utilization, decrease the inactive NetFlow table aging time.

This example displays how to configure the NetFlow table aging time:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# flow platform cache timeout fast time 32 threshold 100
```

Configuring Active Flow Aging

To configure active flow aging, perform this task:

Command	Purpose
Router(config)# flow platform cache timeout active <i>seconds</i>	Configures the aging time for NetFlow table entries regardless of packet activity, which can prevent counter wraparound and inaccurate statistics. - Default: enabled; value: 1920 seconds. - Range for the <i>seconds</i> value: 60–4092.

This example displays how to configure active flow aging:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# flow platform cache timeout active 1920
```

Verifying the NetFlow Table Aging Configuration

To display the NetFlow table aging configuration, perform this task:

Command	Purpose
Router# show platform flow aging	Displays the NetFlow table aging configuration.

This example shows how to display the NetFlow table aging-time configuration:

```
Router# show platform flow aging
Aging scheme  Enabled  Timeout  Packet threshold
-----+-----+-----+-----
          Fast      No       32         100
          Inactive  Yes      300        N/A
          Active   Yes      1920       N/A
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

16

System Event Archive (SEA)

- [Information About the System Event Archive, page 16-1](#)
- [How to Display the SEA Logging System, page 16-2](#)
- [How to Copy the SEA To Another Device, page 16-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Information About the System Event Archive

The primary method of discovering the cause of system failure is system messages. When system messages do not provide the information needed to determine the cause of a failure, you can enable debug traces and attempt to recreate the failure. However, there are several situations in which neither of the above methods provides an optimum solution:

- Reviewing a large number of system messages can be an inefficient method of determining the cause of a failure.
- Debug trace is usually not configured by default.
- You cannot recreate the failure while using debug trace.
- Using debug trace is not an option if the switch on which the failure has occurred is part of your critical network.

The SEA enables each of the CPUs on a switch to report events to the management processor using an out-of-band interface. Each event is logged in nonvolatile memory with a time stamp. You can retrieve the event log by accessing the bootflash on the device, or you can copy the log to another location such as a removable storage device.

The SEA maintains two files in the bootdisk, using up to 32 MB. These files contain the most recent messages recorded to the log:

- sea_log.dat—Applications store the most recent system events in this file.
- sea_console.dat—The most recent console messages are stored in this file.

These files are for system use and should not be removed.

How to Display the SEA Logging System

To display the SEA logging system, perform this task:

Command	Purpose
Router# show logging system [disk size]	Displays the contents of the SEA. (Optional) Use the keyword disk to display the location where the SEA is stored. Use the keyword size to display the current size of the SEA.
Router# clear logging system	Removes the event records stored in the SEA.

The following example shows how to display the SEA:

```
Router# show logging system
SEQ: MM/DD/YY HH:MM:SS MOD/SUB: SEV, COMP, MESSAGE
=====
1: 01/24/07 15:38:40 6/-1 : MAJ, GOLD, syndiagSyncPinnacle failed in slot 6
2: 01/24/07 15:38:40 6/-1 : MAJ, GOLD, queryHyperionSynched[6]: Hyperion out of sync in sw_mode 1
3: 01/24/07 15:38:40 6/-1 : MAJ, GOLD, queryHyperionSynched[6]: Hyperion out of sync in sw_mode 1
4: 01/24/07 15:38:40 6/-1 : MAJ, GOLD, queryHyperionSynched[6]: Hyperion out of sync in sw_mode 1
5: 01/24/07 15:38:40 6/-1 : MAJ, GOLD, queryHyperionSynched[6]: Hyperion out of sync in sw_mode 1
6: 01/24/07 15:38:40 6/-1 : MAJ, GOLD, queryHyperionSynched[6]: Hyperion out of sync in sw_mode 1
7: 01/24/07 15:38:39 6/-1 : MAJ, GOLD, queryHyperionSynched[6]: Hyperion out of sync in sw_mode 1
```

The following example shows how to display the SEA logging system disk:

```
Router# show logging system disk
SEA log disk: bootdisk:
```

The following example shows how to display the current size of the SEA:

```
Router# show logging system size
SEA log size: 33554432 bytes
```

The following example shows how to clear the SEA:

```
Router# clear logging system
Clear logging system operation will take a while.
Do you want to continue? [no]: yes
Router#
```


How to Copy the SEA To Another Device

To copy the SEA to another device, such as a removeable memory device, perform this task:

Command	Purpose
Router# copy logging system file_system	Copies the contents of the SEA to the specified destination file system or process.

The valid values for *file_system* are:

- **bootflash:**
- **disk0:**
- **disk1:**
- **ftp:**
- **http:**
- **https:**
- **rcp:**
- **slavebootflash:**
- **slavedisk0:**
- **slavedisk1:**
- **slavebootdisk:**
- **slavebootflash:**
- **bootdisk:**
- **bootflash:**
- **tftp:**

The following example shows how to copy the SEA to the disk0 file system:

```
Router# copy logging system disk0:
Destination filename [sea_log.dat]?
Copy in
progress...CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
33554432 bytes copied in 112.540 secs (298156 bytes/sec)
```

The following example shows how to copy the SEA using the remote file copy function (rcp):

```
Router# copy logging system rcp:
Address or name of remote host []? 192.0.2.1
Destination username [Router]? username1
Destination filename [sea_log.dat]? /auto/tftpboot-users/username1/sea_log.dat
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

33554432 bytes copied in 48.172 secs (696555 bytes/sec)

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

17

Backplane Traffic Monitoring

- [Prerequisites for Backplane Traffic Monitoring, page 17-1](#)
- [Restrictions for Backplane Traffic Monitoring, page 17-1](#)
- [Information About Traffic Monitoring, page 17-2](#)
- [Default Settings for Backplane Traffic Monitoring, page 17-2](#)
- [How to Configure Backplane Traffic Monitoring, page 17-2](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Backplane Traffic Monitoring

None.

Restrictions for Backplane Traffic Monitoring

The syslog message buffer is limited in size. To reduce false alarms and the number of syslog messages, use the following guidelines:

- Traffic can occur in bursts. If a small number of bursts occur in a monitoring interval, it does not represent a capacity overload issue for the system; the hardware buffers are able to absorb the effects and not cause packet drops. For an example, if you set a monitoring interval to 10 seconds and the threshold to 80 percent, there are a total of 10 traffic utilization readings. Assume only 2 of the readings reached 90 percent and the other 8 readings are 20 percent. If the peak threshold of 90 percent is used to compare with the threshold, an unwanted warning syslog message is generated.

It is better to use the average 34 percent of the 10 readings to compare with the threshold and not generate warning messages in this case. If the peak value comparison is really needed, you can set the interval to 1 second. Setting the interval to 1 second compares the reading directly with the threshold.

- The number of syslog messages that generate syslog messages are from below the threshold and above the threshold.

Information About Traffic Monitoring

Backplane Traffic Monitoring can monitor the backplane and fabric-channel traffic utilization at a configured interval and threshold.

Traffic monitoring allows the switch to monitor the backplane and fabric channel traffic utilization at a configured interval and threshold. Syslog messages are generated when the traffic utilization is above or below the configured threshold. The following examples show several types of syslog messages:

- 00:08:03: %TRAFFIC_UTIL-SP-4-MONITOR_BACKPLANE_REACH_THR: Backplane traffic utilization is 26%, reached threshold(20%) within 10 second interval
- 00:08:13: %TRAFFIC_UTIL-SP-4-MONITOR_BACKPLANE_BELOW_THR: Backplane traffic utilization is 18%, below threshold(20%) within 10 second interval
- 00:08:44: %TRAFFIC_UTIL-SP-4-MONITOR_FABRIC_IG_REACH_THR: Module 1, Channel 0 ingress traffic utilization is 5%, reached threshold(3%) within 30 second interval
- 00:08:44: %TRAFFIC_UTIL-SP-4-MONITOR_FABRIC_EG_REACH_THR: Module 1, Channel 0 egress traffic utilization is 5%, reached threshold(3%) within 30 second interval
- 00:09:14: %TRAFFIC_UTIL-SP-4-MONITOR_FABRIC_IG_BELOW_THR: Module 1, Channel 0 ingress traffic utilization is 1%, below threshold(3%) within 30 second interval
- 00:09:14: %TRAFFIC_UTIL-SP-4-MONITOR_FABRIC_EG_BELOW_THR: Module 1, Channel 0 egress traffic utilization is 1%, below threshold(3%) within 30 second interval

Default Settings for Backplane Traffic Monitoring

- The default threshold is 80 percent.
- Traffic monitor is off by default.

How to Configure Backplane Traffic Monitoring

To configure the Backplane Traffic Monitoring feature, perform one or more of the following tasks:

Command	Purpose
Router(config)# monitor traffic-util backplane interval <i>interval</i> threshold <i>percentage</i>	Configures the backplane utilization traffic monitoring.
Router(config)# monitor traffic-util fabric module { <i>mod-num</i> all } { channel { <i>0</i> <i>1</i> both }} { direction { egress ingress both }} [interval <i>interval</i> threshold <i>percentage</i>]	Configures the fabric channel utilization traffic monitoring.

Command	Purpose
Router(config)# monitor traffic-util fabric logging interval <i>second</i>	Configures the fabric channel utilization traffic monitor SYSLOG interval when the traffic utilization is in the crossed state.
Router(config)# monitor traffic-util backplane logging interval <i>second</i>	Configures the traffic monitor backplane SYSLOG interval when the traffic utilization is in the crossed state.
Router# show catalyst6000 traffic-meter	Displays the percentage of the backplane (shared bus) utilization and traffic monitor status information.

When configuring a range of interfaces, you can enter the *mod-num* as a list or a range. Separate each entry with a comma and each range with a hyphen (-). For example, 1,3,5-9,12.

The following example shows how to enable backplane traffic utilization monitoring:

```
Router(config)# monitor traffic-util backplane logging interval 50 threshold 100
```

The following example shows how to disable backplane traffic utilization monitoring:

```
Router(config)# no monitor traffic-util backplane
```

The following example shows how to specify the fabric channel traffic utilization monitor interval and threshold for a fabric channel on a specific module:

```
Router(config)# monitor traffic-util fabric module 8 channel both direction both interval 50 threshold 60
```

The following example shows how to specify the fabric channel traffic utilization monitor threshold for a specific fabric channel and for egress traffic only:

```
Router(config)# monitor traffic-util fabric module 6 channel 0 direction egress interval 100 threshold 90
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

18

Local SPAN, RSPAN, and ERSPAN

- [Prerequisites for Local SPAN, RSPAN, and ERSPAN, page 18-1](#)
- [Restrictions for Local SPAN, RSPAN, and ERSPAN, page 18-1](#)
- [Information About Local SPAN, RSPAN, and ERSPAN, page 18-7](#)
- [Default Settings for Local SPAN, RSPAN, and ERSPAN, page 18-12](#)
- [How to Configure Local SPAN, RSPAN, and ERSPAN, page 18-12](#)
- [Verifying the SPAN Configuration, page 18-31](#)
- [Configuration Examples for SPAN, page 18-31](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Local SPAN, RSPAN, and ERSPAN

None.

Restrictions for Local SPAN, RSPAN, and ERSPAN

- [Feature Incompatibilities, page 18-2](#)
- [Local SPAN, RSPAN, and ERSPAN Session Limits, page 18-3](#)
- [Local SPAN, RSPAN, and ERSPAN Interface Limits, page 18-3](#)
- [General Restrictions for Local SPAN, RSPAN, and ERSPAN, page 18-3](#)

- [Restrictions for VSPAN, page 18-4](#)
- [Restrictions for RSPAN, page 18-5](#)
- [Restrictions for ERSPAN, page 18-5](#)
- [Restrictions for Distributed Egress SPAN Mode, page 18-7](#)

Feature Incompatibilities

- [Egress SPAN](#) is not supported in egress multicast mode. (CSCsa95965)
- Unknown unicast flood blocking (UUFB) ports cannot be RSPAN or local SPAN egress-only destinations. (CSCsj27695)
- A port-channel interface (an EtherChannel) can be a SPAN source, but you cannot configure active member ports of an EtherChannel as SPAN source ports. Inactive member ports of an EtherChannel can be configured as SPAN sources but they are put into the suspended state and carry no traffic.
- These features are incompatible with SPAN destinations:
 - Private VLANs
 - IEEE 802.1X port-based authentication
 - Port security
 - Spanning Tree Protocol (STP) and related features (PortFast, PortFast BPDU filtering, BPDU Guard, UplinkFast, BackboneFast, EtherChannel Guard, Root Guard, Loop Guard)
 - VLAN trunk protocol (VTP)
 - Dynamic trunking protocol (DTP)
 - IEEE 802.1Q tunneling
- The `dot1ad uni` command (see [Chapter 3, “Ethernet Virtual Connections \(EVCs\)”](#)) imposes these SPAN restrictions:
 - The `dot1ad uni` command does not restrict [egress SPAN](#) configuration.
 - The `dot1ad uni` command does not restrict RSPAN [ingress SPAN](#) source port configuration, but source VLAN filtering does not get applied to traffic from RSPAN ingress SPAN source ports that are configured with the `dot1ad uni` command.
 - Ports that are configured with the `dot1ad uni` command cannot be used as local SPAN or ERSPAN ingress SPAN source ports.



Note

- SPAN destinations can participate in IEEE 802.3Z flow control.
- IP multicast switching using egress packet replication is not compatible with SPAN. In some cases, egress replication can result in multicast packets not being sent to the SPAN destination port. If you are using SPAN and your switching modules are capable of egress replication, enter the `platform ip multicast replication-mode ingress` command to force ingress replication.

Local SPAN, RSPAN, and ERSPAN Session Limits

	Local and Source Sessions		Destination Sessions	
	Local SPAN, RSPAN Source, ERSPAN Source Ingress or Egress or Both	Local SPAN Egress-Only	RSPAN	ERSPAN
Total Sessions				
80	2	14	64	23

Local SPAN, RSPAN, and ERSPAN Interface Limits

	In Each Local SPAN Session	In Each RSPAN Source Session	In Each ERSPAN Source Session	In Each RSPAN Destination Session	In Each ERSPAN Destination Session
Egress or “both” sources	128	128	128	—	—
Ingress sources	128	128	128	—	—
RSPAN and ERSPAN destination session sources	—	—	—	1 RSPAN VLAN	1 IP address
Destinations per session	64	1 RSPAN VLAN	1 IP address	64	64

General Restrictions for Local SPAN, RSPAN, and ERSPAN

- A SPAN destination that is copying traffic from a single egress SPAN source port sends only egress traffic to the network analyzer. If you configure more than one egress SPAN source port, the traffic that is sent to the network analyzer also includes these types of ingress traffic that were received from the egress SPAN source ports:

- Any unicast traffic that is flooded on the VLAN
- Broadcast and multicast traffic

This situation occurs because an egress SPAN source port receives these types of traffic from the VLAN but then recognizes itself as the source of the traffic and drops it instead of sending it back to the source from which it was received. Before the traffic is dropped, SPAN copies the traffic and sends it to the SPAN destination. (CSCds22021)

- Entering additional **monitor session** commands does not clear previously configured SPAN parameters. You must enter the **no monitor session** command to clear configured SPAN parameters.
- Connect a network analyzer to the SPAN destinations.
- Within a SPAN session, all of the SPAN destinations receive all of the traffic from all of the SPAN sources, except when source-VLAN filtering is configured on the SPAN source.
- You can configure destination trunk VLAN filtering to select which traffic is transmitted from the SPAN destination.
- You can configure both Layer 2 LAN ports (LAN ports configured with the **switchport** command) and Layer 3 LAN ports (LAN ports not configured with the **switchport** command) as sources or destinations.

- You cannot mix individual source ports and source VLANs within a single session.
- If you specify multiple ingress source ports, the ports can belong to different VLANs.
- Within a session, you cannot configure both VLANs as SPAN sources and do source VLAN filtering. You can configure VLANs as SPAN sources or you can do source VLAN filtering of traffic from source ports and EtherChannels, but not both in the same session.
- You cannot configure source VLAN filtering for internal VLANs.
- When enabled, local SPAN, RSPAN, and ERSPAN use any previously entered configuration.
- When you specify sources and do not specify a traffic direction (ingress, egress, or both), “both” is used by default.
- SPAN copies Layer 2 Ethernet frames, but SPAN does not copy source trunk port 802.1Q tags. You can configure destinations as trunks to send locally tagged traffic to the traffic analyzer.

**Note**

A destination configured as a trunk tags traffic from a Layer 3 LAN source with the [internal VLAN](#) used by the Layer 3 LAN source.

- Local SPAN sessions, RSPAN source sessions, and ERSPAN source sessions do not copy locally sourced RSPAN VLAN traffic from source trunk ports that carry RSPAN VLANs.
- Local SPAN sessions, RSPAN source sessions, and ERSPAN source sessions do not copy locally sourced ERSPAN GRE-encapsulated traffic from source ports.
- A port or EtherChannel can be a SPAN destination for only one SPAN session. SPAN sessions cannot share destinations.
- SPAN destinations cannot be SPAN sources.
- Destinations never participate in any spanning tree instance. Local SPAN includes BPDUs in the monitored traffic, so any BPDUs seen on the destination are from the source. RSPAN does not support BPDU monitoring.
- All packets forwarded through the switch for transmission from a port that is configured as an egress SPAN source are copied to the SPAN destination, including packets that do not exit the switch through the egress port because STP has put the egress port into the blocking state, or on an egress trunk port because STP has put the VLAN into the blocking state on the trunk port.
- Local-Tx span with fex host port as span destination will not work due to hardware limitation.

Restrictions for VSPAN

**Note**

Local SPAN, RSPAN, and ERSPAN all support [VSPAN](#).

- VSPAN sessions do not support source VLAN filtering.
- For VSPAN sessions with both ingress and egress configured, two packets are forwarded from the destination to the analyzer if the packets get switched on the same VLAN (one as ingress traffic from the ingress port and one as egress traffic from the egress port).

- VSPAN only monitors traffic that leaves or enters Layer 2 ports in the VLAN.
 - If you configure a VLAN as an ingress source and traffic gets routed into the monitored VLAN, the routed traffic is not monitored because it never appears as ingress traffic entering a Layer 2 port in the VLAN.
 - If you configure a VLAN as an egress source and traffic gets routed out of the monitored VLAN, the routed traffic is not monitored because it never appears as egress traffic leaving a Layer 2 port in the VLAN.

Restrictions for RSPAN

- All participating switches must be connected by Layer 2 trunks.
- Any network device that supports RSPAN VLANs can be an RSPAN intermediate device.
- Networks impose no limit on the number of RSPAN VLANs that the networks carry.
- Intermediate network devices might impose limits on the number of RSPAN VLANs that they can support.
- You must configure the RSPAN VLANs in all source, intermediate, and destination network devices. If enabled, the VLAN Trunking Protocol (VTP) can propagate configuration of VLANs numbered 1 through 1024 as RSPAN VLANs. You must manually configure VLANs numbered higher than 1024 as RSPAN VLANs on all source, intermediate, and destination network devices.
- 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.
- RSPAN VLANs can be used only for RSPAN traffic.
- Do not configure a VLAN used to carry management traffic as an RSPAN VLAN.
- Do not assign access ports to RSPAN VLANs. RSPAN puts access ports in an RSPAN VLAN into the suspended state.
- Do not configure any ports in an RSPAN VLAN except trunk ports selected to carry RSPAN traffic.
- MAC address learning is disabled in the RSPAN VLAN.
- You can use output access control lists (ACLs) on the RSPAN VLAN in the RSPAN source switch to filter the traffic sent to an RSPAN destination.
- RSPAN does not support BPDU monitoring.
- Do not configure RSPAN VLANs as sources in VSPAN sessions.
- You can configure any VLAN as an RSPAN VLAN as long as all participating network devices support configuration of RSPAN VLANs and you use the same RSPAN VLAN for each RSPAN session in all participating network devices.

Restrictions for ERSPAN

- For ERSPAN packets, the “protocol type” field value in the GRE header is 0x88BE.
- The payload of a Layer 3 ERSPAN packet is a copied Layer 2 Ethernet frame, excluding any 802.1Q tags.
- ERSPAN adds a 50-byte header to each copied Layer 2 Ethernet frame and replaces the 4-byte cyclic redundancy check (CRC) trailer.

- ERSPAN supports jumbo frames that contain Layer 3 packets of up to 9,202 bytes. If the length of the copied Layer 2 Ethernet frame is greater than 9,170 (9,152-byte Layer 3 packet), ERSPAN truncates the copied Layer 2 Ethernet frame to create a 9,202-byte ERSPAN Layer 3 packet.



Note The Layer 3 IP header in truncated packets retains the nontruncated Layer 3 packet size. The length consistency check between the Layer 2 frame and the Layer 3 packet on ERSPAN destinations that are switches drops truncated ERSPAN packets unless you configure the **no platform verify ip length consistent** global configuration command on the ERSPAN destination switch.

- Regardless of any configured MTU size, ERSPAN creates Layer 3 packets that can be as long as 9,202 bytes. ERSPAN traffic might be dropped by any interface in the network that enforces an MTU size smaller than 9,202 bytes.
- With the default MTU size (1,500 bytes), if the length of the copied Layer 2 Ethernet frame is greater than 1,468 bytes (1,450-byte Layer 3 packet), the ERSPAN traffic is dropped by any interface in the network that enforces the 1,500-byte MTU size.



Note The **mtu** interface command and the **system jumbomtu** command (see the [“Configuring Jumbo Frame Support”](#) section on page 10-6) set the maximum Layer 3 packet size (default is 1,500 bytes, maximum is 9,216 bytes).

- All participating switches must be connected at Layer 3 and the network path must support the size of the ERSPAN traffic.
- ERSPAN does not support packet fragmentation. The “do not fragment” bit is set in the IP header of ERSPAN packets. ERSPAN destination sessions cannot reassemble fragmented ERSPAN packets.
- ERSPAN traffic is subject to the traffic load conditions of the network. You can set the ERSPAN packet IP precedence or DSCP value to prioritize ERSPAN traffic for QoS.
- The only supported destination for ERSPAN traffic is an ERSPAN destination session.
- All ERSPAN source sessions on a switch must use the same origin IP address, configured with the **origin ip address** command (see the [“Configuring ERSPAN Source Sessions”](#) section on page 18-26).
- All ERSPAN destination sessions on a switch must use the same IP address on the same destination interface. You enter the destination interface IP address with the **ip address** command (see the [“Configuring ERSPAN Destination Sessions”](#) section on page 18-28).
- The ERSPAN source session’s destination IP address, which must be configured on an interface on the destination switch, is the source of traffic that an ERSPAN destination session sends to the destinations. You configure the same address in both the source and destination sessions with the **ip address** command.
- The ERSPAN ID differentiates the ERSPAN traffic arriving at the same destination IP address from various different ERSPAN source sessions.

Restrictions for Distributed Egress SPAN Mode

Some switching modules have ASICs that do not support distributed egress SPAN mode for ERSPAN sources.

Enter the **show monitor session egress replication-mode | include Distributed.*Distributed.*Centralized** command to display the slot number of any switching modules that do not support distributed egress SPAN mode for ERSPAN sources.

Enter the **show ASIC-version slot slot_number** command to display the versions of the ASICs on the switching module in the slot where distributed egress SPAN mode is not supported for ERSPAN sources.

Hyperion ASIC revision levels 5.0 and higher and all versions of the Metropolis ASIC support distributed egress SPAN mode for ERSPAN sources. Switching modules with Hyperion ASIC revision levels lower than 5.0 do not support distributed egress SPAN mode for ERSPAN sources.

Information About Local SPAN, RSPAN, and ERSPAN

- [Local SPAN, RSPAN, and ERSPAN Overview, page 18-7](#)
- [Local SPAN, RSPAN, and ERSPAN Sources, page 18-11](#)
- [Local SPAN, RSPAN, and ERSPAN Destinations, page 18-12](#)

Local SPAN, RSPAN, and ERSPAN Overview

- [SPAN Operation, page 18-7](#)
- [Local SPAN Overview, page 18-7](#)
- [RSPAN Overview, page 18-8](#)
- [ERSPAN Overview, page 18-9](#)
- [Traffic Monitored at SPAN Sources, page 18-10](#)

SPAN Operation

SPAN copies traffic from one or more ports, one or more EtherChannels, or one or more VLANs, and sends the copied traffic to one or more destinations for analysis by a network analyzer such as a SwitchProbe device or other Remote Monitoring (RMON) probe. Traffic can also be sent to the processor for packet capture by the Mini Protocol Analyzer, as described in [Chapter 22, “Mini Protocol Analyzer.”](#)

SPAN does not affect the switching of traffic on sources. You must dedicate the destination for SPAN use. The SPAN-generated copies of traffic compete with user traffic for switch resources.

Local SPAN Overview


A local SPAN session is an association of source ports and source VLANs with one or more destinations. You configure a local SPAN session on a single switch. Local SPAN does not have separate source and destination sessions.

Local SPAN sessions do not copy locally sourced RSPAN VLAN traffic from source trunk ports that carry RSPAN VLANs. Local SPAN sessions do not copy locally sourced RSPAN GRE-encapsulated traffic from source ports.

Each local SPAN session can have either ports or VLANs as sources, but not both.

Local SPAN copies traffic from one or more source ports in any VLAN or from one or more VLANs to a destination for analysis (see [Figure 18-1](#)). For example, as shown in [Figure 18-1](#), all traffic on Ethernet port 5 (the source port) is copied to Ethernet port 10. A network analyzer on Ethernet port 10 receives all traffic from Ethernet port 5 without being physically attached to Ethernet port 5.

Figure 18-1 Example SPAN Configuration



RSPAN Overview

RSPAN supports source ports, source VLANs, and destinations on different switches, which provides remote monitoring of multiple switches across your network (see [Figure 18-2](#)). RSPAN uses a Layer 2 VLAN to carry SPAN traffic between switches.


RSPAN consists of an RSPAN source session, an RSPAN VLAN, and an RSPAN destination session. You separately configure RSPAN source sessions and destination sessions on different switches. To configure an RSPAN source session on one switch, you associate a set of source ports or VLANs with an RSPAN VLAN. To configure an RSPAN destination session on another switch, you associate the destinations with the RSPAN VLAN.

The traffic for each RSPAN session is carried as Layer 2 nonroutable traffic over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. All participating switches must be trunk-connected at Layer 2.

RSPAN source sessions do not copy locally sourced RSPAN VLAN traffic from source trunk ports that carry RSPAN VLANs. RSPAN source sessions do not copy locally sourced RSPAN GRE-encapsulated traffic from source ports.

Each RSPAN source session can have either ports or VLANs as sources, but not both.

The RSPAN source session copies traffic from the source ports or source VLANs and switches the traffic over the RSPAN VLAN to the RSPAN destination session. The RSPAN destination session switches the traffic to the destinations.

Figure 18-2 RSPAN Configuration

27389

ERSPAN Overview

ERSPAN supports source ports, source VLANs, and destinations on different switches, which provides remote monitoring of multiple switches across your network (see [Figure 18-3](#)). ERSPAN uses a GRE tunnel to carry traffic between switches.

ERSPAN consists of an ERSPAN source session, routable ERSPAN GRE-encapsulated traffic, and an ERSPAN destination session. You separately configure ERSPAN source sessions and destination sessions on different switches.

To configure an ERSPAN source session on one switch, you associate a set of source ports or VLANs with a destination IP address, ERSPAN ID number, and optionally with a VRF name. To configure an ERSPAN destination session on another switch, you associate the destinations with the source IP address, ERSPAN ID number, and optionally with a VRF name.

ERSPAN source sessions do not copy locally sourced RSPAN VLAN traffic from source trunk ports that carry RSPAN VLANs. ERSPAN source sessions do not copy locally sourced ERSPAN GRE-encapsulated traffic from source ports.

Each ERSPAN source session can have either ports or VLANs as sources, but not both.

The ERSPAN source session copies traffic from the source ports or source VLANs and forwards the traffic using routable GRE-encapsulated packets to the ERSPAN destination session. The ERSPAN destination session switches the traffic to the destinations.

Figure 18-3 *ERSPAN Configuration*

Traffic Monitored at SPAN Sources

- [Monitored Traffic Direction, page 18-10](#)
- [Monitored Traffic Type, page 18-11](#)
- [Duplicate Traffic, page 18-11](#)

Monitored Traffic Direction

You can configure local SPAN sessions, RSPAN source sessions, and ERSPAN source sessions to monitor the following traffic:

- Ingress traffic
 - Called ingress SPAN.
 - Copies traffic received by the sources (ingress traffic).
 - Ingress traffic is sent to the supervisor engine SPAN ASIC to be copied.
- Egress traffic
 - Called egress SPAN.
 - Copies traffic transmitted from the sources (egress traffic).
 - Distributed egress SPAN mode—On some fabric-enabled switching modules, egress traffic can be copied locally by the switching module SPAN ASIC and then sent to the SPAN destinations. See the [“Restrictions for Distributed Egress SPAN Mode” section on page 18-7](#) for information about switching modules that support distributed egress SPAN mode.
 - Centralized egress SPAN mode—On all other switching modules, egress traffic is sent to the supervisor engine SPAN ASIC to be copied and is then sent to the SPAN destinations.
- Both
 - Copies both the received traffic and the transmitted traffic (ingress and egress traffic).
 - Both ingress traffic and egress traffic is sent to the supervisor engine SPAN ASIC to be copied.

Monitored Traffic Type

By default, local SPAN and ERSPAN monitor all traffic, including multicast and bridge protocol data unit (BPDU) frames. RSPAN does not support BPDU monitoring.

Duplicate Traffic

In some configurations, SPAN sends multiple copies of the same source traffic to the destination. For example, in a configuration with a bidirectional SPAN session (both ingress and egress) for two SPAN sources, called s1 and s2, to a SPAN destination, called d1, if a packet enters the switch through s1 and is sent for egress from the switch to s2, ingress SPAN at s1 sends a copy of the packet to SPAN destination d1 and egress SPAN at s2 sends a copy of the packet to SPAN destination d1. If the packet was Layer 2 switched from s1 to s2, both SPAN packets would be the same. If the packet was Layer 3 switched from s1 to s2, the Layer 3 rewrite would alter the source and destination Layer 2 addresses, in which case the SPAN packets would be different.

Local SPAN, RSPAN, and ERSPAN Sources

- [Source Ports and EtherChannels, page 18-11](#)
- [Source VLANs, page 18-11](#)

Source Ports and EtherChannels

A source port or EtherChannel is a port or EtherChannel monitored for traffic analysis. You can configure both Layer 2 and Layer 3 ports and EtherChannels as SPAN sources. SPAN can monitor one or more source ports or EtherChannels in a single SPAN session. You can configure ports or EtherChannels in any VLAN as SPAN sources. Trunk ports or EtherChannels can be configured as sources and mixed with nontrunk sources.

**Note**

SPAN does not copy the encapsulation from trunk sources. You can configure SPAN destinations as trunks to tag the monitored traffic before it is transmitted for analysis.

Source VLANs

A source VLAN is a VLAN monitored for traffic analysis. VLAN-based SPAN (VSPAN) uses a VLAN as the SPAN source. All the ports and EtherChannels in the source VLANs become sources of SPAN traffic.

**Note**

Layer 3 VLAN interfaces on source VLANs are not sources of SPAN traffic. Traffic that enters a VLAN through a Layer 3 VLAN interface is monitored when it is transmitted from the switch through an egress port or EtherChannel that is in the source VLAN.

Local SPAN, RSPAN, and ERSPAN Destinations

A SPAN destination is a Layer 2 port, Layer 3 port, or an EtherChannel, to which local SPAN, RSPAN, or ERSPAN sends traffic for analysis. When you configure a port or EtherChannel as a SPAN destination, it is dedicated for use only by the SPAN feature.

Destination EtherChannels do not support the Port Aggregation Control Protocol (PAgP) or Link Aggregation Control Protocol (LACP) EtherChannel protocols; only the on mode is supported, with all EtherChannel protocol support disabled.

There is no requirement that the member links of a destination EtherChannel be connected to a device that supports EtherChannels. For example, you can connect the member links to separate network analyzers. See [Chapter 23, “EtherChannels,”](#) for more information about EtherChannel.

Destinations, by default, cannot receive any traffic. You can configure Layer 2 destinations to receive traffic from any attached devices.

Destinations, by default, do not transmit anything except SPAN traffic. Layer 2 destinations that you have configured to receive traffic can be configured to learn the Layer 2 address of any devices attached to the destination and transmit traffic that is addressed to the devices.

You can configure trunks as destinations, which allows trunk destinations to transmit encapsulated traffic. You can use allowed VLAN lists to configure destination trunk VLAN filtering.

Default Settings for Local SPAN, RSPAN, and ERSPAN

- Local SPAN: disabled
- RSPAN: disabled
- ERSPAN: disabled
- Default operating mode for egress SPAN sessions: centralized

How to Configure Local SPAN, RSPAN, and ERSPAN

- [Configuring a Destination as an Unconditional Trunk \(Optional\), page 18-13](#)
- [Configuring Destination Trunk VLAN Filtering \(Optional\), page 18-13](#)
- [Configuring Destination Port Permit Lists \(Optional\), page 18-15](#)
- [Configuring the Egress SPAN Mode \(Optional\), page 18-15](#)
- [Configuring Local SPAN, page 18-16](#)
- [Configuring RSPAN, page 18-20](#)
- [Configuring ERSPAN, page 18-26](#)
- [Configuring Source VLAN Filtering in Global Configuration Mode, page 18-30](#)

Configuring a Destination as an Unconditional Trunk (Optional)

To configure the destination as a trunk so that the monitored traffic is tagged as it leaves the destination, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface {type slot/port port-channel number}	Selects the interface to configure.
Step 3	Router(config-if)# switchport	Configures the interface for Layer 2 switching (required only if the interface is not already configured for Layer 2 switching).
Step 4	Router(config-if)# switchport trunk encapsulation dot1q	Configures the encapsulation, which configures the interface as an 802.1Q trunk.
Step 5	Router(config-if)# switchport mode trunk	Configures the interface to trunk unconditionally.

This example shows how to configure a port as an unconditional IEEE 802.1Q trunk:

```
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport
Router(config-if)# switchport trunk encapsulation dot1q
Router(config-if)# switchport mode trunk
```

Configuring Destination Trunk VLAN Filtering (Optional)



Note

- In addition to filtering VLANs on a trunk, you can also apply the allowed VLAN list to access ports.
- Destination trunk VLAN filtering is applied at the destination. Destination trunk VLAN filtering does not reduce the amount of traffic being sent from the SPAN sources to the SPAN destinations.

When a destination is a trunk, you can use the list of VLANs allowed on the trunk to filter the traffic transmitted from the destination. (CSCeb01318)

Destination trunk VLAN filtering removes the restriction that, within a SPAN session, all destinations receive all the traffic from all the sources. Destination trunk VLAN filtering allows you to select, on a per-VLAN basis, the traffic that is transmitted from each destination trunk to the network analyzer.

To configure destination trunk VLAN filtering on a destination trunk, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type slot/port</i>	Selects the destination trunk port to configure.
Step 3	Router(config-if)# switchport trunk allowed vlan { add except none remove } <i>vlan</i> [, <i>vlan</i> [, <i>vlan</i> [, ...]]	Configures the list of VLANs allowed on the trunk.

- The *vlan* parameter is either a single VLAN number from 1 through 4094, or a range of VLANs described by two VLAN numbers, the lesser one first, separated by a dash. Do not enter any spaces between comma-separated *vlan* parameters or in dash-specified ranges.
- All VLANs are allowed by default.
- To remove all VLANs from the allowed list, enter the **switchport trunk allowed vlan none** command.
- To add VLANs to the allowed list, enter the **switchport trunk allowed vlan add** command.
- You can modify the allowed VLAN list without removing the SPAN configuration.

This example shows the configuration of a local SPAN session that has several VLANs as sources and several trunk ports as destinations, with destination trunk VLAN filtering that filters the SPAN traffic so that each destination trunk port transmits the traffic from one VLAN:

```

interface GigabitEthernet1/1
description SPAN destination interface for VLAN 10
no ip address
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 10
switchport mode trunk
switchport nonegotiate
!
interface GigabitEthernet1/2
description SPAN destination interface for VLAN 11
no ip address
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 11
switchport mode trunk
switchport nonegotiate
!
interface GigabitEthernet1/3
description SPAN destination interface for VLAN 12
no ip address
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 12
switchport mode trunk
switchport nonegotiate
!
interface GigabitEthernet1/4
description SPAN destination interface for VLAN 13
no ip address
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 13
switchport mode trunk
switchport nonegotiate

```

```

!
monitor session 1 source vlan 10 - 13
monitor session 1 destination interface Gi1/1 - 4

```

Configuring Destination Port Permit Lists (Optional)

To prevent accidental configuration of ports as destinations, you can create a permit list of the ports that are valid for use as destinations. With a destination port permit list configured, you can only configure the ports in the permit list as destinations.

To configure a destination port permit list, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor permit-list	Enables use of the destination port permit list.
Step 3	Router(config)# monitor permit-list destination interface <i>type slot/port[-port] [, type slot/port - port]</i>	Configures a destination port permit list or adds to an existing destination port permit list.

This example shows how to configure a destination port permit list that includes Gigabit Ethernet ports 5/1 through 5/4 and 6/1:

```

Router# configure terminal
Router(config)# monitor permit-list
Router(config)# monitor permit-list destination interface gigabitethernet 5/1-4,
gigabitethernet 6/1

```

This example shows how to verify the configuration:

```

Router(config)# do show monitor permit-list
SPAN Permit-list      :Admin Enabled
Permit-list ports     :Gi5/1-4,Gi6/1

```

Configuring the Egress SPAN Mode (Optional)

Centralized egress SPAN mode is the default. See the [“Restrictions for Distributed Egress SPAN Mode” section on page 18-7](#) for information about switching modules that support distributed egress SPAN mode.

To configure the egress SPAN mode, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session egress replication-mode distributed	Enables distributed egress SPAN mode. Note Enter the no monitor session egress replication-mode distributed command to enable centralized egress SPAN mode.
Step 3	Router(config)# end	Exits configuration mode.

This example shows how to enable distributed egress SPAN mode:

```
Router# configure terminal
Router(config)# monitor session egress replication-mode distributed
Router(config)# end
```

This example shows how to disable distributed egress SPAN mode:

```
Router# configure terminal
Router(config)# monitor session egress replication-mode centralized
Router(config)# end
```

This example shows how to display the configured egress SPAN mode:

```
Router# show monitor session egress replication-mode | include Configured
Configured mode : Centralized
```

Configuring Local SPAN

- [Configuring Local SPAN \(SPAN Configuration Mode\), page 18-16](#)
- [Configuring Local SPAN \(Global Configuration Mode\), page 18-18](#)

Configuring Local SPAN (SPAN Configuration Mode)



Note

To tag the monitored traffic as it leaves a destination, you must configure the destination to trunk unconditionally before you configure it as a destination (see the [“Configuring a Destination as an Unconditional Trunk \(Optional\)”](#) section on page 18-13).

To configure a local SPAN session in SPAN configuration mode, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>local_SPAN_session_number</i> type [local local-tx]	Configures a local SPAN session number and enters local SPAN session configuration mode. Note • Enter the local keyword to configure ingress or egress or both SPAN sessions. • Enter the local-tx keyword to configure egress-only SPAN sessions.
Step 3	Router(config-mon-local)# description <i>session_description</i>	(Optional) Describes the local SPAN session.

	Command	Purpose
Step 4	Router(config-mon-local)# source { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i> } [rx tx both]	Associates the local SPAN session number with the source ports or VLANs, and selects the traffic direction to be monitored. Note - When you enter the local-tx keyword, the rx and both keywords are not available and the tx keyword is required. - To make best use of the available SPAN sessions, it is always preferable to configure local-tx sessions instead of local sessions with the tx keyword.
Step 5	Router(config-mon-local)# filter <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i>	(Optional) Configures source VLAN filtering when the local SPAN source is a trunk port.
Step 6	Router(config-mon-local)# destination { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> } [ingress [learning]]	Associates the local SPAN session number with the destinations.
Step 7	Router(config-mon-local)# no shutdown	Activates the local SPAN session. Note The no shutdown command and shutdown commands are not supported for local-tx egress-only SPAN sessions.
Step 8	Router(config-mon-local)# end	Exits configuration mode.

- session_description* can be up to 240 characters and cannot contain special characters. The description can contain spaces.



Note You can enter 240 characters after the **description** command.

- local_span_session_number* can range from 1 to 80.
- single_interface* is as follows:
 - interface** *type slot/port*; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - interface** **port-channel** *number*



Note Destination port channel interfaces must be configured with the **channel-group group_num mode on** command and the **no channel-protocol** command. See the “[How to Configure EtherChannels](#)” section on page 23-7.

- interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- interface_range* is **interface** *type slot/first_port - last_port*.
- mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...

- *single_vlan* is the ID number of a single VLAN.
- *vlan_list* is *single_vlan* , *single_vlan* , *single_vlan* ...
- *vlan_range* is *first_vlan_ID* - *last_vlan_ID*.
- *mixed_vlan_list* is, in any order, *single_vlan* , *vlan_range* , ...
- Enter the **ingress** keyword to configure destinations to receive traffic from attached devices.
- Enter the **learning** keyword to enable MAC address learning from the destinations, which allows the switch to transmit traffic that is addressed to devices attached to the destinations.

When configuring destinations with the **ingress** and **learning** keywords, note the following:

- Configure the destinations for Layer 2 switching. See the [“How to Configure LAN Interfaces for Layer 2 Switching” section on page 21-5](#).
- If the destination is a trunk and the attached device transmits untagged traffic back to the switch, use 802.1Q trunking with the native VLAN configured to accept the traffic from the attached device.
- Do not configure the destinations with Layer 3 addresses. Use a VLAN interface to route traffic to and from devices attached to destinations.
- Destinations are held in the down state. To route the traffic to and from attached devices, configure an additional active Layer 2 port in the VLAN to keep the VLAN interface up.

This example shows how to configure session 1 to monitor ingress traffic from Gigabit Ethernet port 1/1 and configure Gigabit Ethernet port 1/2 as the destination:

```
Router(config)# monitor session 1 type local
Router(config-mon-local)# source interface gigabitethernet 1/1 rx
Router(config-mon-local)# destination interface gigabitethernet 1/2
```

For additional examples, see the [“Configuration Examples for SPAN” section on page 18-31](#).

Configuring Local SPAN (Global Configuration Mode)



Note

- To tag the monitored traffic as it leaves a destination, you must configure the destination to trunk unconditionally before you configure it as a destination (see the [“Configuring a Destination as an Unconditional Trunk \(Optional\)” section on page 18-13](#)).
- You can configure up to two local SPAN sessions in global configuration mode.
- You can use SPAN configuration mode for all SPAN configuration tasks.
- You must use SPAN configuration mode to configure the supported maximum number of SPAN sessions.

Local SPAN does not use separate source and destination sessions. To configure a local SPAN session, configure local SPAN sources and destinations with the same session number. To configure a local SPAN session, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>local_span_session_number</i> source { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i> } [rx tx both]	Associates the local SPAN source session number with the source ports or VLANs and selects the traffic direction to be monitored.
Step 3	Router(config)# monitor session <i>local_span_session_number</i> destination { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> } [ingress [learning]]	Associates the local SPAN session number and the destinations.

- *local_span_session_number* can range from 1 to 80.
- *single_interface* is as follows:
 - **interface type slot/port**; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - **interface port-channel number**

**Note**

Destination port channel interfaces must be configured with the **channel-group group_num mode on** command and the **no channel-protocol** command. See the “[How to Configure EtherChannels](#)” section on page 23-7.

- *interface_list* is *single_interface* , *single_interface* , *single_interface* ...

**Note**

In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- *interface_range* is **interface type slot/first_port - last_port**.
- *mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- *single_vlan* is the ID number of a single VLAN.
- *vlan_list* is *single_vlan* , *single_vlan* , *single_vlan* ...
- *vlan_range* is *first_vlan_ID* - *last_vlan_ID*.
- *mixed_vlan_list* is, in any order, *single_vlan* , *vlan_range* , ...
- Enter the **ingress** keyword to configure destinations to receive traffic from attached devices.
- Enter the **learning** keyword to enable MAC address learning from the destinations, which allows the switch to transmit traffic that is addressed to devices attached to the destinations.

When configuring destinations with the **ingress** and **learning** keywords, note the following:

- Configure the destinations for Layer 2 switching. See the “[How to Configure LAN Interfaces for Layer 2 Switching](#)” section on page 21-5.
- If the destination is a trunk and the attached device transmits untagged traffic back to the switch, use 802.1Q trunking with the native VLAN configured to accept the traffic from the attached device.

- Do not configure the destinations with Layer 3 addresses. Use a VLAN interface to route traffic to and from devices attached to destinations.
- Destinations are held in the down state. To route the traffic to and from attached devices, configure an additional active Layer 2 port in the VLAN to keep the VLAN interface up.

This example shows how to configure Gigabit Ethernet port 5/1 as a bidirectional source for session 1:

```
Router(config)# monitor session 1 source interface gigabitethernet 5/1
```

This example shows how to configure Gigabit Ethernet port 5/48 as the destination for SPAN session 1:

```
Router(config)# monitor session 1 destination interface gigabitethernet 5/48
```

For additional examples, see the “Configuration Examples for SPAN” section on page 18-31.

Configuring RSPAN

- [Configuring RSPAN VLANs, page 18-20](#)
- [Configuring RSPAN Sessions \(SPAN Configuration Mode\), page 18-20](#)
- [Configuring RSPAN Sessions \(Global Configuration Mode\), page 18-23](#)

Configuring RSPAN VLANs

To configure a VLAN as an RSPAN VLAN, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# vlan <i>vlan_ID</i> [{- ,} <i>vlan_ID</i>]	Creates or modifies an Ethernet VLAN, a range of Ethernet VLANs, or several Ethernet VLANs specified in a comma-separated list (do not enter space characters).
Step 3	Router(config-vlan)# remote-span	Configures the VLAN as an RSPAN VLAN.
Step 4	Router(config-vlan)# end	Updates the VLAN database and returns to privileged EXEC mode.

Configuring RSPAN Sessions (SPAN Configuration Mode)

- [Configuring RSPAN Source Sessions in SPAN Configuration Mode, page 18-21](#)
- [Configuring RSPAN Destination Sessions in SPAN Configuration Mode, page 18-22](#)

Configuring RSPAN Source Sessions in SPAN Configuration Mode

To configure an RSPAN source session in SPAN configuration mode, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>RSPAN_source_session_number</i> type rspan-source	Configures an RSPAN source session number and enters RSPAN source session configuration mode for the session.
Step 3	Router(config-mon-rspan-src)# description <i>session_description</i>	(Optional) Describes the RSPAN source session.
Step 4	Router(config-mon-rspan-src)# source { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i> } [rx tx both]	Associates the RSPAN source session number with the source ports or VLANs, and selects the traffic direction to be monitored.
Step 5	Router(config-mon-rspan-src)# filter <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i>	(Optional) Configures source VLAN filtering when the RSPAN source is a trunk port.
Step 6	Router(config-mon-rspan-src)# destination remote vlan <i>rspan_vlan_ID</i>	Associates the RSPAN source session number session number with the RSPAN VLAN.
Step 7	Router(config-mon-rspan-src)# no shutdown	Activates the RSPAN source session.
Step 8	Router(config-mon-rspan-src)# end	Exits configuration mode.

- *session_description* can be up to 240 characters and cannot contain special characters. The description can contain spaces.



Note You can enter 240 characters after the **description** command.

- *RSPAN_source_span_session_number* can range from 1 to 80.
- *single_interface* is as follows:
 - **interface type slot/port**; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - **interface port-channel number**
- *interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- *interface_range* is **interface type slot/first_port - last_port**.
- *mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- *single_vlan* is the ID number of a single VLAN.
- *vlan_list* is *single_vlan* , *single_vlan* , *single_vlan* ...
- *vlan_range* is *first_vlan_ID - last_vlan_ID*.
- *mixed_vlan_list* is, in any order, *single_vlan* , *vlan_range* , ...

- See the [“Configuring RSPAN VLANs”](#) section on page 18-20 for information about the RSPAN VLAN ID.

This example shows how to configure session 1 to monitor bidirectional traffic from Gigabit Ethernet port 1/1:

```
Router(config)# monitor session 1 type rspan-source
Router(config-mon-rspan-src)# source interface gigabitethernet 1/1
Router(config-mon-rspan-src)# destination remote vlan 2
```

For additional examples, see the [“Configuration Examples for SPAN”](#) section on page 18-31.

Configuring RSPAN Destination Sessions in SPAN Configuration Mode



Note

- To tag the monitored traffic, you must configure the port to trunk unconditionally before you configure it as a destination (see the [“Configuring a Destination as an Unconditional Trunk \(Optional\)”](#) section on page 18-13).
- You can configure an RSPAN destination session on the RSPAN source session switch to monitor RSPAN traffic locally.

To configure an RSPAN destination session, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>RSPAN_destination_session_number</i> type rspan-destination	Configures an RSPAN destination session number and enters RSPAN destination session configuration mode for the session.
Step 3	Router(config-mon-rspan-dst)# description <i>session_description</i>	(Optional) Describes the RSPAN destination session.
Step 4	Router(config-mon-rspan-dst)# source remote vlan <i>rspan_vlan_ID</i>	Associates the RSPAN destination session number with the RSPAN VLAN.
Step 5	Router(config-mon-rspan-dst)# destination { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> } [ingress [learning]]	Associates the RSPAN destination session number with the destinations.
Step 6	Router(config-mon-rspan-dst)# end	Exits configuration mode.

- *RSPAN_destination_span_session_number* can range from 1 to 80.
- *single_interface* is as follows:
 - **interface** *type slot/port*; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - **interface** **port-channel** *number*



Note

Destination port channel interfaces must be configured with the **channel-group** *group_num* **mode on** command and the **no channel-protocol** command. See the [“How to Configure EtherChannels”](#) section on page 23-7.

- *interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- *interface_range* is **interface** *type slot/first_port - last_port*.
- *mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- Enter the **ingress** keyword to configure destinations to receive traffic from attached devices.
- Enter the **learning** keyword to enable MAC address learning from the destinations, which allows the switch to transmit traffic that is addressed to devices attached to the destinations.

When configuring destinations with the **ingress** and **learning** keywords, note the following:

- Configure the destinations for Layer 2 switching. See the [“How to Configure LAN Interfaces for Layer 2 Switching”](#) section on page 21-5.
- If the destination is a trunk and the attached device transmits untagged traffic back to the switch, use 802.1Q trunking with the native VLAN configured to accept the traffic from the attached device.
- Do not configure the destinations with Layer 3 addresses. Use a VLAN interface to route traffic to and from devices attached to destinations.
- Destinations are held in the down state. To route the traffic to and from attached devices, configure an additional active Layer 2 port in the VLAN to keep the VLAN interface up.
- The **no shutdown** command and **shutdown** commands are not supported for RSPAN destination sessions.

This example shows how to configure RSPAN VLAN 2 as the source for session 1 and Gigabit Ethernet port 1/2 as the destination:

```
Router(config)# monitor session 1 type rspan-destination
Router(config-rspan-dst)# source remote vlan 2
Router(config-rspan-dst)# destination interface gigabitethernet 1/2
```

For additional examples, see the [“Configuration Examples for SPAN”](#) section on page 18-31.

Configuring RSPAN Sessions (Global Configuration Mode)

- [Configuring RSPAN Source Sessions in Global Configuration Mode, page 18-24](#)
- [Configuring RSPAN Destination Sessions in Global Configuration Mode, page 18-25](#)

Configuring RSPAN Source Sessions in Global Configuration Mode

To configure an RSPAN source session in global configuration mode, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>RSPAN_source_session_number</i> source { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i> } [rx tx both]	Associates the RSPAN source session number with the source ports or VLANs, and selects the traffic direction to be monitored.
Step 3	Router(config)# monitor session <i>RSPAN_source_session_number</i> destination remote vlan <i>rspan_vlan_ID</i>	Associates the RSPAN source session session number with the RSPAN VLAN.

- To configure RSPAN VLANs, see the [“Configuring RSPAN VLANs” section on page 18-20](#).
- *RSPAN_source_span_session_number* can range from 1 to 80.
- *single_interface* is as follows:
 - **interface type slot/port**; type is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - **interface port-channel number**
- *interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note

In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- *interface_range* is **interface type slot/first_port - last_port**.
- *mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- *single_vlan* is the ID number of a single VLAN.
- *vlan_list* is *single_vlan* , *single_vlan* , *single_vlan* ...
- *vlan_range* is *first_vlan_ID - last_vlan_ID*.
- *mixed_vlan_list* is, in any order, *single_vlan* , *vlan_range* , ...
- See the [“Configuring RSPAN VLANs” section on page 18-20](#) for information about the RSPAN VLAN ID.

This example shows how to configure Gigabit Ethernet port 5/2 as the source for session 2:

```
Router(config)# monitor session 2 source interface gigabitethernet 5/2
```

This example shows how to configure RSPAN VLAN 200 as the destination for session 2:

```
Router(config)# monitor session 2 destination remote vlan 200
```

For additional examples, see the [“Configuration Examples for SPAN” section on page 18-31](#).

Configuring RSPAN Destination Sessions in Global Configuration Mode



Note

- To tag the monitored traffic, you must configure the port to trunk unconditionally before you configure it as a destination (see the [“Configuring a Destination as an Unconditional Trunk \(Optional\)”](#) section on page 18-13).
- You can configure an RSPAN destination session on the RSPAN source session switch to monitor RSPAN traffic locally.

To configure an RSPAN destination session in global configuration mode, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>RSPAN_destination_session_number</i> source remote vlan <i>rspan_vlan_ID</i>	Associates the RSPAN destination session number with the RSPAN VLAN.
Step 3	Router(config)# monitor session <i>RSPAN_destination_session_number</i> destination { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> } [ingress [learning]]	Associates the RSPAN destination session number with the destinations.

- RSPAN_destination_span_session_number* can range from 1 to 80.
- See the [“Configuring RSPAN VLANs”](#) section on page 18-20 for information about the RSPAN VLAN ID.
- single_interface* is as follows:
 - interface** *type slot/port*; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - interface** **port-channel** *number*



Note

Destination port channel interfaces must be configured with the **channel-group** *group_num* **mode on** command and the **no channel-protocol** command. See the [“How to Configure EtherChannels”](#) section on page 23-7.

- interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note

In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- interface_range* is **interface** *type slot/first_port - last_port*.
- mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- Enter the **ingress** keyword to configure destinations to receive traffic from attached devices.
- Enter the **learning** keyword to enable MAC address learning from the destinations, which allows the switch to transmit traffic that is addressed to devices attached to the destinations.

When configuring destinations with the **ingress** and **learning** keywords, note the following:

- Configure the destinations for Layer 2 switching. See the [“How to Configure LAN Interfaces for Layer 2 Switching”](#) section on page 21-5.
- If the destination is a trunk and the attached device transmits untagged traffic back to the switch, use 802.1Q trunking with the native VLAN configured to accept the traffic from the attached device.
- Do not configure the destinations with Layer 3 addresses. Use a VLAN interface to route traffic to and from devices attached to destinations.
- Destinations are held in the down state. To route the traffic to and from attached devices, configure an additional active Layer 2 port in the VLAN to keep the VLAN interface up.

This example shows how to configure RSPAN VLAN 200 as the source for session 3:

```
Router(config)# monitor session 3 source remote vlan 200
```

This example shows how to configure Gigabit Ethernet port 5/47 as the destination for session 3:

```
Router(config)# monitor session 3 destination interface gigabitethernet 5/47
```

For additional examples, see the [“Configuration Examples for SPAN”](#) section on page 18-31.

Configuring ERSPAN

- [Configuring ERSPAN Source Sessions, page 18-26](#)
- [Configuring ERSPAN Destination Sessions, page 18-28](#)

Configuring ERSPAN Source Sessions

To configure an ERSPAN source session, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>ERSPAN_source_session_number</i> type erspan-source	Configures an ERSPAN source session number and enters ERSPAN source session configuration mode for the session.
Step 3	Router(config-mon-erspan-src)# description <i>session_description</i>	(Optional) Describes the ERSPAN source session.
Step 4	Router(config-mon-erspan-src)# source { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i> } [rx tx both]	Associates the ERSPAN source session number with the source ports or VLANs, and selects the traffic direction to be monitored.
Step 5	Router(config-mon-erspan-src)# filter <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i>	(Optional) Configures source VLAN filtering when the ERSPAN source is a trunk port.
Step 6	Router(config-mon-erspan-src)# destination	Enters ERSPAN source session destination configuration mode.

	Command	Purpose
Step 7	Router(config-mon-erspan-src-dst)# ip address <i>ip_address</i>	Configures the ERSPAN flow destination IP address, which must also be configured on an interface on the destination switch and be entered in the ERSPAN destination session configuration (see the “Configuring ERSPAN Destination Sessions” section on page 18-28, Step 6).
Step 8	Router(config-mon-erspan-src-dst)# erspan-id <i>ERSPAN_flow_id</i>	Configures the ID number used by the source and destination sessions to identify the ERSPAN traffic, which must also be entered in the ERSPAN destination session configuration (see the “Configuring ERSPAN Destination Sessions” section on page 18-28, Step 7).
Step 9	Router(config-mon-erspan-src-dst)# origin ip address <i>ip_address</i> [force]	Configures the IP address used as the source of the ERSPAN traffic.
Step 10	Router(config-mon-erspan-src-dst)# ip ttl <i>ttl_value</i>	(Optional) Configures the IP time-to-live (TTL) value of the packets in the ERSPAN traffic.
Step 11	Router(config-mon-erspan-src-dst)# ip prec <i>ipp_value</i>	(Optional) Configures the IP precedence value of the packets in the ERSPAN traffic.
Step 12	Router(config-mon-erspan-src-dst)# ip dscp <i>dscp_value</i>	(Optional) Configures the IP DSCP value of the packets in the ERSPAN traffic.
Step 13	Router(config-mon-erspan-src-dst)# vrf <i>vrf_name</i>	(Optional) Configures the VRF name to use instead of the global routing table.
Step 14	Router(config-mon-erspan-src)# no shutdown	Activates the ERSPAN source session.
Step 15	Router(config-mon-erspan-src-dst)# end	Exits configuration mode.

- *session_description* can be up to 240 characters and cannot contain special characters. The description can contain spaces.



Note You can enter 240 characters after the **description** command.

- *ERSPAN_source_span_session_number* can range from 1 to 80.
- *single_interface* is as follows:
 - **interface** *type slot/port*; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - **interface** **port-channel** *number*



Note Port channel interfaces must be configured with the **channel-group** *group_num* **mode on** command and the **no channel-protocol** command. See the [“How to Configure EtherChannels”](#) section on page 23-7.

- *interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- *interface_range* is **interface** type *slot/first_port - last_port*.
- *mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- *single_vlan* is the ID number of a single VLAN.
- *vlan_list* is *single_vlan* , *single_vlan* , *single_vlan* ...
- *vlan_range* is *first_vlan_ID - last_vlan_ID*.
- *mixed_vlan_list* is, in any order, *single_vlan* , *vlan_range* , ...
- *ERSPAN_flow_id* can range from 1 to 1023.
- All ERSPAN source sessions on a switch must use the same source IP address. Enter the **origin ip address ip_address force** command to change the origin IP address configured in all ERSPAN source sessions on the switch.
- *ttl_value* can range from 1 to 255.
- *ipp_value* can range from 0 to 7.
- *dscp_value* can range from 0 to 63.

This example shows how to configure session 3 to monitor bidirectional traffic from Gigabit Ethernet port 4/1:

```
Router(config)# monitor session 3 type erspan-source
Router(config-mon-erspan-src)# source interface gigabitethernet 4/1
Router(config-mon-erspan-src)# destination
Router(config-mon-erspan-src-dst)# ip address 10.1.1.1
Router(config-mon-erspan-src-dst)# origin ip address 20.1.1.1
Router(config-mon-erspan-src-dst)# erspan-id 101
```

For additional examples, see the “Configuration Examples for SPAN” section on page 18-31.

Configuring ERSPAN Destination Sessions



Note

You cannot monitor ERSPAN traffic locally.

To configure an ERSPAN destination session, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>ERSPAN_destination_session_number</i> type erspan-destination	Configures an ERSPAN destination session number and enters ERSPAN destination session configuration mode for the session.
Step 3	Router(config-mon-erspan-dst)# description <i>session_description</i>	(Optional) Describes the ERSPAN destination session.
Step 4	Router(config-mon-erspan-dst)# destination { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> } [ingress [learning]]	Associates the ERSPAN destination session number with the destinations.
Step 5	Router(config-mon-erspan-dst)# source	Enters ERSPAN destination session source configuration mode.

	Command	Purpose
Step 6	Router(config-mon-erspan-dst-src) # ip address <i>ip_address</i> [force]	Configures the ERSPAN flow destination IP address. This must be an address on a local interface and match the address that you entered in the “ Configuring ERSPAN Source Sessions ” section on page 18-26, Step 7.
Step 7	Router(config-mon-erspan-dst-src) # erspan-id <i>ERSPAN_flow_id</i>	Configures the ID number used by the destination and destination sessions to identify the ERSPAN traffic. This must match the ID that you entered in the “ Configuring ERSPAN Source Sessions ” section on page 18-26, Step 8.
Step 8	Router(config-mon-erspan-dst-src) # vrf <i>vrf_name</i>	(Optional) Configures the VRF name used instead of the global routing table.
Step 9	Router(config-mon-erspan-dst) # no shutdown	Activates the ERSPAN destination session.
Step 10	Router(config-mon-erspan-dst-src) # end	Exits configuration mode.

- *ERSPAN_destination_span_session_number* can range from 1 to 80.
- *single_interface* is as follows:
 - **interface** *type slot/port*; *type* is **fastethernet**, **gigabitethernet**, or **tengigabitethernet**.
 - **interface** **port-channel** *number*



Note Destination port channel interfaces must be configured with the **channel-group** *group_num* **mode on** command and the **no channel-protocol** command. See the “[How to Configure EtherChannels](#)” section on page 23-7.

- *interface_list* is *single_interface* , *single_interface* , *single_interface* ...



Note In lists, you must enter a space before and after the comma. In ranges, you must enter a space before and after the dash.

- *interface_range* is **interface** *type slot/first_port - last_port*.
- *mixed_interface_list* is, in any order, *single_interface* , *interface_range* , ...
- All ERSPAN destination sessions on a switch must use the same IP address on the same destination interface. Enter the **ip address** *ip_address* **force** command to change the IP address configured in all ERSPAN destination sessions on the switch.



Note You must also change all ERSPAN source session destination IP addresses (see the “[Configuring ERSPAN Source Sessions](#)” section on page 18-26, Step 7).

- *ERSPAN_flow_id* can range from 1 to 1023.
- Enter the **ingress** keyword to configure destinations to receive traffic from attached devices.

- Enter the **learning** keyword to enable MAC address learning from the destinations, which allows the switch to transmit traffic that is addressed to devices attached to the destinations.

When configuring destinations with the **ingress** and **learning** keywords, note the following:

- Configure the destinations for Layer 2 switching. See the “[How to Configure LAN Interfaces for Layer 2 Switching](#)” section on page 21-5.
- If the destination is a trunk and the attached device transmits untagged traffic back to the switch, use 802.1Q trunking with the native VLAN configured to accept the traffic from the attached device.
- Do not configure the destinations with Layer 3 addresses. Use a VLAN interface to route traffic to and from devices attached to destinations.
- Destinations are held in the down state. To route the traffic to and from attached devices, configure an additional active Layer 2 port in the VLAN to keep the VLAN interface up.

This example shows how to configure an ERSPAN destination session to send ERSPAN ID 101 traffic arriving at IP address 10.1.1.1 to Gigabit Ethernet port 2/1:

```
Router(config)# monitor session 3 type erspan-destination
Router(config-erspan-dst)# destination interface gigabitethernet 2/1
Router(config-erspan-dst)# source
Router(config-erspan-dst-src)# ip address 10.1.1.1
Router(config-erspan-dst-src)# erspan-id 101
```

For additional examples, see the “[Configuration Examples for SPAN](#)” section on page 18-31.

Configuring Source VLAN Filtering in Global Configuration Mode



Note

- To configure source VLAN filtering in SPAN configuration mode, see these sections:
 - [Configuring Local SPAN \(SPAN Configuration Mode\)](#), page 18-16
 - [Configuring RSPAN Source Sessions in SPAN Configuration Mode](#), page 18-21
 - [Configuring ERSPAN](#), page 18-26
- Source VLAN filtering reduces the amount of traffic that is sent from SPAN sources to SPAN destinations.

Source VLAN filtering monitors specific VLANs when the source is a trunk port.

To configure source VLAN filtering when the local SPAN or RSPAN source is a trunk port, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# monitor session <i>session_number</i> filter <i>single_vlan</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i>	Configures source VLAN filtering when the local SPAN or RSPAN source is a trunk port.

- *single_vlan* is the ID number of a single VLAN.
- *vlan_list* is *single_vlan* , *single_vlan* , *single_vlan* ...

- *vlan_range* is *first_vlan_ID* - *last_vlan_ID*.
- *mixed_vlan_list* is, in any order, *single_vlan* , *vlan_range* , ...

This example shows how to monitor VLANs 1 through 5 and VLAN 9 when the source is a trunk port:

```
Router(config)# monitor session 2 filter vlan 1 - 5 , 9
```

Verifying the SPAN Configuration

To verify the configuration, enter the **show monitor session** command.

This example shows how to verify the configuration of session 2:

```
Router# show monitor session 2
Session 2
-----
Type : Remote Source Session

Source Ports:
    RX Only:      Gi3/1
Dest RSPAN VLAN: 901
Router#
```

This example shows how to display the full details of session 2:

```
Router# show monitor session 2 detail
Session 2
-----
Type : Remote Source Session

Source Ports:
    RX Only:      Gi1/1-3
    TX Only:      None
    Both:         None
Source VLANs:
    RX Only:      None
    TX Only:      None
    Both:         None
Source RSPAN VLAN: None
Destination Ports: None
Filter VLANs:    None
Dest RSPAN VLAN: 901
```

Configuration Examples for SPAN

This example shows the configuration of RSPAN source session 2:

```
Router(config)# monitor session 2 source interface gigabitethernet1/1 - 3 rx
Router(config)# monitor session 2 destination remote vlan 901
```

This example shows how to clear the configuration for sessions 1 and 2:

```
Router(config)# no monitor session range 1-2
```

This example shows the configuration of an RSPAN source session with multiple sources:

```
Router(config)# monitor session 2 source interface gigabitethernet 5/15 , 7/3 rx
Router(config)# monitor session 2 source interface gigabitethernet 1/2 tx
Router(config)# monitor session 2 source interface port-channel 102
Router(config)# monitor session 2 source filter vlan 2 - 3
```

```
Router(config)# monitor session 2 destination remote vlan 901
```

This example shows how to remove sources for a session:

```
Router(config)# no monitor session 2 source interface gigabitethernet 5/15 , 7/3
```

This example shows how to remove options for sources for a session:

```
Router(config)# no monitor session 2 source interface gigabitethernet 1/2
```

```
Router(config)# no monitor session 2 source interface port-channel 102 tx
```

This example shows how to remove source VLAN filtering for a session:

```
Router(config)# no monitor session 2 filter vlan 3
```

This example shows the configuration of RSPAN destination session 8:

```
Router(config)# monitor session 8 source remote vlan 901
```

```
Router(config)# monitor session 8 destination interface gigabitethernet 1/2 , 2/3
```

This example shows the configuration of ERSPAN source session 12:

```
monitor session 12 type erspan-source
description SOURCE_SESSION_FOR_VRF_GRAY
source interface Gi8/48 rx
destination
  erspan-id 120
  ip address 10.8.1.2
  origin ip address 32.1.1.1
vrf gray
```

This example shows the configuration of ERSPAN destination session 12:

```
monitor session 12 type erspan-destination
description DEST_SESSION_FOR_VRF_GRAY
destination interface Gi4/48
source
  erspan-id 120
  ip address 10.8.1.2
vrf gray
```

This example shows the configuration of ERSPAN source session 13:

```
monitor session 13 type erspan-source
source interface Gi6/1 tx
destination
  erspan-id 130
  ip address 10.11.1.1
  origin ip address 32.1.1.1
```

This example shows the configuration of ERSPAN destination session 13:

```
monitor session 13 type erspan-destination
destination interface Gi6/1
source
  erspan-id 130
  ip address 10.11.1.1
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

19

SNMP IfIndex Persistence

- [Prerequisites for SNMP IfIndex Persistence, page 19-1](#)
- [Restrictions for SNMP IfIndex Persistence, page 19-1](#)
- [Information About SNMP IfIndex Persistence, page 19-1](#)
- [Default Settings for SNMP IfIndex Persistence, page 19-2](#)
- [How to Configure SNMP IfIndex Persistence, page 19-2](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for SNMP IfIndex Persistence

None.

Restrictions for SNMP IfIndex Persistence

None.

Information About SNMP IfIndex Persistence

The SNMP ifIndex persistence feature provides an interface index (ifIndex) value that is retained and used when the switch reboots. The ifIndex value is a unique identifying number associated with a physical or logical interface.

There is no requirement in the relevant RFCs that the correspondence between particular ifIndex values and their interfaces be maintained when the switch reboots, but many applications (for example, device inventory, billing, and fault detection) require maintenance of this correspondence.

You can poll the switch at regular intervals to correlate the interfaces to the ifIndexes, but it is not practical to poll constantly. The SNMP ifIndex persistence feature provides permanent ifIndex values, which eliminates the need to poll interfaces.

The following definitions are based on RFC 2233, “The Interfaces Group MIB using SMIV2.” The following terms are values in the Interfaces MIB (IF-MIB):

- **ifIndex**—A unique number (greater than zero) that identifies each interface for SNMP identification of that interface.
- **ifName**—The text-based name of the interface, for example, “ethernet 3/1.”
- **ifDescr**—A description of the interface. Recommended information for this description includes the name of the manufacturer, the product name, and the version of the interface hardware and software.

Default Settings for SNMP IfIndex Persistence

SNMP ifIndex persistence is disabled by default.

How to Configure SNMP IfIndex Persistence

- [Enabling SNMP IfIndex Persistence Globally, page 19-2](#)
- [Enabling and Disabling SNMP IfIndex Persistence on Specific Interfaces, page 19-3](#)

Enabling SNMP IfIndex Persistence Globally

To globally enable SNMP ifIndex persistence, perform this task:

Command	Purpose
Router(config)# snmp-server ifindex persist	Globally enables SNMP ifIndex persistence.

In the following example, SNMP ifIndex persistence is enabled for all interfaces:

```
router(config)# snmp-server ifindex persist
```

Disabling SNMP IfIndex Persistence Globally

To globally disable SNMP ifIndex persistence after enabling it, perform this task:

Command	Purpose
Router(config)# no snmp-server ifindex persist	Globally disables SNMP ifIndex persistence.

In the following example, SNMP ifIndex persistence is disabled for all interfaces:

```
router(config)# no snmp-server ifindex persist
```

Enabling and Disabling SNMP IfIndex Persistence on Specific Interfaces

To enable SNMP ifIndex persistence only on a specific interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface { vlan <i>vlan_ID</i> } { type slot/port } { port-channel <i>port_channel_number</i> }	Selects an interface to configure.
Step 2	Router(config-if)# snmp ifindex persist	Enables SNMP ifIndex persistence on the specified interface.
Step 3	Router(config-if)# exit	Exits interface configuration mode.



Note

The **[no] snmp ifindex persistence** interface command cannot be used on subinterfaces. A command applied to an interface is automatically applied to all the subinterfaces associated with that interface.

In the following example, SNMP ifIndex persistence is enabled for Ethernet interface 3/1 only:

```
router(config)# interface ethernet 3/1
router(config-if)# snmp ifindex persist
router(config-if)# exit
```

In the following example, SNMP ifIndex persistence is disabled for Ethernet interface 3/1 only:

```
router(config)# interface ethernet 3/1
router(config-if)# no snmp ifindex persist
router(config-if)# exit
```

Clearing SNMP IfIndex Persistence Configuration from a Specific Interface

To clear the interface-specific SNMP ifIndex persistence setting and configure the interface to use the global configuration setting, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface. Note that the syntax of the interface command will vary depending on the platform you are using.
Step 2	Router(config-if)# snmp ifindex clear	Clears any interface-specific SNMP ifIndex persistence configuration for the specified interface and returns to the global configuration setting.
Step 3	Router(config-if)# exit	Exits interface configuration mode.

In the following example, any previous setting for SNMP ifIndex persistence on Ethernet interface 3/1 is removed from the configuration. If SNMP ifIndex persistence is globally enabled, SNMP ifIndex persistence will be enabled for Ethernet interface 3/1. If SNMP ifIndex persistence is globally disabled, SNMP ifIndex persistence will be disabled for Ethernet interface 3/1.

```
router(config)# interface ethernet 3/1
router(config-if)# snmp ifindex clear
router(config-if)# exit
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

20

Top-N Reports

- [Prerequisites for Top-N Reports, page 20-1](#)
- [Restrictions for Top-N Reports, page 20-1](#)
- [Information About Top-N Reports, page 20-1](#)
- [Default Settings for Top-N Reports, page 20-2](#)
- [How to Use Top-N Reports, page 20-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Top-N Reports

None.

Restrictions for Top-N Reports

None.

Information About Top-N Reports

- [Top-N Reports Overview, page 20-2](#)
- [Top-N Reports Operation, page 20-2](#)

Top-N Reports Overview

Top-N reports allows you to collect and analyze data for each physical port on a switch. When Top-N reports start, they obtain statistics from the appropriate hardware counters and then go into sleep mode for a user-specified interval. When the interval ends, the reports obtain the current statistics from the same hardware counters, compare the current statistics from the earlier statistics, and store the difference. The statistics for each port are sorted by one of the statistic types that are listed in [Table 20-1](#).

Table 20-1 Valid Top-N Statistic Types

Statistic Type	Definition
broadcast	Number of input/output broadcast packets
bytes	Number of input/output bytes
errors	Number of input errors
multicast	Number of input/output multicast packets
overflow	Number of buffer overflows
packets	Number of input/output packets
utilization	Utilization



Note

When calculating the port utilization, Top-N reports bundles the Tx and Rx lines into the same counter and also looks at the full-duplex bandwidth when calculating the percentage of utilization. For example, a Gigabit Ethernet port would be 2000-Mbps full duplex.

Top-N Reports Operation

When you enter the **collect top** command, processing begins and the system prompt reappears immediately. When processing completes, the reports are not displayed immediately on the screen; the reports are saved for later viewing. The Top-N reports notify you when the reports are complete by sending a syslog message to the screen.

To view the completed reports, enter the **show top counters interface report** command. Only completed reports are displayed. For reports that are not completed, there is a short description of the process information.

To terminate a Top-N reports process, enter the **clear top counters interface report** command. Pressing **Ctrl-C** does not terminate Top-N reports processes. The completed reports remain available for viewing until you remove them by entering the **clear top counters interface report {all | report_num}** command.

Default Settings for Top-N Reports

None.

How to Use Top-N Reports

- [Enabling Top-N Reports Creation, page 20-3](#)
- [Displaying Top-N Reports, page 20-3](#)
- [Clearing Top-N Reports, page 20-4](#)

Enabling Top-N Reports Creation

To enable Top-N reports creation, perform this task:

Command	Purpose
Router# collect top [<i>number_of_ports</i>] counters interface { <i>type</i> all layer-2 layer-3 } [sort-by <i>statistic_type</i>] [interval <i>seconds</i>]	Enables Top-N reports creation.

When enabling Top-N reports creation, note the following information:

- You can specify the number of busiest ports for which to create reports (the default is 20).
- You can specify the statistic type by which ports are determined to be the busiest (the default is utilization). The supported values for *statistic_type* are **broadcast**, **bytes**, **errors**, **multicast**, **overflow**, **packets**, and **utilization**.
- You can specify the interval over which statistics are collected (range: 0 through 999; the default is 30 seconds).
- Except for a utilization report (configured with the **sort-by utilization** keywords), you can specify an interval of zero to create a report that displays the current counter values instead of a report that displays the difference between the start-of-interval counter values and the end-of-interval counter values.

This example shows how to enable Top-N reports creation for an interval of 76 seconds for the four ports with the highest utilization:

```
Router# collect top 4 counters interface all sort-by utilization interval 76
TopN collection started.
```

Displaying Top-N Reports

To display Top-N reports, perform this task:

Command	Purpose
Router# show top counters interface report [<i>report_num</i>]	Displays Top-N reports.
	Note To display information about all the reports, do not enter a <i>report_num</i> value.

Top-N reports statistics are not displayed in these situations:

- If a port is not present during the first poll.
- If a port is not present during the second poll.

- If a port's speed or duplex changes during the polling interval.
- If a port's type changes from Layer 2 to Layer 3 during the polling interval.
- If a port's type changes from Layer 3 to Layer 2 during the polling interval.

This example shows how to display information about all the Top-N reports:

```
Router# show top counters interface report
Id Start Time                               Int N   Sort-By   Status   Owner
-----
1  08:18:25 UTC Tue Nov 23 2004 76  20  util     done     console
2  08:19:54 UTC Tue Nov 23 2004 76  20  util     done     console
3  08:21:34 UTC Tue Nov 23 2004 76  20  util     done     console
4  08:26:50 UTC Tue Nov 23 2004 90  20  util     done     console
```



Note

Reports for which statistics are still being obtained are shown with a status of pending.

This example shows how to display a specific Top-N report:

```
Router# show top counters interface report 1
Started By      : console
Start Time     : 08:18:25 UTC Tue Nov 23 2004
End Time       : 08:19:42 UTC Tue Nov 23 2004
Port Type      : All
Sort By        : util
Interval       : 76 seconds
Port   Band  Util Bytes      Packets      Broadcast  Multicast  In-  Buf-
      width  (Tx + Rx) (Tx + Rx) (Tx + Rx) (Tx + Rx) err  ovflw
-----
Gi2/5   100   50  726047564  11344488  11344487    1           0    0
Gi2/48  100   35  508018905  7937789   0           43          0    0
Gi2/46  100   25  362860697  5669693   0           43          0    0
Gi2/47  100   22  323852889  4762539   4762495     43          0    0
```

Clearing Top-N Reports

To clear Top-N reports, perform one of these tasks:

Command	Purpose
Router# clear top counters interface report	Clears all the Top-N reports that have a status of done.
Router# clear top counters interface report <i>[report_num]</i>	Clears Top-N report number <i>report_num</i> regardless of status.

This example shows how to remove all reports that have a status of done:

```
Router# clear top counters interface report
04:00:06: %TOPN_COUNTERS-5-DELETED: TopN report 1 deleted by the console
04:00:06: %TOPN_COUNTERS-5-DELETED: TopN report 2 deleted by the console
04:00:06: %TOPN_COUNTERS-5-DELETED: TopN report 3 deleted by the console
04:00:06: %TOPN_COUNTERS-5-DELETED: TopN report 4 deleted by the console
```

This example shows how to remove a report number 4:

```
Router# clear top counters interface report 4
04:52:12: %TOPN_COUNTERS-5-KILLED: TopN report 4 killed by the console
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

21

Layer 2 Traceroute Utility

- [Prerequisites for the Layer 2 Traceroute Utility, page 21-1](#)
- [Restrictions for the Layer 2 Traceroute Utility, page 21-1](#)
- [Information About the Layer 2 Traceroute Utility, page 21-2](#)
- [How to Use the Layer 2 Traceroute Utility, page 21-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for the Layer 2 Traceroute Utility

None.

Restrictions for the Layer 2 Traceroute Utility

- Cisco Discovery Protocol (CDP) must be enabled on all the devices in the network. For the Layer 2 traceroute utility to function properly, do not disable CDP. If any devices in the Layer 2 path are transparent to CDP, the Layer 2 traceroute utility cannot identify these devices on the path.
- A switch is defined as reachable from another switch when you can test connectivity by using the **ping** privileged EXEC command. All devices in the Layer 2 path must be mutually reachable. To verify the ping connectivity you need to use the IP address that the CDP advertises on its Layer 2 interfaces.
- 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 Layer 2 path from the source device to the destination device. All devices 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 Layer 2 traceroute utility 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 Layer 2 traceroute utility uses the associated MAC address and identifies the Layer 2 path.
 - If an ARP entry does not exist, the Layer 2 traceroute utility 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 utility terminates at that hop and displays an error message.
- The Layer 2 traceroute utility is not supported in Token Ring VLANs.

Information About the Layer 2 Traceroute Utility

The Layer 2 traceroute utility identifies the Layer 2 path that a packet takes from a source device to a destination device. Layer 2 traceroute supports only unicast source and destination MAC addresses. The utility determines the path by using the MAC address tables of the switches in the path. When the Layer 2 traceroute utility detects a device in the path that does not support Layer 2 traceroute, it continues to send Layer 2 trace queries and allows them to time out.

The Layer 2 traceroute utility can only identify the path from the source device to the destination device. The utility cannot identify the path that a packet takes from the source host to the source device or from the destination device to the destination host.

How to Use the Layer 2 Traceroute Utility

To display the Layer 2 path that a packet takes from a source device to a destination device, perform one of these tasks in privileged EXEC mode:

Command	Purpose
Router# traceroute mac [interface type interface_number] <i>source_mac_address</i> [interface type interface_number] <i>destination_mac_address</i> [vlan vlan_id] [detail]	Uses MAC addresses to trace the path that packets take through the network.
Router# traceroute mac ip { <i>source_ip_address</i> <i>source_hostname</i> } { <i>destination_ip_address</i> <i>destination_hostname</i> } [detail]	Uses IP addresses to trace the path that packets take through the network.

These examples show how to use the **traceroute mac** and **traceroute mac ip** commands to display the physical path a packet takes through the network to reach its destination:

```
Router# traceroute mac 0000.0201.0601 0000.0201.0201
```

```
Source 0000.0201.0601 found on con6[WS-C2950G-24-EI] (2.2.6.6)
con6 (2.2.6.6) :Fa0/1 => Fa0/3
con5          (2.2.5.5      ) : Fa0/3 => Gi0/1
con1          (2.2.1.1      ) : Gi0/1 => Gi0/2
con2          (2.2.2.2      ) : Gi0/2 => Fa0/1
Destination 0000.0201.0201 found on con2[WS-C3550-24] (2.2.2.2)
Layer 2 trace completed
```

```
Router#
```

```
Router# traceroute mac 0001.0000.0204 0001.0000.0304 detail
```

```
Source 0001.0000.0204 found on VAYU[WS-C6509] (2.1.1.10)
1 VAYU / WS-C6509 / 2.1.1.10 :
    Gi6/1 [full, 1000M] => Po100 [auto, auto]
2 PANI / WS-C6509 / 2.1.1.12 :
    Po100 [auto, auto] => Po110 [auto, auto]
3 BUMI / WS-C6509 / 2.1.1.13 :
    Po110 [auto, auto] => Po120 [auto, auto]
4 AGNI / WS-C6509 / 2.1.1.11 :
    Po120 [auto, auto] => Gi8/12 [full, 1000M] Destination 0001.0000.0304
found on AGNI[WS-C6509] (2.1.1.11) Layer 2 trace completed.
Router#
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

22

Mini Protocol Analyzer

- [Prerequisites for the Mini Protocol Analyzer, page 22-1](#)
- [Restrictions for the Mini Protocol Analyzer, page 22-1](#)
- [Information About the Mini Protocol Analyzer, page 22-2](#)
- [How to Configure the Mini Protocol Analyzer, page 22-2](#)
- [Configuration Examples for the Mini Protocol Analyzer, page 22-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for the Mini Protocol Analyzer

None.

Restrictions for the Mini Protocol Analyzer

The PFC and any DFCs provide hardware support for analysis of IPv4 traffic. Analysis of low volume IPv6 traffic is supported in software.

Information About the Mini Protocol Analyzer

The Mini Protocol Analyzer captures network traffic from a SPAN session and stores the captured packets in a local memory buffer. Using the provided filtering options, you can limit the captured packets to:

- Packets from selected VLANs, ACLs, or MAC addresses.
- Packets of a specific EtherType.
- Packets of a specified packet size.

You can start and stop the capture using immediate commands, or you can schedule the capture to begin at a specified date and time.

The captured data can be displayed on the console, stored to a local file system, or exported to an external server using normal file transfer protocols. The format of the captured file is libpcap, which is supported by many packet analysis and sniffer programs. Details of this format can be found at the following URL:

<http://www.tcpdump.org/>

By default, only the first 68 bytes of each packet are captured.

How to Configure the Mini Protocol Analyzer

- [Configuring a Capture Session, page 22-2](#)
- [Filtering the Packets to be Captured, page 22-4](#)
- [Starting and Stopping a Capture, page 22-5](#)
- [Displaying and Exporting the Capture Buffer, page 22-7](#)

Configuring a Capture Session

To configure a capture session using the Mini Protocol Analyzer, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# [no] monitor session <i>number</i> type capture	Configures a SPAN session number with packets directed to the processor for capture. Enters capture session configuration mode. The session number range is 1 to 80. The no prefix removes the session.
Step 3	Router(config-mon-capture)# buffer-size <i>buf_size</i>	(Optional) Sets the size in KB of the capture buffer. The range is 32-65535 KB; the default is 2048 KB.
Step 4	Router(config-mon-capture)# description <i>session_description</i>	(Optional) Describes the capture session. The description can be up to 240 characters and cannot contain special characters. If the description contains spaces, it must be enclosed in quotation marks("").

	Command	Purpose
Step 5	Router(config-mon-capture)# rate-limit <i>pps</i>	(Optional) Sets a limit on the number of packets per second (<i>pps</i>) that can be captured. The range is 10-100000 packets per seconds; the default is 10000 packets per second.
Step 6	Router(config-mon-capture)# source {{ interface { <i>single_interface</i> <i>interface_list</i> <i>interface_range</i> <i>mixed_interface_list</i> } port-channel <i>channel_id</i> } { vlan { <i>vlan_ID</i> <i>vlan_list</i> <i>vlan_range</i> <i>mixed_vlan_list</i> }} [rx tx both]	Associates the capture session with source ports or VLANs, and selects the traffic direction to be monitored. The default traffic direction is both.
Step 7	Router(config-mon-capture)# exit	Exits the capture session configuration mode.

- Only one capture session is supported; multiple simultaneous capture sessions cannot be configured.
- The **source interface** command argument is either a single interface, or a range of interfaces described by two interface numbers (the lesser one first, separated by a dash), or a comma-separated list of interfaces and ranges.



Note When configuring a source interface list, you must enter a space before and after the comma. When configuring a source interface range, you must enter a space before and after the dash.

- The **source vlan** command argument is either a single VLAN number from 1 through 4094 (except reserved VLANs), or a range of VLANs described by two VLAN numbers (the lesser one first, separated by a dash), or a list of VLANs and ranges.



Note When configuring a source VLAN list, do not enter a space before or after the comma. When configuring a source VLAN range, do not enter a space before or after the dash. Note that this requirement differs from the requirement for source interface lists and ranges.

- Data capture does not begin when the capture session is configured. The capture is started by the **monitor capture start** or **monitor capture schedule** command described in the [“Starting and Stopping a Capture”](#) section on page 22-5.
- Although the capture buffer is linear by default, it can be made circular as a run-time option in the **monitor capture start** or **monitor capture schedule** command.
- When no hardware rate limit registers are available, the capture session is disabled.
- The source VLAN cannot be changed if a VLAN filter is configured. Remove any VLAN filters before changing the source VLAN.

Filtering the Packets to be Captured

To filter the packets to be captured by the Mini Protocol Analyzer, perform this task in capture session configuration mode:

	Command	Purpose
Step 1	Router(config-mon-capture) # filter access-group {acl_number acl_name}	(Optional) Captures only packets from the specified ACL.
Step 2	Router(config-mon-capture) # filter vlan {vlan_ID vlan_list vlan_range mixed_vlan_list}	(Optional) Captures only packets from the specified source VLAN or VLANs.
Step 3	Router(config-mon-capture) # filter ethertype type	(Optional) Captures only packets of the specified EtherType. The <i>type</i> can be specified in decimal, hex, or octal. Note Configure <i>type</i> as 0x86dd to filter IPv6 traffic. Analysis of low volume IPv6 traffic is supported in software.
Step 4	Router(config-mon-capture) # filter length min_len [max_len]	(Optional) Captures only packets whose size is between <i>min_len</i> and <i>max_len</i> , inclusive. If <i>max_len</i> is not specified, only packets of exactly size <i>min_len</i> will be captured. The range for <i>min_len</i> is 0 to 9216 bytes and the range for <i>max_len</i> is 1 to 9216 bytes.
Step 5	Router(config-mon-capture) # filter mac-address mac_addr	(Optional) Captures only packets from the specified MAC address.
Step 6	Router(config-mon-capture) # end	Exits the configuration mode.

- Several options are provided for filtering the packets to be captured. Filtering by ACL and VLAN is performed in hardware before any rate-limiting is applied; all other filters are executed in software. Software filtering can decrease the capture rate.
- The **filter vlan** argument is either a single VLAN number from 1 through 4094 (except reserved VLANs), or a range of VLANs described by two VLAN numbers (the lesser one first, separated by a dash), or a list of VLANs and ranges.



Note When configuring a filter VLAN list, you must enter a space before and after the comma. When configuring a filter VLAN range, you must enter a space before and after the dash. Note that this requirement differs from the requirement for source VLAN lists and ranges described in the preceding section.

- To enter an EtherType as a decimal number, enter the number (1 to 65535) with no leading zero. To enter a hexadecimal number, precede four hexadecimal characters with the prefix 0x. To enter an octal number, enter numeric digits (0 to 7) with a leading zero. For example, the 802.1Q EtherType can be entered in decimal notation as 33024, in hexadecimal as 0x8100, or in octal as 0100400.

- Enter a MAC address as three 2-byte values in dotted hexadecimal format. An example is 0123.4567.89ab.
- The **no** keyword removes the filter.

**Note**

After removing a VLAN filter using the **no** keyword, you must exit configuration mode, reenter the capture configuration mode, and issue the **source vlan** command before making other capture configuration changes.

- When you configure a VLAN filter, the capture source or destination must be a VLAN. When you configure a port filter, the capture source or destination must be a port.

Starting and Stopping a Capture

The commands to start and stop a capture are not stored as configuration settings. These commands are executed from the console in EXEC mode. You can start a capture immediately or you can set a future date and time for the capture to start. The capture ends when one of the following conditions occurs:

- A stop or clear command is entered from the console.
- The capture buffer becomes full, unless it is configured as a circular buffer.
- The optionally specified number of seconds has elapsed.
- The optionally specified number of packets has been captured.

When the capture stops, the SPAN session is ended and no further capture session packets are forwarded to the processor.

When starting a packet capture, you have the option to override some configured settings.

To start, stop, or cancel a capture, perform this task:

	Command	Purpose
Step 1	Router# monitor capture [buffer size <i>buf_size</i>] [length <i>cap_len</i>] [linear circular] [filter <i>acl_number</i> <i>acl_name</i>] { start [for count (packets seconds)] schedule at <i>time date</i> }	Starts a capture with optional run-time configuration changes. The capture can start immediately or it can start at a specified time and date. • The buffer size option overrides the configured or default capture buffer size. • The length option determines the number of bytes that will be captured from each packet. The range for <i>cap_len</i> is 0 to 9216 bytes; the default is 68 bytes. A value of 0 causes the entire packet to be captured. • The circular option specifies that the capture buffer will overwrite earlier entries once it fills. The linear option specifies that the capture will stop when the buffer fills. The default is linear. • The filter option applies the specified ACL. • The for option specifies that the capture will end after the specified number of seconds has elapsed or the specified number of packets has been captured.
Step 2	Router# monitor capture stop	Stops the capture.
Step 3	Router# monitor capture clear [filter]	Clears any run-time configuration settings, clears any pending scheduled capture, and clears the capture buffer. The filter option clears only the run-time filter settings.

When using these commands, note the following information:

- The format for *time* and *date* is hh:mm:ss dd mmm yyyy. The hour is specified in 24-hour notation, and the month is specified by a three-letter abbreviation. For example, to set a capture starting time of 7:30 pm on October 31, 2006, use the notation 19:30:00 31 oct 2006. The time zone is GMT.
- When you specify a capture filter ACL in the start command, the new ACL will not override any configured ACLs. The new ACL will execute in software.

Displaying and Exporting the Capture Buffer

To display the captured packets or information about the capture session, or to export the captured packets for analysis, perform this task:

	Command	Purpose
Step 1	Router# show monitor capture	Displays the capture session configuration.
Step 2	Router# show monitor capture status	Displays the capture session state, mode, and packet statistics.
Step 3	Router# show monitor capture buffer [<i>start</i> [<i>end</i>]] [detail][dump [<i>nowrap</i> [<i>dump_length</i>]]] [acl <i>acl_number</i> <i>acl_name</i>]]	Displays the capture buffer contents. - The <i>start</i> and <i>end</i> parameters specify packet number indices in the capture buffer. When a <i>start</i> index is specified with no <i>end</i> index, only the single packet at the <i>start</i> index is displayed. When both the <i>start</i> and <i>end</i> indices are specified, all packets between these indices are displayed. The range is 1 to 4294967295. - The detail option adds expanded and formatted protocol and envelope information for each packet, including the packet arrival time. - The dump option displays the hexadecimal contents of the packet. If <i>nowrap</i> is specified with <i>dump_length</i>, one line of hexadecimal packet content of <i>dump_length</i> characters will be displayed for each packet. If <i>dump_length</i> is not specified, a line of 72 characters will be displayed. The range of <i>dump_length</i> is 14 to 256. - The acl option causes the display of only those packets that match the specified ACL.
Step 4	Router# show monitor capture buffer [<i>start</i> [<i>end</i>]] brief [acl <i>acl_number</i> <i>acl_name</i>]]	Displays only packet header information.
Step 5	Router# monitor capture export buffer url	Copies the contents of the capture buffer to the specified file system or file transfer mechanism.

Configuration Examples for the Mini Protocol Analyzer

- [General Configuration Examples, page 22-8](#)
- [Filtering Configuration Examples, page 22-9](#)
- [Operation Examples, page 22-10](#)
- [Display Examples, page 22-10](#)

General Configuration Examples

This example shows how to minimally configure the Mini Protocol Analyzer:

```
Router#
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# monitor session 1 type capture
Router(config-mon-capture)# end
```

```
Router# show mon cap
Capture instance [1] :
=====
Capture Session ID : 1
Session status      : up
rate-limit value    : 10000
redirect index      : 0x807
buffer-size         : 2097152
capture state       : OFF
capture mode        : Linear
capture length      : 68
```

Router#

This example shows how to configure the buffer size, session description, and rate limit:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# monitor session 1 type capture
Router(config-mon-capture)# buffer-size 4096
Router(config-mon-capture)# description "Capture from ports, no filtering."
Router(config-mon-capture)# rate-limit 20000
Router(config-mon-capture)# end
```

```
Router#
Router# show monitor capture
Capture instance [1] :
=====
Capture Session ID : 1
Session status      : up
rate-limit value    : 20000
redirect index      : 0x807
buffer-size         : 4194304
capture state       : OFF
capture mode        : Linear
capture length      : 68
```

Router#

This example shows how to configure the source as a mixed list of ports:

```
Router(config-mon-capture)# source interface gig 3/1 - 3 , gig 3/5
```

This example shows how to configure the source as a mixed list of VLANs:

```
Router(config-mon-capture)# source vlan 123,234-245
```

Filtering Configuration Examples

This example shows how to configure for capturing packets with the following attributes:

- The packets belong to VLANs 123 or 234 through 245
- The packets are of 802.1Q EtherType (hexadecimal 0x8100, decimal 33024)
- The packet size is exactly 8192 bytes
- The source MAC address is 01:23:45:67:89:ab
- The packets conform to ACL number 99

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# monitor session 1 type capture
Router(config-mon-capture)# source vlan 123,234-245
Router(config-mon-capture)# filter ethertype 0x8100
Router(config-mon-capture)# filter length 8192
Router(config-mon-capture)# filter mac-address 0123.4567.89ab
Router(config-mon-capture)# filter access-group 99
Router(config-mon-capture)# end
```

```
Router# show monitor capture
Capture instance [1] :
=====
Capture Session ID : 1
Session status      : up
rate-limit value    : 20000
redirect index      : 0x7E07
Capture vlan        : 1019
buffer-size         : 4194304
capture state       : OFF
capture mode        : Linear
capture length      : 68
Sw Filters          :
    ethertype       : 33024
    src mac         : 0123.4567.89ab
    Hw acl          : 99
```

```
Router# show monitor session 1
Session 1
-----
Type                : Capture Session
Description          : capture from ports
Source VLANs        :
    Both             : 123,234-245
Capture buffer size  : 4096 KB
Capture rate-limit   :
    value            : 20000
Capture filters      :
    ethertype        : 33024
    src mac          : 0123.4567.89ab
    acl              : 99
```

```
Egress SPAN Replication State:
Operational mode      : Centralized
Configured mode       : Distributed (default)
```

```
Router#
```

This example shows how to capture packets whose size is less than 128 bytes:

```
Router(config-mon-capture)# filter length 0 128
```

This example shows how to capture packets whose size is more than 256 bytes:

```
Router(config-mon-capture)# filter length 256 9216
```

Operation Examples

This example shows how to start and stop a capture:

```
Router# monitor capture start
Router# monitor capture stop
Router#
```

This example shows how to start a capture to end after 60 seconds:

```
Router# monitor capture start for 60 seconds
Router#
```

This example shows how to start a capture at a future date and time:

```
Router# monitor capture schedule at 11:22:33 30 jun 2008
capture will start at : <11:22:33 UTC Mon Jun 30 2008> after 32465825 secs
Router#
```

This example shows how to start a capture with options to override the buffer size and to change to a circular buffer:

```
Router# monitor capture buffer size 65535 circular start
Router#
```

This example shows how to export the capture buffer to an external server and a local disk:

```
Router# monitor capture export buffer tftp://server/user/capture_file.cap
Router# monitor capture export buffer disk1:capture_file.cap
```

Display Examples

- [Displaying the Configuration, page 22-10](#)
- [Displaying the Capture Session Status, page 22-11](#)
- [Displaying the Capture Buffer Contents, page 22-12](#)

Displaying the Configuration

To display the capture session configuration, enter the **show monitor capture** command.

```
Router# show monitor capture
Capture instance [1] :
=====
Capture Session ID : 1
Session status      : up
rate-limit value    : 10000
redirect index       : 0x807
buffer-size          : 2097152
capture state        : OFF
capture mode         : Linear
```

```
capture length      : 68
```

This example shows how to display more details using the **show monitor session *n*** command:

```
Router# show monitor session 1
Session 1
-----
Type                : Capture Session
Source Ports        :
    Both            : Gi3/1-3,Gi3/5
Capture buffer size  : 32 KB
Capture filters      : None

Egress SPAN Replication State:
Operational mode     : Centralized
Configured mode      : Distributed (default)
```

This example shows how to display the full details using the **show monitor session *n* detail** command:

```
Router# show monitor session 1 detail
Session 1
-----
Type                : Capture Session
Description          : -
Source Ports        :
    RX Only         : None
    TX Only         : None
    Both            : Gi3/1-3,Gi3/5
Source VLANs        :
    RX Only         : None
    TX Only         : None
    Both            : None
Source RSPAN VLAN   : None
Destination Ports   : None
Filter VLANs        : None
Dest RSPAN VLAN     : None
Source IP Address    : None
Source IP VRF        : None
Source ERSPAN ID     : None
Destination IP Address : None
Destination IP VRF   : None
Destination ERSPAN ID : None
Origin IP Address    : None
IP QOS PREC          : 0
IP TTL               : 255
Capture dst_cpu_id   : 1
Capture vlan         : 0
Capture buffer size  : 32 KB
Capture rate-limit
    value            : 10000
Capture filters      : None

Egress SPAN Replication State:
Operational mode     : Centralized
Configured mode      : Distributed (default)
```

Displaying the Capture Session Status

To display the capture session status, enter the **show monitor capture status** command.

```
Router# show monitor capture status
capture state        : ON
```

```

capture mode          : Linear
Number of packets
    received : 253
    dropped  : 0
    captured : 90

```

Displaying the Capture Buffer Contents

To display the capture session contents, enter the **show monitor capture buffer** command. These examples show the resulting display using several options of this command:

```

Router# show monitor capture buffer
1      IP: s=10.12.0.5 , d=224.0.0.10, len 60
2      346  0180.c200.000e 0012.44d8.5000 88CC 020707526F7
3      60   0180.c200.0000 0004.c099.06c5 0026 42420300000
4      60   ffff.ffff.ffff 0012.44d8.5000 0806 00010800060
5      IP: s=7.0.84.23 , d=224.0.0.5, len 116
6      IP: s=10.12.0.1 , d=224.0.0.10, len 60

```

```

Router# show monitor capture buffer detail
1      Arrival time : 09:44:30 UTC Fri Nov 17 2006
      Packet Length : 74 , Capture Length : 68
      Ethernet II : 0100.5e00.000a 0008.a4c8.c038 0800
      IP: s=10.12.0.5 , d=224.0.0.10, len 60, proto=88
2      Arrival time : 09:44:31 UTC Fri Nov 17 2006
      Packet Length : 346 , Capture Length : 68
346    0180.c200.000e 0012.44d8.5000 88CC 020707526F757463031

```

```

Router# show monitor capture buffer dump
1      IP: s=10.12.0.5 , d=224.0.0.10, len 60
08063810:          0100 5E00000A          ..^...
08063820: 0008A4C8 C0380800 45C0003C 00000000  ..$H@8..E@.<....
08063830: 0258CD8F 0A0C0005 E000000A 0205EE6A  .XM.....`.....nj
08063840: 00000000 00000000 00000000 00000064  .....d
08063850: 0001000C 01000100 0000000F 0004          .....
2      346  0180.c200.000e 0012.44d8.5000 88CC 020707526F757465720415
3      60   0180.c200.0000 0004.c099.06c5 0026 42420300000000000800000
4      60   ffff.ffff.ffff 0012.44d8.5000 0806 0001080006040001001244
5      IP: s=7.0.84.23 , d=224.0.0.5, len 116
0806FCB0:          0100 5E000005          ..^...
0806FCC0: 0015C7D7 AC000800 45C00074 00000000  ..GW,...E@.t....
0806FCD0: 01597D55 07005417 E0000005 0201002C  .Y}U..T.`.....,
0806FCE0: 04040404 00000000 00000002 00000010  .....
0806FCF0: 455D8A10 FFFF0000 000A1201 0000          E].....

```

```

Router# show monitor capture buffer dump nowrap
1      74   0100.5e00.000a 0008.a4c8.c038 0800 45C0003C0000000
2      346  0180.c200.000e 0012.44d8.5000 88CC 020707526F7574
3      60   0180.c200.0000 0004.c099.06c5 0026 424203000000000
4      60   ffff.ffff.ffff 0012.44d8.5000 0806 000108000604000

```


**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

23

Restrictions for PFC QoS

- [General Guidelines](#), page 23-2
- [PFC and DFC Guidelines](#), page 23-4
- [Class Map Command Restrictions](#), page 23-5
- [Policy Map Class Command Restrictions](#), page 23-5
- [Supported Granularity for CIR and PIR Rate Values](#), page 23-5
- [Supported Granularity for CIR and PIR Token Bucket Sizes](#), page 23-6
- [IP Precedence and DSCP Values](#), page 23-7

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:

http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html

- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

General Guidelines

- With Release 15.0(1)SY1 and later releases, you can increase the supported number of QoS TCAM entries:

	TCAM Banks	QoS TCAM Entries	
		PFC4 Mode	PFC4XL Mode
Release 15.0(1)SY - Not configurable in Release 15.0(1)SY. - Default for 15.0(1)SY1 and later releases. - Command in Release 15.0(1)SY1 and later releases: <code>platform hardware acl reserve qos-banks 1</code>	1	16K	64K
Release 15.0(1)SY1 and later releases. Command: <code>platform hardware acl reserve qos-banks 2</code>	2	32K	128K

Changes to the supported number of QoS TCAM entries take effect after a reload. Enter the **show platform hardware acl global-config** command to display the QoS TCAM entry configuration:

```
Router# show platform hardware acl global-config | include [Bb]anks
Reserved QoS Banks:
  Current 1 banks
  Latest set 1 banks
  After next reload 1 banks
Router#
```

Enter the **show platform hardware pfc mode** command to display the PFC mode.

- PFC QoS supports IGMP, MLD, and PIM traffic.
- The **match ip precedence** and **match ip dscp** commands filter only IPv4 traffic.
- The **match precedence** and **match dscp** commands filter IPv4 and IPv6 traffic.
- The **set ip dscp** and **set ip precedence** commands are saved in the configuration file as **set dscp** and **set precedence** commands.
- PFC QoS supports the **set dscp** and **set precedence** policy map class commands for IPv4 and IPv6 traffic.
- The flowmask requirements of QoS, NetFlow, and NetFlow data export (NDE) might conflict, especially if you configure microflow policing.
- With egress ACL support for remarked DSCP and VACL capture both configured on an interface, VACL capture might capture two copies of each packet, and the second copy might be corrupt.
- You cannot configure egress ACL support for remarked DSCP on tunnel interfaces.
- Egress ACL support for remarked DSCP supports IP unicast traffic.
- Egress ACL support for remarked DSCP is not relevant to multicast traffic. PFC QoS applies ingress QoS changes to multicast traffic before applying [egress QoS](#).
- NetFlow and NetFlow data export (NDE) do not support interfaces where egress ACL support for remarked DSCP is configured.

- When egress ACL support for remarked DSCP is configured on any interface, you must configure an interface-specific flowmask to enable NetFlow and NDE support on interfaces where egress ACL support for remarked DSCP is not configured. Enter either the **platform flow ip interface-destination-source** or the **platform flow ip interface-full** global configuration mode command.
- Interface counters are not accurate on interfaces where egress ACL support for remarked DSCP is configured.
- You cannot apply microflow policing to traffic that has been permitted by egress ACL support for remarked DSCP.
- Traffic that has been permitted by egress ACL support for remarked DSCP cannot be tagged as MPLS traffic. (The traffic can be tagged as MPLS traffic on another network device.)
- When you apply both ingress policing and egress policing to the same traffic, both the input policy and the output policy must either mark down traffic or drop traffic. PFC QoS does not support ingress markdown with egress drop or ingress drop with egress markdown. (CSCea23571)
- If traffic is both aggregate and microflow policed, then the aggregate and microflow policers must both be in the same policy-map class and each must use the same **conform-action** and **exceed-action** keyword option: **drop**, **set-dscp-transmit**, **set-prec-transmit**, or **transmit**.
- You cannot configure PFC QoS features on tunnel interfaces.
- PFC QoS does not rewrite the payload ToS byte in tunnel traffic.
- PFC QoS filters only by ACLs, dscp values, or IP precedence values.
- For these commands, PFC QoS applies identical configuration to all LAN ports controlled by the same application-specific integrated circuit (ASIC):
 - **rcv-queue cos-map**
 - **wrr-queue cos-map**
- Except for WS-X6904-40G-2T, WS-X6908-10GE, WS-X6816-10T-2T, WS-X6716-10T, WS-X6816-10G-2T, WS-X6716-10GE, WS-X6704-10GE, WS-X6848-SFP-2T, WS-X6748-SFP, WS-X6824-SFP-2T, WS-X6724-SFP, WS-X6848-TX-2T, WS-X6748-GE-TX modules, PFC QoS applies identical configuration to all LAN ports controlled by the same application-specific integrated circuit (ASIC) for these commands:
 - **rcv-queue random-detect**
 - **rcv-queue queue-limit**
 - **wrr-queue queue-limit**
 - **wrr-queue bandwidth**
 - **priority-queue cos-map**
 - **wrr-queue threshold**
 - **rcv-queue threshold**
 - **wrr-queue random-detect**
 - **wrr-queue random-detect min-threshold**
 - **wrr-queue random-detect max-threshold**
- Configure these commands only on physical ports. Do not configure these commands on logical interfaces:
 - **priority-queue cos-map**
 - **wrr-queue cos-map**
 - **wrr-queue random-detect**
 - **wrr-queue random-detect max-threshold**
 - **wrr-queue random-detect min-threshold**

- **wrr-queue threshold**
- **wrr-queue queue-limit**
- **wrr-queue bandwidth**
- **rcv-queue cos-map**
- **rcv-queue bandwidth**
- **rcv-queue random-detect**
- **rcv-queue random-detect max-threshold**
- **rcv-queue random-detect min-threshold**
- **rcv-queue queue-limit**
- **rcv-queue cos-map**
- **rcv-queue threshold**

**Note**

IP multicast switching using egress packet replication is not compatible with QoS. In some cases, egress replication can result in the incorrect COS or DSCP marking of packets. If you are using QoS and your switching modules are capable of egress replication, enter the **platform ip multicast replication-mode ingress** command to force ingress replication.

PFC and DFC Guidelines

- The PFC and DFCs support QoS for IPv6 unicast and multicast traffic.
- To display information about IPv6 PFC QoS, enter the **show platform qos ipv6** command.
- The QoS features implemented in the port ASICs (queue architecture and dequeuing algorithms) support IPv4 and IPv6 traffic.
- The PFC and DFCs support IPv6 named extended ACLs and named standard ACLs.
- The PFC and DFCs support the **match protocol ipv6** command.
- With egress ACL support for remarked DSCP configured, the PFC and DFCs do not provide hardware-assistance for these features:
 - Cisco IOS reflexive ACLs
 - Network Address Translation (NAT)
- You cannot apply microflow policing to ARP traffic.
- The PFC and DFCs do not apply egress policing to traffic that is being bridged to the RP.
- The PFC and DFCs do not apply egress policing or egress DSCP mutation to multicast traffic from the RP.
- PFC QoS does not rewrite the ToS byte in bridged multicast traffic.
- The PFC and DFCs support up to 1022 aggregate policers, but some PFC QoS commands other than the **police** command will be included in this count. By default, any policy using a **set** or **trust** command will be included in the aggregate policer count. You can disable the addition of the **set** or **trust** commands to the aggregate policer count by entering the **no platform qos marking statistics** command, but you will then be unable to collect statistics for the classmaps associated with these commands. You can view the aggregate policer count in the QoS Policer Resources section of the output of the **show platform hardware capacity qos** command.

Class Map Command Restrictions

- PFC QoS supports a single **match** command in **class-map match-all** class maps, except that the **match protocol** command can be configured in a class map with the **match dscp** or **match precedence** command.
- PFC QoS supports multiple **match** commands in **class-map match-any** class maps.



Note PFC QoS supports a maximum of 9 commands that match DSCP or IP precedence values in class maps and ACLs.

- PFC QoS does not support these class map commands:
 - match classmap**
 - match destination-address**
 - match input-interface**
 - match source-address**

Policy Map Class Command Restrictions

PFC QoS does not support these policy map class commands:

- set qos-group**
- service-policy**

Supported Granularity for CIR and PIR Rate Values

CIR and PIR Rate Value Range	Granularity
32768 to 2097152 (2 Mbs)	32768 (32 Kb)
2097153 to 4194304 (4 Mbs)	65536 (64 Kb)
4194305 to 8388608 (8 Mbs)	131072 (128 Kb)
8388609 to 16777216 (16 Mbs)	262144 (256 Kb)
16777217 to 33554432 (32 Mbs)	524288 (512 Kb)
33554433 to 67108864 (64 Mbs)	1048576 (1 Mb)
67108865 to 134217728 (128 Mbs)	2097152 (2 Mb)
134217729 to 268435456 (256 Mbs)	4194304 (4 Mb)
268435457 to 536870912 (512 Mbs)	8388608 (8 Mb)
536870913 to 1073741824 (1 Gbs)	16777216 (16 Mb)
1073741825 to 2147483648 (2 Gbs)	33554432 (32 Mb)
2147483649 to 4294967296 (4 Gbs)	67108864 (64 Mb)
4294967297 to 8589934592 (8 Gbs)	134217728 (128 Mb)
8589934593 to 17179869184 (16 Gbs)	268435456 (256 Mb)

CIR and PIR Rate Value Range	Granularity
17179869185 to 34359738368 (32 Gbs)	536870912 (512 Mb)
34359738369 to 68719476736 (64 Gbs)	1073741824 (1024 Mb)

Within each range, PFC QoS programs the PFC with rate values that are multiples of the granularity values.

Supported Granularity for CIR and PIR Token Bucket Sizes

CIR and PIR Token Bucket Size Range	Granularity
1 to 32768 (32 KB)	1024 (1 KB)
32769 to 65536 (64 KB)	2048 (2 KB)
65537 to 131072 (128 KB)	4096 (4 KB)
131073 to 262144 (256 KB)	8196 (8 KB)
262145 to 524288 (512 KB)	16392 (16 KB)
524289 to 1048576 (1 MB)	32768 (32 KB)
1048577 to 2097152 (2 MB)	65536 (64 KB)
2097153 to 4194304 (4 MB)	131072 (128 KB)
4194305 to 8388608 (8 MB)	262144 (256 KB)
8388609 to 16777216 (16 MB)	524288 (512 KB)
16777217 to 33554432 (32 MB)	1048576 (1 MB)

Within each range, PFC QoS programs the PFC with token bucket sizes that are multiples of the granularity values.

IP Precedence and DSCP Values

3-bit IP Precedence	6 most significant bits of ToS						6-bit DSCP
	8	7	6	5	4	3	
0	0	0	0	0	0	0	0
	0	0	0	0	0	1	1
	0	0	0	0	1	0	2
	0	0	0	0	1	1	3
	0	0	0	1	0	0	4
	0	0	0	1	0	1	5
	0	0	0	1	1	0	6
	0	0	0	1	1	1	7
1	0	0	1	0	0	0	8
	0	0	1	0	0	1	9
	0	0	1	0	1	0	10
	0	0	1	0	1	1	11
	0	0	1	1	0	0	12
	0	0	1	1	0	1	13
	0	0	1	1	1	0	14
	0	0	1	1	1	1	15
2	0	1	0	0	0	0	16
	0	1	0	0	0	1	17
	0	1	0	0	1	0	18
	0	1	0	0	1	1	19
	0	1	0	1	0	0	20
	0	1	0	1	0	1	21
	0	1	0	1	1	0	22
	0	1	0	1	1	1	23
3	0	1	1	0	0	0	24
	0	1	1	0	0	1	25
	0	1	1	0	1	0	26
	0	1	1	0	1	1	27
	0	1	1	1	0	0	28
	0	1	1	1	0	1	29
	0	1	1	1	1	0	30
	0	1	1	1	1	1	31

3-bit IP Precedence	6 most significant bits of ToS						6-bit DSCP
	8	7	6	5	4	3	
4	1	0	0	0	0	0	32
	1	0	0	0	0	1	33
	1	0	0	0	1	0	34
	1	0	0	0	1	1	35
	1	0	0	1	0	0	36
	1	0	0	1	0	1	37
	1	0	0	1	1	0	38
	1	0	0	1	1	1	39
5	1	0	1	0	0	0	40
	1	0	1	0	0	1	41
	1	0	1	0	1	0	42
	1	0	1	0	1	1	43
	1	0	1	1	0	0	44
	1	0	1	1	0	1	45
	1	0	1	1	1	0	46
	1	0	1	1	1	1	47
6	1	1	0	0	0	0	48
	1	1	0	0	0	1	49
	1	1	0	0	1	0	50
	1	1	0	0	1	1	51
	1	1	0	1	0	0	52
	1	1	0	1	0	1	53
	1	1	0	1	1	0	54
	1	1	0	1	1	1	55
7	1	1	1	0	0	0	56
	1	1	1	0	0	1	57
	1	1	1	0	1	0	58
	1	1	1	0	1	1	59
	1	1	1	1	0	0	60
	1	1	1	1	0	1	61
	1	1	1	1	1	0	62
	1	1	1	1	1	1	63


Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

24

PFC QoS Overview

This chapter provides an overview of quality of service (QoS) as implemented on the Policy Feature Card (PFC) and Distributed Forwarding Cards (DFCs) in Cisco IOS Release 15.4SY. The term “PFC QoS” refers to hardware accelerated QoS. PFC QoS is implemented on various switch components in addition to the PFC and any DFCs.

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- For information about AutoQoS, see [Chapter 28, “AutoQoS.”](#)
- For information about QoS and MPLS, see [Chapter 29, “MPLS QoS.”](#)
- QoS in Cisco IOS Release 15.4SY (PFC QoS) uses some Cisco IOS modular QoS CLI (MQC). Because PFC QoS is implemented in hardware, it supports only a subset of the MQC syntax.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

QoS makes network performance more predictable and bandwidth utilization more effective. Without QoS, 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. QoS minimizes the chance that priority traffic would be dropped.

PFC QoS has these capabilities:

- Classification
- Marking
- Policing
- Congestion avoidance
- Mutation

You can configure two types of policies:

- Policies that define any combination of egress DSCP mutation, egress EXP mutation, classification, marking, and policing.
- Policies that define queueing.

Queueing configuration and classification, marking, and policing configuration are all optional.

Queueing configuration and classification, marking, and policing configuration have no common parameters or commands.

You configure QoS policies with the following commands:

- **class-map**—Defines a traffic class based on packet-matching criteria. Class maps are referenced in policy maps.
- **table-map**—Defines mapping from one set of traffic field values to another set of traffic fields values. Table maps are referenced in policy maps. Not used in queueing policies.
- **policy-map**—Defines or uses any combination of the following:
 - Class maps for classification
 - Table maps for marking and mutation
 - Policy trust mode
 - Policing
 - Queueing (cannot be combined with classification, marking, trust mode, or policing)
- **service-policy**—Applies policies to interfaces.

These chapters describe QoS:

- [Restrictions for PFC QoS, page 23-1](#)
- [Classification, Marking, and Policing, page 25-1](#)
- [Policy-Based Queueing, page 26-1](#)
- [QoS Global and Interface Options, page 27-1](#)



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

25

Classification, Marking, and Policing

- [Information About Classification, Marking, and Policing Policies, page 25-1](#)
- [How to Configure Classification, Marking, and Policing Policies, page 25-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Information About Classification, Marking, and Policing Policies

- [Classification, Marking, and Policing Policy Overview, page 25-1](#)
- [Traffic Classification, page 25-2](#)
- [Traffic Marking, page 25-3](#)
- [Information about Policing, page 25-4](#)

Classification, Marking, and Policing Policy Overview

The classification, marking, and policing configuration is defined by the policies attached to an interface. Classification, marking, and policing are unaffected by any QoS commands configured on ingress interfaces or by queueing policies.

Traffic Classification

Traffic classification allows the network to recognize traffic as it falls into classes that you have configured. Network traffic must be classified to apply specific QoS to it.

Classification can be inclusive (for example, all of the traffic in a Layer 2 VLAN or all of the traffic passing through an interface) or classification can be very specific (for example, you can use a class map with **match** commands that recognize specific aspects of the traffic).

You can classify and apply QoS (for example, marking) and then, on another interface or network device, classify again based on the marked value and apply other QoS.

The PFC and any DFCs supports a single **match** command in **class-map match-all** class maps, except that the **match protocol** command can be configured in a class map with the **match dscp** or **match precedence** command.

The PFC and any DFCs supports multiple **match** commands in **class-map match-any** class maps.



Note

PFC QoS supports a maximum of 9 commands that match DSCP or IP precedence values in class maps and ACLs.

Class maps can use the **match** commands listed in [Table 25-1](#) to configure a traffic class that is based on the match criteria.

Table 25-1 Traffic Classification Class Map match Commands and Match Criteria

match Commands	Direction	Match Criteria
match access-group { <i>access_list_number</i> name <i>access_list_name</i> }	Both	Access control list (ACL). Note Use ACLs to match the following: —CoS value —VLAN ID —Packet length
match any	Both	Any match criteria.
match cos	Ingress	CoS value.
match discard-class	Both	Discard class value.
match dscp Note The match protocol command can be configured in a class map with the match dscp command.	Both	DSCP value.
match l2 miss	Ingress	Layer 2 traffic flooded in a VLAN because it is addressed to a currently unlearned MAC-Layer destination address.
match mpls experimental topmost	Both	MPLS EXP value in the topmost label.
match precedence Note The match protocol command can be configured in a class map with the match precedence command.	Both	IP precedence values.

Table 25-1 Traffic Classification Class Map match Commands and Match Criteria (continued)

match Commands	Direction	Match Criteria
match protocol {arp ip ipv6}	Both	Protocol.
Note The match protocol command can be configured in a class map with the match dscp or match precedence command.		
match qos-group	Both	QoS group ID.

The PFC and any DFCs supports these ACL types for use with the **match access group** command:

Protocol	Numbered ACLs	Extended ACLs	Named ACLs
IPv4	Yes: 1 to 99 1300 to 1999	Yes: 100 to 199 2000 to 2699	Yes
IPv6	Not applicable	Yes (named)	Yes
MAC Layer	Not applicable	Not applicable	Yes
ARP	Not applicable	Not applicable	Yes

Traffic Marking



Note

Policing can also mark traffic.

Marking network traffic allows you to set or modify the attributes for a specific class of traffic, which allows class-based QoS features to recognize traffic classes based on the marking. There are two traffic marking methods:

- You can apply a configured value with a policy-map **set** command. [Table 25-2](#) lists the available policy-map **set** commands and the corresponding attribute.

Table 25-2 Configured-Value Policy-Map Commands

set Commands	Traffic Attribute	Ingress?	Egress?
set cos <i>cos_value</i>	Layer 2 CoS value	Y	Y
set dscp <i>dscp_value</i>	Layer 3 DSCP value	Y	Y
set precedence <i>precedence_value</i>	Layer 3 IP precedence value	Y	Y
set dscp tunnel <i>dscp_value</i>	Layer 3 DSCP value in the tunnel header of a Generic Routing Encapsulation (GRE) tunneled packet	Y	Y
set discard-class <i>discard_value</i>	Discard-class value	Y	Y
set qos-group <i>group_id</i>	QoS group ID	Y	Y
set mpls experimental imposition <i>exp_value</i>	MPLS experimental (EXP) field on all imposed label entries	Y	Y
set mpls experimental topmost <i>exp_value</i>	MPLS experimental (EXP) field on all topmost label entries	Y	Y

- You can apply a map to received values with a policy-map **set** command. [Table 25-3](#) lists the available policy-map **set** commands and the corresponding attribute.

Table 25-3 *Mapped-Value Policy-Map Commands*

set Commands	Map Name	Traffic Attribute	Ingress?	Egress?
set dscp cos	dscp-cos-map	Layer 3 DSCP value	Y	N
set dscp precedence	dscp-precedence-map	Layer 3 DSCP value	Y	N

Information about Policing

- [Overview of Policing, page 25-4](#)
- [Per-Interface Policers, page 25-5](#)
- [Aggregate Policers, page 25-5](#)
- [Microflow Policers, page 25-6](#)

Overview of Policing

You can configure policers to do the following:

- Mark traffic
- Limit bandwidth utilization and mark traffic

Policers can be applied to ingress and egress interfaces. Ingress policing is applied first, followed by egress policing.

Policing can rate limit incoming and outgoing traffic so that it adheres to the traffic forwarding rules defined by the QoS configuration. Sometimes these configured rules for how traffic should be forwarded through the system are referred to as a contract. If the traffic does not adhere to this contract, it is marked down to a lower DSCP value or dropped.

Policing does not buffer out-of-profile packets. As a result, policing does not affect transmission delay. In contrast, traffic shaping works by buffering out-of-profile traffic, which moderates traffic bursts. (PFC QoS does not support shaping.)

The PFC and DFCs support ingress and egress PFC QoS, which includes ingress and egress policing. Policers act on ingress traffic per-port or per-VLAN. For egress traffic, the policers act per-VLAN only.

Policing uses the Layer 2 frame size. You specify the bandwidth utilization limit as a committed information rate (CIR). You can also specify a higher peak information rate (PIR). Packets that exceed a rate are “out of profile” or “nonconforming.”

In each policer, you specify if out-of-profile packets are to be dropped or to have a new DSCP value applied to them (applying a new DSCP value is called “markdown”). Because out-of-profile packets do not retain their original priority, they are not counted as part of the bandwidth consumed by in-profile packets.

If you configure a PIR, the PIR out-of-profile action cannot be less severe than the CIR out-of-profile action. For example, if the CIR out-of-profile action is to mark down the traffic, then the PIR out-of-profile action cannot be to transmit the traffic.

For all policers, PFC QoS uses a configurable global table that maps the internal DSCP value to a marked-down DSCP value. When markdown occurs, PFC QoS gets the marked-down DSCP value from the table. You cannot specify marked-down DSCP values in individual policers.

**Note**

- By default, the markdown table is configured so that no markdown occurs: the marked-down DSCP values are equal to the original DSCP values. To enable markdown, configure the table appropriately for your network.
- When you apply both ingress policing and egress policing to the same traffic, both the input policy and the output policy must either mark down traffic or drop traffic. PFC QoS does not support ingress markdown with egress drop or ingress drop with egress markdown.

Per-Interface Policers

PFC QoS applies the bandwidth limits specified in a per-interface policer to matched traffic. For example, if you configure a per-interface policer to allow 1 Mbps for all matched TFTP traffic, it limits the TFTP traffic to 1 Mbps.

You define per-interface policers in a policy map class with the **police** command. If you attach per-interface policers to multiple ingress ports, each one polices the matched traffic on each ingress port separately.

Aggregate Policers

- [Aggregate Policer Overview, page 25-5](#)
- [Distributed Aggregate Policers, page 25-5](#)
- [Nondistributed Aggregate Policers, page 25-6](#)

Aggregate Policer Overview

PFC QoS applies the bandwidth limits specified in an aggregate policer cumulatively to all flows in matched traffic. For example, if you configure an aggregate policer to allow 1 Mbps for all TFTP traffic flows on VLAN 1 and VLAN 3, it limits the TFTP traffic for all flows combined on VLAN 1 and VLAN 3 to 1 Mbps.

You create named aggregate policers with the **platform qos aggregate-policer** command. If you attach a named aggregate policer to multiple ingress ports, it polices the matched traffic from all the ingress ports to which it is attached.

You can configure up to 1,023 aggregate policers; the configured aggregate policers can be applied to interfaces to configure up to 16,384 policer instances.

Distributed Aggregate Policers

When distributed aggregate policing is enabled, aggregate policers synchronize policing on interfaces supported by different DFC-equipped switching modules or the PFC. Distributed aggregate policing applies to the first 4,096 aggregate policer instances of these types:

- Aggregate policers applied to VLAN, tunnel, or port channel interfaces.
- Shared aggregate policers.
- Aggregate policers in egress policies.

With distributed aggregate policing enabled, aggregate policers in excess of the hardware-supported capacity function as nondistributed aggregate policers.

Nondistributed Aggregate Policers

Nondistributed aggregate policing works independently on each DFC-equipped switching module and independently on the PFC, which supports any non-DFC-equipped switching modules. Aggregate policing does not combine flow statistics from different DFC-equipped switching modules. You can display aggregate policing statistics for each DFC-equipped switching module and for the PFC and any non-DFC-equipped switching modules supported by the PFC.

Each PFC or DFC polices independently, which might affect QoS features being applied to traffic that is distributed across the PFC and any DFCs. Examples of these QoS feature are:

- Policers applied to a port channel interface.
- Policers applied to a switched virtual interface.
- Egress policers applied to either a Layer 3 interface or an SVI.

Policies affected by this restriction deliver an aggregate rate that is the sum of all the independent policing rates.

Microflow Policers

PFC QoS applies the bandwidth limit specified in a microflow policer separately to each flow in matched traffic. For example, if you configure a microflow policer to limit the TFTP traffic to 1 Mbps on VLAN 1 and VLAN 3, then 1 Mbps is allowed for each flow in VLAN 1 and 1 Mbps for each flow in VLAN 3. In other words, if there are three flows in VLAN 1 and four flows in VLAN 3, the microflow policer allows each of these flows 1 Mbps.

You can configure PFC QoS to apply the bandwidth limits in a microflow policer as follows:

- You can create microflow policers with up to 127 different rate and burst parameter combinations.
- You create microflow policers in a policy map class with the **police flow** command.
- You can configure a microflow policer to use only source addresses, which applies the microflow policer to all traffic from a source address regardless of the destination addresses.
- You can configure a microflow policer to use only destination addresses, which applies the microflow policer to all traffic to a destination address regardless of the source addresses.
- For MAC-Layer microflow policing, PFC QoS considers MAC-Layer traffic with the same protocol and the same source and destination MAC-Layer addresses to be part of the same flow, including traffic with different EtherTypes. You can configure MAC ACLs to filter IPX traffic.
- You cannot apply microflow policing to ARP traffic.
- Release 15.1(1)SY1 and later releases support Egress Microflow Destination-Only Policing. [Egress policing is per-VLAN](#), applied on either a Layer 3 interface or an SVI. With releases earlier than Release 15.1(1)SY1, policies attached with the **output** keyword do not support microflow policing.

You can include both an aggregate policer and a microflow policer in each policy map class to police a flow based on both its own bandwidth utilization and on its bandwidth utilization combined with that of other flows.



Note

If traffic is both aggregate and microflow policed, then the aggregate and microflow policers must both be in the same policy-map class and each must use the same **conform-action** and **exceed-action** keyword option: **drop**, **set-dscp-transmit**, **set-prec-transmit**, or **transmit**.

For example, you could create a microflow policer with a bandwidth limit suitable for individuals in a group, and you could create a named aggregate policer with bandwidth limits suitable for the group as a whole. You could include both policers in policy map classes that match the group's traffic. The combination would affect individual flows separately and the group aggregately.

For policy map classes that include both an aggregate and a microflow policer, PFC QoS responds to an out-of-profile status from either policer and, as specified by the policer, applies a new DSCP value or drops the packet. If both policers return an out-of-profile status, then if either policer specifies that the packet is to be dropped, it is dropped; otherwise, PFC QoS applies a marked-down DSCP value.

How to Configure Classification, Marking, and Policing Policies

- [Enabling Distributed Aggregate Policing, page 25-7](#)
- [Configuring a Class Map, page 25-8](#)
- [Configuring a Policy Map, page 25-9](#)
- [Attaching a Policy Map to an Interface, page 25-17](#)
- [Configuring Dynamic Per-Session Attachment of a Policy Map, page 25-19](#)

Enabling Distributed Aggregate Policing

To enable distributed aggregate policing, perform this task:

Command	Purpose
Router(config)# platform qos police distributed {strict loose}	Enables distributed aggregate policing. • The strict keyword prevents application of aggregate policers to interfaces that would exceed available hardware acceleration resources. • The loose keyword allows application of aggregate policers as nondistributed to interfaces that would exceed available hardware acceleration resources.

This example shows how to enable strict distributed aggregate policing globally:

```
Router(config)# platform qos police distributed strict
Router(config)#
```

This example shows how to disable distributed aggregate policing globally:

```
Router(config)# no platform qos police distributed
Router(config)#
```

Configuring a Class Map

- [Creating a Class Map, page 25-8](#)
- [Configuring Filtering in a Class Map, page 25-8](#)
- [Verifying Class Map Configuration, page 25-8](#)

Creating a Class Map

To create a class map, perform this task:

Command	Purpose
Router(config)# class-map [match-all match-any] <i>class_name</i>	Creates a class map. Note If you do not enter a match keyword, the default is match-all .

Configuring Filtering in a Class Map

To configure filtering in a class map, enter match commands as described in [Table 25-1, “Traffic Classification Class Map match Commands and Match Criteria”](#).

Verifying Class Map Configuration

To verify class map configuration, perform this task:

	Command	Purpose
Step 1	Router (config-cmap)# end	Exits configuration mode.
Step 2	Router# show class-map <i>class_name</i>	Verifies the configuration.

This example shows how to create a class map named **ipp5** and how to configure filtering to match traffic with IP precedence 5:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# class-map ipp5
Router(config-cmap)# match ip precedence 5
Router(config-cmap)# end
```

This example shows how to verify the configuration:

```
Router# show class-map ipp5
Class Map match-all ipp5 (id 1)
  Match ip precedence 5
```

Configuring a Policy Map

- [Policy Map Overview, page 25-9](#)
- [Creating a Policy Map, page 25-9](#)
- [Policy Map Class Configuration Guidelines and Restrictions, page 25-9](#)
- [Creating a Policy Map Class and Configuring Filtering, page 25-9](#)
- [Configuring Policy Map Class Actions, page 25-10](#)
- [Verifying Policy Map Configuration, page 25-16](#)

Policy Map Overview

Policy maps can contain one or more policy map classes, each with different policy map commands. Configure a separate policy map class in the policy map for each type of traffic that an interface receives. Put all commands for each type of traffic in the same policy map class. PFC QoS does not attempt to apply commands from more than one policy map class to matched traffic.

Creating a Policy Map

To create a policy map, perform this task:

Command	Purpose
Router(config)# policy-map <i>policy_name</i>	Creates a policy map.

Policy Map Class Configuration Guidelines and Restrictions

- PFC QoS does not support the **class** *class_name* **destination-address**, **class** *class_name* **input-interface**, **class** *class_name* **qos-group**, and **class** *class_name* **source-address** policy map commands.
- PFC QoS supports the **class default** policy map command.
- PFC QoS does not detect the use of unsupported commands until you attach a policy map to an interface.
- PFC QoS does not support multiple ACL matches per class.

Creating a Policy Map Class and Configuring Filtering

To create a policy map class and configure it to filter with a class map, perform this task:

Command	Purpose
Router(config-pmap)# class <i>class_name</i>	Creates a policy map class and configures it to filter with a class map. Note PFC QoS supports class maps that contain a single match command.

Configuring Policy Map Class Actions

- [Policy Map Class Action Restrictions, page 25-10](#)
- [Configuring Policy Map Class Marking, page 25-11](#)
- [Configuring the Policy Map Class Trust State, page 25-11](#)
- [Configuring Policy Map Class Policing, page 25-11](#)

Policy Map Class Action Restrictions

- Policy maps can contain one or more policy map classes.
- Put all commands for each type of traffic in the same policy map class.
- PFC QoS only applies commands from one policy map class to traffic. After traffic has matched the filtering in one policy map class, QoS does not apply the filtering configured in other policy map classes.
- For hardware-switched traffic, PFC QoS does not support the **bandwidth**, **priority**, **queue-limit**, or **random-detect** policy map class commands. You can configure these commands because they can be used for software-switched traffic.
- PFC QoS does not support the **set qos-group** policy map class commands.
- PFC QoS supports the **set ip dscp** and **set ip precedence** policy map class commands for IPv4 traffic.
 - You can use the **set ip dscp** and **set ip precedence** commands on non-IP traffic to mark the internal DSCP value, which is the basis of the egress Layer 2 CoS value.
 - The **set ip dscp** and **set ip precedence** commands are saved in the configuration file as **set dscp** and **set precedence** commands.
- PFC QoS supports the **set dscp** and **set precedence** policy map class commands for IPv4 and IPv6 traffic.
- You cannot do all three of the following in a policy map class:
 - Mark traffic with the **set** commands
 - Configure the trust state
 - Configure policing

In a policy map class, you can either mark traffic with the **set** commands or do one or both of the following:

- Configure the trust state
- Configure policing



Note

When configure policing, you can mark traffic with policing keywords.

Configuring Policy Map Class Marking

PFC QoS supports policy map class marking for all traffic with **set** policy map class commands. To configure policy map class marking, perform this task:

Command	Purpose
Router(config-pmap-c)# set { dscp <i>dscp_value</i> precedence <i>ip_precedence_value</i> }	Configures the policy map class to mark matched traffic with the configured DSCP or IP precedence value.

Configuring the Policy Map Class Trust State



Note

You cannot attach a policy map that configures a trust state with the **service-policy output** command.

To configure the policy map class trust state, perform this task:

Command	Purpose
Router(config-pmap-c)# trust { cos dscp ip-precedence }	Configures the policy map class trust state, which selects the value that PFC QoS uses as the source of the initial internal DSCP value.

When configuring the policy map class trust state, note the following information:

- Enter the **no trust** command to use the trust state configured on the ingress port (this is the default).
- With the **cos** keyword, PFC QoS sets the internal DSCP value from received or ingress port CoS.
- With the **dscp** keyword, PFC QoS uses received DSCP.
- With the **ip-precedence** keyword, PFC QoS sets DSCP from received IP precedence.

Configuring Policy Map Class Policing

- [Policy Map Class Policing Restrictions, page 25-12](#)
- [Using a Named Aggregate Policer, page 25-12](#)
- [Configuring a Per-Interface Policer, page 25-13](#)
- [Configuring a Per-Interface Microflow Policer, page 25-14](#)

Policy Map Class Policing Restrictions

- PFC QoS does not support the **set-qos-transmit** policer keyword.
- PFC QoS does not support the **set-dscp-transmit** or **set-prec-transmit** keywords as arguments to the **exceed-action** keyword.
- PFC QoS does not detect the use of unsupported keywords until you attach a policy map to an interface.
- Policing with the **conform-action transmit** keywords sets the port trust state of matched traffic to trust DSCP or to the trust state configured by a **trust** command in the policy map class.

Using a Named Aggregate Policer

To use a named aggregate policer, perform this task:

Command	Purpose
Router(config-pmap-c) # police aggregate <i>aggregate_name</i>	Configures the policy map class to use a previously defined named aggregate policer.

- When distributed aggregate policing is enabled, note the following information:
 - Distributed aggregate policers synchronize policing on interfaces supported by different DFC-equipped switching modules or the PFC. Distributed aggregate policing applies to the first 4,096 aggregate policers of these types:
 - Aggregate policers applied to VLAN, tunnel, or port channel interfaces.
 - Shared aggregate policers.
 - Aggregate policers in egress policies.
 - Distributed aggregate policers in excess of the hardware-supported capacity function as nondistributed aggregate policers.
- When distributed aggregate policing is not enabled, note the following information:
 - Aggregate policing works independently on each DFC-equipped switching module and independently on the PFC, which supports any non-DFC-equipped switching modules. Aggregate policing does not combine flow statistics from different DFC-equipped switching modules. You can display aggregate policing statistics for each DFC-equipped switching module and for the PFC and any non-DFC-equipped switching modules supported by the PFC.
 - Each PFC or DFC polices independently, which might affect QoS features being applied to traffic that is distributed across the PFC and any DFCs. Examples of these QoS feature are:
 - Policers applied to a port channel interface.
 - Policers applied to a switched virtual interface.
 - Egress policers applied to either a Layer 3 interface or an SVI.
 - Policers affected by this restriction deliver an aggregate rate that is the sum of all the independent policing rates.

Configuring a Per-Interface Policer

To configure a per-interface policer, perform this task:

Command	Purpose
Router(config-pmap-c)# police <i>bits_per_second</i> <i>normal_burst_bytes</i> [<i>maximum_burst_bytes</i>] [pir <i>peak_rate_bps</i>] [[conform-action { drop set-dscp-transmit <i>dscp_value</i> set-prec-transmit <i>ip_precedence_value</i> transmit }] exceed-action { drop policed-dscp transmit }] violate-action { drop policed-dscp transmit }]	Creates a per-interface policer and configures the policy-map class to use it.

- When you apply both ingress policing and egress policing to the same traffic, both the input policy and the output policy must either mark down traffic or drop traffic. PFC QoS does not support ingress markdown with egress drop or ingress drop with egress markdown.
- Policing uses the Layer 2 frame size.
- See the [“Restrictions for PFC QoS” section on page 23-1](#) for information about rate and burst size granularity.
- The valid range of values for the CIR *bits_per_second* parameter is as follows:
 - Minimum—32 kilobits per second, entered as 32000
 - Maximum—256 gigabits per second, entered as 256000000000
- The *normal_burst_bytes* parameter sets the CIR token bucket size.
- The *maximum_burst_bytes* parameter sets the PIR token bucket size.
- When configuring the size of a token bucket, note the following information:
 - Because the token bucket must be large enough to hold at least one frame, configure the token bucket size to be larger than the maximum size of the traffic being policed.
 - For TCP traffic, configure the token bucket size as a multiple of the TCP window size, with a minimum value at least twice as large as the maximum size of the traffic being policed.
 - The *maximum_burst_bytes* parameter must be set larger than the *normal_burst_bytes* parameter.
 - To sustain a specific rate, set the token bucket size to be at least the rate value divided by 2000.
 - The minimum token bucket size is 1 byte, entered as 1.
 - The maximum token bucket size is 512 megabytes, entered as 512000000.
- The valid range of values for the **pir** *bits_per_second* parameter is as follows:
 - Minimum—32 kilobits per second, entered as 32000 (the value cannot be smaller than the CIR *bits_per_second* parameters)
 - Maximum—256 gigabits per second, entered as 256000000000
- (Optional) You can specify a conform action for matched in-profile traffic as follows:
 - The default conform action is **transmit**, which sets the policy map class trust state to *trust DSCP* unless the policy map class contains a **trust** command.
 - To set PFC QoS labels in untrusted traffic, you can enter the **set-dscp-transmit** keyword to mark matched untrusted traffic with a new DSCP value or enter the **set-prec-transmit** keyword to mark matched untrusted traffic with a new IP precedence value. The **set-dscp-transmit** and **set-prec-transmit** keywords are only supported for IP traffic. PFC QoS sets egress ToS and CoS from the configured value.
 - You can enter the **drop** keyword to drop all matched traffic.

- Ensure that aggregate and microflow policers that are applied to the same traffic each specify the same conform-action behavior.
- (Optional) For traffic that exceeds the CIR, you can specify an exceed action as follows:
 - For marking without policing, you can enter the **transmit** keyword to transmit all matched out-of-profile traffic.
 - The default exceed action is **drop**, except with a *maximum_burst_bytes* parameter (**drop** is not supported with a *maximum_burst_bytes* parameter).



Note When the exceed action is **drop**, PFC QoS ignores any configured violate action.

- You can enter the **policed-dscp-transmit** keyword to cause all matched out-of-profile traffic to be marked down as specified in the markdown map.



Note When you create a policer that does not use the **pir** keyword and the *maximum_burst_bytes* parameter is equal to the *normal_burst_bytes* parameter (which is the case if you do not enter the *maximum_burst_bytes* parameter), the **exceed-action policed-dscp-transmit** keywords cause PFC QoS to mark traffic down as defined by the **policed-dscp max-burst** markdown map.

- (Optional) for traffic that exceeds the PIR, you can specify a violate action as follows:
 - For marking without policing, you can enter the **transmit** keyword to transmit all matched out-of-profile traffic.
 - The default violate action is equal to the exceed action.
 - You can enter the **policed-dscp-transmit** keyword to cause all matched out-of-profile traffic to be marked down as specified in the markdown map.

This example shows how to create a policy map named **max-pol-ipp5** that uses the class-map named **ipp5**, which is configured to trust received IP precedence values and is configured with a maximum-capacity aggregate policer and with a microflow policer:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# policy-map max-pol-ipp5
Router(config-pmap)# class ipp5
Router(config-pmap-c)# trust ip-precedence
Router(config-pmap-c)# police 2000000000 2000000 conform-action set-prec-transmit 6
exceed-action policed-dscp-transmit
Router(config-pmap-c)# police flow 10000000 10000 conform-action set-prec-transmit 6
exceed-action policed-dscp-transmit
Router(config-pmap-c)# end
```

Configuring a Per-Interface Microflow Policer

To configure a per-interface microflow policer, perform this task:

Command	Purpose
Router(config-pmap-c)# police flow [mask { src-only dest-only full-flow }] bits_per_second normal_burst_bytes [[[conform-action { drop set-dscp-transmit <i>dscp_value</i> set-prec-transmit <i>ip_precedence_value</i> transmit }] exceed-action { drop policed-dscp transmit }] violate-action { drop policed-dscp transmit }]	Creates a per-interface microflow policer and configures the policy-map class to use it.

- When you apply both ingress policing and egress policing to the same traffic, both the input policy and the output policy must either mark down traffic or drop traffic. PFC QoS does not support ingress markdown with egress drop or ingress drop with egress markdown.
- Policing uses the Layer 2 frame size.
- See the [“Restrictions for PFC QoS” section on page 23-1](#) for information about rate and burst size granularity.
- You can enter the **mask src-only** keywords to base flow identification only on source addresses, which applies the microflow policer to all traffic from each source address. PFC QoS supports the **mask src-only** keywords for both IP traffic and MAC traffic.
- You can enter the **mask dest-only** keywords to base flow identification only on destination addresses, which applies the microflow policer to all traffic to each source address. PFC QoS supports the **mask dest-only** keywords for both IP traffic and MAC traffic. Release 15.1(1)SY1 and later releases support Egress Microflow Destination-Only Policing. [Egress policing is per-VLAN](#), applied on either a Layer 3 interface or an SVI.
- By default and with the **mask full-flow** keywords, PFC QoS bases IP flow identification on source IP address, destination IP address, the Layer 3 protocol, and Layer 4 port numbers.
- PFC QoS considers MAC-Layer traffic with the same protocol and the same source and destination MAC-Layer addresses to be part of the same flow, including traffic with different EtherTypes.
- Microflow policers do not support the *maximum_burst_bytes* parameter, the **pir bits_per_second** keyword and parameter, or the **violate-action** keyword.



Note The flowmask requirements of microflow policing, NetFlow, and NetFlow data export (NDE) might conflict.

- The valid range of values for the CIR *bits_per_second* parameter is as follows:
 - Minimum—32 kilobits per second, entered as 32000
 - Maximum—256 gigabits per second, entered as 256000000000
- The *normal_burst_bytes* parameter sets the CIR token bucket size.
- When configuring the size of a token bucket, note the following information:
 - Because the token bucket must be large enough to hold at least one frame, configure the token bucket size to be larger than the maximum size of the traffic being policed.
 - For TCP traffic, configure the token bucket size as a multiple of the TCP window size, with a minimum value at least twice as large as the maximum size of the traffic being policed.
 - The *maximum_burst_bytes* parameter must be set larger than the *normal_burst_bytes* parameter.
 - To sustain a specific rate, set the token bucket size to be at least the rate value divided by 2000.
 - The minimum token bucket size is 1 byte, entered as 1.
 - The maximum token bucket size is 512 megabytes, entered as 512000000.

- (Optional) You can specify a conform action for matched in-profile traffic as follows:
 - The default conform action is **transmit**, which sets the policy map class trust state to *trust DSCP* unless the policy map class contains a **trust** command.
 - To set PFC QoS labels in untrusted traffic, you can enter the **set-dscp-transmit** keyword to mark matched untrusted traffic with a new DSCP value or enter the **set-prec-transmit** keyword to mark matched untrusted traffic with a new IP precedence value. The **set-dscp-transmit** and **set-prec-transmit** keywords are only supported for IP traffic. PFC QoS sets egress ToS and CoS from the configured value.
 - You can enter the **drop** keyword to drop all matched traffic.
 - Ensure that aggregate and microflow policers that are applied to the same traffic each specify the same conform-action behavior.
- (Optional) For traffic that exceeds the CIR, you can specify an exceed action as follows:
 - For marking without policing, you can enter the **transmit** keyword to transmit all matched out-of-profile traffic.
 - The default exceed action is **drop**, except with a *maximum_burst_bytes* parameter (**drop** is not supported with a *maximum_burst_bytes* parameter).



Note When the exceed action is **drop**, PFC QoS ignores any configured violate action.

- You can enter the **policed-dscp-transmit** keyword to cause all matched out-of-profile traffic to be marked down as specified in the markdown map.



Note When you create a policer that does not use the **pir** keyword and the *maximum_burst_bytes* parameter is equal to the *normal_burst_bytes* parameter (which is the case if you do not enter the *maximum_burst_bytes* parameter), the **exceed-action** **policed-dscp-transmit** keywords cause PFC QoS to mark traffic down as defined by the **policed-dscp max-burst** markdown map.

Verifying Policy Map Configuration

To verify policy map configuration, perform this task:

	Command	Purpose
Step 1	Router(config-pmap-c) # end	Exits policy map class configuration mode.
		Note Enter additional class commands to create additional classes in the policy map.
Step 2	Router# show policy-map <i>policy_name</i>	Verifies the configuration.

This example shows how to verify the configuration:

```
Router# show policy-map max-pol-ipp5
Policy Map max-pol-ipp5
  class ipp5

  class ipp5
    police flow 10000000 10000 conform-action set-prec-transmit 6 exceed-action
    policed-dscp-transmit
```

```

trust precedence
  police 2000000000 2000000 2000000 conform-action set-prec-transmit 6 exceed-action
  policed-dscp-transmit

Router#

```

Attaching a Policy Map to an Interface

To attach a policy map to an interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{ vlan <i>vlan_ID</i> } { <i>type slot/port[.subinterface]</i> } { port-channel <i>number[.subinterface]</i> }}	Selects the interface to configure.
Step 2	Router(config-if)# service-policy [input output] <i>policy_map_name</i>	Attaches a policy map to the interface.
Step 3	Router(config-if)# end	Exits configuration mode.

- Do not attach a service policy to a port that is a member of an EtherChannel.
- PFC QoS supports the **output** keyword only on Layer 3 interfaces (either LAN ports configured as Layer 3 interfaces or VLAN interfaces). You can attach both an input and an output policy map to a Layer 3 interface.
- VLAN-based or port-based PFC QoS on Layer 2 ports is not relevant to policies attached to Layer 3 interfaces with the **output** keyword.
- Release 15.1(1)SY1 and later releases support Egress Microflow Destination-Only Policing. [Egress policing is per-VLAN](#), applied on either a Layer 3 interface or an SVI. With releases earlier than Release 15.1(1)SY1, policies attached with the **output** keyword do not support microflow policing.
- You cannot attach a policy map that configures a trust state with the **service-policy output** command.
- Filtering based on IP precedence or DSCP in policies attached with the **output** keyword uses the received IP precedence or DSCP values. Filtering based on IP precedence or DSCP in policies attached with the **output** keyword is not based on any IP precedence or DSCP changes made by ingress QoS.
- A shared aggregate policer cannot be applied in both ingress and egress directions.
- When distributed aggregate policing is enabled, aggregate policers synchronize policing on interfaces supported by different DFC-equipped switching modules or the PFC. Distributed aggregate policing applies to the first 4,096 aggregate policer instances of these types:
 - Aggregate policers applied to VLAN, tunnel, or port channel interfaces.
 - Shared aggregate policers.
 - Aggregate policers in egress policies.

With distributed aggregate policing enabled, aggregate policers in excess of the hardware-supported capacity function as nondistributed aggregate policers.

- Nondistributed aggregate policing works independently on each DFC-equipped switching module and independently on the PFC, which supports any non-DFC-equipped switching modules. Aggregate policing does not combine flow statistics from different DFC-equipped switching modules. You can display aggregate policing statistics for each DFC-equipped switching module and for the PFC and any non-DFC-equipped switching modules supported by the PFC.

Each PFC or DFC polices independently, which might affect QoS features being applied to traffic that is distributed across the PFC and any DFCs. Examples of these QoS feature are:

- Policers applied to a port channel interface.
- Policers applied to a switched virtual interface.
- Egress policers applied to either a Layer 3 interface or an SVI.

Policers affected by this restriction deliver an aggregate rate that is the sum of all the independent policing rates.

- For nonaggregate policers, each PFC or DFC polices independently, which might affect QoS features being applied to traffic that is distributed across the PFC and any DFCs. Examples of these QoS feature are:

- Policers applied to a port channel interface.
- Policers applied to a switched virtual interface.
- Egress policers applied to either a Layer 3 interface or an SVI.

Policers affected by this restriction deliver an aggregate rate that is the sum of all the independent policing rates.

- When you apply both ingress policing and egress policing to the same traffic, both the input policy and the output policy must either mark down traffic or drop traffic. PFC QoS does not support ingress markdown with egress drop or ingress drop with egress markdown.

This example shows how to attach the policy map named **pmap1** to gigabit Ethernet port 5/36:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/36
Router(config-if)# service-policy input pmap1
Router(config-if)# end
```

This example shows how to verify the configuration:

```
Router# show policy-map interface gigabitethernet 5/36
gigabitethernet5/36
  service-policy input: pmap1
    class-map: cmap1 (match-all)
      0 packets, 0 bytes
      5 minute rate 0 bps
      match: ip precedence 5
    class cmap1
      police 8000 8000 conform-action transmit exceed-action drop
    class-map: cmap2 (match-any)
      0 packets, 0 bytes
      5 minute rate 0 bps
      match: ip precedence 2
        0 packets, 0 bytes
        5 minute rate 0 bps
    class cmap2
      police 8000 10000 conform-action transmit exceed-action drop
Router#
```

Configuring Dynamic Per-Session Attachment of a Policy Map

- [Policy Map Dynamic Per-Session Attachment Prerequisites](#), page 25-19
- [Defining and Associating Policy Maps](#), page 25-19

Policy Map Dynamic Per-Session Attachment Prerequisites

- Define the ingress and [egress QoS](#) policy maps to be assigned when users are authenticated.
- Configure identity policies to specify the policy maps to be assigned.
- In the user profiles on the RADIUS server, configure the Cisco vendor-specific attributes (VSAs) to specify which ingress and egress QoS policy maps will be assigned to each user.

Defining and Associating Policy Maps

To define the policy maps and associate them with an identity policy, follow these steps:

	Command	Purpose
Step 1	Router(config)# policy-map <i>in_policy_name</i>	Configures an ingress QoS policy map.
Step 2	Router(config-pmap)# class <i>class_map_name</i> ...	Configures policy map class.
Step 3	Router(config-pmap-c)# exit	Exits policy map class configuration submenu.
Step 4	Router(config)# policy-map <i>out_policy_name</i>	Configures an egress QoS policy map.
Step 5	Router(config-pmap)# class <i>class_map_name</i> ...	Configures policy map class.
Step 6	Router(config-pmap-c)# exit	Exits policy map class configuration submenu.
Step 7	Router(config)# identity policy <i>policy1</i>	Creates an identity policy, and enters identity policy configuration submenu.
Step 8	Router(config-identity-policy)# service-policy type qos input <i>in_policy_name</i>	Associates the ingress QoS policy map with this identity.
Step 9	Router(config-identity-policy)# service-policy type qos output <i>out_policy_name</i>	Associates the egress QoS policy map with this identity.
Step 10	Router(config-identity-policy)# end	Exits identity policy configuration submenu and returns to privileged EXEC mode.

To remove the identity policy, use the **no identity policy *policy_name*** command.

After the policy maps have been defined, configure the Cisco AV pair attributes in each user profile on the RADIUS server using the policy map names:

- cisco-avpair = "ip:sub-policy-In=*in_policy_name*"
- cisco-avpair = "ip:sub-policy-Out=*out_policy_name*"

To set the Cisco AV pair attributes on the RADIUS server, perform the following task:

Command or Action	Purpose
sub-policy-In = <i>in_policy_name</i> sub-policy-Out = <i>out_policy_name</i>	Enters the two Cisco AV pairs for service policy on the RADIUS server in the user file. When the switch requests the policy name, this information in the user file is supplied. A RADIUS user file contains an entry for each user that the RADIUS server will authenticate. Each entry, which is also referred to as a <i>user profile</i>, establishes an attribute the user can access. In this example, you have configured a service policy that attaches a QoS policy map to the interface and specifies the direction (inbound for data packets traveling into the interface or outbound for data packets leaving the interface). The policy map applied in the inbound direction is <i>example_in_qos</i> and the outbound policy map is <i>example_out_qos</i>.

This example shows the configuration in the user file on the RADIUS server:

```
userid      Password ="cisco"
  Service-Type = Framed,
  Framed-Protocol = PPP,
  cisco-avpair = "sub-policy-In=example_in_qos",
  cisco-avpair = "sub-policy-Out=example_out_qos"
```

This example shows the output of the **show epm session summary** command when a session is active:

```
Router# show epm session summary

EPM Session Information
-----
Total sessions seen so far : 5
Total active sessions      : 1
Session IP Address         : 192.0.2.1
-----
```

This example shows the output of the **show epm session ip ip_addr** command when a session is active on the interface with IP address 192.0.2.1:

```
Router# show epm session ip 192.0.2.1

Admission feature      : AUTHPROXY
AAA Policies           :
Input Service Policy   : in_policy_name
Output Service Policy  : out_policy_name
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

26

Policy-Based Queueing

- [Prerequisites for Policy-Based Queueing, page 26-1](#)
- [Restrictions for Policy-Based Queueing, page 26-2](#)
- [Information About Policy-Based Queueing, page 26-4](#)
- [How to Configure Policy-Based Queueing, page 26-10](#)
- [Configuration Examples for Policy-Based Queueing, page 26-18](#)
- [Queueing on FEX, page 26-45](#)

**Note**

- Queueing is optional. Use the commands described in this section to configure queueing on ports that serve congested links.
- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Policy-Based Queueing

None.

Restrictions for Policy-Based Queueing

- Ports on the WS-X6904-40G-2T switching modules can be configured with two priority queues (2p6q4t). Priority queue 1 has higher priority than priority queue 2.
- The WS-X6904-40G-2T egress queues support QoS shaping.
 - Shaping buffers traffic that exceeds the total egress port bandwidth. You can configure shaping on the egress queues and for egress traffic on the port.
 - You configure shaping as a percentage of the total egress port bandwidth.
 - Shaping can be configured only on the egress queues of WS-X6904-40G-2T ports. You cannot configure shaping on a port-channel interface.
 - Ports with shaping configured can be members of an EtherChannel, but shaping can cause the egress traffic rate to vary significantly and cause significant differences between the member ports.
 - You configure shaping with the **shape average percent** command. The **show running-configuration** command displays the **shape average percent** command, but the **show policy-map type lan-queueing** command displays the shaping configuration with the **cir percentage** command under the heading “Average Rate Traffic Shaping” heading.

See the “[Configuring Policy-Based Queueing on WS-X6904-40G-2T Ports](#)” section on page 26-12.

- DSCP-based queueing is supported on 8q4t, 1p7q2t, 2p6q4t, and 1p7q4t ports. The Supervisor Engine 2T-10GE ports are 8q4t/1p7q4t with the **platform qos 10g-only** global configuration command configured, which disables the Gigabit Ethernet ports on the supervisor engine.
- The presence of a **match dscp** command or a **match precedence** command in a class map that is used in a queueing policy attached to a port enables DSCP-based queueing on the port in the direction that the queueing policy is attached.
- CoS-based queueing is always used for non-IP traffic, IP multicast traffic, and IP unknown unicast flood traffic.
- Class maps that are used in queueing policies can contain any combination and number of **match dscp**, **match precedence**, or **match cos** commands.
- You can attach one input and one output queueing policy to an interface, in addition to policies that configure marking or policing. To support per-queue shaping on WS-X6904-40G-2T egress nonprioritynonpriority queues, the output queueing policy on WS-X6904-40G-2T port egress queues supports child policies.
- To support migration from Cisco IOS Release 12.2SX configurations, Cisco IOS Release 15.4SY supports global configuration mode and interface configuration mode queueing commands.
 - When you attach an ingress or egress queueing policy to a port, all interface configuration mode queueing commands on the port are deleted.
 - An attached ingress or egress queueing policy supercedes the effect of any configured global configuration mode queueing commands.
 - If you attach a queueing policy in only one direction, the queueing configuration of the other direction is either the defaults or is defined by any configured global configuration mode queueing commands.
 - You cannot configure any interface configuration mode queueing commands on a port that has a queueing policy attached.

- Policy-based queueing is supported on all modules that Cisco IOS Release 15.4SY supports with the Supervisor Engine 2T-10GE.
- Queueing policies are specific to particular port types because a queueing policy cannot contain any commands that are not supported by the port to which it will be attached (see the [“Configuration Examples for Policy-Based Queueing”](#) section on page 26-18). Commands not supported on a port are not ignored. You cannot successfully apply a policy with unsupported commands to a port.
- For clarity, configure queueing policy names that correspond to the port type it supports; for example **1q2t_1q8t_ingress**.
- Queueing policies can contain multiple class maps; each class map configures a queue.
- A policy-map class defined by a **class** command that uses a class map and that contains the policy-map class **priority** command configures the priority queue. (The priority queue is the highest-numbered queue.) The class-map filters the QoS values (CoS or DSCP) it is configured with to the priority queue.
- Enabling SRR on a port disables the priority queue.
- A policy-map class defined by a **class** command that uses a class map and that does not contain the policy-map class **priority** command configures the highest-numbered nonpriority queue. (If the port has a priority queue, the highest-numbered nonpriority queue is numbered one less than the priority queue. If the port does not have a priority queue, the highest-numbered queue is a nonpriority queue.) Subsequent such commands define the configuration of the remaining nonpriority queues in reverse numerical order. You cannot skip a queue, but configuration of all queues is not required. The class-map filters the QoS values it is configured with to the nonpriority queue being configured.
- A **class** command that uses the **class-default** keywords configures queue 1. The **class-default** keywords filter all remaining QoS values to queue 1.
- For each nonpriority queue, policy-map class **queue-buffers** or **random-detect** commands assign QoS values (CoS or DSCP) to thresholds within a queue. Thresholds are configured in numerical order. You cannot skip a threshold, but configuration of all threshold is not required. QoS values that are to be applied to the thresholds must be from the group of values that the class-map filters to the queue.
 - The first **queue-buffers** or **random-detect** command assigns QoS values to the first threshold, and configures the percentage value applied to it.
 - Subsequent **queue-buffers** or **random-detect** commands that are configured with the same percentage value assign additional QoS values to the first threshold.
 - The next **queue-buffers** or **random-detect** command that is configured with a different percentage value assigns QoS values to the numerically next threshold, and configures the percentage value applied to it.
 - Subsequent **queue-buffers** or **random-detect** commands that are configured with the same percentage value assign additional QoS values to the numerically next threshold.
 - Each **queue-buffers** or **random-detect** command with a different percentage defines the next unconfigured threshold and any subsequent commands that repeat a percentage value assign additional QoS values to a configured threshold.
 - All unconfigured thresholds are at 100%.
 - All unassigned QoS labels are assigned to the highest-numbered threshold.

Information About Policy-Based Queueing

- [Port-Based Queue Types, page 26-4](#)
- [Queueing Policies, page 26-8](#)

Port-Based Queue Types

- [Ingress and Egress Buffers and Queues, page 26-4](#)
- [Ingress Queue Types, page 26-5](#)
- [Egress Queue Types, page 26-6](#)
- [Module to Queue Type Mappings, page 26-7](#)

Ingress and Egress Buffers and Queues

**Note**

The WS-X6904-40G-2T egress queues support QoS [shaping](#).

The Ethernet port ASICs have buffers that are divided into a fixed number of queues. When congestion avoidance is enabled, PFC QoS uses the traffic's Layer 2 CoS value or, on some port types, the Layer 3 DSCP values, to assign traffic to the queues. The buffers and queues store frames temporarily as they transit the switch. PFC QoS allocates the port ASIC memory as buffers for each queue on each port.

The Ethernet ports support the following types of queues:

- Nonpriority queues
- Priority queues

The Ethernet ports support the following types of scheduling algorithms between queues:

- Shaped round robin (SRR)—SRR allows a queue to use only the allocated bandwidth.
- Deficit weighted round robin (DWRR)—DWRR keeps track of any lower-priority queue under-transmission caused by traffic in a higher-priority queue and compensates in the next round.
- Weighted Round Robin (WRR)—WRR does not explicitly reserve bandwidth for the queues. Instead, the amount of bandwidth assigned to each queue is user configurable. The percentage or weight allocated to a queue defines the amount of bandwidth allocated to the queue.
- Priority queueing—Strict priority queueing allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued, giving delay-sensitive data preferential treatment over other traffic. The switch services traffic in the strict-priority transmit queue before servicing the nonpriority queues. After transmitting a packet from a nonpriority queue, the switch checks for traffic in the strict-priority queue. If the switch detects traffic in the strict-priority queue, it suspends its service of the nonpriority queue and completes service of all traffic in the strict-priority queue before returning to the nonpriority queue.

The Ethernet ports provide congestion avoidance with these types of thresholds within a queue:

- Weighted Random Early Detection (WRED)—On ports with WRED drop thresholds, frames with a given QoS label are admitted to the queue based on a random probability designed to avoid buffer congestion. The probability of a frame with a given QoS label being admitted to the queue or discarded depends on the weight and threshold assigned to that QoS label.

For example, if CoS 2 is assigned to queue 1, threshold 2, and the threshold 2 levels are 40 percent (low) and 80 percent (high), then frames with CoS 2 will not be dropped until queue 1 is at least 40 percent full. As the queue depth approaches 80 percent, frames with CoS 2 have an increasingly higher probability of being discarded rather than being admitted to the queue. Once the queue is over 80 percent full, all CoS 2 frames are dropped until the queue is less than 80 percent full. The frames the switch discards when the queue level is between the low and high thresholds are picked out at random, rather than on a per-flow basis or in a FIFO manner. This method works well with protocols such as TCP that can adjust to periodic packet drops by backing off and adjusting their transmission window size.

- **Tail-drop thresholds**—On ports with tail-drop thresholds, frames with a given QoS label are admitted to the queue until the drop threshold associated with that QoS label is exceeded; subsequent frames of that QoS label are discarded until the threshold is no longer exceeded. For example, if CoS 1 is assigned to queue 1, threshold 2, and the threshold 2 watermark is 60 percent, then frames with CoS 1 will not be dropped until queue 1 is 60 percent full. All subsequent CoS 1 frames will be dropped until the queue is less than 60 percent full. With some port types, you can configure the nonpriority receive queue to use both a tail-drop and a WRED-drop threshold by mapping a CoS value to the queue or to the queue and a threshold. The switch uses the tail-drop threshold for traffic carrying CoS values mapped only to the queue. The switch uses WRED-drop thresholds for traffic carrying CoS values mapped to the queue and a threshold. All LAN ports of the same type use the same drop-threshold configuration.



Note

- You can enable DSCP-based queues and thresholds on 8q4t, 1p7q2t, and 1p7q4t ports (see the [“Module to Queue Type Mappings”](#) section on page 26-7), either in a queueing policy or with legacy interface commands (see the [“Legacy Configuration Procedures for DSCP-Based Queue Mapping”](#) section on page 36-14).
- DSCP-based queueing is supported on 8q4t, 1p7q2t, and 1p7q4t ports. The Supervisor Engine 2T-10GE ports are 8q4t/1p7q4t with the **platform qos 10g-only** global configuration command configured. To configure DSCP-based queue mapping on Supervisor Engine 2T ports, you must enter **shutdown** interface configuration mode commands for the Supervisor Engine 2T Gigabit Ethernet ports, and then enter the **platform qos 10g-only** global configuration command, which disables the Gigabit Ethernet ports on the Supervisor Engine 2T.

The switch provides congestion avoidance with the combination of multiple queues and the scheduling algorithms associated with each queue.

Ingress Queue Types

To see the queue structure of a LAN port, enter the **show queueing interface** *type slot/port* | **include type** command. The command displays one of the following architectures:

- **1q2t** indicates one nonpriority queue with one configurable tail-drop threshold and one nonconfigurable tail-drop threshold.
- **2q4t** indicates two nonpriority queues, each with four configurable tail-drop thresholds.
- **2q8t** indicates two nonpriority queues, each with eight configurable tail-drop thresholds.
- **8q4t** indicates eight nonpriority queues, each with four thresholds, each configurable as either WRED-drop or tail-drop, with support for DSCP-based queueing.
- **8q8t** indicates eight nonpriority queues, each with eight thresholds, each configurable as either WRED-drop or tail-drop.

- **1p1q4t** indicates:
 - One strict-priority queue
 - One nonpriority queue with four configurable tail-drop thresholds.
- **1p7q2t** indicates the following:
 - One strict-priority queue
 - Seven nonpriority queues, each with two thresholds, each threshold configurable as either WRED-drop or tail-drop
 - Supports DSCP-based queueing
- **2p6q4t** indicates the following:
 - Two strict-priority queues
 - Six nonpriority queues, each with four thresholds, each threshold configurable as either WRED-drop or tail-drop
 - Supports DSCP-based queueing
- **1p7q4t** indicates the following:
 - One strict-priority queue
 - Seven nonpriority queues, each with four thresholds, each threshold configurable as either WRED-drop or tail-drop
 - Supports DSCP-based queueing

Egress Queue Types

To see the queue structure of an egress LAN port, enter the **show queueing interface** *type slot/port* | **include type** command. The command displays one of the following architectures:

- **1p3q8t** indicates the following:
 - One strict-priority queue
 - Three nonpriority queues, each with eight thresholds, each threshold configurable as either WRED-drop or tail-drop
- **2p6q4t** indicates the following:
 - Two strict-priority queues
 - Six nonpriority queues, each with four thresholds, each threshold configurable as either WRED-drop or tail-drop
 - Supports DSCP-based queueing
 - [Supports per-queue shaping and per-port shaping](#)
- **1p7q4t** indicates the following:
 - One strict-priority queue
 - Seven nonpriority queues, each with four thresholds, each threshold configurable as either WRED-drop or tail-drop
 - Supports DSCP-based queueing
 - [On WS-X6904-40G-2T, supports per-queue shaping and per-port shaping](#)

- **1p7q8t** indicates the following:
 - One strict-priority queue
 - Seven nonpriority queues, each with eight thresholds, each threshold configurable as either WRED-drop or tail-drop

Module to Queue Type Mappings

- [Table 26-1—Supervisor Engine Module QoS Queue Structures](#)
- [Table 26-2—40-Gigabit Ethernet Modules](#)
- [Table 26-3—10-Gigabit Ethernet Modules](#)
- [Table 26-4—Gigabit and 10/100/1000 Ethernet Modules](#)

Table 26-1 Supervisor Engine Module QoS Queue Structures

Supervisor Engines	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
VS-S2T-10G-XL, VS-S2T-10G							
With Gigabit Ethernet ports enabled	2q4t	WRR	1p3q4t	DWRR or SRR		128 MB	112 MB
	Does not support DSCP-based queueing.						
With Gigabit Ethernet ports disabled	8q4t	WRR	1p7q4t	DWRR or SRR			
	• Supports DSCP-based queueing.						

Table 26-2 40-Gigabit Ethernet Modules

Modules	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-X6904-40G-2TXL, WS-X6904-40G-2T (supports DSCP-based queueing; supports egress queue shaping)	1p7q4t or 2p6q4t	DWRR	1p7q4t or 2p6q4t	DWRR	See the Catalyst 6900 Series data sheet .		
	Note On WS-X6904-40G-2T ports, the bandwidth command must be configured if you configure any nondefault values for any other queueing commands on the port. (CSCtz05347)						

Table 26-3 10-Gigabit Ethernet Modules

Modules	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-X6908-10GE (supports DSCP-based queueing)	8q4t	DWRR	1p7q4t	DWRR SRR	200 MB	108 MB	90 MB

Table 26-3 10-Gigabit Ethernet Modules (continued)

Modules	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-X6816-10T-2T, WS-X6716-10T, WS-X6816-10G-2T, WS-X6716-10GE (supports DSCP-based queueing)							
Performance mode	8q4t	DWRR	1p7q4t	DWRR SRR	198 MB	108 MB per port	90 MB per port
Oversubscription mode	1p7q2t	DWRR	1p7q4t	DWRR SRR	91 MB	90 MB per port	1 MB per port group
WS-X6704-10GE	8q8t	WRR	1p7q8t	DWRR	16 MB	2 MB	14 MB

Table 26-4 Gigabit and 10/100/1000 Ethernet Modules

Modules	Ingress Queue and Drop Thresholds	Ingress Queue Scheduler	Egress Queue and Drop Thresholds	Egress Queue Scheduler	Total Buffer Size	Ingress Buffer Size	Egress Buffer Size
WS-X6848-TX-2T, WS-X6748-GE-TX, WS-X6848-SFP-2T, WS-X6748-SFP, WS-X6824-SFP-2T, WS-X6724-SFP							
	2q8t	WRR	1p3q8t	DWRR	1.3 MB	166 KB	1.2 MB

Queueing Policies

Queueing policies use class-maps with **match** commands (see [Table 26-5](#)) and policy maps with scheduling and congestion management commands (see [Table 26-6](#)).

Table 26-5 Queueing Policy Class Map match Commands and Match Criteria

match Commands	Match Criteria
match cos <i>cos_list</i>	CoS values.
match dscp <i>dscp_list</i>	DSCP values.
match precedence <i>precedence_list</i>	Precedence values.

Note

- The presence of a **match dscp** command or a **match precedence** command in a class map that is used in a queueing policy attached to a port configures DSCP-based queueing in the direction of the policy (ingress or egress) on that port.
- CoS-based queueing is always used for non-IP traffic, IP multicast traffic, and IP unknown unicast flood traffic.
- Class maps that are used in queueing policies can contain both **match dscp** commands and **match cos** commands.

Table 26-6 Queueing Policy-Map Class Commands

Queueing Commands	Description
<code>bandwidth [remaining] percent percentage</code>	Allocates bandwidth between nonpriority queues. The remaining keyword is required on ports that have a priority queue.
<code>shape average percent percentage</code>	On WS-X6904-40G-2T ports, configures shaping on egress priority queues. Enables SRR on other ports, which allocates limited bandwidth between nonpriority egress queues (see “SRR” in the “Module to Queue Type Mappings” section on page 26-7).
<code>priority</code>	Applies a policy-map class to a priority queue.
<code>priority level {1 2} [percent percentage]</code>	(Required to configure multiple priority queues or shaping.) On WS-X6904-40G-2T ports, applies a policy-map class to one of the priority queues. Priority queue 1 has higher priority than priority queue 2. The optional percent percentage keyword and argument configure shaping on egress priority queues.
<code>queue-buffers ratio weight</code>	Sets the queue size.
<code>queue-limit multiple-type-based</code>	Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
<code>queue-limit cos cos_value percent percent_of_qsize</code>	Applies one CoS value to a tail-drop threshold and configures the threshold percentage.
<code>queue-limit dscp dscp_value percent percent_of_qsize</code>	Applies one DSCP value to a tail-drop threshold and configures the threshold percentage.
<code>queue-limit precedence precedence_value percent percent_of_qsize</code>	Applies one precedence value to a tail-drop threshold and configures the threshold percentage.
<code>queue-limit cos values cos_list percent percent_of_qsize</code>	Applies multiple CoS values to a tail-drop threshold and configures the threshold percentage.
<code>queue-limit dscp values dscp_list percent percent_of_qsize</code>	Applies multiple DSCP values to a tail-drop threshold and configures the threshold percentage.
<code>queue-limit precedence values precedence_list percent percent_of_qsize</code>	Applies multiple precedence values to a tail-drop threshold and configures the threshold percentage.

Table 26-6 Queueing Policy-Map Class Commands (continued)

Queueing Commands	Description
<code>random-detect cos-based [aggregate]</code>	Enables application of CoS values to a WRED-drop threshold. The aggregate keyword enables use of the values keyword.
<code>random-detect dscp-based [aggregate]</code>	Enables application of DSCP values to a WRED-drop threshold. The aggregate keyword enables use of the values keyword.
<code>random-detect precedence-based [aggregate]</code>	Enables application of precedence values to a WRED-drop threshold. The aggregate keyword enables use of the values keyword.
<code>random-detect multiple-type-based [aggregate]</code>	Enables application of CoS, precedence, and DSCP values to a WRED-drop threshold. The aggregate keyword enables use of the values keyword.
<code>random-detect cos cos_value percent min_percent max_percent</code>	Applies one CoS value to a WRED-drop threshold and configures the threshold percentages.
<code>random-detect dscp dscp_value percent min_percent max_percent</code>	Applies one DSCP value to a WRED-drop threshold and configures the threshold percentages.
<code>random-detect precedence precedence_value percent min_percent max_percent</code>	Applies one DSCP value to a WRED-drop threshold and configures the threshold percentages.
<code>random-detect cos values cos_list percent min_percent max_percent</code>	Applies multiple CoS values to a WRED-drop threshold and configures the threshold percentages; requires the aggregate keyword.
<code>random-detect dscp values dscp_list percent min_percent max_percent</code>	Applies multiple DSCP values to a WRED-drop threshold and configures the threshold percentages; requires the aggregate keyword.
<code>random-detect precedence values precedence_list percent min_percent max_percent</code>	Applies multiple precedence values to a WRED-drop threshold and configures the threshold percentages.

How to Configure Policy-Based Queueing

- [Configuring a Queueing Policy Class Map, page 26-11](#)
- [Verifying a Queueing Policy Class Map, page 26-11](#)
- [Configuring Queueing Policy Maps, page 26-11](#)
- [Verifying a Queueing Policy Map, page 26-17](#)
- [Attaching a Queueing Policy Map to an Interface, page 26-17](#)



Note

See the “[Configuration Examples for Policy-Based Queueing](#)” section on [page 26-18](#) for detailed information about which queueing commands are supported by each queue type.

Configuring a Queueing Policy Class Map

To configure a queueing policy class map, perform this task:

	Command	Purpose
Step 1	Router(config)# class-map type lan-queueing match-any <i>class_name</i>	Creates a class map. Create a class map for the thresholds on each type of queue that you are configuring. Give the class map a name that allows you to easily associate it with the queue type and threshold for which you configure it.
Step 2	Router (config-cmap)# match cos <i>cos_value1</i> [<i>cos_value2</i> ... [<i>cos_valueN</i>]]	(Optional) Configures the queueing policy class map to filter based on CoS values. You can enter multiple commands.
Step 3	Router (config-cmap)# match dscp <i>dscp_value1</i> [<i>dscp_value2</i> ... [<i>dscp_valueN</i>]]	(Optional) Configures the queueing policy class map to filter based on DSCP values and enables DSCP-based queueing on the port in the direction that the queueing policy is attached. You can enter multiple commands.
Step 4	Router (config-cmap)# match precedence <i>precedence_value1</i> [<i>precedence_value2</i> ... [<i>precedence_valueN</i>]]	(Optional) Configures the queueing policy class map to filter based on precedence values and enables DSCP-based queueing on the port in the direction that the queueing policy is attached. You can enter multiple commands.
Step 5	Router (config-cmap)# end	Exits configuration mode.

This example shows how to create a class map named **cos5** and how to configure filtering to match traffic with CoS 5:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# class-map cos5
Router(config-cmap)# match cos 5
Router(config-cmap)# end
```

Verifying a Queueing Policy Class Map

To verify the queueing policy class map, perform this task:

Command	Purpose
Router# show class-map <i>class_name</i>	Verifies the configuration.

Configuring Queueing Policy Maps

- [Configuring Policy-Based Queueing on WS-X6904-40G-2T Ports, page 26-12](#)
- [Creating a Queueing Policy, page 26-13](#)
- [Configuring a Priority Queue, page 26-14](#)

- [Configuring Nonpriority Queues, page 26-14](#)
- [Configuring Thresholds, page 26-15](#)

**Note**

- You can attach one input queueing policy map and one output queueing policy map to an interface.
- Queueing policy maps contain one policy-map class for each queue.
- Each policy-map class configures a queue.

Configuring Policy-Based Queueing on WS-X6904-40G-2T Ports

- [Configuring WS-X6904-40G-2T Nonpriority Egress Queue Shaping \(Except class-default\), page 26-12](#)
- [Configuring WS-X6904-40G-2T Multiple Priority Queues or Egress Priority-Queue Shaping, page 26-13](#)
- [Configuring WS-X6904-40G-2T Class-Default Egress Shaping, page 26-13](#)

**Note**

Use the procedures in these sections to configure the WS-X6904-40G-2T nonpriority ingress queues and single priority queues (ingress or egress) without shaping:

- [Creating a Queueing Policy, page 26-13](#)
- [Configuring a Priority Queue, page 26-14](#)
- [Configuring Nonpriority Queues, page 26-14](#)

Configure thresholds as described in the “[Configuring Thresholds](#)” section on [page 26-15](#).

Configuring WS-X6904-40G-2T Nonpriority Egress Queue Shaping (Except class-default)

The first **class** *class_map_name* command that you enter configures the highest numbered nonpriority queue. Subsequent **class** *class_map_name* commands configure the remaining nonpriority queues in reverse numerical order (highest-numbered to queue #2). To configure shaping on a nonpriority egress queue, perform this task:

	Command	Purpose
Step 1	Router(config)# policy-map type lan-queueing <i>child_policy_name</i>	Creates a child policy map for the nonpriority queues.
Step 2	Router(config-pmap)# class <i>class_map_name</i>	Creates a policy map class for a nonpriority egress queue numbered greater than one.
Step 3	Router(config-pmap-c)# bandwidth remaining percent <i>percentage</i>	Allocates DWRR bandwidth.
Step 4	Router(config-pmap-c)# shape average percent <i>percentage</i>	Configures shaping.
Step 5	Router(config-pmap-c)# queue-buffers ratio <i>weight</i>	(Optional) Sets queue buffer size.
Step 6	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring WS-X6904-40G-2T Multiple Priority Queues or Egress Priority-Queue Shaping

To enable the WS-X6904-40G-2T 2p6q4t port architecture (two priority queues) or to configure egress priority-queue shaping, perform this task:

	Command	Purpose
Step 1	Router(config-pmap)# class <i>class_map_name</i>	Creates a policy map class.
Step 2	Router(config-pmap-c)# priority level {1 2} [percent <i>percentage</i>]	Applies the class map to one of the priority queues. Priority queue 1 has higher priority than priority queue 2. Optionally, for egress queues, configures shaping .
Step 3	Router(config-pmap-c)# queue-buffers ratio <i>weight</i>	(Optional) Sets queue buffer size.
Step 4	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring WS-X6904-40G-2T Class-Default Egress Shaping

The **class class-default** command configures queue #1. To configure class-default egress shaping, perform this task:

	Command	Purpose
Step 1	Router(config)# policy-map type lan-queueing <i>parent_policy_name</i>	Creates the parent policy map for the port.
Step 2	Router(config-pmap)# class class-default	Configures the class-default policy map class.
Step 3	Router(config-pmap-c)# shape average percent <i>percentage</i> or Router(config-pmap-c)# shape average <value (<i>absolute number</i>)>	Configures shaping.
Step 4	Router(config-pmap-c)# queue-buffers ratio <i>weight</i>	(Optional) Sets queue buffer size.
Step 5	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Creating a Queueing Policy

To create a queueing policy, perform this task:

	Command	Purpose
Step 1	Router(config)# policy-map type lan-queueing <i>policy_name</i>	Creates a policy map.
Step 2	Router(config-pmap)# end	(Optional) Exits policy map class configuration mode.

Configuring a Priority Queue

1p1q4t, 1p1q8t, 1p7q2t, 1p3q8t, 1p7q8t, and 1p7q4t ports have priority queues. To configure the priority queue, perform this task:

	Command	Purpose
Step 1	Router(config-pmap)# class <i>class_map_name</i>	Creates a policy map class.
Step 2	Router(config-pmap-c)# priority	Applies the class map to the priority queue. Note The priority queue is not supported if SRR is enabled.
Step 3	Router(config-pmap-c)# queue-buffers ratio <i>weight</i>	(Optional) Sets queue buffer size.
Step 4	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring Nonpriority Queues

The first **class** *class_map_name* command that you enter that is not followed by the **priority** keyword configures the highest numbered nonpriority queue. Subsequent **class** *class_map_name* commands configure the remaining nonpriority queues in reverse numerical order (highest-numbered to queue #2). The **class** **class-default** command configures queue #1. To configure a nonpriority queue, perform this task:

	Command	Purpose
Step 1	Router(config-pmap)# class { <i>class_map_name</i> class-default }	Creates a policy map class. Note - Enter a class map name to configure nonpriority queues numbered greater than one. - Enter the class-default keyword to configure queue #1.
Step 2	Router(config-pmap-c)# bandwidth [remaining] percent <i>percentage</i>	Allocates WRR or DWRR bandwidth. Note - On WS-X6904-40G-2T ports, the bandwidth command must be configured if you configure any nondefault values for any other queueing commands on the port. (CSCtz05347) - Not required on ports that have one queue. - The remaining keyword is required on ports that have a priority queue.
Step 3	Router(config-pmap-c)# shape average percent <i>percentage</i>	(Egress queues only) Enables SRR (not supported on WS-X6904-40G-2T ports), which allocates limited bandwidth between nonpriority egress queues. - See “SRR” in the “Module to Queue Type Mappings” section on page 26-7). - Configuring SRR disables the priority queue.

	Command	Purpose
Step 4	Router(config-pmap-c)# <i>queue-buffers ratio weight</i>	(Optional) Sets queue buffer size. Note Not required on ports that have one queue.
Step 5	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring Thresholds

- [Threshold Configuration Guidelines and Restrictions, page 26-15](#)
- [Configuring a Threshold as Tail-Drop with CoS-Based Queueing, page 26-15](#)
- [Configuring a Threshold as WRED-Drop with CoS-Based Queueing, page 26-16](#)
- [Configuring a Threshold as Tail-Drop with DSCP-Based Queueing, page 26-16](#)
- [Configuring a Threshold as WRED-Drop with DSCP-Based Queueing, page 26-17](#)

Threshold Configuration Guidelines and Restrictions

- To configure tail drop thresholds, enter **queue-limit** commands.
- For ports that support configuration as either tail-drop or WRED drop:
 - Enter **queue-limit** commands to configure a threshold as tail drop.
 - Enter **random-detect** commands to configure a threshold as WRED drop.
- The first **queue-limit cos**, **queue-limit dscp**, **random-detect cos**, or **random-detect dscp** command that you enter configures threshold #1.
- Subsequent commands with different threshold percentage values configure the remaining thresholds in numerical order (from threshold #2 to the highest-numbered threshold).
- Subsequent commands with an already configured threshold percentage value apply addition QoS values to the threshold indicated by the percentage value.
- Ports configured for DSCP-based queueing use CoS-based queueing for non-IP traffic, IP multicast traffic, and IP unknown unicast flood traffic. In queueing policies that configure DSCP-based queueing, configure CoS-based queueing to provide specific QoS for non-IP traffic, IP multicast traffic, and IP unknown unicast flood traffic.

Configuring a Threshold as Tail-Drop with CoS-Based Queueing

To configure a threshold as tail-drop with CoS-based queueing, perform this task:

	Command	Purpose
Step 1	Router(config-pmap-c)# queue-limit cos { <i>one_value</i> values <i>value_list</i> } percent <i>percentage</i>	Applies CoS to a tail-drop threshold and configures the threshold percentage.
Step 2	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring a Threshold as WRED-Drop with CoS-Based Queueing

To configure a threshold as WRED-drop with CoS-based queueing, perform this task:

	Command	Purpose
Step 1	Router(config-pmap-c)# random-detect cos-based [aggregate]	Enables application of CoS values to a WRED-drop threshold. Enter the aggregate keyword to configure multiple CoS values on the threshold with the values keyword.
Step 2	Router(config-pmap-c)# random-detect cos {one_value values value_list} percent min_% max_%	Applies CoS to a WRED-drop threshold and configures the threshold percentage.
Step 3	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring a Threshold as Tail-Drop with DSCP-Based Queueing

To configure a threshold as tail-drop with DSCP-based queueing, perform this task:

	Command	Purpose
Step 1	Router(config-pmap-c)# queue-limit multiple-type-based	(Optional) Enables application of CoS, DSCP, or precedence values to a tail-drop threshold.
Step 2	Router(config-pmap-c)# queue-limit {cos dscp precedence} {one_value values value_list} percent percentage	Applies QoS values to a tail-drop threshold and configures the threshold percentage.
Step 3	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Configuring a Threshold as WRED-Drop with DSCP-Based Queueing

To configure a threshold as WRED-drop with DSCP-based queueing, perform this task:

	Command	Purpose
Step 1	Router(config-pmap-c)# random-detect { dscp-based precedence-based multiple-type-based } [aggregate]	Enables application of QoS values to a WRED-drop threshold. Note - Enter the dscp-based keyword to enable application of DSCP values to a WRED-drop threshold. - Enter the precedence-based keyword to enable application of precedence values to a WRED-drop threshold. - Enter the multiple-type-based keyword to enable application of both CoS, DSCP, or precedence values to a WRED-drop threshold. - Enter the aggregate keyword to configure multiple QoS values on the threshold with the values keyword.
Step 2	Router(config-pmap-c)# random-detect { cos dscp precedence } { one_value values <i>value_list</i> } percent <i>min_%</i> <i>max_%</i>	Applies QoS values to a WRED-drop threshold and configures the threshold percentage.
Step 3	Router(config-pmap-c)# end	(Optional) Exits policy map class configuration mode.

Verifying a Queueing Policy Map

Use the **show policy-map** *policy_name* to verify the configuration.

Attaching a Queueing Policy Map to an Interface

To attach a queueing policy to an interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 2	Router(config-if)# service-policy type lan-queueing [input output] <i>policy_map_name</i>	Attaches a queueing policy to the interface.
Step 3	Router(config-if)# end	Exits configuration mode.

Use the **show policy-map interface** command to verify the configuration.

Configuration Examples for Policy-Based Queueing

- [Queueing Policy Sample Configuration, page 26-18](#)
- [Queueing Policy Commands Supported by Each Queue Type, page 26-19](#)
- [Queueing Policy Commands Sample Configurations for Each Queue Type, page 26-33](#)

Queueing Policy Sample Configuration

Without comments:

```
policy-map type lan-queueing p1
  class cos5
    priority
  class cos123
    bandwidth remaining percent 25
    queue-limit cos 2 percent 20
    queue-limit cos 3 percent 30
  class class-default
    queue-limit cos 6 percent 60
```

With comments:

```
policy-map type lan-queueing p1 ! For 1p3q8t
  class cos5 ! Configured to filter CoS 5
!The filtering configured in the class map selects the values that go to the queue
!
    priority ! Applies the class map to the priority queue (#4)
!
!
!First non-priority class applies to highest-numbered non-priority queue (#3)
  class cos123 ! Configured to filter CoS 1, 2, and 3
!The filtering configured in the class map selects the values that go to the queue
!
!
!'remaining' keyword required on ports that have a priority queue
    bandwidth remaining percent 25
!
!
!First queue-limit command assigns CoS 2 to threshold #1 and configures it at 20%
    queue-limit cos 2 percent 20
!Any other queue-limit command with the same percentage
!applies additional configuration to this threshold
!
!Next queue-limit command with different percentage value configures the next threshold
!Assigns CoS 3 to threshold #2 and configures it at 30%
    queue-limit cos 3 percent 30
!Any other queue-limit command with the same percentage
!applies additional configuration to this threshold
!
!Thresholds 3-8 are unconfigured
!All unconfigured thresholds are at 100%
!No explicit configuration provided for CoS 1: defaults to last threshold
!
!End of queue 3 configuration
!
!Queue 2 is unconfigured
!
  class class-default ! applies to queue #1
!'class-default' gets all remaining CoS values:
```

```
!0, 4, 6, and 7
!  
!  
!Threshold 1 is explicitly configured:  
    queue-limit cos 6 percent 50  
!  
!Remaining thresholds (2-8) are not configured by the queueing policy  
!and cannot be configured by anything else  
!No explicit configuration provided for CoS 0, 4, and 7:  
!CoS values not explicitly configured default to the last threshold
```

Queueing Policy Commands Supported by Each Queue Type

- [1q2t, 1q8t Ingress Queue Supported Commands, page 26-20](#)
- [2q8t Ingress Queue Supported Commands, page 26-21](#)
- [8q4t Ingress Queue Supported Commands, page 26-22](#)
- [8q8t Ingress Queue Supported Commands, page 26-23](#)
- [1p1q4t Ingress Queue Supported Commands, page 26-24](#)
- [1p1q8t Ingress Queue Supported Commands, page 26-25](#)
- [1p7q2t Ingress Queue Supported Commands, page 26-26](#)
- [1p3q8t Egress Queue Supported Commands, page 26-27](#)
- [1p7q8t Egress Queue Supported Commands, page 26-28](#)
- [1p7q4t, 2p6q4t Ingress or Egress Queue Supported Commands, page 26-29](#)
- [WS-X6904-40G-2T 1p7q4t, 2p6q4t Egress Queue Supported Commands, page 26-31](#)

1q2t, 1q8t Ingress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queuing policy_map_name ! For 1q2t, 1q8t
```

- Nonpriority queue 1 policy commands:

```
class class-default ! Receives all CoS values.
    ! bandwidth percent percentage ! WRR or DWRR bandwidth allocation.
    ! queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration (repeat to configure sequential thresholds):

```
! queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values
! to a tail-drop threshold.

queue-limit cos {one_value | values value_list} percent percentage

! queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage

! random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword

! random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword

! random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword

! random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values
! to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword

! random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.

! random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.

! random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

2q8t Ingress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 2q8t, 8q8t
```

- Class commands for nonpriority queues numbered higher than 1—Configures queues in reverse numerical order; repeat to configure the next queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
```

- Class command for queue #1:

```
class class-default ! Receives all remaining CoS values.
```

- Nonpriority queue configuration commands:

```
bandwidth percent percentage ! WRR or DWRR bandwidth allocation.
```

```
queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds in each queue:

```
! queue-limit multiple-type-based
```

```
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
```

```
queue-limit cos {one_value | values value_list} percent percentage
```

```
! Applies CoS to a tail-drop threshold and configures the threshold percentage
```

```
! queue-limit dscp {one_value | values value_list} percent percentage
```

```
! Applies one DSCP value to a tail-drop threshold
```

```
! and configures the threshold percentage
```

```
! random-detect cos-based [aggregate]
```

```
! Enables application of CoS values to a WRED-drop threshold.
```

```
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
! random-detect dscp-based [aggregate]
```

```
! Enables application of DSCP values to a WRED-drop threshold.
```

```
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
! random-detect precedence-based [aggregate]
```

```
! Enables application of precedence values to a WRED-drop threshold.
```

```
! random-detect multiple-type-based [aggregate]
```

```
! Enables application of CoS, precedence, and DSCP values
```

```
! to a WRED-drop threshold.
```

```
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
! random-detect cos {one_value | values value_list} percent min_% max_%
```

```
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
```

```
! random-detect dscp {one_value | values value_list} percent min_% max_%
```

```
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

```
! random-detect {precedence one_value | values precedence value_list} percent min_% max_%
```

```
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

8q4t Ingress Queue Supported Commands



Note

- Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 8q4t
```

- Class commands for nonpriority queues numbered higher than 1—Configures queues in reverse numerical order; repeat to configure the next queue:

```
class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
```

- Class command for queue #1:

```
class class-default ! Receives all remaining QoS values (CoS, DSCP, precedence).
```

- Nonpriority queue configuration commands:

```
bandwidth percent percentage ! WRR or DWRR bandwidth allocation.
```

```
queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds in each queue:

```
queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
```

```
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
```

```
queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
```

```
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
```

```
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
```

```
random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

```
random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

8q8t Ingress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop or WRED-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 1p1q8t
```

- Priority queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
  priority ! Applies the class map to the priority queue
```

- Nonpriority queue policy commands:

```
class class-default ! Receives all remaining CoS values.
  ! bandwidth remaining percent percentage ! WRR or DWRR bandwidth allocation.
  ! queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds:

```
! queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
! queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values
! to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

1p1q4t Ingress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queuing policy_map_name ! For 1p1q4t
```

- Priority queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
  priority ! Applies the class map to the priority queue
```

- Nonpriority queue policy commands:

```
class class-default ! Receives all remaining CoS values.
  ! bandwidth remaining percent percentage ! WRR or DWRR bandwidth allocation.
  ! queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds:

```
! queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values
! to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
! queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
! random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values
! to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```


1p1q8t Ingress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop or WRED-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 1p1q8t
```

- Priority queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
  priority ! Applies the class map to the priority queue
```

- Nonpriority queue policy commands:

```
class class-default ! Receives all remaining CoS values.
  ! bandwidth remaining percent percentage ! WRR or DWRR bandwidth allocation.
  ! queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds:

```
! queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
! queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values
! to a WRED-drop threshold.
!
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

1p7q2t Ingress Queue Supported Commands



Note

- Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 1p7q2t
```

- Priority queue:

```
class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
  priority ! Applies the class map to the priority queue
```

- Class commands for nonpriority queues numbered higher than 1—Configures queues in reverse numerical order; repeat to configure the next queue:

```
class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
```

- Class command for queue #1:

```
class class-default ! Receives all remaining QoS values (CoS, DSCP, precedence).
```

- Nonpriority queue configuration commands:

```
bandwidth percent percentage ! WRR or DWRR bandwidth allocation.
!
queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds in each queue:

```
queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos {one_value | values value_list} percent min_% max_%
```

```

! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
random-detect {precedence one_value | values precedence value_list} percent min_%
max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.

```

1p3q8t Egress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop and WRED-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 1p3q8t
```

- Priority queue:

```

class class_map_name ! Receives CoS values filtered by class_map_name.
  priority ! Applies the class map to the priority queue

```

- Class commands for nonpriority queues numbered higher than 1—Configures queues in reverse numerical order; repeat to configure the next queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
```

- Class command for queue #1:

```
class class-default ! Receives all remaining CoS values.
```

- Nonpriority queue configuration commands:

```

bandwidth percent percentage ! WRR or DWRR bandwidth allocation.
!
! shape average percent percentage ! SRR bandwidth allocation.
!
queue-buffers ratio weight ! Queue buffer size.

```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds:

```

! queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
! queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.

```

```

! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values
! to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect {precedence one_value | values precedence value_list}
percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.

```

1p7q8t Egress Queue Supported Commands



Note

- Supports CoS-based queueing with tail-drop and WRED-drop thresholds.
- Unsupported commands are included in this section as comments.

```
policy-map type lan-queueing policy_map_name ! For 1p7q8t
```

- Priority queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
priority ! Applies the class map to the priority queue
```

- Class commands for nonpriority queues numbered higher than 1—Configures queues in reverse numerical order; repeat to configure the next queue:

```
class class_map_name ! Receives CoS values filtered by class_map_name.
```

- Class command for queue #1:

```
class class-default ! Receives all remaining CoS values.
```

- Nonpriority queue configuration commands:

```
bandwidth percent percentage ! WRR or DWRR bandwidth allocation.
!
! shape average percent percentage ! SRR bandwidth allocation.
!
queue-buffers ratio weight ! Queue buffer size.
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds:

```
! queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
! queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword

```

```

!
! random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
! random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values
! to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
! random-detect {precedence one_value | values precedence value_list}
percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.

```

1p7q4t, 2p6q4t Ingress or Egress Queue Supported Commands



Note

- Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds. Supports SRR (except WS-X6904-40G-2T) or DWRR dequeueing.
- WS-X6904-40G-2T supports [shaping](#). See the “[WS-X6904-40G-2T 1p7q4t, 2p6q4t Egress Queue Supported Commands](#)” section on page 26-31.
- Unsupported commands are included in this section as comments.

- Priority queue except WS-X6904-40G-2T with multiple priority queues:

```

class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
priority ! Applies the class map to the priority queue
! Not supported if SRR mode is enabled.
!
queue-buffers ratio weight ! Queue buffer size.
!

```

- WS-X6904-40G-2T priority queues:

```

class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
priority [level { 1 | 2}]
! Applies the class map to one of the priority queues
!
queue-buffers ratio weight ! Queue buffer size.
!

```

- Class commands for nonpriority queues numbered higher than 1—Configures queues in reverse numerical order; repeat to configure the next queue:

```

class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.

```

- Class command for queue #1:

```
class class-default ! Receives all remaining QoS values (CoS, DSCP, precedence).
```

- Nonpriority queue configuration commands:

```
shape average percent percentage
! Enables SRR on nonpriority egress queues.
```

```
bandwidth remaining percent percentage
! DWRR bandwidth allocation.
```



Note On WS-X6904-40G-2T ports, the **bandwidth** command must be configured if you configure any nondefault values for any other queueing commands on the port. (CSCtz05347)

```
!
queue-buffers ratio weight ! Queue buffer size.
!
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds in each queue:

```
queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
```

WS-X6904-40G-2T 1p7q4t, 2p6q4t Egress Queue Supported Commands



Note

- Supports CoS-based, DSCP-based, and precedence-based queuing with tail-drop and WRED-drop thresholds. Supports DWRR dequeuing.
- Unsupported commands are included in this section as comments.
- You configure shaping with a child policy map for the nonpriority queues that you apply to a class-default parent policy map.
- Without shaping, you configure all policy-map class commands in one policy map.

- (For shaping) Child policy map command for nonpriority queues numbered higher than 1:

```
policy-map type lan-queuing child_policy_name
```

- Policy map class commands for nonpriority queues numbered higher than 1—Configures nonpriority queues in reverse numerical order; repeat to configure the next queue:

```
class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
```

- Configuration commands for nonpriority queues numbered higher than 1:

```
shape average percent percentage
! Configures shaping on egress nonpriority queues.
```

```
bandwidth remaining percent percentage
! DWRR bandwidth allocation.
```



Note

On WS-X6904-40G-2T ports, the **bandwidth** command must be configured if you configure any nondefault values for any other queuing commands on the port. (CSCtz05347)

```
!
queue-buffers ratio weight ! Queue buffer size.
!
```

- Nonpriority queue threshold configuration—Repeat to configure sequential thresholds in each queue:

```
queue-limit multiple-type-based
! Enables application of CoS, precedence, and DSCP values to a tail-drop threshold.
!
queue-limit cos {one_value | values value_list} percent percentage
! Applies CoS to a tail-drop threshold and configures the threshold percentage
!
queue-limit dscp {one_value | values value_list} percent percentage
! Applies one DSCP value to a tail-drop threshold
! and configures the threshold percentage
!
random-detect cos-based [aggregate]
! Enables application of CoS values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect dscp-based [aggregate]
! Enables application of DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect precedence-based [aggregate]
! Enables application of precedence values to a WRED-drop threshold.
```

```

! The 'aggregate' keyword allows use of the 'values' keyword
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect multiple-type-based [aggregate]
! Enables application of CoS, precedence, and DSCP values to a WRED-drop threshold.
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos {one_value | values value_list} percent min_% max_%
! Applies CoS to a WRED-drop threshold and configures the threshold percentages.
!
random-detect dscp {one_value | values value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.
!
random-detect {precedence one_value | values precedence value_list} percent min_% max_%
! Applies DSCP to a WRED-drop threshold and configures the threshold percentages.

```

- Single priority queue without shaping:

```

class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
priority ! Applies the class map to the priority queue
!
queue-buffers ratio weight ! Queue buffer size.
!

```

- Multiple priority queues without shaping:

```

class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
priority level { 1 | 2}
! Applies the class map to one of the priority queues
!
queue-buffers ratio weight ! Queue buffer size.
!

```

- Egress priority queues (single or multiple) with shaping:

```

class class_map_name
! Receives QoS values (CoS, DSCP, precedence) values filtered by class_map_name.
priority level { 1 | 2} percent percentage
! Applies the class map to one of the priority queues
! Configures shaping on the queue
!
queue-buffers ratio weight ! Queue buffer size.
!

```

- (For shaping) Parent policy map command for queue #1 (class-default):

```

policy-map type lan-queuing parent_policy_name

```

- Class command for queue #1 (class-default):

```

class class-default ! Receives all remaining QoS values (CoS, DSCP, precedence).

shape average percent percentage
! Configures shaping on egress nonpriority queues.

bandwidth remaining percent percentage
! DWRR bandwidth allocation.

```



Note

On WS-X6904-40G-2T ports, the **bandwidth** command must be configured if you configure any nondefault values for any other queueing commands on the port. (CSCtz05347)

```

!

```



```

queue-buffers ratio weight ! Queue buffer size.
!
service-policy child_service_policy
! If shaping is configured on the other nonpriority queues

```

To configure class-default thresholds, see the [“Nonpriority queue threshold configuration” information](#).

Queueing Policy Commands Sample Configurations for Each Queue Type

- [1q2t Ingress Queue Sample Configuration, page 26-33](#)
- [1q8t Ingress Queue Sample Configuration, page 26-33](#)
- [2q8t Ingress Queue Sample Configuration, page 26-34](#)
- [8q4t, 8q8t Ingress Queue Sample Configuration \(CoS-Based Queueing\), page 26-35](#)
- [8q4t Ingress Queue Sample Configuration \(DSCP-Based Queueing\), page 26-35](#)
- [1p1q4t Ingress Queue Sample Configuration, page 26-38](#)
- [1p1q8t Ingress Queue Sample Configuration, page 26-38](#)
- [1p3q8t Egress Queue Sample Configuration, page 26-39](#)
- [1p7q8t Egress Queue Sample Configuration, page 26-40](#)
- [1p7q4t Ingress or Egress Queue Sample Configuration \(CoS-Based Queueing\), page 26-41](#)
- [1p7q4t Ingress or Egress Queue Sample Configuration \(DSCP-Based Queueing\), page 26-42](#)



Note

These sample configurations approximate the default queueing enabled by the **auto qos default** and **platform qos queueing-only** global configuration commands.

1q2t Ingress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop thresholds.

```

policy-map type lan-queueing p_map_1q2t
  class class-default ! Receives all CoS values.
  !
  ! Configures threshold #1:
  queue-limit cos values 0 1 3 4 percent 80
  ! Applies CoS values to threshold 1 and configures the threshold percentage
  !
  ! Other thresholds unconfigured; default to 100%
  ! Remaining CoS values are not explicitly configured:
  ! default to threshold 8 at 100%

```

1q8t Ingress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop thresholds.

```

policy-map type lan-queueing p_map_1q8t
  class class-default ! Receives all CoS values.
  !
  ! Configures threshold #1:
  queue-limit cos 0 percent 50
  ! Applies CoS 0 to threshold 1 and configures the threshold percentage
  !
  ! Configures threshold #2:
  queue-limit cos values 1 2 3 4 percent 60
  ! Applies CoS 1, 2, 3, 4 to threshold 2 and configures the threshold percentage
  !
  ! Configures threshold #3:
  queue-limit cos values 6 7 percent 80
  ! Applies CoS 6 and 7 to threshold 3 and configures the threshold percentage
  !
  ! Other thresholds unconfigured; default to 100%
  ! CoS 5 is not explicitly configured: defaults to threshold 8 at 100%

```

2q8t Ingress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop thresholds.

```

class-map type lan-queueing match-any c_map_cos_5
  match cos 5
!
policy-map type lan-queueing p_map_2q8t ! For 2q8t
!
! Configures queue #2:
class c_map_cos_5
! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_cos_5.
!
  bandwidth percent 10 ! WRR bandwidth allocation.
  queue-buffers ratio 20 ! Queue buffer size.
  !
  queue-limit cos 5 percent 100
  ! Applies CoS 5 to threshold 1 and configures the threshold percentage
!
! Configures queue #1:
class class-default ! Receives all remaining CoS values.
!
  bandwidth percent 90 ! WRR bandwidth allocation.
  !
  queue-buffers ratio 20 ! Queue buffer size.
  !
  ! Configures threshold #1:
  queue-limit cos values 0 1 percent 70
  ! Applies CoS 0 1 to threshold 1 and configures the threshold percentage
  !
  ! Configures threshold #2:
  queue-limit cos values 2 3 percent 80
  ! Applies CoS 2 3 to threshold 2 and configures the threshold percentage
  !
  ! Configures threshold #3:
  queue-limit cos 4 percent 90
  ! Applies CoS 4 to threshold 3 and configures the threshold percentage
  !
  ! CoS 6 and 7 default to threshold 4 at 100%

```

8q4t, 8q8t Ingress Queue Sample Configuration (CoS-Based Queueing)



Note

- Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds. This sample configures CoS-based queueing.

```
class-map type lan-queueing match-any c_map_cos_5
  match cos 5
!
policy-map type lan-queueing p_map_8q4t_cos_8q8t ! For 8q4t CoS-based queueing and 8q8t

  ! Configures queue #8:
  class c_map_cos_5 ! Receives CoS values values filtered by c_map_cos_5.
  !
    bandwidth percent 90 ! WRR bandwidth allocation.
    queue-buffers ratio 20 ! Queue buffer size.
    !
    ! CoS 5 defaults to threshold 4 at 100%
  !
  ! Configures queue #1:
  class class-default
  ! Receives all remaining CoS values.
  !
    bandwidth percent 10 ! WRR bandwidth allocation.
    queue-buffers ratio 80 ! Queue buffer size.
    !
    ! Configures threshold #1:
    random-detect cos-based aggregate
    ! The 'aggregate' keyword allows use of the 'values' keyword
    !
    random-detect cos values 0 1 percent 40 70
    !
    ! Configures threshold #2:
    random-detect cos-based aggregate
    ! The 'aggregate' keyword allows use of the 'values' keyword
    !
    random-detect cos values 2 3 percent 40 80
    !
    ! Configures threshold #3:
    random-detect cos-based
    !
    random-detect cos value 4 percent 50 90
    !
    ! Configures threshold #4:
    random-detect cos-based aggregate
    ! The 'aggregate' keyword allows use of the 'values' keyword
    !
    random-detect cos values 6 7 percent 50 100
```

8q4t Ingress Queue Sample Configuration (DSCP-Based Queueing)



Note

- Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds. This sample configures DSCP-based queueing.

```

class-map type lan-queueing match-any c_map_cos_5_dscp_40_46
    match cos 5
    !
    match dscp 40 46
    ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_48-63
    match dscp 48 49 50 51 52 53 54 55
    match dscp 56 57 58 59 60 61 62 63
    ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_32_34-38
    match dscp 32 34 35 36 37 38
    ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_24_26_28_30
    match dscp 24 26 28 30
    ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_18_20_22
    match dscp 18 20 22
    ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_10_12_14
    match dscp 10 12 14
    ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
policy-map type lan-queueing p_map_8q4t_dscp ! For 8q4t DSCP-based queueing

    ! Configures queue #8:
    class c_map_cos_5_dscp_40_46
    ! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_cos_5.
    !
        bandwidth percent 90 ! WRR bandwidth allocation.
        queue-buffers ratio 20 ! Queue buffer size.
        !
        ! CoS 5 and DSCP 40, 46 default to threshold 4 at 100%
    !
    ! Configures queue #7:
    class c_map_dscp_48-63
    ! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_48-63.
    !
        bandwidth remaining percent 10 ! WRR bandwidth allocation.
        queue-buffers ratio 10 ! Queue buffer size.
        !
        ! DSCP 48-63 default to threshold 4 at 100%
    !
    ! Configures queue #6:
    class c_map_dscp_32_34-38
    ! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_32_34-38.
    !
        bandwidth remaining percent 10 ! WRR bandwidth allocation.
        queue-buffers ratio 10 ! Queue buffer size.
        !
        ! DSCP 32, 34-38 default to threshold 4 at 100%
    !
    ! Configures queue #5:
    class c_map_dscp_24_26_28_30
    ! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_24_26_28_30.
    !
        bandwidth remaining percent 10 ! WRR bandwidth allocation.
        queue-buffers ratio 10 ! Queue buffer size.

```

```

!
! DSCP 24, 26, 28, 30 default to threshold 4 at 100%
!
! Configures queue #4:
class c_map_dscp_18_20_22
! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_18_20_22.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1:
random-detect dscp-based
random-detect dscp 20 percent 70 100
random-detect dscp 22 percent 70 100
random-detect dscp 18 percent 70 100
!
! Configures queue #3:
class c_map_dscp_10_12_14
! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_10_12_14.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1:
random-detect dscp-based
random-detect dscp 14 percent 70 100
random-detect dscp 12 percent 70 100
random-detect dscp 10 percent 70 100
!
! Configures queue #1:
class class-default
! Receives all remaining QoS values (CoS, DSCP, precedence).
!
bandwidth percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 80 ! Queue buffer size.
!
! Configures threshold #1:
random-detect cos-based aggregate
random-detect cos values 0 1 percent 40 70
!
! Configures threshold #2:
random-detect cos-based aggregate
random-detect cos values 2 3 percent 40 80
!
! Configures threshold #3:
random-detect cos-based
random-detect cos value 4 percent 50 90
!
! Configures threshold #4:
random-detect cos-based aggregate
random-detect cos values 6 7 percent 50 100
! DSCP values default to this threshold
! 0-9, 11, 13, 15-17, 19, 21, 23, 25, 27, 29, 31, 33, 39, 41-45, 47

```

1p1q4t Ingress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop thresholds.

```
class-map type lan-queueing match-any c_map_cos_5
  match cos 5
!
policy-map type lan-queueing p_map_1p1q4t ! For 1p1q4t
!
  ! Configures the priority queue:
  class c_map_cos_5 ! Receives CoS values filtered by c_map_cos_5.
    priority ! Applies the class map to the priority queue
  !
  ! Configures queue #1:
  class class-default ! Receives all remaining CoS values.
  !
    ! Configures threshold #1
    queue-limit cos values 0 1 percent 70
    ! Applies CoS 0 1 to threshold 1 and configures the threshold percentage
    !
    ! Configures threshold #2
    queue-limit cos values 2 3 percent 80
    ! Applies CoS 2 3 to threshold 2 and configures the threshold percentage
    !
    ! Configures threshold #3
    queue-limit cos 4 percent 90
    ! Applies CoS 4 to threshold 3 and configures the threshold percentage
    !
    ! CoS 6 and 7 default to threshold 4 at 100%
```

1p1q8t Ingress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop or WRED-drop thresholds.

```
class-map type lan-queueing match-any c_map_cos_5
  match cos 5
!
policy-map type lan-queueing p_map_1p1q8t ! For 1p1q8t
!
  ! Configures the priority queue:
  class c_map_cos_5 ! Receives CoS values filtered by c_map_cos_5.
    priority ! Applies the class map to the priority queue
  !
  class class-default ! Receives all remaining CoS values.
  !
    random-detect cos-based
    ! Enables application of CoS values to a WRED-drop threshold.
    !
    ! Configures threshold #1
    random-detect cos 0 percent 40 70
    ! Applies CoS to WRED-drop threshold 1 and configures the threshold percentages.
    !
    ! Configures threshold #2
    random-detect cos 1 percent 40 70
    ! Applies CoS to WRED-drop threshold 2 and configures the threshold percentages.
    !
```

```

! Configures threshold #3
random-detect cos 2 percent 50 80
! Applies CoS to WRED-drop threshold 3 and configures the threshold percentages.
!
! Configures threshold #4
random-detect cos 3 percent 50 80
! Applies CoS to WRED-drop threshold 4 and configures the threshold percentages.
!
! Configures threshold #5
random-detect cos 4 percent 60 90
! Applies CoS to WRED-drop threshold 5 and configures the threshold percentages.
!
! Configures threshold #6
random-detect cos 6 percent 60 90
! Applies CoS to WRED-drop threshold 6 and configures the threshold percentages.
!
! Configures threshold #7
random-detect cos 7 percent 70 100
! Applies CoS to WRED-drop threshold 7 and configures the threshold percentages.

```

1p3q8t Egress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop and WRED-drop thresholds.

```

class-map type lan-queueing match-any c_map_cos_2_3_4
    match cos 2 3 4
!
class-map type lan-queueing match-any c_map_cos_5
    match cos 5
!
class-map type lan-queueing match-any c_map_cos_6_7
    match cos 6 7
!
policy-map type lan-queueing p_map_1p3q8t
!
    ! Configures the priority queue:
    class c_map_cos_5 ! Receives CoS values filtered by c_map_cos_5.
        priority ! Applies the class map to the priority queue
    !
    ! Configures queue #3 for 1p3q8t:
    class c_map_cos_6_7
    ! Receives CoS values values filtered by c_map_cos_6_7.
    !
        bandwidth remaining percent 40 ! WRR bandwidth allocation.
        queue-buffers ratio 15 ! Queue buffer size.
    !
    ! Configures threshold #1:
    random-detect cos-based aggregate
    ! The 'aggregate' keyword allows use of the 'values' keyword
    !
    random-detect cos values 6 7 percent 70 100
!
! Configures queue #2 for 1p3q8t:
class c_map_cos_2_3_4
! Receives CoS values values filtered by c_map_cos_2_3_4.
!
    bandwidth remaining percent 30 ! WRR bandwidth allocation.
    queue-buffers ratio 20 ! Queue buffer size.
    !
    ! Configures threshold #1:

```

```

random-detect cos-based
!
random-detect cos 2 percent 40 70
! Applies CoS to WRED-drop threshold 1 and configures the threshold percentages.
!
! Configures threshold #2:
random-detect cos-based aggregate
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos values 3 4 percent 70 100
! Applies CoS to WRED-drop threshold 2 and configures the threshold percentages.
!
! Configures queue #1:
class class-default
! Receives all remaining CoS values.
!
    bandwidth remaining percent 25 ! WRR bandwidth allocation.
    queue-buffers ratio 20 ! Queue buffer size.
    !
    ! Configures threshold #1:
    random-detect cos-based
    !
    random-detect cos 0 percent 40 70
    ! Applies CoS to WRED-drop threshold 1 and configures the threshold percentages.
    !
    ! Configures threshold #2:
    random-detect cos-based
    !
    random-detect cos 1 percent 70 100
    ! Applies CoS to WRED-drop threshold 2 and configures the threshold percentages.

```

1p7q8t Egress Queue Sample Configuration



Note

Supports CoS-based queueing with tail-drop and WRED-drop thresholds.

```

class-map type lan-queuing match-any c_map_cos_2_3_4
    match cos 2 3 4
!
class-map type lan-queuing match-any c_map_cos_5
    match cos 5
!
class-map type lan-queuing match-any c_map_cos_6_7
    match cos 6 7
!
policy-map type lan-queuing p_map_1p7q8t
!
    ! Configures the priority queue:
    class c_map_cos_5 ! Receives CoS values filtered by c_map_cos_5.
        priority ! Applies the class map to the priority queue
    !
    ! Configures queue #7:
    class c_map_cos_6_7
    ! Receives CoS values values filtered by c_map_cos_6_7.
    !
        bandwidth remaining percent 40 ! WRR bandwidth allocation.
        queue-buffers ratio 15 ! Queue buffer size.
        !
        ! Configures threshold #1:
        random-detect cos-based aggregate
        ! The 'aggregate' keyword allows use of the 'values' keyword

```



```

!
    random-detect cos values 6 7 percent 70 100
!
! Configures queue #6 for 1p7q8t:
class c_map_cos_2_3_4
! Receives CoS values values filtered by c_map_cos_2_3_4.
!
    bandwidth remaining percent 30 ! WRR bandwidth allocation.
    queue-buffers ratio 20 ! Queue buffer size.
    !
    ! Configures threshold #1:
    random-detect cos-based
    !
    random-detect cos 2 percent 40 70
    ! Applies CoS to WRED-drop threshold 1 and configures the threshold percentages.
    !
    ! Configures threshold #2:
    random-detect cos-based aggregate
    ! The 'aggregate' keyword allows use of the 'values' keyword
    !
    random-detect cos values 3 4 percent 70 100
    ! Applies CoS to WRED-drop threshold 2 and configures the threshold percentages.
!
! Configures queue #1:
class class-default
! Receives all remaining CoS values.
!
    bandwidth remaining percent 25 ! WRR bandwidth allocation.
    queue-buffers ratio 50 ! Queue buffer size.
    !
    ! Configures threshold #1:
    random-detect cos-based
    !
    random-detect cos 0 percent 40 70
    ! Applies CoS to WRED-drop threshold 1 and configures the threshold percentages.
    !
    ! Configures threshold #2:
    random-detect cos-based
    !
    random-detect cos 1 percent 70 100
    ! Applies CoS to WRED-drop threshold 2 and configures the threshold percentages.

```

1p7q4t Ingress or Egress Queue Sample Configuration (CoS-Based Queueing)



Note

Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds. Supports SRR (except WS-X6904-40G-2T) or DWRR dequeuing. This sample configures CoS-based queueing.

```

class-map type lan-queueing match-any c_map_cos_5
    match cos 5
!
policy-map type lan-queueing p_map_1p7q4t_cos

    ! Configures the priority queue:
    class c_map_cos_5 ! Receives CoS values filtered by c_map_cos_5.
        priority ! Applies the class map to the priority queue
    !
    ! Configures queue #1:
    class class-default
    ! Receives all remaining CoS values.

```

```
!
bandwidth remaining percent 85 ! WRR bandwidth allocation.
```



Note On WS-X6904-40G-2T ports, the **bandwidth** command must be configured if you configure any nondefault values for any other queueing commands on the port. (CSCtz05347)

```
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1:
random-detect cos-based aggregate
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos values 0 1 percent 40 70
!
! Configures threshold #2:
random-detect cos-based aggregate
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos values 2 3 percent 40 80
!
! Configures threshold #3:
random-detect cos-based
!
random-detect cos value 4 percent 50 90
!
! Configures threshold #4:
random-detect cos-based aggregate
! The 'aggregate' keyword allows use of the 'values' keyword
!
random-detect cos values 6 7 percent 50 100
```

1p7q4t Ingress or Egress Queue Sample Configuration (DSCP-Based Queueing)



Note

Supports CoS-based, DSCP-based, and precedence-based queueing with tail-drop and WRED-drop thresholds. Supports SRR (except WS-X6904-40G-2T) or DWRR dequeuing. This sample configures DSCP-based queueing.

```
class-map type lan-queueing match-any c_map_cos_5_dscp_40_46
  match cos 5
  !
  match dscp 40 46
  ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_48-63
  match dscp 48 49 50 51 52 53 54 55
  match dscp 56 57 58 59 60 61 62 63
  ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_32_34-38
  match dscp 32 34 35 36 37 38
  ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_24_26_28_30
  match dscp 24 26 28 30
  ! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_18_20_22
```

```

match dscp 18 20 22
! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
class-map type lan-queueing match-any c_map_dscp_10_12_14
match dscp 10 12 14
! NOTE: Enables DSCP-based queueing on the port in the direction of the queueing policy.
!
policy-map type lan-queueing p_map_1p7q4t_dscp

```

```

! Configures the priority queue:
class c_map_cos_5_dscp_40_46 ! Receives CoS values filtered by c_map_cos_5_dscp_40_46.
priority ! Applies the class map to the priority queue
!
! Configures queue #7:
class c_map_dscp_48-63
! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_48-63.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.

```



Note On WS-X6904-40G-2T ports, the **bandwidth** command must be configured if you configure any nondefault values for any other queueing commands on the port. (CSCtz05347)

```

queue-buffers ratio 10 ! Queue buffer size.
!
! DSCP 48-63 default to threshold 4 at 100%
!
! Configures queue #6:
class c_map_dscp_32_34-38
! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_32_34-38.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! DSCP 32, 34-38 default to threshold 4 at 100%
!
! Configures queue #5:
class c_map_dscp_24_26_28_30
! Receives QoS values (CoS, DSCP, precedence) values
! filtered by c_map_dscp_24_26_28_30.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1
queue-limit dscp values 24 30 percent 100
!
! Configures threshold #2
queue-limit dscp 28 percent 100
!
! Configures threshold #3
queue-limit dscp 26 percent 100
!
! Configures queue #4:
class c_map_dscp_18_20_22
! Receives QoS values (CoS, DSCP, precedence) values filtered by c_map_dscp_18_20_22.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1:
random-detect dscp-based

```

```

random-detect dscp 20 percent 70 100
!
! Configures threshold #2:
random-detect dscp-based
random-detect dscp 22 percent 70 100
!
! Configures threshold #3:
random-detect dscp-based
random-detect dscp 18 percent 70 100
!
! Configures queue #3:
class c_map_dscp_10_12_14
! Receives QoS values (CoS, DSCP, precedence) values filtered by ! c_map_dscp_10_12_14.
!
bandwidth remaining percent 10 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1:
random-detect dscp-based
random-detect dscp 14 percent 70 100
random-detect dscp 12 percent 70 100
random-detect dscp 10 percent 70 100
!
! Configures queue #1:
class class-default
! Receives all remaining QoS values (CoS, DSCP, precedence).
!
bandwidth remaining percent 25 ! WRR bandwidth allocation.
queue-buffers ratio 10 ! Queue buffer size.
!
! Configures threshold #1:
random-detect cos-based aggregate
random-detect cos values 0 1 percent 40 70
!
! Configures threshold #2:
random-detect cos-based aggregate
random-detect cos values 2 3 percent 40 80
!
! Configures threshold #3:
random-detect cos-based
random-detect cos value 4 percent 50 90
!
! Configures threshold #4:
! Remaining DSCP values default to this threshold
! 0-9, 11, 13, 15-17, 19, 21, 23, 25, 27, 29, 31, 33, 39, 41-45, 47
!
random-detect cos-based aggregate
random-detect cos values 6 7 percent 50 100

```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Queuing on FEX

QoS on FEX host ports is enabled by default with default queuing configuration. Queuing on FEX provides the facility to configure your own queuing model on FEX ports. Queuing is supported in the egress direction only. The queuing features are configurable through the lan-queueing policies and do not support any legacy queuing commands. Priority queue configuration is optional. Queuing on FEX is supported in DSCP mode only. When applied to an interface the configuration will be propagated to all the interfaces of a member FEX, in case of a stack.

DoD certification mandates the queuing configuration for supported features. These features are as follows:

- Priority queuing
- Queue bandwidth
- DSCP to queue map
- Queue buffer

Priority Queuing

This example shows how to configure priority queuing:

```
Router# show policy-map type lan-queueing ltest
Policy Map type lan-queueing ltest
Class ltest
  bandwidth remaining 10 (%)
Class ltest1
  priority
Class class-default
```

Queue Bandwidth

This example shows how to configure queue bandwidth percentage:

```
Router# show policy-map type lan-queueing ltest
Policy Map type lan-queueing ltest
Class ltest
  bandwidth remaining 10 (%)
  Average Rate Traffic Shaping
  cir 30%
Class ltest1
  priority
Class class-default
*Dec 15 08:39:42.535: 020: 00 00 00 64 00 04 01 64 00 00 00 64 00 05 01 64
...d...d...d...d
*Dec 15 08:39:42.535: 030: 00 00 00 64
```

DSCP to Queue Map

The ability to map a DSCP to a particular queue is provided.

Queue Buffer

This example shows how to configure buffer per queue:

```
Router# show policy-map type lan-queuing ltest
Policy Map type lan-queuing ltest
Class ltest
  bandwidth remaining 10 (%)
  queue-buffers ratio 30
Class ltest1
  priority
  queue-buffers ratio 40
Class class-default
```

27

QoS Global and Interface Options

- [How to Configure the Ingress LAN Port CoS Value, page 27-2](#)
- [How to Configure Egress DSCP Mutation, page 27-3](#)
- [How to Configure Ingress CoS Mutation on IEEE 802.1Q Tunnel Ports, page 27-4](#)
- [How to Configure DSCP Value Maps, page 27-7](#)
- [How to Configure Trusted Boundary with Cisco Device Verification, page 27-10](#)
- [Legacy Configuration Procedures for Queueing-Only Mode, page 27-11](#)
- [Legacy Configuration Procedures for VLAN-Based PFC QoS on Layer 2 LAN Ports, page 27-12](#)
- [Legacy Configuration Procedures for Port Trust State, page 27-13](#)
- [Legacy Configuration Procedures for DSCP-Based Queue Mapping, page 27-14](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

How to Configure the Ingress LAN Port CoS Value



Note

A service policy applied to a port overrides any commands configured on the port.

Whether or not PFC QoS uses the CoS value applied with the **platform qos cos** command depends on the trust state of the port and the trust state of the traffic received through the port. The **platform qos cos** command does not configure the trust state of the port or the trust state of the traffic received through the port.

To use the CoS value applied with the **platform qos cos** command as the basis of internal DSCP:

- On a port that receives only untagged ingress traffic, configure the ingress port as trusted or configure a trust CoS policy map that matches the ingress traffic.
- On a port that receives tagged ingress traffic, configure a trust CoS policy map that matches the ingress traffic.
- The original ingress CoS value remains known.
 - By default, for IPv4 and IPv6 traffic, the ingress CoS value is overwritten by the DSCP value.
 - By default, for other traffic that is not tagged, the ingress CoS value is used, rather than the configured port CoS value.
 - Use the **platform qos cos override** interface command to use the value configured with the **platform qos cos** interface command instead of the original ingress CoS value.

You can configure the CoS value that PFC QoS assigns to untagged frames from ingress LAN ports configured as trusted and to all frames from ingress LAN ports configured as untrusted.

To configure the CoS value for an ingress LAN port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} {port-channel number}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos cos port_cos	Configures the ingress LAN port CoS value.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to configure the CoS value 5 on Fast Ethernet port 5/24 and verify the configuration:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface fastethernet 5/24
Router(config-if)# platform qos cos 5
Router(config-if)# end
Router# show queueing interface fastethernet 5/24 | include Default COS
Default COS is 5
Router#
```


How to Configure Egress DSCP Mutation

- [Configuring Named DSCP Mutation Maps, page 27-3](#)
- [Attaching an Egress DSCP Mutation Map to an Interface, page 27-4](#)

Configuring Named DSCP Mutation Maps

To configure a named DSCP mutation map, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos map dscp-mutation <i>map_name dscp1 [dscp2 [dscp3 [dscp4 [dscp5 [dscp6</i> <i>[dscp7 [dscp8]]]]]] to mutated_dscp</i>	Configures a named DSCP mutation map.
Step 2	Router(config)# end	Exits configuration mode.

- You can enter up to 8 DSCP values that map to a mutated DSCP value.
- You can enter multiple commands to map additional DSCP values to a mutated DSCP value.
- You can enter a separate command for each mutated DSCP value.

This example shows how to map DSCP 30 to mutated DSCP value 8:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# platform qos map dscp-mutation mutmap1 30 to 8
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show platform qos maps dscp-mutation
DSCP mutation map mutmap1: (dscp= d1d2)
  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 :   08 31 32 33 34 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
<...Output Truncated...>
Router#
```



Note

In the DSCP mutation map displays, the marked-down DSCP values are shown in the body of the matrix; the first digit of the original DSCP value is in the column labeled d1 and the second digit is in the top row. In the example shown, DSCP 30 maps to DSCP 08.

Attaching an Egress DSCP Mutation Map to an Interface

To attach an egress DSCP mutation map to an interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{ vlan <i>vlan_ID</i> } { type <i>slot/port</i> [. <i>subinterface</i>]} { port-channel <i>number</i> [. <i>subinterface</i>]}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos dscp-mutation <i>mutation_map_name</i>	Attaches an egress DSCP mutation map to the interface.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to attach the egress DSCP mutation map named mutmap1 to Fast Ethernet port 5/36:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/36
Router(config-if)# platform qos dscp-mutation mutmap1
Router(config-if)# end
```

How to Configure Ingress CoS Mutation on IEEE 802.1Q Tunnel Ports

- [Ingress CoS Mutation Configuration Guidelines and Restrictions, page 27-4](#)
- [Configuring Ingress CoS Mutation Maps, page 27-6](#)
- [Applying Ingress CoS Mutation Maps to IEEE 802.1Q Tunnel Ports, page 27-6](#)



Note

- IEEE 802.1Q tunnel ports configured to trust received CoS support ingress CoS mutation (see the [“Applying Ingress CoS Mutation Maps to IEEE 802.1Q Tunnel Ports”](#) section on page 27-6 for the list of supported modules).

When you configure ingress CoS mutation on an IEEE 802.1Q tunnel port that you have configured to trust received CoS, PFC QoS uses the mutated CoS value instead of the received CoS value in the ingress drop thresholds and for any trust CoS marking and policing.

Ingress CoS Mutation Configuration Guidelines and Restrictions

- Ports that are not configured as IEEE 802.1Q tunnel ports do not support ingress CoS mutation.
- Ports that are not configured to trust received CoS do not support ingress CoS mutation.
- Ingress CoS mutation does not change the CoS value carried by the customer frames. When the customer traffic exits the 802.1Q tunnel, the original CoS is intact.
- These switching modules support ingress CoS mutation:
 - WS-X6704-10GE

- WS-X6848-SFP-2T, WS-X6748-SFP
 - WS-X6824-SFP-2T, WS-X6724-SFP
 - WS-X6848-TX-2T, WS-X6748-GE-TX
- Ingress CoS mutation configuration applies to all ports in a port group. The port groups are:
 - WS-X6704-10GE:
4 ports, 4 port groups, 1 port in each group
 - WS-X6848-SFP-2T, WS-X6748-SFP:
48 ports, 4 port groups: ports 1–12, 13–24, 25–36, and 37–48
 - WS-X6824-SFP-2T, WS-X6724-SFP:
24 ports, 2 port groups: ports 1–12 and 13–24
 - WS-X6848-TX-2T, WS-X6748-GE-TX:
48 ports, 4 port groups: ports 1–12, 13–24, 25–36, and 37–48
- To avoid ingress CoS mutation configuration failures, only create EtherChannels where all member ports support ingress CoS mutation or where no member ports support ingress CoS mutation. Do not create EtherChannels with mixed support for ingress CoS mutation.
- If you configure ingress CoS mutation on a port that is a member of an EtherChannel, the ingress CoS mutation is applied to the port-channel interface.
- You can configure ingress CoS mutation on port-channel interfaces.
- With ingress CoS mutation configured on a port-channel interface, the following occurs:
 - The ingress CoS mutation configuration is applied to the port groups of all member ports of the EtherChannel. If any member port cannot support ingress CoS mutation, the configuration fails.
 - If a port in the port group is a member of a second EtherChannel, the ingress CoS mutation configuration is applied to the second port-channel interface and to the port groups of all member ports of the second EtherChannel. If any member port of the second EtherChannel cannot support ingress CoS mutation, the configuration fails on the first EtherChannel. If the configuration originated on a nonmember port in a port group that has a member port of the first EtherChannel, the configuration fails on the nonmember port.
 - The ingress CoS mutation configuration propagates without limit through port groups, member ports, and port-channel interfaces, regardless of whether or not the ports are configured to trust CoS or are configured as IEEE 802.1Q tunnel ports.
- An EtherChannel where you want to configure ingress CoS mutation must not have member ports that are in port groups containing member ports of other EtherChannels that have member ports that do not support ingress CoS mutation. (This restriction extends without limit through all port-group-linked member ports and port-channel-interface-linked ports.)
- A port where you want to configure ingress CoS mutation must not be in a port group that has a member port of an EtherChannel that has members that do not support ingress CoS mutation. (This restriction extends without limit through all port-group-linked member ports and port-channel-interface-linked ports.)
- There can be only be one ingress CoS mutation configuration applied to all port-group-linked member ports and port-channel-interface-linked ports.

Configuring Ingress CoS Mutation Maps

To configure an ingress CoS mutation map, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos map cos-mutation <i>mutation_map_name mutated_cos1 mutated_cos2</i> <i>mutated_cos3 mutated_cos4 mutated_cos5</i> <i>mutated_cos6 mutated_cos7 mutated_cos8</i>	Configures an ingress CoS mutation map. You must enter 8 mutated CoS values to which PFC QoS maps ingress CoS values 0 through 7.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure a CoS mutation map named testmap:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# platform qos map cos-mutation testmap 4 5 6 7 0 1 2 3
Router(config)# end
Router#
```

This example shows how to verify the map configuration:

```
Router(config)# show platform qos maps cos-mutation
COS mutation map testmap
cos-in   :   0   1   2   3   4   5   6   7
-----
cos-out  :   4   5   6   7   0   1   2   3
Router#
```

Applying Ingress CoS Mutation Maps to IEEE 802.1Q Tunnel Ports

To attach an ingress CoS mutation map to an IEEE 802.1Q tunnel port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} <i>{port-channel number}}</i>	Selects the interface to configure.
Step 2	Router(config-if)# platform qos cos-mutation <i>mutation_map_name</i>	Attaches an ingress CoS mutation map to the interface.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to attach the ingress CoS mutation map named testmap to Gigabit Ethernet port 1/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 1/1
Router(config-if)# platform qos cos-mutation testmap
Router(config-if)# end
Router# show platform qos maps cos-mutation
COS mutation map testmap
cos-in   :   0   1   2   3   4   5   6   7
-----
cos-out  :   4   5   6   7   0   1   2   3

testmap is attached on the following interfaces
Gi1/1
Router#
```

How to Configure DSCP Value Maps

- [Mapping Received CoS Values to Internal DSCP Values, page 27-7](#)
- [Mapping Received IP Precedence Values to Internal DSCP Values, page 27-7](#)
- [Configuring DSCP Markdown Values, page 27-8](#)
- [Mapping Internal DSCP Values to Egress CoS Values, page 27-9](#)

Mapping Received CoS Values to Internal DSCP Values

To configure the mapping of received CoS values to the DSCP value that PFC QoS uses internally on the PFC, perform this task:

	Command	Purpose
Step 1	Router(config)# table-map cos-discard-class-map <i>dscp1 dscp2 dscp3 dscp4 dscp5 dscp6 dscp7 dscp8</i>	Configures the received CoS to internal DSCP map. You must enter 8 DSCP values to which PFC QoS maps CoS values 0 through 7.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the received CoS to internal DSCP map:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# table-map cos-discard-class-map 0 1 2 3 4 5 6 7
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show platform qos maps cos-discard-class
Router#
```

Mapping Received IP Precedence Values to Internal DSCP Values

To configure the mapping of received IP precedence values to the DSCP value that PFC QoS uses internally on the PFC, perform this task:

	Command	Purpose
Step 1	Router(config)# table-map precedence-discard-class-map <i>dscp1 dscp2 dscp3 dscp4 dscp5 dscp6 dscp7 dscp8</i>	Configures the received IP precedence to internal DSCP map. You must enter 8 internal DSCP values to which PFC QoS maps received IP precedence values 0 through 7.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure the received IP precedence to internal DSCP map:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# table-map precedence-discard-class-map 0 1 2 3 4 5 6 7
Router(config)# end
```

This example shows how to verify the configuration:

```
Router# show platform qos maps precedence-discard-class
Router#
```

Configuring DSCP Markdown Values

To configure the mapping of DSCP markdown values used by policers, perform this task:

	Command	Purpose
Step 1	Router(config)# table-map policed-discard-class {normal-burst max-burst} dscp1 [dscp2 [dscp3 [dscp4 [dscp5 [dscp6 [dscp7 [dscp8]]]]]] to markdown_dscp	Configures a DSCP markdown map.
Step 2	Router(config)# end	Exits configuration mode.

When configuring a DSCP markdown map, note the following information:

- You can enter the **normal-burst** keyword to configure the markdown map used by the **exceed-action policed-dscp-transmit** keywords.
- You can enter the **max-burst** keyword to configure the markdown map used by the **violate-action policed-dscp-transmit** keywords.



Note When you create a policer that does not use the **pir** keyword, and the *maximum_burst_bytes* parameter is equal to the *normal_burst_bytes* parameter (which occurs if you do not enter the *maximum_burst_bytes* parameter), the **exceed-action policed-dscp-transmit** keywords cause PFC QoS to mark traffic down as defined by the **policed-dscp max-burst** markdown map.

- To avoid out-of-sequence packets, configure the markdown maps so that conforming and nonconforming traffic uses the same queue.
- You can enter up to 8 DSCP values that map to a marked-down DSCP value.
- You can enter multiple commands to map additional DSCP values to a marked-down DSCP value.
- You can enter a separate command for each marked-down DSCP value.



Note Configure marked-down DSCP values that map to CoS values consistent with the markdown penalty.

This example shows how to map DSCP 1 to marked-down DSCP value 0:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# table-map policed-discard-class normal-burst 1 to 0
Router(config)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show table-map policed-discard-class-normal-burst-map
Normal Burst Policed-dscp map:                                     (dscp= d1d2)
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 :    50 51 52 53 54 55 56 57 58 59
6 :    60 61 62 63

```

Router#



Note

In the Policed-dscp displays, the marked-down DSCP values are shown in the body of the matrix; the first digit of the original DSCP value is in the column labeled d1 and the second digit is in the top row. In the example shown, DSCP 41 maps to DSCP 41.

Mapping Internal DSCP Values to Egress CoS Values

To configure the mapping of the DSCP value that PFC QoS uses internally on the PFC to the CoS value used for egress LAN port scheduling and congestion avoidance, perform this task:

	Command	Purpose
Step 1	Router(config)# table-map discard-class-cos-map <i>dscp1 [dscp2 [dscp3 [dscp4 [dscp5 [dscp6 [dscp7 [dscp8]]]]]] to cos_value</i>	Configures the internal DSCP to egress CoS map.
Step 2	Router(config)# end	Exits configuration mode.

- You can enter up to 8 DSCP values that PFC QoS maps to a CoS value.
- You can enter multiple commands to map additional DSCP values to a CoS value.
- You can enter a separate command for each CoS value.

This example shows how to configure internal DSCP values 0, 8, 16, 24, 32, 40, 48, and 54 to be mapped to egress CoS value 0:

```

Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# table-map discard-class-cos-map 0 8 16 24 32 40 48 54 to 0
Router(config)# end
Router#

```

This example shows how to verify the configuration:

```

Router# show table-map discard-class-cos-map
Dscp-cos map:                                     (dscp= d1d2)
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 :    06 06 06 06 00 06 07 07 07 07 07
6 :    07 07 07 07
Router#

```



Note

In the `Dscp-cos map` display, the CoS values are shown in the body of the matrix; the first digit of the DSCP value is in the column labeled d1 and the second digit is in the top row. In the example shown, DSCP values 41 through 47 all map to CoS 05.

How to Configure Trusted Boundary with Cisco Device Verification

The trusted boundary with Cisco device verification feature configures Ethernet LAN ports to use CDP to detect whether or not a Cisco IP phone is attached to the port.

- If CDP detects a Cisco IP phone, QoS applies a configured **mls qos trust dscp**, **mls qos trust ip-precedence**, or **mls qos trust cos** interface command.
- If CDP does not detect a Cisco IP phone, QoS ignores any configured nondefault trust state.

To configure trusted boundary with Cisco device verification, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} {port-channel number}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos trust device cisco-phone	Configures trusted boundary with Cisco device verification.
Step 3	Router(config-if)# end	Exits configuration mode.

When configuring trusted boundary with Cisco device verification, CDP must be enabled on the port to use trusted boundary with Cisco device verification.

This example shows how to configure trusted boundary with Cisco device verification on Gigabit Ethernet port 1/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 1/1
Router(config-if)# mls qos trust device cisco-phone
Router(config-if)# end
Router#
```

This example shows how to verify the configuration on a port configured to trust CoS, but that does not have a Cisco IP phone attached:

```
Router# show queueing interface gigabitethernet 1/1 | include [Tt]rust
Trust boundary enabled
  Port is untrusted
  Extend trust state: not trusted [COS = 0]
Router#
```


Legacy Configuration Procedures for Queueing-Only Mode

**Note**

You can configure the queueing-only functionality with service policies.

To enable queueing-only mode on the switch, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos queueing-only	Enables queueing-only mode on the switch.
Step 2	Router(config)# end	Exits configuration mode.

When you enable queueing-only mode, the following actions occur:

- Except on ports configured with service policies, disables policing and marking (preserves all ingress QoS labels).
- Configures ingress queueing on ports to which an ingress queueing policy is not attached. Configures egress queueing on ports to which an egress queueing policy is not attached.
- Configures all ports to trust Layer 2 CoS.

**Note**

The switch applies the port CoS value to untagged ingress traffic and to traffic that is received through ports that cannot be configured to trust CoS.

This example shows how to enable queueing-only mode:

```
Router# configure terminal
Router(config)# platform qos queueing-only
Router(config)# end
Router#
```

Legacy Configuration Procedures for VLAN-Based PFC QoS on Layer 2 LAN Ports



Note

- You can attach policy maps to Layer 3 interfaces for application of PFC QoS to egress traffic. VLAN-based or port-based PFC QoS on Layer 2 ports is not relevant to application of PFC QoS to egress traffic on Layer 3 interfaces.
- By default, PFC QoS uses policy maps attached to LAN ports. For ports configured as Layer 2 LAN ports with the **switchport** keyword, you can configure PFC QoS to use policy maps attached to a VLAN. Ports not configured with the **switchport** keyword are not associated with a VLAN.

To enable VLAN-based PFC QoS on a Layer 2 LAN port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} {port-channel number}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos vlan-based	Enables VLAN-based PFC QoS on a Layer 2 LAN port or a Layer 2 EtherChannel.
Step 3	Router(config-if)# end	Exits configuration mode.

- The configured port trust state does not affect marking when the **platform qos vlan-based** interface command is configured.
- A service policy attached to the Layer 3 VLAN interface defines QoS for ports where the **platform qos vlan-based** interface command is configured.
- Service policies attached to ports configured with the **platform qos vlan-based** interface command are ignored.
- Configuring a Layer 2 LAN port for VLAN-based PFC QoS preserves the policy map port configuration. The **no platform qos vlan-based** port command reenables any previously configured port commands.

This example shows how to enable VLAN-based PFC QoS on Fast Ethernet port 5/42:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/42
Router(config-if)# platform qos vlan-based
Router(config-if)# end
```

This example shows how to verify the configuration:

```
Router# show platform qos | begin QoS is vlan-based
QoS is vlan-based on the following interfaces:
Fa5/42
<...Output Truncated...>
```

Legacy Configuration Procedures for Port Trust State

To configure a port to which a service policy is not attached as untrusted, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} {port-channel number}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos trust none remark	Configures the port as untrusted and marks all non-MPLS traffic.
Step 3	Router(config-if)# end	Exits configuration mode.

To configure a port to which a service policy is not attached to trust CoS or IP precedence, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} {port-channel number}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos trust [precedence cos]	Configures the trust state of the port. Note DSCP is trusted by default.
Step 3	Router(config-if)# end	Exits configuration mode.



Note

The trust state of a port is unrelated to enables ingress queueing. To avoid dropping traffic because of inconsistent CoS values, configure ports to trust CoS only when the received traffic carries CoS values that you know to be consistent with network policy.

This example shows how to configure Gigabit Ethernet port 1/1 with the **trust cos** keywords:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 1/1
Router(config-if)# platform qos trust cos
Router(config-if)# end
Router#
```

This example shows how to verify the configuration:

```
Router# show queueing interface gigabitethernet 1/1 | include trust
Trust state: trust COS
Router#
```

Legacy Configuration Procedures for DSCP-Based Queue Mapping

- [Configuring Ingress DSCP-Based Queue Mapping, page 27-15](#)
- [Mapping DSCP Values to Standard Transmit-Queue Thresholds, page 27-16](#)
- [Mapping DSCP Values to the Transmit Strict-Priority Queue, page 27-18](#)



Note

- Do not use the procedures in this section if you have policy-based queueing configured.
- You can enable DSCP-based queues and thresholds on 8q4t, 1p7q2t, and 1p7q4t ports (see the [“Module to Queue Type Mappings”](#) section on page 35-7)
- DSCP-based queueing is supported on 8q4t, 1p7q2t, and 1p7q4t ports. The Supervisor Engine 6T-10GE ports are 8q4t/1p7q4t with the **platform qos 10g-only** global configuration command configured. To configure DSCP-based queue mapping on Supervisor Engine 6T ports, you must enter **shutdown** interface configuration mode commands for the Supervisor Engine 6T Gigabit Ethernet ports, and then enter the **platform qos 10g-only** global configuration command, which disables the Gigabit Ethernet ports on the Supervisor Engine 6T.
- In releases where [CSCts82932](#) is not resolved, do not use the default DSCP-based queue mapping for 8q4t ingress queues unless you configure supporting bandwidth and queue limits.

Enabling DSCP-Based Queue Mapping

To enable DSCP-based queue mapping, perform this task:

	Command	Purpose
Step 1	Router(config)# interface tengigabitethernet <i>slot/port</i>	Selects the interface to configure.
Step 2	Router(config-if)# platform qos queue-mode mode-dscp	Enables DSCP-based queue mapping.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to enable DSCP-based queue mapping on 10-Gigabit Ethernet port 6/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface tengigabitethernet 6/1
Router(config-if)# platform qos queue-mode mode-dscp
Router(config-if)# end
```

This example shows how to verify the configuration:

```
Router# show queueing interface tengigabitethernet 6/1 | include Queueing Mode
Queueing Mode In Tx direction: mode-dscp
Queueing Mode In Rx direction: mode-dscp
```

Configuring Ingress DSCP-Based Queue Mapping

- [Enabling DSCP-Based Queue Mapping, page 27-14](#)
- [Mapping DSCP Values to Standard Receive-Queue Thresholds, page 27-15](#)



Note

Ingress DSCP-to-queue mapping is supported only on ports configured to trust DSCP.

Mapping DSCP Values to Standard Receive-Queue Thresholds

To map DSCP values to the standard receive-queue thresholds, perform this task:

	Command	Purpose
Step 1	Router(config)# interface tengigabitethernet slot/port	Selects the interface to configure.
Step 2	Router(config-if)# rcv-queue dscp-map queue_# threshold_# dscp1 [dscp2 [dscp3 [dscp4 [dscp5 [dscp6 [dscp7 [dscp8]]]]]]]	Maps DSCP values to the standard receive queue thresholds.
Step 3	Router(config-if)# end	Exits configuration mode.

When mapping DSCP values, note the following information:

- You can enter up to 8 DSCP values that map to a queue and threshold.
- You can enter multiple commands to map additional DSCP values to the queue and threshold.
- You must enter a separate command for each queue and threshold.

This example shows how to map the DSCP values 0 and 1 to threshold 1 in the standard receive queue for 10-Gigabit Ethernet port 6/1 port 6/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface tengigabitethernet 6/1
Router(config-if)# rcv-queue dscp-map 1 1 0 1
Router(config-if)# end
Router#
```

**Note**

The receive queue mapping is shown in the second `queue thresh dscp-map` displayed by the **show queueing interface** command.

This example shows how to verify the configuration:

```
Router# show queueing interface tengigabitethernet 1/1 | begin queue thresh dscp-map
<...Output Truncated...>
queue thresh dscp-map
-----
1      1      0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25 27 29 31 33 39 41 42 43 44 45
47
1      2
1      3
1      4
2      1      14
2      2      12
2      3      10
2      4
3      1      22
3      2      20
3      3      18
3      4
4      1      24 30
4      2      28
4      3      26
4      4
5      1      32 34 35 36 37 38
5      2
5      3
5      4
6      1      48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
6      2
6      3
6      4
7      1
7      2
7      3
7      4
8      1      40 46
8      2
8      3
8      4
<...Output Truncated...>
Router#
```

Mapping DSCP Values to Standard Transmit-Queue Thresholds

To map DSCP values to standard transmit-queue thresholds, perform this task:

	Command	Purpose
Step 1	Router(config)# interface tengigabitethernet <i>slot/port</i>	Selects the interface to configure.
Step 2	Router(config-if)# wrr-queue dscp-map <i>transmit_queue_# threshold_# dscp1 [dscp2 [dscp3 [dscp4 [dscp5 [dscp6 [dscp7 [dscp8]]]]]]]</i>	Maps DSCP values to a standard transmit-queue threshold.
Step 3	Router(config-if)# end	Exits configuration mode.

- You can enter up to 8 DSCP values that map to a queue and threshold.
- You can enter multiple commands to map additional DSCP values to the queue and threshold.
- You must enter a separate command for each queue and threshold.

This example shows how to map the DSCP values 0 and 1 to standard transmit queue 1/threshold 1 for 10-Gigabit Ethernet port 6/1 port 6/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface tengigabitethernet 6/1
Router(config-if)# wrr-queue dscp-map 1 1 0 1
Router(config-if)# end
Router#
```

**Note**

The eighth queue is the strict priority queue in the output of the **show queueing interface** command.

This example shows how to verify the configuration:

```
Router# show queueing interface tengigabitethernet 6/1 | begin queue thresh dscp-map
queue thresh dscp-map
-----
1      1      0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25 27 29 31 33 39 41 42 43 44 45
47
1      2
1      3
1      4
2      1      14
2      2      12
2      3      10
2      4
3      1      22
3      2      20
3      3      18
3      4
4      1      24 30
4      2      28
4      3      26
4      4
5      1      32 34 35 36 37 38
5      2
5      3
5      4
6      1      48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
6      2
6      3
6      4
7      1
7      2
7      3
7      4
8      1      40 46
<...Output Truncated...>
Router#
```

Mapping DSCP Values to the Transmit Strict-Priority Queue

To map DSCP values to the transmit strict-priority queue, perform this task:

	Command	Purpose
Step 1	Router(config)# interface tengigabitethernet <i>slot/port</i>	Selects the interface to configure.
Step 2	Router(config-if)# priority-queue dscp-map <i>queue_# dscp1 [dscp2 [dscp3 [dscp4 [dscp5 [dscp6 [dscp7 [dscp8]]]]]]]</i>	Maps DSCP values to the transmit strict-priority queue. You can enter multiple priority-queue dscp-map commands to map more than 8 DSCP values to the strict-priority queue.
Step 3	Router(config-if)# end	Exits configuration mode.

- The queue number is always 1.
- You can enter up to 8 DSCP values to map to the queue.
- You can enter multiple commands to map additional DSCP values to the queue.

This example shows how to map DSCP value 7 to the strict-priority queue on 10 Gigabit Ethernet port 6/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface tengigabitethernet 6/1
Router(config-if)# priority-queue dscp-map 1 7
Router(config-if)# end
Router#
```



Note

The strict priority queue is queue 8 in the output of the **show queueing interface** command.

This example shows how to verify the configuration:

```
Router# show queueing interface tengigabitethernet 6/1 | begin queue thresh dscp-map
queue thresh dscp-map
-----
<...Output Truncated...>
      8      1      7 40 46
<...Output Truncated...>
Router#
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

28

AutoQoS

- [Prerequisites for AutoQoS, page 28-1](#)
- [Restrictions for AutoQoS, page 28-1](#)
- [Information About AutoQoS, page 28-2](#)
- [Default Settings for AutoQoS, page 28-4](#)
- [How to Configure AutoQoS, page 28-4](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for AutoQoS

None.

Restrictions for AutoQoS

- AutoQoS generates commands for the port and adds the generated commands to the running configuration.
- The generated QoS commands are applied as if you were entering them from the CLI. An existing configuration might cause the application of the generated commands to fail or an existing configuration might be overridden by the generated commands. These actions occur without warning. If the generated commands are successfully applied, any configuration that was not overridden remains in the running configuration. Any commands that were overridden might still exist in the startup-config file.

- Some of the generated commands are the type of PFC QoS commands that are applied to [all ports controlled by a port ASIC](#). When one of these commands is applied, PFC QoS displays the messages caused by application of the command to all the ports controlled by the port ASIC. Depending on the module, these commands are applied to as many as 48 ports. See the “Number of port groups” and “Port ranges per port group” listed for each module in the *Release Notes for Cisco IOS Release 15.2SY*.
- You might not be able to configure support for Cisco IP phones and the other autoQoS options on ports that are controlled by the same port ASIC because of conflicting port trust state requirements.
- If application of the generated commands fails, the previous running configuration is restored.
- Enable autoQoS before you configure other QoS commands. If necessary, you can modify the QoS configuration after the autoQoS configuration completes.
- AutoQoS cannot attach a policy map to an interface if there is already a policy map attached.
- Do not modify a policy map or class map that includes AUTOQOS in its name.
- You cannot configure autoQoS on the following:
 - Port-channel interfaces
 - VLAN interfaces (also known as switch virtual interfaces or SVIs)
 - Tunnel interfaces
 - Loopback interfaces
 - Subinterfaces on any type of interface

Information About AutoQoS

- [AutoQoS Support for a Cisco IP Phone, page 28-2](#)
- [AutoQoS Support for Cisco IP Communicator, page 28-3](#)
- [AutoQoS Support for Marked Traffic, page 28-3](#)

**Note**

AutoQoS is a macro that applies the recommended Architecture for Voice, Video, and Integrated Data (AVVID) QoS settings to a port.

AutoQoS Support for a Cisco IP Phone

Cisco IP phones are usually connected directly to ports. Optionally, you can attach a PC to the phone and use the phone as a hop to the switch.

The traffic that comes from the phone can be marked with an 802.1Q or 802.1p tag. The tag contains a VLAN ID and CoS value. When you configure the port to trust the CoS value that comes from the phone, the switch uses the CoS value to prioritize the phone traffic.

There is a three-port switch built into Cisco IP phones that forwards the traffic that comes from the PC, the phone, and the switch port. Cisco IP phones have trust and classification capabilities that you need to configure (see the [“How to Configure Cisco IP Phone Support” section on page 19-5](#)).

AutoQoS supports Cisco IP phones with the **auto qos voip cisco-phone** interface configuration command. When you enter the **auto qos voip cisco-phone** interface configuration command on a port that is configured to support an IP phone and to which an IP phone is connected, the autoQoS feature does the following:

- If QoS was not already enabled, enables QoS globally.
- If VLAN-based QoS was configured for the port, reverts to the default port-based QoS (done for all ports on switching modules with **1p1q0t/1p3q1t** ports).
- Sets the port trust state to trust CoS.
- Creates and applies a trust-CoS QoS policy to ports on switching modules with non-Gigabit Ethernet **1q4t/2q2t** ports, which do not support port trust.

AutoQoS Support for Cisco IP Communicator

The Cisco IP Communicator program runs on a PC and emulates a Cisco IP phone. The Cisco IP Communicator marks its voice traffic with a DSCP value instead of a CoS value. When you configure the port to trust the DSCP value that comes from the Cisco IP Communicator, the switch uses the DSCP value to prioritize the Cisco IP Communicator traffic.

AutoQoS supports the Cisco IP Communicator program with the **auto qos voip cisco-softphone** interface configuration command. When you enter the **auto qos voip cisco-softphone** interface configuration command on a port that is connected to a device running the Cisco IP Communicator program, the autoQoS feature does the following:

- If QoS was not already enabled, enables QoS globally.
- If VLAN-based QoS was configured for the port, reverts to the default port-based QoS (done for all ports on switching modules with **1p1q0t/1p3q1t** ports).
- If a trust state was configured for the port, reverts to the default untrusted state.
- Creates and applies ingress policers to trust DSCP 46 and remark DSCP 26 packets to DSCP 24. Packets with other DSCP values or out-of-profile packets are remarked with DSCP 0.

AutoQoS Support for Marked Traffic

Ports that connect to the interior of your network might receive traffic that has already been marked with QoS labels that are consistent with your network QoS policies, and which do not need to be changed. You can use the QoS trust feature to process the marked traffic using the received QoS values.

AutoQoS supports marked traffic with the **auto qos voip trust** interface configuration command. When you enter the **auto qos voip trust** interface configuration command, the autoQoS feature does the following:

- If QoS was not already enabled, enables QoS globally.
- If VLAN-based QoS was configured for the port, reverts to the default port-based QoS (done for all ports on switching modules with **1p1q0t/1p3q1t** ports).
- If the port is configured with the **switchport** command, sets the port trust state to trust CoS.

- If the port is not configured with the **switchport** command, sets the port trust state to trust DSCP.
- Creates and applies a trust-CoS or trust-DSCP QoS policy to ports on switching modules with non-Gigabit Ethernet **1q4t/2q2t** ports, which do not support port trust.

Default Settings for AutoQoS

None.

How to Configure AutoQoS

- [Configuring AutoQoS Support for a Cisco IP Phone, page 28-5](#)
- [Configuring AutoQoS Support for Cisco IP Communicator, page 28-5](#)
- [Configuring AutoQoS Support for Marked Traffic, page 28-6](#)



Note

AutoQoS generates commands that are appropriate for the QoS port architecture of the port on which you enter an **auto qos voip** command. For each of the different **auto qos voip** commands, autoQoS generates different QoS commands for each of these QoS port architectures:

- 1p1q0t/1p3q1t
- 1p1q4t/1p2q2t
- 1p1q4t/1p3q8t
- 1p1q8t/1p2q1t
- 1q2t/1p2q2t
- 1q2t/1p3q8t
- 1q4t/2q2t
- 1q8t/1p3q8t
- 1q8t/1p7q8t
- 2q8t/1p3q8t
- 8q4t/1p7q4t
- 8q8t/1p7q8t

The procedures in the following sections include the commands that you need to enter to display the generated commands, but the specific commands that autoQoS generates are not listed in this document.

Configuring AutoQoS Support for a Cisco IP Phone



Note

Complete the configuration procedures in the [“How to Configure Cisco IP Phone Support”](#) section on [page 19-5](#) before you configure autoQoS for a Cisco IP phone.

To configure autoQoS for a Cisco IP phone, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 3	Router(config-if)# auto qos voip cisco-phone	Configures autoQoS for a Cisco IP phone.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

When configuring autoQoS for a Cisco IP phone, note the following information:

- To disable autoQoS on an interface, use the **no auto qos voip** interface configuration command.



Note

The **no auto qos voip** interface configuration command does not delete the received CoS to internal DSCP map created by autoQoS.

- You might see messages that instruct you to configure other ports to trust CoS. You must do so to enable the autoQoS generated commands.

This example shows how to enable autoQoS on Gigabit Ethernet interface 1/1:

```
Router(config)# interface gigabitethernet 1/1
Router(config-if)# auto qos voip cisco-phone
```

Displays the generated [received CoS to internal DSCP map](#).

```
Router# show running-config | include qos map cos-dscp
```

Configuring AutoQoS Support for Cisco IP Communicator

To configure autoQoS for Cisco IP Communicator, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 3	Router(config-if)# auto qos voip cisco-softphone	Configures autoQoS for Cisco IP Communicator.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

- To disable autoQoS on an interface, use the **no auto qos voip** interface configuration command.



Note The **no auto qos voip** interface configuration command does not delete the policy, class, and DSCP markdown maps created by autoQoS.

- You cannot configure support for Cisco IP Communicator on ports that are configured with the **switchport** keyword.
- PFC QoS supports 1023 aggregate policers and each use of the **auto qos voip cisco-softphone** command on a port uses two aggregate policers.

This example shows how to enable autoQoS on Gigabit Ethernet interface 1/1:

```
Router(config)# interface gigabitethernet 1/1
Router(config-if)# auto qos voip cisco-softphone
```

Displays the configured autoQoS commands.

```
Router# show auto qos interface type slot/port
```

Displays the policy map and policers created by autoQoS.

```
Router# show policy-map AUTOQOS-CISCO-SOFT-PHONE
```

Displays the class maps created by autoQoS.

```
Router# show class-map AUTOQOS-CISCO-SOFTPHONE-SIGNAL
Router# show class-map AUTOQOS-CISCO-SOFTPHONE-DATA
```

Displays the [DSCP markdown maps](#) created by autoQoS.

```
Router# show running-config | include qos map policed-dscp
```

Configuring AutoQoS Support for Marked Traffic

To configure autoQoS for marked traffic, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface type slot/port	Selects the interface to configure.
Step 3	Router(config-if)# auto qos voip trust	Configures autoQoS for marked traffic.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

When configuring autoQoS to trust marked traffic, note the following information:

- To disable autoQoS on an interface, use the **no auto qos voip** interface configuration command.



Note The **no auto qos voip** interface configuration command does not delete the [received CoS to internal DSCP map](#) created by autoQoS.

- For ports configured with the **switchport** command, you might see messages that instruct you to configure other ports to trust CoS. You must do so to enable the autoQoS generated commands.

This example shows how to enable autoQoS on Gigabit Ethernet interface 1/1:

```
Router(config)# interface gigabitethernet 1/1  
Router(config-if)# auto qos voip trust
```

Displays the configured autoQoS commands.

```
Router# show auto qos interface type slot/port
```

For ports configured with the **switchport** command, displays the generated [received CoS to internal DSCP map](#).

```
Router# show running-config | include qos map cos-dscp
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

29

MPLS QoS

- Terminology, page 29-2
- MPLS QoS Features, page 29-2
- MPLS QoS Overview, page 29-4
- MPLS QoS, page 29-5
- MPLS QoS Default Configuration, page 29-13
- MPLS QoS Commands, page 29-15
- MPLS QoS Restrictions, page 29-15
- How to Configure MPLS QoS, page 29-16
- MPLS DiffServ Tunneling Modes, page 29-26
- How to Configure Short Pipe Mode, page 29-30
- How to Configure Uniform Mode, page 29-34

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- MPLS QoS extends to MPLS traffic the PFC QoS features described in [Chapter 24, “PFC QoS Overview.”](#)
- This chapter provides supplemental information on MPLS QoS features. Be sure that you understand the PFC QoS features before you read this chapter.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Terminology

- *Class of Service* (CoS) refers to three bits in an 802.1Q header that are used to indicate the priority of the Ethernet frame as it passes through a switched network. The CoS bits in the 802.1Q header are commonly referred to as the 802.1p bits. To maintain QoS when a packet traverses both Layer 2 and Layer 3 domains, the type of service (ToS) and CoS values can be mapped to each other.
- *Classification* is the process used for selecting traffic to be marked for QoS.
- *Differentiated Services Code Point* (DSCP) is the first six bits of the ToS byte in the IP header. DSCP is only present in an IP packet.
- *E-LSP* is a label switched path (LSP) on which nodes infer the QoS treatment for MPLS packets exclusively from the experimental (EXP) bits in the MPLS header. Because the QoS treatment is inferred from the EXP (both class and drop precedence), several classes of traffic can be multiplexed onto a single LSP (use the same label). A single LSP can support up to eight classes of traffic because the EXP field is a 3-bit field. The maximum number of classes would be less after reserving some values for control plane traffic or if some of the classes have a drop precedence associated with them.
- *EXP bits* define the QoS treatment (per-hop behavior) that a node should give to a packet. It is the equivalent of the DiffServ Code Point (DSCP) in the IP network. A DSCP defines a class and drop precedence. The EXP bits are generally used to carry all the information encoded in the IP DSCP. In some cases, however, the EXP bits are used exclusively to encode the dropping precedence.
- *Frames* carry traffic at Layer 2. Layer 2 frames carry Layer 3 packets.
- *IP precedence* is the three most significant bits of the ToS byte in the IP header.
- *QoS tags* are prioritization values carried in Layer 3 packets and Layer 2 frames. A Layer 2 CoS label can have a value ranging between zero for low priority and seven for high priority. A Layer 3 IP precedence label can have a value ranging between zero for low priority and seven for high priority. IP precedence values are defined by the three most significant bits of the 1-byte ToS byte. A Layer 3 DSCP label can have a value between 0 and 63. DSCP values are defined by the six most significant bits of the 1-byte IP ToS field.
- *LERs* (label edge routers) are devices that impose and dispose of labels upon packets; also referred to as Provider Edge (PE) routers.
- *LSRs* (label switching routers) are devices that forward traffic based upon labels present in a packet; also referred to as Provider (P) routers.
- *Marking* is the process of setting a Layer 3 DSCP value in a packet. Marking is also the process of choosing different values for the MPLS EXP field to mark packets so that they have the priority that they require during periods of congestion.
- *Packets* carry traffic at Layer 3.
- *Policing* is limiting bandwidth used by a flow of traffic. Policing can mark or drop traffic.

MPLS QoS Features

- [MPLS Experimental Field, page 29-3](#)
- [Trust, page 29-3](#)
- [Classification, page 29-3](#)
- [Policing and Marking, page 29-3](#)

- [Preserving IP ToS, page 29-4](#)
- [EXP Mutation, page 29-4](#)
- [MPLS DiffServ Tunneling Modes, page 29-4](#)

MPLS Experimental Field

Setting the MPLS experimental (EXP) field value satisfies the requirement of service providers who do not want the value of the IP precedence field modified within IP packets transported through their networks.

By choosing different values for the MPLS EXP field, you can mark packets so that packets have the priority that they require during periods of congestion.

By default, the IP precedence value is copied into the MPLS EXP field during imposition. You can mark the MPLS EXP bits with an MPLS QoS policy.

Trust

For received Layer 3 MPLS packets, the PFC usually trusts the EXP value in the received topmost label. None of the following have any effect on MPLS packets:

- Interface trust state
- Port CoS value
- Policy-map **trust** command

For received Layer 2 MPLS packets, the PFC can either trust the EXP value in the received topmost label or apply port trust or policy trust to the MPLS packets for CoS and egress queueing purposes.

Classification

Classification is the process that selects the traffic to be marked. Classification accomplishes this by partitioning traffic into multiple priority levels, or classes of service. Traffic classification is the primary component of class-based QoS provisioning. The PFC makes classification decisions based on the EXP bits in the received topmost label of received MPLS packets (after a policy is installed). See the [“Configuring a Class Map to Classify MPLS Packets” section on page 29-17](#) for information.

Policing and Marking

Policing causes traffic that exceeds the configured rate to be discarded or marked down to a higher drop precedence. Marking is a way to identify packet flows to differentiate them. Packet marking allows you to partition your network into multiple priority levels or classes of service.

The MPLS QoS policing and marking features that you can implement depend on the received traffic type and the forwarding operation applied to the traffic. See [“Configuring a Policy Map” section on page 29-20](#) for information.

Preserving IP ToS

The PFC automatically preserves the IP ToS during all MPLS operations including imposition, swapping, and disposition. You do not need to enter a command to save the IP ToS.

EXP Mutation

You can configure up to eight egress EXP mutation maps to mutate the internal EXP value before it is written as the egress EXP value. You can attach egress EXP mutation maps to these interface types:

- LAN port subinterfaces
- Layer 3 VLAN interfaces
- Layer 3 LAN ports

You cannot attach EXP mutation maps to Layer 2 LAN ports (ports configured with the **switchport** command).

For configuration information, see the [“Configuring MPLS QoS Egress EXP Mutation” section on page 29-24](#).

MPLS DiffServ Tunneling Modes

The PFC uses MPLS DiffServ tunneling modes. Tunneling provides QoS transparency from one edge of a network to the other edge of the network. See the [“MPLS DiffServ Tunneling Modes” section on page 29-26](#) for information.

MPLS QoS Overview

MPLS QoS enables network administrators to provide differentiated types of service across an MPLS network. Differentiated service satisfies a range of requirements by supplying for each transmitted packet the service specified for that packet by its QoS. Service can be specified in different ways, for example, using the IP precedence bit settings in IP packets.

Specifying the QoS in the IP Precedence Field

When you send IP packets from one site to another, the IP precedence field (the first three bits of the DSCP field in the header of an IP packet) specifies the QoS. Based on the IP precedence marking, the packet is given the treatment configured for that quality of service. If the service provider network is an MPLS network, then the IP precedence bits are copied into the MPLS EXP field at the edge of the network. However, the service provider might want to set QoS for an MPLS packet to a different value determined by the service offering.

In that case, the service provider can set the MPLS EXP field. The IP header remains available for the customer’s use; the QoS of an IP packet is not changed as the packet travels through the MPLS network.


For more information, see the [“MPLS DiffServ Tunneling Modes” section on page 29-26](#).

MPLS QoS

- [MPLS Topology Overview, page 29-5](#)
- [LERs at the Input Edge of an MPLS Network, page 29-6](#)
- [LSRs in the Core of an MPLS Network, page 29-6](#)
- [LERs at the Output Edge of an MPLS Network, page 29-7](#)
- [LERs at the EoMPLS Edge, page 29-7](#)
- [LERs at the IP Edge \(MPLS, MPLS VPN\), page 29-8](#)
- [LSRs at the MPLS Core, page 29-11](#)

MPLS Topology Overview

Figure 29-1 MPLS Network Connecting Two Sites of a Customer's IP Network



- Networks are bidirectional, but for the purpose of this overview, the packets move left to right.
- CE1—Customer equipment 1
- PE1—Service provider ingress label edge router (LER)
- P1—Label switch router (LSR) within the core of the network of the service provider
- P2—LSR within the core of the network of the service provider
- PE2—service provider egress LER
- CE2—Customer equipment 2
- PE1 and PE2 are at the boundaries between the MPLS network and the IP network.

MPLS QoS supports IP QoS. For MPLS packets, the EXP value is mapped into an internal DSCP so that the PFC can apply non-MPLS QoS marking and policing.

For both the ingress and egress policies, MPLS QoS marking and policing decisions are made on a per-interface basis at an ingress PFC. The ingress interfaces are physical ports, subinterfaces, or VLANs.

The QoS policy ACLs are programmed in QoS TCAM separately for ingress and egress lookup. The ternary content addressable memory (TCAM) egress lookup takes place after the IP forwarding table (FIB) and NetFlow lookups are completed.

The results of each QoS TCAM lookup yield an index into RAM that contains policer configuration and policing counters. Additional RAM contains the microflow policer configuration; the microflow policing counters are maintained in the respective NetFlow entries that match the QoS ACL.

The results of ingress and egress aggregate and microflow policing are combined into a final policing decision. The out-of-profile packets can be either dropped or marked down in the DSCP.

LERs at the Input Edge of an MPLS Network



Note

Incoming labels are aggregate or nonaggregate. The aggregate label indicates that the arriving MPLS or MPLS VPN packet must be switched through an IP lookup to find the next hop and the outgoing interface. The nonaggregate label indicates that the packet contains the IP next hop information.

This section describes how edge LERs can operate at either the ingress or the egress side of an MPLS network.

At the ingress side of an MPLS network, LERs process packets as follows:

1. Layer 2 or Layer 3 traffic enters the edge of the MPLS network at the edge LER (PE1).
2. The PFC receives the traffic from the input interface and uses the 802.1p bits or the IP ToS bits to determine the EXP bits and to perform any classification, marking, and policing. For classification of incoming IP packets, the input service policy can also use access control lists (ACLs).
3. For each incoming packet, the PFC performs a lookup on the IP address to determine the next-hop router.
4. The appropriate label is pushed (imposition) into the packet, and the EXP value resulting from the QoS decision is copied into the MPLS EXP field in the label header.
5. The PFC forwards the labeled packets to the appropriate output interface for processing.
6. The PFC also forwards the 802.1p bits or the IP ToS bits to the output interface.
7. At the output interface, the labeled packets are differentiated by class for marking or policing. For LAN interfaces, egress classification is still based on IP, not on MPLS.
8. The labeled packets (marked by EXP) are sent to the core MPLS network.

LSRs in the Core of an MPLS Network

This section describes how LSRs used at the core of an MPLS network process packets:

1. Incoming MPLS-labeled packets (and 802.1p bits or IP ToS bits) from an edge LER (or other core device) arrive at the core LSR.
2. The PFC receives the traffic from the input interface and uses the EXP bits to perform classification, marking, and policing.
3. The PFC or DFCs perform a table lookup to determine the next-hop LSR.
4. An appropriate label is placed (swapped) into the packet and the MPLS EXP bits are copied into the label header.
5. The PFC forwards the labeled packets to the appropriate output interface for processing.
6. The PFC also forwards the 802.1p bits or the IP ToS bits to the output interface.
7. The outbound packet is differentiated by the MPLS EXP field for marking or policing.

8. The labeled packets (marked with EXP) are sent to another LSR in the core MPLS network or to an LER at the output edge.

**Note**

Within the service provider network, there is no IP precedence field for the queueing algorithm to use because the packets are MPLS packets. The packets remain MPLS packets until they arrive at PE2, the provider edge router.

LERs at the Output Edge of an MPLS Network

At the egress side of an MPLS network, LERs process packets as follows:

1. MPLS-labeled packets (and 802.1p bits or IP ToS bits) from a core LSR arrive at the egress LER (PE2) from the MPLS network backbone.
2. The PFC pops the MPLS labels (disposition) from the packets. Aggregate labels are classified using the original 802.1p bits or the IP ToS bits. Nonaggregate labels are classified with the EXP value by default.
3. For aggregate labels, the PFC performs a lookup on the IP address to determine the packet's destination; the PFC then forwards the packet to the appropriate output interface for processing. For non-aggregate labels, forwarding is based on the label. By default, non-aggregate labels are popped at the penultimate-hop router (next to last), not the egress PE router.
4. The PFC also forwards the 802.1p bits or the IP ToS bits to the output interface.
5. The packets are differentiated according to the 802.1p bits or the IP ToS bits and treated accordingly.

**Note**

The MPLS EXP bits allow you to specify the QoS for an MPLS packet. The IP precedence and DSCP bits allow you to specify the QoS for an IP packet.

LERs at the EoMPLS Edge

This section summarizes the Ethernet over MPLS (EoMPLS) QoS features that function on the LERs. EoMPLS QoS support is similar to IP-to-MPLS QoS:

- For EoMPLS, if the port is untrusted, the CoS trust state is automatically configured for VC type 4 (VLAN mode), not for VC type 5 (port mode). 802.1q CoS preservation across the tunnel is similar.
- Packets received on tunnel ingress are treated as untrusted for EoMPLS interfaces, except for VC Type 4 where trust CoS is automatically configured on the ingress port and policy marking is not applied.
- If the ingress port is configured as trusted, packets received on an EoMPLS interface are never marked by QoS policy in the original IP packet header (marking by IP policy works on untrusted ports).
- 802.1p CoS is preserved from entrance to exit, if available through the 802.1q header.
- After exiting the tunnel egress, queueing is based on preserved 802.1p CoS if 1p tag has been tunneled in the EoMPLS header (VC type 4); otherwise, queueing is based on the CoS derived from the QoS decision.

LERs at the IP Edge (MPLS, MPLS VPN)

This section provides information about QoS features for LERs at the ingress (CE-to-PE) and egress (PE-to-CE) edges for MPLS and MPLS VPN networks. Both MPLS and MPLS VPN support general MPLS QoS features. See the [“MPLS VPN” section on page 29-11](#) for additional MPLS VPN-specific QoS information.

IP to MPLS

- [IP to MPLS Overview, page 29-8](#)
- [Classification for IP-to-MPLS, page 29-8](#)
- [Classification for IP-to-MPLS Mode MPLS QoS, page 29-9](#)
- [Classification at IP-to-MPLS Ingress Port, page 29-9](#)
- [Classification at IP-to-MPLS Egress Port, page 29-9](#)

IP to MPLS Overview

The PFC provides the following MPLS QoS capabilities at the IP-to-MPLS edge:

- Assigning an EXP value based on the **platform qos trust** or **policy-map** command
- Marking an EXP value using a policy
- Policing traffic using a policy

This section provides information about the MPLS QoS classification that the PFC supports at the IP-to-MPLS edge. Additionally, this section provides information about the capabilities provided by the ingress and egress interface modules. For Ethernet to MPLS, the ingress interface, MPLS QoS, and egress interface features are similar to corresponding features for IP to MPLS.

Classification for IP-to-MPLS

The PFC ingress and egress policies for IP traffic classify traffic on the original received IP using **match** commands for IP precedence, IP DSCP, and IP ACLs. Egress policies do not classify traffic on the imposed EXP value nor on a marking done by an ingress policy.

After the PFC applies the port trust and QoS policies, it assigns the internal DSCP. The PFC then assigns the EXP value based on the internal DSCP-to-EXP global map for the labels that it imposes. If more than one label is imposed, the EXP value is the same in each label. The PFC preserves the original IP ToS when the MPLS labels are imposed.

The PFC assigns the egress CoS based on the internal DSCP-to-CoS global map. If the default internal DSCP-to-EXP and the internal DSCP-to-CoS maps are consistent, then the egress CoS has the same value as the imposed EXP.

If the ingress port receives both IP-to-IP and IP-to-MPLS traffic, classification should be used to separate the two types of traffic. For example, if the IP-to-IP and IP-to-MPLS traffic have different destination address ranges, you can classify traffic on the destination address, and then apply IP ToS policies to the IP-to-IP traffic and apply a policy (that marks or sets the EXP value in the imposed MPLS header) to the IP-to-MPLS traffic. See the following two examples:

- A PFC policy to mark IP ToS sets the internal DSCP—If it is applied to all traffic, then for IP-to-IP traffic, the egress port will rewrite the CoS (derived from the internal DSCP) to the IP ToS byte in the egress packet. For IP-to-MPLS traffic, the PFC will map the internal DSCP to the imposed EXP value.
- A PFC policy to mark MPLS EXP sets the internal DSCP—If it is applied to all traffic, then for IP-to-IP traffic, the egress port rewrites the IP ToS according to the ingress IP policy (or trust). The CoS is mapped from the ToS. For IP-to-MPLS traffic, the PFC will map the internal DSCP to the imposed EXP value.

Classification for IP-to-MPLS Mode MPLS QoS

MPLS QoS at the ingress to PE1 supports:

- Matching on IP precedence or DSCP values or filtering with an access group
- The **set mpls experimental imposition** and **police** commands

MPLS QoS at the egress of PE1 supports the **mpls experimental topmost** command.

Classification at IP-to-MPLS Ingress Port

Classification for IP-to-MPLS is the same as for IP-to-IP. LAN port classification is based on the received Layer 2 802.1Q CoS value.

Classification at IP-to-MPLS Egress Port

LAN port classification is based on the received EXP value and the egress CoS values is mapped from that value.

If the egress port is a trunk, the LAN ports copy the egress CoS into the egress 802.1Q field.

MPLS to IP

- [MPLS to IP Overview, page 29-9](#)
- [Classification for MPLS-to-IP, page 29-10](#)
- [Classification for MPLS-to-IP MPLS QoS, page 29-10](#)
- [Classification at MPLS-to-IP Ingress Port, page 29-10](#)
- [Classification at MPLS-to-IP Egress Port, page 29-11](#)

MPLS to IP Overview

MPLS QoS supports these capabilities at the MPLS-to-IP edge:

- Option to propagate EXP value into IP DSCP on exit from an MPLS domain per egress interface
- Option to use IP service policy on the MPLS-to-IP egress interface

This section provides information about the MPLS-to-IP MPLS QoS classification. Additionally, this section provides information about the capabilities provided by the ingress and egress modules.

For MPLS to Ethernet, the ingress interface, MPLS QoS, and egress interface features are similar to corresponding features for MPLS to IP except for the case of EoMPLS decapsulation where egress IP policy cannot be applied (packets can be classified as MPLS only).

Classification for MPLS-to-IP

The PFC assigns the internal DSCP (internal priority that the PFC assigns to each frame) based on the QoS result. The QoS result is affected by the following:

- Default trust EXP value
- Label type (per-prefix or aggregate)
- Number of VPNs
- Explicit NULL use
- QoS policy

There are three different classification modes:

- Regular MPLS classification—For nonaggregate labels, in the absence of MPLS recirculation, the PFC classifies the packet based on MPLS EXP ingress or egress policy. The PFC queues the packet based on COS derived from EXP-to-DSCP-to-CoS mapping. The underlying IP DSCP is either preserved after egress decapsulation, or overwritten from the EXP (through the EXP-to-DSCP map).
- IP classification for aggregate label hits in VPN CAM—The PFC does one of the following:
 - Preserves the underlying IP ToS
 - Rewrites the IP ToS by a value derived from the EXP-to-DSCP global map
 - Changes the IP ToS to any value derived from the egress IP policy

In all cases, egress queueing is based on the final IP ToS from the DSCP-to-CoS map.

- IP classification with aggregate labels not in VPN CAM—After recirculation, the PFC differentiates the MPLS-to-IP packets from the regular IP-to-IP packets based on the ingress reserved VLAN specified in the MPLS decapsulation adjacency. The reserved VLAN is allocated per VRF both for VPN and non-VPN cases. The ingress ToS after recirculation can be either the original IP ToS value, or derived from the original EXP value. The egress IP policy can overwrite this ingress ToS to an arbitrary value.



Note

For information about recirculation, see the [“Recirculation” section on page 37-4](#).

For incoming MPLS packets on the PE-to-CE ingress, the PFC supports MPLS classification only. Ingress IP policies are not supported. PE-to-CE traffic from the MPLS core is classified or policed on egress as IP.

Classification for MPLS-to-IP MPLS QoS

MPLS QoS at the ingress to PE2 supports matching on the EXP value and the **police** command.

MPLS QoS at the egress of PE2 supports matching on IP precedence or DSCP values or filtering with an access group and the **police** command.

Classification at MPLS-to-IP Ingress Port

LAN port classification is based on the EXP value. The **match mpls experimental** command matches on the EXP value in the received topmost label.

Classification at MPLS-to-IP Egress Port

Classification for MPLS-to-IP is the same as it is for IP-to-IP.

The LAN interface classification is based on the egress CoS.

If the egress port is a trunk, the LAN ports copy the egress CoS into the egress 802.1Q field.



Note

For MPLS to IP, egress IP ACL or QoS is not effective on the egress interface if the egress interface has MPLS IP (or tag IP) enabled. The exception is a VPN CAM hit, in which case the packet is classified on egress as IP.

MPLS VPN

The following PE MPLS QoS features are supported for MPLS VPN:

- Classification, policing, or marking of CE-to-PE IP traffic through the VPN subinterface
- Per-VPN QoS (per-port, per-VLAN, or per-subinterface)

For customer edge (CE)-to-PE traffic, or for CE-to-PE-to-CE traffic, the subinterface support allows you to apply IP QoS ingress or egress policies to subinterfaces and to physical interfaces. Per-VPN policing is also provided for a specific interface or subinterface associated with a given VPN on the CE side.

In situations when there are multiple interfaces belonging to the same VPN, you can perform per-VPN policing aggregation using the same shared policer in the ingress or egress service policies for all similar interfaces associated with the same PFC.

For aggregate VPN labels, the EXP propagation in recirculation case may not be supported because MPLS adjacency does not know which egress interface the final packet will use.



Note

For information on recirculation, see the [“Recirculation” section on page 37-4](#).

The PFC propagates the EXP value if all interfaces in the VPN have EXP propagation enabled.

The following PE MPLS QoS features are supported:

- General MPLS QoS features for IP packets
- Classification, policing, or marking of CE-to-PE IP traffic through the VPN subinterface
- Per-VPN QoS (per-port, per-VLAN, or per-subinterface)

LSRs at the MPLS Core

This section provides information about MPLS QoS features for LSRs at the core (MPLS-to-MPLS) for MPLS and MPLS VPN networks. Ingress features, egress interface, and PFC features for Carrier Supporting Carrier (CsC) QoS features are similar to those used with MPLS to MPLS described in the next section. A difference between CsC and MPLS to MPLS is that with CsC labels can be imposed inside the MPLS domain.

MPLS to MPLS

- [Classification for MPLS-to-MPLS, page 29-12](#)
- [Classification for MPLS-to-MPLS MPLS QoS, page 29-13](#)

- [Classification at MPLS-to-MPLS Ingress Port, page 29-13](#)
- [Classification at MPLS-to-MPLS Egress Port, page 29-13](#)

MPLS to MPLS Overview

MPLS QoS at the MPLS core supports the following:

- Per-EXP policing based on a service policy
- Copying the input topmost EXP value into the newly imposed EXP value
- Optional EXP mutation (changing of EXP values on an interface edge between two neighboring MPLS domains) on the egress boundary between MPLS domains
- Microflow policing based on individual label flows for a particular EXP value
- Optional propagation of topmost EXP value into the underlying EXP value when popping the topmost label from a multi-label stack.

The following section provides information about MPLS-to-MPLS MPLS QoS classification. Additionally, the section provides information about the capabilities provided by the ingress and egress modules.

Classification for MPLS-to-MPLS

For received MPLS packets, the PFC ignores the port trust state, the ingress CoS, and any policy-map **trust** commands. Instead, the PFC trusts the EXP value in the topmost label.



Note

The MPLS QoS ingress and egress policies for MPLS traffic classify traffic on the EXP value in the received topmost label when you enter the **match mpls experimental** command.

MPLS QoS maps the EXP value to the internal DSCP using the EXP-to-DSCP global map. What the PFC does next depends on whether it is swapping labels, imposing a new label, or popping a label:

- Swapping labels—When swapping labels, the PFC preserves the EXP value in the received topmost label and copies it to the EXP value in the outgoing topmost label. The PFC assigns the egress CoS using the internal DSCP-to-CoS global map. If the DSCP global maps are consistent, then the egress CoS is based on the EXP in the outgoing topmost label.

The PFC can mark down out-of-profile traffic using the **police** command's **exceed** and **violate** actions. It does not mark in-profile traffic, so the **conform** action must be transmitted and the **set** command cannot be used. If the PFC is performing a markdown, it uses the internal DSCP as an index into the internal DSCP markdown map. The PFC maps the result of the internal DSCP markdown to an EXP value using the internal DSCP-to-EXP global map. The PFC rewrites the new EXP value to the topmost outgoing label and does not copy the new EXP value to the other labels in the stack. The PFC assigns the egress CoS using the internal DSCP-to-CoS global map. If the DSCP maps are consistent, then the egress CoS is based on the EXP value in the topmost outgoing label.

- Imposing an additional label—When imposing a new label onto an existing label stack, the PFC maps the internal DSCP to the EXP value in the imposed label using the internal DSCP-to-EXP map. It then copies the EXP value in the imposed label to the underlying swapped label. The PFC assigns the egress CoS using the internal DSCP-to-CoS global map. If the DSCP maps are consistent, the egress CoS is based on the EXP value in the imposed label.

The PFC can mark in-profile and mark down out-of-profile traffic. After it marks the internal DSCP, the PFC uses the internal DSCP-to-EXP global map to map the internal DSCP to the EXP value in the newly imposed label. The PFC then copies the EXP in the imposed label to the underlying swapped label. The PFC assigns the egress CoS using the internal DSCP-to-CoS global map. Therefore, the egress CoS is based on the EXP in the imposed label.

- Popping a label—When popping a label from a multi-label stack, the PFC preserves the EXP value in the exposed label. The PFC assigns the egress CoS using the internal DSCP-to-CoS global map. If the DSCP maps are consistent, then the egress CoS is based on the EXP value in the popped label.
- If EXP propagation is configured for the egress interface, the PFC maps the internal DSCP to the EXP value in the exposed label using the DSCP-to-EXP global map. The PFC assigns the egress CoS using the internal DSCP-to-CoS global map. If the DSCP maps are consistent, the egress CoS is based on the EXP value in the exposed label.

Classification for MPLS-to-MPLS MPLS QoS

MPLS QoS at the ingress to P1 or P2 supports the following:

- Matching with the **mpls experimental topmost** command
- The **set mpls experimental imposition, police**, and **police** with **set imposition** commands

MPLS QoS at the egress of P1 or P2 supports matching with the **mpls experimental topmost** command.

Classification at MPLS-to-MPLS Ingress Port

LAN port classification is based on the egress CoS from the PFC. The **match mpls experimental** command matches on the EXP value in the received topmost label.

Classification at MPLS-to-MPLS Egress Port

LAN port classification is based on the egress CoS value from the PFC. The **match mpls experimental** command matches on the egress CoS; it does not match on the EXP in the topmost label. If the egress port is a trunk, the LAN ports copy the egress CoS into the egress 802.1Q field.

MPLS QoS Default Configuration

Feature	Default Value
PFC QoS global enable state	With all other PFC QoS parameters at default values, default EXP is mapped from IP precedence. With PFC QoS enabled and all other PFC QoS parameters at default values, PFC QoS sets Layer 3 DSCP to zero (untrusted ports only), Layer 2 CoS to zero, the imposed EXP to zero in all traffic transmitted from LAN ports (default is untrusted). For trust CoS, the default EXP value is mapped from CoS; for trust DSCP, the default EXP value is mapped from IP precedence.
PFC QoS port enable state	Enabled when PFC QoS is globally enabled
Port CoS value	0
Microflow policing	Enabled

Feature	Default Value
IntraVLAN microflow policing	Disabled
Port-based or VLAN-based PFC QoS	Port-based
EXP to DSCP map (DSCP set from EXP values)	EXP 0 = DSCP 0 EXP 1 = DSCP 8 EXP 2 = DSCP 16 EXP 3 = DSCP 24 EXP 4 = DSCP 32 EXP 5 = DSCP 40 EXP 6 = DSCP 48 EXP 7 = DSCP 56
IP precedence to DSCP map (DSCP set from IP precedence values)	IP precedence 0 = DSCP 0 IP precedence 1 = DSCP 8 IP precedence 2 = DSCP 16 IP precedence 3 = DSCP 24 IP precedence 4 = DSCP 32 IP precedence 5 = DSCP 40 IP precedence 6 = DSCP 48 IP precedence 7 = DSCP 56
DSCP to EXP map (EXP set from DSCP values)	DSCP 0–7 = EXP 0 DSCP 8–15 = EXP 1 DSCP 16–23 = EXP 2 DSCP 24–31 = EXP 3 DSCP 32–39 = EXP 4 DSCP 40–47 = EXP 5 DSCP 48–55 = EXP 6 DSCP 56–63 = EXP 7
Marked-down DSCP from DSCP map	Marked-down DSCP value equals original DSCP value (no mark down)
EXP mutation map	No mutation map by default
Policers	None
Policy maps	None
MPLS flow mask in NetFlow table	Label + EXP value

Feature	Default Value
MPLS core QoS	There are four possibilities at the MPLS core QoS: - Swapping—Incoming EXP field is copied to outgoing EXP field. - Swapping + imposition—Incoming EXP field is copied to both the swapped EXP field and the imposed EXP field. Note If there is a service policy with a set for EXP field, its EXP field will be placed into the imposed label and also into the swapped label. - Disposition of topmost label—Exposed EXP field is preserved. - Disposition of only label—Exposed IP DSCP is preserved.
MPLS to IP edge QoS	Preserve the exposed IP DSCP

MPLS QoS Commands

MPLS QoS supports the following MPLS QoS commands:

- **match mpls experimental topmost**
- **set mpls experimental imposition**
- **police**
- **platform qos map exp-dscp**
- **platform qos map dscp-exp**
- **platform qos map exp-mutation**
- **platform qos exp-mutation**
- **show platform qos mpls**
- **no platform qos mpls trust exp**



Note

For information about supported non-MPLS QoS commands, see [Chapter 24, “PFC QoS Overview.”](#)

The following commands are not supported:

- **set qos-group**
- **set discard-class**

MPLS QoS Restrictions

When configuring MPLS QoS, follow these guidelines and restrictions:

- For IP-to-MPLS or EoMPLS imposition when the received packet is an IP packet:

- When QoS is disabled, the EXP value is based on the received IP ToS.
 - When QoS is queuing only, the EXP value is based on the received IP ToS.
- For EoMPLS imposition when the received packet is a non-IP packet:
 - When QoS is disabled, the EXP value is based on the ingress CoS.
 - When QoS is queuing only, the EXP value is based on the received IP ToS.
- For MPLS-to-MPLS operations:
 - Swapping when QoS is disabled, the EXP value is based on the original EXP value (in the absence of EXP mutation).
 - Swapping when QoS is queuing only, the EXP value is based on the original EXP value (in the absence of EXP mutation).
 - Imposing additional label when QoS is disabled, the EXP value is based on the original EXP value (in the absence of EXP mutation).
 - Imposing an additional label when QoS is queuing only, the EXP value is based on the original EXP value (in the absence of EXP mutation).
 - Popping one label when QoS is disabled, the EXP value is based on the underlying EXP value.
 - Popping one label when QoS is queuing only, the EXP value is based on the underlying EXP value.
- EXP value is irrelevant to MPLS-to-IP disposition.
- The **no platform qos rewrite ip dscp** command is incompatible with MPLS. The default **platform qos rewrite ip dscp** command must remain enabled in order for the PFC to assign the correct EXP value for the labels that it imposes.
- The **no platform qos mpls trust exp** command allows you to treat MPLS packets similarly to Layer 2 packets for CoS and egress queueing purposes by applying port trust or policy trust instead of the default EXP value.
- The **bandwidth percent** command is not supported on Ten-Gigabit Ethernet interfaces and subinterfaces.

How to Configure MPLS QoS

- [Enabling Queueing-Only Mode, page 29-17](#)
- [Configuring a Class Map to Classify MPLS Packets, page 29-17](#)
- [Configuring a Policy Map, page 29-20](#)
- [Displaying a Policy Map, page 29-24](#)
- [Configuring MPLS QoS Egress EXP Mutation, page 29-24](#)
- [Configuring EXP Value Maps, page 29-25](#)

Enabling Queueing-Only Mode

To enable queueing-only mode, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos queueing-only	Enables queueing-only mode.
Step 2	Router(config)# end	Exits configuration mode.

When you enable queueing-only mode, the router does the following:

- Disables marking and policing globally
- Configures all ports to trust Layer 2 CoS



Note

The switch applies the port CoS value to untagged ingress traffic and to traffic that is received through ports that cannot be configured to trust CoS.

This example shows how to enable queueing-only mode:

```
Router# configure terminal
Router(config)# platform qos queueing-only
Router(config)# end
Router#
```

Restrictions and Usage Guidelines

If QoS is disabled (**no platform qos**) for the PFC, the EXP value is determined as follows:

- For IP-to-MPLS or EoMPLS imposition when the received packet is an IP packet when QoS is queuing only (**platform qos queueing-only**), the EXP value is based on the received IP ToS.
- For EoMPLS imposition when the received packet is a non-IP packet when QoS is queuing only, the EXP value is based on the received IP ToS.
- For MPLS-to-MPLS operations:
 - Swapping when QoS is queuing only, the EXP value is based on the original EXP value (in the absence of EXP mutation).
 - Imposing additional label when QoS is queuing only, the EXP value is based on the original EXP value (in the absence of EXP mutation).
 - Popping one label when QoS is queuing only, the EXP value is based on the underlying EXP value.
- The EXP value is irrelevant to MPLS-to-IP disposition.

Configuring a Class Map to Classify MPLS Packets

You can use the **match mpls experimental topmost** command to define traffic classes inside the MPLS domain by packet EXP values. This allows you to define service policies to police the EXP traffic on a per-interface basis by using the **police** command.

To configure a class map, perform this task beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# class-map <i>class_name</i>	Specifies the class map to which packets will be matched.
Step 2	Router(config-cmap)# match mpls experimental topmost <i>value</i>	Specifies the packet characteristics that will be matched to the class.
Step 3	Router(config-cmap)# exit	Exits class-map configuration mode.

This example shows that all packets that contain MPLS experimental value 3 are matched by the traffic class named exp3:

```

Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# class-map exp3
Router(config-cmap)# match mpls experimental topmost 3
Router(config-cmap)# exit
Router(config)# policy-map exp3
Router(config-pmap)# class exp3
Router(config-pmap-c)# police 1000000 8000000 conform-action transmit exceed-action drop
Router(config-pmap-c)# exit
Router(config-pmap)# end
Router# show class exp3
Class Map match-all exp3 (id 61)
  Match mpls experimental topmost 3
Router# show policy-map exp3
Policy Map exp3
Class exp3
  police cir 1000000 bc 8000000 be 8000000 conform-action transmit exceed-action drop
Router# show running-config interface gigabitethernet 3/27
Building configuration...

Current configuration : 173 bytes
!
interface GigabitEthernet3/27
  ip address 47.0.0.1 255.0.0.0
  tag-switching ip
end

Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 3/27
Router(config-if)# service-policy input exp3
Router(config-if)#
Router#
Enter configuration commands, one per line. End with CNTL/Z.
Router# show running-config interface gigabitethernet 3/27
Building configuration...

Current configuration : 173 bytes
!
interface GigabitEthernet3/27
  ip address 47.0.0.1 255.0.0.0
  tag-switching ip
  service-policy input exp3
end

Router#
1w4d: %SYS-5-CONFIG_I: Configured from console by console
Router# show platform qos mpls
QoS Summary [MPLS]:          (* - shared aggregates, Mod - switch module)

```

Int	Mod	Dir	Class-map	DSCP	Agg Id	Trust	Fl Id	AgForward-By	AgPoliced-By
Gi3/27	5	In	exp3	0	2	dscp	0	0	0
All	5	-	Default	0	0*	No	0	3466140423	0

```
Router# show policy-map interface gigabitethernet 3/27
GigabitEthernet3/27
```

```
Service-policy input: exp3
```

```
class-map: exp3 (match-all)
  Match: mpls experimental topmost 3
  police :
    1000000 bps 8000000 limit 8000000 extended limit
  Earl in slot 5 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps
```

```
Class-map: class-default (match-any)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any
```

```
Router# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router(config)# interface gigabitethernet 3/27
```

```
Router(config-if)# service-policy output ip2tag
```

```
Router(config-if)# end
```

```
Router# show platform qos ip
```

```
QoS Summary [IPv4]: (* - shared aggregates, Mod - switch module)
```

Int	Mod	Dir	Class-map	DSCP	Agg Id	Trust	Fl Id	AgForward-By	AgPoliced-By
Vl300	5	In	x	44	1	No	0	0	0
Gi3/27	5	Out	iptcp	24	2	--	0	0	0
All	5	-	Default	0	0*	No	0	3466610741	0

Restrictions and Usage Guidelines

- The **match mpls experimental** command specifies the name of an EXP field value to be used as the match criterion against which packets are checked to determine if they belong to the class specified by the class map.
- To use the **match mpls experimental** command, you must first enter the **class-map** command to specify the name of the class whose match criteria you want to establish. After you identify the class, you can use the **match mpls experimental** command to configure its match criteria.
- If you specify more than one command in a class map, only the last command entered applies. The last command overrides the previously entered commands.

Configuring a Policy Map

You can attach only one policy map to an interface. Policy maps can contain one or more policy map classes, each with different policy map commands.

Configure a separate policy map class in the policy map for each type of traffic that an interface receives. Put all commands for each type of traffic in the same policy map class. MPLS QoS does not attempt to apply commands from more than one policy map class to matched traffic.

Configuring a Policy Map to Set the EXP Value on All Imposed Labels

To set the value of the MPLS EXP field on all imposed label entries, use the **set mpls experimental imposition** command in QoS policy-map class configuration mode. To disable the setting, use the **no** form of this command.



Note

The **set mpls experimental imposition** command replaces the **set mpls experimental** command.

	Command	Purpose
Step 1	Router(config)# policy-map <i>policy_name</i>	Creates a policy map.
Step 2	Router(config-pmap)# class-map <i>name</i> [match-all match-any]	Accesses the QoS class-map configuration mode to configure QoS class maps.
Step 3	Router(config-pmap-c)# set mpls experimental imposition { <i>mpls-exp-value</i> <i>from-field</i> [table <i>table-map-name</i>]}	Sets the value of the MPLS experimental (EXP) field on all imposed label entries.
Step 4	Router(config-pmap-c)# exit	Exits class-map configuration mode.

The following example sets the MPLS EXP imposition value according to the DSCP value defined in the MPLS EXP value 3:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# access-1 101 p tcp any any
Router(config)# class-map iptcp
Router(config-cmap)# match acc 101
Router(config-cmap)# exit
Router(config)#
Router(config-cmap)# policy-map ip2tag
Router(config-pmap)# class iptcp
Router(config-pmap-c)# set mpls exp imposition 3
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)#
Router#
1w4d: %SYS-5-CONFIG_I: Configured from console by console
Router#
Router# show policy-map ip2tag
  Policy Map ip2tag
    Class iptcp
      set mpls experimental imposition 3
Router# show class iptcp
  Class Map match-all iptcp (id 62)
    Match access-group101
```

```

Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# interface gigabitethernet 3/27
Router(config-if)# ser in ip2tag
Router(config-if)#
Routers
1w4d: %SYS-5-CONFIG_I: Configured from console by console
Router# show pol ip2tag
    Policy Map ip2tag
      Class iptcp
        set mpls experimental imposition 3
Router# show class-map iptcp
    Class Map match-all iptcp (id 62)
      Match access-group 101

Router# show access-1 101
Extended IP access list 101
    10 permit tcp any any
Router# show policy-map interface gigabitethernet 3/27
GigabitEthernet3/27

Service-policy input: ip2tag

    class-map: iptcp (match-all)
      Match: access-group 101
      set mpls experimental 3:
      Earl in slot 5 :
        0 bytes
        5 minute offered rate 0 bps
        aggregate-forwarded 0 bytes

    class-map: class-default (match-any)
      Match: any

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any

```

This example shows how to verify the configuration:

```

Router# show policy map ip2tag
    Policy Map ip2tag
      Class iptcp
        set mpls experimental imposition 3

```

EXP Value Imposition Guidelines and Restrictions

When setting the EXP value on all imposed labels, follow these guidelines and restrictions:

- Use the **set mpls experimental imposition** command during label imposition. This command sets the MPLS EXP field on all imposed label entries.
- The **set mpls experimental imposition** command is supported only on input interfaces (imposition).
- The **set mpls experimental imposition** command does not mark the EXP value directly; instead, it marks the internal DSCP that is mapped to EXP through the internal DSCP-to-EXP global map.
- It is important to note that classification (based on the original received IP header) and marking (done to the internal DSCP) do not distinguish between IP-to-IP traffic and IP-to-MPLS traffic. The commands that you use to mark IP ToS and mark EXP have the same result as when you mark the internal DSCP.

- To set the pushed label entry value to a value different from the default value during label imposition, use the **set mpls experimental imposition** command.
- You optionally can use the **set mpls experimental imposition** command with the IP precedence, DSCP field, or QoS IP ACL to set the value of the MPLS EXP field on all imposed label entries.
- When imposing labels onto the received IP traffic with the PFC, you can mark the EXP field using the **set mpls experimental imposition** command.

Configuring a Policy Map Using the Police Command

Policing is a function in the PFC hardware that provides the ability to rate limit a particular traffic class to a specific rate. The PFC supports aggregate policing and microflow policing.

Aggregate policing meters all traffic that ingresses into a port, regardless of different source, destination, protocol, source port, or destination port. Microflow policing meters all traffic that ingresses into a port, on a per flow (per source, destination, protocol, source port, and destination port). For additional information on aggregate and microflow policing, see [Chapter 25, “Classification, Marking, and Policing.”](#)

	Command	Purpose
Step 1	Router(config)# policy-map <i>policy_name</i>	Creates a policy map.
Step 2	Router(config-pmap)# class-map <i>name</i> [match-all match-any]	Accesses the QoS class map configuration mode to configure QoS class maps.
Step 3	Router(config-pmap-c)# police { <i>aggregate name</i> }	Adds the class to a shared aggregate policer.
Step 4	Router(config-pmap-c)# police <i>bps burst_normal burst_max conform-action action exceed-action action violate-action action</i>	Creates a per-class-per-interface policer.
Step 5	Router(config-pmap-c)# police flow { <i>bps [burst_normal] [conform-action action] [exceed-action action]</i> }	Creates an ingress flow policer. (Not supported in egress policy.)
Step 6	Router(config-pmap-c)# exit	Exits class-map configuration mode.

This is an example of creating a policy map with a policer:

```
Router(config)# policy-map ip2tag
Router(config-pmap)# class iptcp
Router(config-pmap-c)# no set mpls exp topmost 3
Router(config-pmap-c)# police 1000000 1000000 c set-mpls-exp?
set-mpls-exp-imposition-transmit

Router(config-pmap-c)# police 1000000 1000000 c set-mpls-exp-imposit 3 e d
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface gigabitethernet 3/27
Router(config-if)# ser in ip2tag
Router(config-if)#
```

This is an example of verifying the configuration:

```
Router# show pol ip2tag
Policy Map ip2tag
Class iptcp
  police cir 1000000 bc 1000000 be 1000000 conform-action
set-mpls-exp-imposition-transmit 3 exceed-action drop
Router# show running-config interface gigabitethernet 3/27
```

```

Building configuration...

Current configuration : 202 bytes
!
interface GigabitEthernet3/27
  logging event link-status
  service-policy input ip2tag
end

Router# show policy interface gigabitethernet 3/27
GigabitEthernet3/27

Service-policy input: ip2tag

class-map: iptcp (match-all)
  Match: access-group 101
  police :
    1000000 bps 1000000 limit 1000000 extended limit
  Earl in slot 5 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: set-mpls-exp-imposition-transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

class-map: class-default (match-any)
  Match: any

Class-map: class-default (match-any)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any

```

Restrictions and Usage Guidelines

The following restrictions and guidelines apply when using the **police** command to configure a policy map:

- With MPLS, the **exceed-action** *action* command and the **violate-action** *action* command work similarly to IP usage. The packet may get dropped or the EXP value is marked down.
- With MPLS, the **set-dscp transmit** *action* command and the **set-prec-transmit** *action* command set the internal DSCP that is mapped into the CoS bits, which affects queueing, however, they do not change the EXP value, except for imposition.
- When swapping labels for received MPLS traffic with the PFC, you can mark down out-of-profile traffic using the **police** command **exceed-action policed-dscp-transmit** and **violate-action policed-dscp-transmit** keywords. The PFC does not mark in-profile traffic; when marking down out-of-profile traffic, the PFC marks the outgoing topmost label. The PFC does not propagate the marking down through the label stack.
- With MPLS, the flow key is based on the label and EXP value; there is no flowmask option. Otherwise, flow key operation is similar to IP-to-IP.
- You can use the **police** command to set the pushed label entry value to a value different from the default value during label imposition.
- When imposing labels onto the received IP traffic with the PFC, you can mark the EXP field using the **conform-action set-mpls-exp-imposition-transmit** keywords.
- During IP-to-MPLS imposition, IP ToS marking is not supported. If you configure a policy to mark IP ToS, the PFC marks the EXP value.

Displaying a Policy Map

You can display a policy map with an interface summary for MPLS QoS classes or with the configuration of all classes configured for all service policies on the specified interface.

Displaying the Configuration of All Classes

To display the configuration of all classes configured for all service policies on the specified interface, perform this task:

Command	Purpose
Router# show policy interface <i>interface_type</i> <i>interface_number</i>	Displays the configuration of all classes configured for all policy maps on the specified interface.

This example shows the configurations for all classes on Gigabit Ethernet interface 3/27:

```
Router# show policy interface gigabitethernet 3/27
GigabitEthernet3/27

Service-policy input: ip2tag

  class-map: iptcp (match-all)
    Match: access-group 101
    police :
      1000000 bps 1000000 limit 1000000 extended limit
    Earl in slot 5 :
      0 bytes
      5 minute offered rate 0 bps
      aggregate-forwarded 0 bytes action: set-mpls-exp-imposition-transmit
      exceeded 0 bytes action: drop
      aggregate-forward 0 bps exceed 0 bps

  class-map: class-default (match-any)
    Match: any

  Class-map: class-default (match-any)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: any
```

Configuring MPLS QoS Egress EXP Mutation

- [Configuring Named EXP Mutation Maps, page 29-25](#)
- [Attaching an Egress EXP Mutation Map to an Interface, page 29-25](#)

Configuring Named EXP Mutation Maps

To configure a named EXP mutation map, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos map exp-mutation <i>name mutated_exp1 mutated_exp2 mutated_exp3</i> <i>mutated_exp4 mutated_exp5 mutated_exp6</i> <i>mutated_exp7 mutated_exp8</i>	Configures a named EXP mutation map.
Step 2	Router(config)# end	Exits configuration mode.

When configuring a named EXP mutation map, note the following information:

- You can enter up to eight input EXP values that map to a mutated EXP value.
- You can enter multiple commands to map additional EXP values to a mutated EXP value.
- You can enter a separate command for each mutated EXP value.
- You can configure 15 ingress EXP mutation maps to mutate the internal EXP value before it is written as the ingress EXP value. You can attach ingress EXP mutation maps to any interface that PFC QoS supports.
- PFC QoS derives the egress EXP value from the internal DSCP value. If you configure ingress EXP mutation, PFC QoS does not derive the ingress EXP value from the mutated EXP value.

Attaching an Egress EXP Mutation Map to an Interface

To attach an egress EXP mutation map to an interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port[.subinterface]} {port-channel number[.subinterface]}}	Selects the interface to configure.
Step 2	Router(config-if)# platform qos exp-mutation <i>exp-mutation-table-name</i>	Attaches an egress EXP mutation map to the interface.
Step 3	Router(config-if)# end	Exits configuration mode.

This example shows how to attach the egress EXP mutation map named mutemap2:

```
Router(config)# interface gigabitethernet 3/26
Router(config-if)# platform qos exp-mutation mutemap2
Router(config-if)# end
```

Configuring EXP Value Maps

- [Configuring an Ingress-EXP to Internal-DSCP Map, page 29-26](#)
- [Configuring a Named Egress-DSCP to Egress-EXP Map, page 29-26](#)

Configuring an Ingress-EXP to Internal-DSCP Map

To configure an ingress-EXP to internal-DSCP map, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos map exp-dscp values	Configures the ingress-EXP value to internal-DSCP map. You must enter eight DSCP values corresponding to the EXP values. Valid values are 0 through 63.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure an ingress-EXP to internal-DSCP map:

```
Router(config)# platform qos map exp-dscp 43 43 43 43 43 43 43 43
Router(config)#
```

This example shows how to verify the configuration:

```
Router(config)# show platform qos map exp-dscp
Exp-dscp map:
  exp:   0   1   2   3   4   5   6   7
-----
  dscp: 43 43 43 43 43 43 43 43
```

Configuring a Named Egress-DSCP to Egress-EXP Map

To configure a named egress-DSCP to egress-EXP map, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos map dscp-exp dscp_values to exp_values	Configures a named egress-DSCP to egress-EXP map. You can enter up to eight DSCP values at one time to a single EXP value. Valid values are 0 through 7.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to configure a named egress-DSCP to egress-EXP map:

```
Router(config)# platform qos map dscp-exp 20 25 to 3
Router(config)#
```

MPLS DiffServ Tunneling Modes

Tunneling provides QoS the ability to be transparent from one edge of a network to the other edge of the network. A tunnel starts where there is label imposition. A tunnel ends where there is label disposition; that is, where the label is removed from the stack, and the packet goes out as an MPLS packet with a different per-hop behavior (PHB) layer underneath or as an IP packet with the IP PHB layer.

For the PFC, there are two ways to forward packets through a network:

- **Short Pipe mode**—In Short Pipe mode, the egress PE router uses the original packet marking instead of the marking used by the intermediate provider (P) routers. EXP marking does not propagate to the packet ToS byte.

For a description of this mode, see the [“Short Pipe Mode” section on page 29-27](#).

For the configuration information, see the “How to Configure Short Pipe Mode” section on page 29-30.

- Uniform mode—In Uniform mode, the marking in the IP packet may be manipulated to reflect the service provider’s QoS marking in the core. This mode provides consistent QoS classification and marking throughout the network including CE and core routers. EXP marking is propagated to the underlying ToS byte.

For a description, see the “Uniform Mode” section on page 29-28.

For the configuration procedure, see the “How to Configure Uniform Mode” section on page 29-34.

Both tunneling modes affect the behavior of edge and penultimate label switching routers (LSRs) where labels are put onto packets and removed from packets. They do not affect label swapping at intermediate routers. A service provider can choose different types of tunneling modes for each customer.

For additional information, see “MPLS DiffServ Tunneling Modes” at this URL:

http://www.cisco.com/en/US/docs/ios-xml/ios/mp_te_diffserv/configuration/15-mt/mp-diffserv-tun-mode.html.

Short Pipe Mode

Short pipe mode is used when the customer and service provider are in different DiffServ domains. It allows the service provider to enforce its own DiffServ policy while preserving customer DiffServ information, which provides a DiffServ transparency through the service provider network.


QoS policies implemented in the core do not propagate to the packet ToS byte. The classification based on MPLS EXP value ends at the customer-facing egress PE interface; classification at the customer-facing egress PE interface is based on the original IP packet header and not the MPLS header.



Note

The presence of an egress IP policy (based on the customer’s PHB marking and not on the provider’s PHB marking) automatically implies the Short Pipe mode.

Figure 29-2 Short Pipe Mode Operation with VPNs



Short Pipe mode functions as follows:

1. CE1 transmits an IP packet to PE1 with an IP DSCP value of 1.
2. PE1 sets the MPLS EXP field to 5 in the imposed label entries.
3. PE1 transmits the packet to P1.
4. P1 sets the MPLS EXP field value to 5 in the swapped label entry.
5. P1 transmits the packet to P2.
6. P2 pops the IGP label entry.
7. P2 transmits the packet to PE2.
8. PE2 pops the BGP label.
9. PE2 transmits the packet to CE2, but does QoS based on the IP DSCP value.

For additional information, see “MPLS DiffServ Tunneling Modes” at this URL:


http://www.cisco.com/en/US/docs/ios-xml/ios/mp_te_diffserv/configuration/15-mt/mp-diffserv-tun-mode.html

Short Pipe Mode Restrictions

Short Pipe mode is not supported if the MPLS-to-IP egress interface is EoMPLS (the adjacency has the end of marker (EOM) bit set).

Uniform Mode

In Uniform mode, packets are treated uniformly in the IP and MPLS networks; that is, the IP precedence value and the MPLS EXP bits always correspond to the same PHB. Whenever a router changes or recolors the PHB of a packet, that change must be propagated to all encapsulation markings. The propagation is performed by a router only when a PHB is added or exposed due to label imposition or disposition on any router in the packet's path. The color must be reflected everywhere at all levels. For example, if a packet's QoS marking is changed in the MPLS network, the IP QoS marking reflects that change.

Figure 29-3 Uniform Mode Operation

*In both the MPLS-to-MPLS and the MPLS-to-IP cases, the PHBs of the topmost popped label is copied into the new top label or the IP DSCP if no label remains

The procedure varies according to whether IP precedence bit markings or DSCP markings are present.

The following actions occur if there are IP precedence bit markings:

1. IP packets arrive in the MPLS network at PE1, the service provider edge router.
2. A label is copied onto the packet.
3. If the MPLS EXP field value is recolored (for example, if the packet becomes out-of-rate because too many packets are being transmitted), that value is copied to the IGP label. The value of the BGP label is not changed.
4. At the penultimate hop, the IGP label is removed. That value is copied into the next lower level label.
5. When all MPLS labels have been removed from the packet that is sent out as an IP packet, the IP precedence or DSCP value is set to the last changed EXP value in the core.

The following is an example when there are IP precedence bit markings:

1. At CE1 (customer equipment 1), the IP packet has an IP precedence value of 3.
2. When the packet arrives in the MPLS network at PE1 (the service provider edge router), the IP precedence value of 3 is copied to the imposed label entries of the packet.
3. The MPLS EXP field in the IGP label header might be changed within the MPLS core (for example, at P1) by a mark down.



Note

Because the IP precedence bits are 3, the BGP label and the IGP label also contain 3 because in Uniform mode, the labels always are identical. The packet is treated uniformly in the IP and MPLS networks.

Uniform Mode Restrictions

If the egress IP ACLs or service policies are configured on the MPLS-to-IP exit point, the Uniform mode is always enforced because of recirculation.

MPLS DiffServ Tunneling Restrictions and Usage Guidelines

The MPLS DiffServ tunneling restrictions and usage guidelines are as follows:

- One label-switched path (LSP) can support up to eight classes of traffic (that is, eight PHBs) because the MPLS EXP field is a 3-bit field.
- MPLS DiffServ tunneling modes support E-LSPs. An E-LSP is an LSP on which nodes determine the QoS treatment for MPLS packet exclusively from the EXP bits in the MPLS header.

The following features are supported with the MPLS differentiated service (DiffServ) tunneling modes:

- MPLS per-hop behavior (PHB) layer management. (Layer management is the ability to provide an additional layer of PHB marking to a packet.)
- Improved scalability of the MPLS layer management by control on managed customer edge (CE) routers.
- MPLS can tunnel a packet's QoS (that is, the QoS is transparent from edge to edge). With QoS transparency, the IP marking in the IP packet is preserved across the MPLS network.
- The MPLS EXP field can be marked differently and separately from the PHB marked in the IP precedence or DSCP field.

How to Configure Short Pipe Mode

- [Ingress PE Router—Customer Facing Interface, page 29-30](#)
- [Configuring Ingress PE Router—P Facing Interface, page 29-32](#)
- [Configuring the P Router—Output Interface, page 29-33](#)
- [Configuring the Egress PE Router—Customer Facing Interface, page 29-34](#)



Note

- The steps that follow show one way, but not the only way, to configure Short Pipe mode.
- The Short Pipe mode on the egress PE (or PHP) is automatically configured when you attach to the interface an egress service policy that includes an IP class.

Ingress PE Router—Customer Facing Interface

This procedure configures a policy map to set the MPLS EXP field in the imposed label entries.

To set the EXP value, the ingress LAN port must be untrusted.

For MPLS and VPN, the ingress PE supports all ingress PFC IP policies.

To configure a policy map to set the MPLS EXP field in the imposed label entries, perform this task:

	Command	Purpose
Step 1	Router(config)# access-list <i>ipv4_acl_number_or_name</i> permit any	Creates an IPv4 access list.
Step 2	Router(config)# class-map <i>class_name</i>	Creates a class map.
Step 3	Router(config-cmap)# match access-group <i>ipv4_acl_number_or_name</i>	Configures the class map to filter with the ACL created in step 1.
Step 4	Router(config)# policy-map <i>policy_map_name</i>	Creates a named QoS policy.
Step 5	Router(config-pmap)# class <i>class_name</i>	Configures the policy to use the class map created in step 2.
Step 6	Router(config-pmap-c)# police <i>bits_per_second</i> [<i>normal_burst_bytes</i>] conform-action set-mpls-exp-transmit <i>exp_value</i> exceed-action drop	Configures policing, including the following: - Action to take on packets that conform to the rate limit specified in the service level agreement (SLA). - Action to take on packets that exceed the rate limit specified in the SLA. The <i>exp_value</i> sets the MPLS EXP field.
Step 7	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 8	Router(config-if)# no platform qos trust	Configures the interface as untrusted.
Step 9	Router(config-if)# service-policy input <i>policy_map_name</i>	Attaches the policy map created in step 4 to the interface as an input service policy.

Configuration Example

This example shows how to configure a policy map to set the MPLS EXP field in the imposed label entries:

```
Router(config)# access-list 1 permit any
Router(config)# class-map CUSTOMER-A
Router(config-cmap)# match access-group 1
Router(config)# policy-map set-MPLS-PHB
Router(config-pmap)# class CUSTOMER-A
Router(config-pmap-c)# police 50000000 conform-action set-mpls-exp-transmit 4
exceed-action drop
Router(config)# interface gigabitethernet 3/1
Router(config-if)# no platform qos trust
Router(config)# interface gigabitethernet 3/1.31
Router(config-if)# service-policy input set-MPLS-PHB
```

Configuring Ingress PE Router—P Facing Interface

This procedure classifies packets based on their MPLS EXP field and provides appropriate discard and scheduling treatments.

To classify packets based on their MPLS EXP field and provide appropriate discard and scheduling treatments, perform this task:

	Command	Purpose
Step 1	Router(config)# class-map <i>class_name</i>	Specifies the class map to which packets will be mapped (matched). Creates a traffic class.
Step 2	Router(config-c-map)# match mpls experimental <i>exp_list</i>	Specifies the MPLS EXP field values used as a match criteria against which packets are checked to determine if they belong to the class.
Step 3	Router(config)# policy-map <i>name</i>	Configures the QoS policy for packets that match the class or classes.
Step 4	Router(config-p-map)# class <i>class_name</i>	Associates the traffic class with the service policy.
Step 5	Router(config-p-map-c)# bandwidth { <i>bandwidth_kbps</i> percent <i>percent</i> }	Specifies the minimum bandwidth guarantee to a traffic class. You can specify the minimum bandwidth guarantee in kilobits per second or by percent of the overall bandwidth.
Step 6	Router(config-p-map)# class class-default	Specifies the default class so that you can configure or modify its policy.
Step 7	Router(config-p-map-c)# random-detect	Enables a WRED drop policy for a traffic class that has a bandwidth guarantee.
Step 8	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 9	Router(config-if)# service-policy output <i>name</i>	Attaches a QoS policy to an interface and specifies that policies should be applied on packets leaving the interface.



Note

The **bandwidth** command and **random-detect** command are not supported on LAN ports.



Note

The **bandwidth percent** command is not supported on Ten-Gigabit Ethernet interfaces and subinterfaces.

Configuration Example

This example shows how to classify packets based on their MPLS EXP field and provide appropriate discard and scheduling treatments:

```
Router(config)# class-map MPLS-EXP-4
Router(config-c-map)# match mpls experimental 4
Router(config)# policy-map output-qos
Router(config-p-map)# class MPLS-EXP-4
Router(config-p-map-c)# bandwidth percent 40
Router(config-p-map)# class class-default
Router(config-p-map-c)# random-detect
Router(config)# interface pos 4/1
Router(config-if)# service-policy output output-qos
```


Configuring the P Router—Output Interface

To classify packets based on their MPLS EXP field and provide appropriate discard and scheduling treatments, perform this task:

	Command	Purpose
Step 1	Router(config)# class-map <i>class_name</i>	Specifies the class map to which packets will be mapped (matched). Creates a traffic class.
Step 2	Router(config-c-map)# match mpls experimental <i>exp_list</i>	Specifies the MPLS EXP field values used as a match criteria against which packets are checked to determine if they belong to the class.
Step 3	Router(config)# policy-map <i>name</i>	Configures the QoS policy for packets that match the class or classes.
Step 4	Router(config-p-map)# class <i>class_name</i>	Associates the traffic class with the service policy.
Step 5	Router(config-p-map-c)# bandwidth { <i>bandwidth_kbps</i> percent <i>percent</i> }	Specifies the minimum bandwidth guarantee to a traffic class. You can specify the minimum bandwidth guarantee in kilobits per second or by percent of the overall bandwidth.
Step 6	Router(config-p-map)# class class-default	Specifies the default class so that you can configure or modify its policy.
Step 7	Router(config-p-map-c)# random-detect	Applies WRED to the policy based on the IP precedence or the MPLS EXP field value.
Step 8	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 9	Router(config-if)# service-policy output <i>name</i>	Attaches a QoS policy to an interface and specifies that policies should be applied on packets leaving the interface.



Note

The **bandwidth** command and **random-detect** command are not supported on LAN ports.



Note

The **bandwidth percent** command is not supported on Ten-Gigabit Ethernet interfaces and subinterfaces.

Configuration Example

This example shows how to classify packets based on their MPLS EXP field and provide appropriate discard and scheduling treatments:

```
Router(config)# class-map MPLS-EXP-4
Router(config-c-map)# match mpls experimental 4
Router(config)# policy-map output-qos
Router(config-p-map)# class MPLS-EXP-4
Router(config-p-map-c)# bandwidth percent 40
Router(config-p-map)# class class-default
Router(config-p-map-c)# random-detect
Router(config)# interface pos 2/1
Router(config-if)# service-policy output output-qos
```

Configuring the Egress PE Router—Customer Facing Interface

To classify a packet based on its IP DSCP value and provide appropriate discard and scheduling treatments, perform this task:

	Command	Purpose
Step 1	Router(config)# class-map <i>class_name</i>	Specifies the class map to which packets will be mapped (matched). Creates a traffic class.
Step 2	Router(config-c-map)# match ip dscp <i>dscp_values</i>	Uses the DSCP values as the match criteria.
Step 3	Router(config)# policy-map <i>name</i>	Configures the QoS policy for packets that match the class or classes.
Step 4	Router(config-p-map)# class <i>class_name</i>	Associates the traffic class with the service policy.
Step 5	Router(config-p-map-c)# bandwidth { <i>bandwidth_kbps</i> percent <i>percent</i> }	Specifies the minimum bandwidth guarantee to a traffic class. You can specify the minimum bandwidth guarantee in kilobits per second or by percent of the overall bandwidth.
Step 6	Router(config-p-map)# class class-default	Specifies the default class so that you can configure or modify its policy.
Step 7	Router(config-p-map-c)# random-detect dscp-based	Enables a WRED drop policy for a traffic class that has a bandwidth guarantee.
Step 8	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 9	Router(config-if)# service-policy <i>output name</i>	Attaches a QoS policy to an interface and specifies that policies should be applied on packets leaving the interface.



Note

The **bandwidth percent** command is not supported on Ten-Gigabit Ethernet interfaces and subinterfaces.

Configuration Example

This example shows how to classify a packet based on its IP DSCP value and provide appropriate discard and scheduling treatments:

```
Router(config)# class-map IP-PREC-4
Router(config-c-map)# match ip precedence 4
Router(config)# policy-map output-qos
Router(config-p-map)# class IP-PREC-4
Router(config-p-map-c)# bandwidth percent 40
Router(config-p-map)# class class-default
Router(config-p-map-c)# random-detect
Router(config)# interface gigabitethernet 3/2.32
Router(config-if)# service-policy output output-qos
```

How to Configure Uniform Mode

- [Configuring the Ingress PE Router—Customer Facing Interface, page 29-35](#)
- [Configuring the Ingress PE Router—P Facing Interface, page 29-36](#)

- [Configuring the Egress PE Router—Customer Facing Interface, page 29-37](#)

**Note**

The steps that follow show one way, but not the only way, to configure the Uniform mode.

Configuring the Ingress PE Router—Customer Facing Interface

For Uniform mode, setting the trust state to IP precedence or IP DSCP allows the PFC to copy the IP PHB into the MPLS PHB.

**Note**

This description applies to PFC QoS for LAN ports.

To configure a policy map to set the MPLS EXP field in imposed label entries, perform this task:

	Command	Purpose
Step 1	Router(config)# access-list <i>ipv4_acl_number_or_name</i> permit any	Creates an IPv4 access list.
Step 2	Router(config)# class-map <i>class_name</i>	Creates a class map.
Step 3	Router(config-cmap)# match access-group <i>ipv4_acl_number_or_name</i>	Configures the class map to filter with the ACL created in Step 1 .
Step 4	Router(config)# policy-map <i>policy_map_name</i>	Creates a named QoS policy.
Step 5	Router(config-pmap)# class <i>class_name</i>	Configures the policy to use the class map created in step 2 .
Step 6	Router(config-pmap-c)# police <i>bits_per_second</i> [<i>normal_burst_bytes</i>] conform-action transmit exceed-action drop	Configures policing, including the following: • Action to take on packets that conform to the rate limit specified in the SLA. • Action to take on packets that exceed the rate limit specified in the SLA.
Step 7	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 8	Router(config-if)# platform qos trust dscp	Configures received DSCP as the basis of the internal DSCP for all the port's ingress traffic.
Step 9	Router(config-if)# service-policy input <i>policy_map_name</i>	Attaches the policy map created in step 4 to the interface as an input service policy.

Configuration Example

This example shows how to configure a policy map to set the MPLS EXP field in imposed label entries:

```
Router(config)# access-list 1 permit any
Router(config)# class-map CUSTOMER-A
Router(config-cmap)# match access-group 1
Router(config)# policy-map SLA-A
Router(config-pmap)# class CUSTOMER-A
Router(config-pmap-c)# police 50000000 conform-action transmit exceed-action drop
Router(config)# interface gigabitethernet 3/1
Router(config-if)# platform qos trust dscp
Router(config)# interface gigabitethernet 3/1.31
Router(config-if)# service-policy input SLA-A
```

Configuring the Ingress PE Router—P Facing Interface

To classify packets based on their MPLS EXP field and provide appropriate discard and scheduling treatments, perform this task:

	Command	Purpose
Step 1	Router(config)# class-map <i>class_name</i>	Specifies the class map to which packets will be mapped (matched). Creates a traffic class.
Step 2	Router(config-c-map)# match mpls experimental <i>exp_list</i>	Specifies the MPLS EXP field values used as a match criteria against which packets are checked to determine if they belong to the class.
Step 3	Router(config)# policy-map <i>name</i>	Configures the QoS policy for packets that match the class or classes.
Step 4	Router(config-p-map)# class <i>class_name</i>	Associates the traffic class with the service policy.
Step 5	Router(config-p-map-c)# bandwidth { <i>bandwidth_kbps</i> percent <i>percent</i> }	Specifies the minimum bandwidth guarantee to a traffic class. You can specify the minimum bandwidth guarantee in kilobits per second or by percent of the overall bandwidth.
Step 6	Router(config-p-map)# class class-default	Specifies the default class so that you can configure or modify its policy.
Step 7	Router(config-p-map-c)# random-detect	Enables a WRED drop policy for a traffic class that has a bandwidth guarantee.
Step 8	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 9	Router(config-if)# service-policy output <i>name</i>	Attaches a QoS policy to an interface and specifies that policies should be applied on packets leaving the interface.



Note

The **bandwidth** command and **random-detect** command are not supported on LAN ports.



Note

The **bandwidth percent** command is not supported on Ten-Gigabit Ethernet interfaces and subinterfaces.

Configuration Example

This example shows how to classify packets based on their MPLS EXP field and provide appropriate discard and scheduling treatments:

```
Router(config)# class-map MPLS-EXP-3
Router(config-c-map)# match mpls experimental 3
Router(config)# policy-map output-qos
Router(config-p-map)# class MPLS-EXP-3
Router(config-p-map-c)# bandwidth percent 40
Router(config-p-map)# class class-default
Router(config-p-map-c)# random-detect
Router(config)# interface pos 4/1
Router(config-if)# service-policy output output-qos
```

Configuring the Egress PE Router—Customer Facing Interface

To configure the egress PE router at the customer-facing interface, perform this task:

	Command	Purpose
Step 1	Router(config)# class-map <i>class_name</i>	Specifies the class map to which packets will be mapped (matched). Creates a traffic class.
Step 2	Router(config-c-map)# match ip precedence <i>precedence-value</i>	Identifies IP precedence values as match criteria.
Step 3	Router(config)# policy-map <i>name</i>	Configures the QoS policy for packets that match the class or classes.
Step 4	Router(config-p-map)# class <i>class_name</i>	Associates the traffic class with the service policy.
Step 5	Router(config-p-map-c)# bandwidth { <i>bandwidth_kbps</i> percent <i>percent</i> }	Specifies the minimum bandwidth guarantee to a traffic class. You can specify the minimum bandwidth guarantee in kilobits per second or by percent of the overall bandwidth.
Step 6	Router(config-p-map)# class class-default	Specifies the default class so that you can configure or modify its policy.
Step 7	Router(config-p-map-c)# random-detect	Applies WRED to the policy based on the IP precedence or the MPLS EXP field value.
Step 8	Router(config)# interface <i>type slot/port</i>	Selects an interface to configure.
Step 9	Router(config-if) mpls propagate-cos	Enables propagation of EXP value into the underlying IP DSCP at the MPLS domain exit LER egress port.
Step 10	Router(config-if)# service-policy output <i>name</i>	Attaches a QoS policy to an interface and specifies that policies should be applied on packets coming into the interface.



Note

The **bandwidth** command and **random-detect** command are not supported on LAN ports.



Note

The **bandwidth percent** command is not supported on Ten-Gigabit Ethernet interfaces and subinterfaces.

Configuration Example

This example shows how to configure the egress PE router at the customer-facing interface:

```
Router(config)# class-map IP-PREC-4
Router(config-c-map)# match ip precedence 4
Router(config)# policy-map output-qos
Router(config-p-map)# class IP-PREC-4
Router(config-p-map-c)# bandwidth percent 40
Router(config-p-map)# class class-default
Router(config-p-map-c)# random-detect
Router(config)# interface gigabitethernet 3/2.32
Router(config-if) mpls propagate-cos
Router(config-if)# service-policy output output-qos
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

30

PFC QoS Statistics Data Export

- [Prerequisites for PFC QoS Statistics Data Export, page 30-1](#)
- [Restrictions for PFC QoS Statistics Data Export, page 30-1](#)
- [Information About PFC QoS Statistics Data Export, page 30-1](#)
- [Default Settings for PFC QoS Statistics Data Export, page 30-2](#)
- [How to Configure PFC QoS Statistics Data Export, page 30-2](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for PFC QoS Statistics Data Export

None.

Restrictions for PFC QoS Statistics Data Export

None.

Information About PFC QoS Statistics Data Export

The PFC QoS statistics data export feature generates per-LAN-port and per-aggregate policer utilization information and forwards this information in UDP packets to traffic monitoring, planning, or accounting applications. You can enable PFC QoS statistics data export on a per-LAN-port or on a per-aggregate

policer basis. The statistics data generated per port consists of counts of the input and output packets and bytes. The aggregate policer statistics consist of counts of allowed packets and counts of packets exceeding the policed rate.

The PFC QoS statistics data collection occurs periodically at a fixed interval, but you can configure the interval at which the data is exported. PFC QoS statistics collection is enabled by default, and the data export feature is disabled by default for all ports and all configured aggregate policers.

**Note**

The PFC QoS statistics data export feature is completely separate from NetFlow Data Export and does not interact with it.

Default Settings for PFC QoS Statistics Data Export

Feature	Default Value
Global PFC QoS data export	Disabled
Per port PFC QoS data export	Disabled
Per named aggregate policer PFC QoS data export	Disabled
Per class map policer PFC QoS data export	Disabled
PFC QoS data export time interval	300 seconds
Export destination	Not configured
PFC QoS data export field delimiter	Pipe character ()

How to Configure PFC QoS Statistics Data Export

- [Enabling PFC QoS Statistics Data Export Globally, page 30-2](#)
- [Enabling PFC QoS Statistics Data Export for a Port, page 30-3](#)
- [Enabling PFC QoS Statistics Data Export for a Named Aggregate Policers, page 30-4](#)
- [Enabling PFC QoS Statistics Data Export for a Class Map, page 30-5](#)
- [Setting the PFC QoS Statistics Data Export Time Interval, page 30-6](#)
- [Configuring PFC QoS Statistics Data Export Destination Host and UDP Port, page 30-7](#)
- [Setting the PFC QoS Statistics Data Export Field Delimiter, page 30-8](#)

Enabling PFC QoS Statistics Data Export Globally

To enable PFC QoS statistics data export globally, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos statistics-export	Enables PFC QoS statistics data export globally.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable PFC QoS statistics data export globally and verify the configuration:

```
Router# configure terminal
Router(config)# platform qos statistics-export
Router(config)# end
% Warning: Export destination not set.
% Use 'platform qos statistics-export destination' command to configure the export
destination
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
Export Status : enabled
Export Interval : 300 seconds
Export Delimiter : |
Export Destination : Not configured
Router#
```



Note

You must enable PFC QoS statistics data export globally for other PFC QoS statistics data export configuration to take effect.

Enabling PFC QoS Statistics Data Export for a Port

To enable PFC QoS statistics data export for a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Selects the interface to configure.
Step 2	Router(config-if)# platform qos statistics-export	Enables PFC QoS statistics data export for the port.
Step 3	Router(config)# end	Exits configuration mode.

This example shows how to enable PFC QoS statistics data export on GigabitEthernet port 5/24 and verify the configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/24
Router(config-if)# platform qos statistics-export
Router(config-if)# end
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
Export Status : enabled
Export Interval : 300 seconds
Export Delimiter : |
Export Destination : Not configured

QoS Statistics Data Export is enabled on following ports:
-----
GigabitEthernet5/24
Router#
```

When enabled on a port, PFC QoS statistics data export contains the following fields, separated by the delimiter character:

- Export type (“1” for a port)
- Slot/port
- Number of ingress packets

- Number of ingress bytes
- Number of egress packets
- Number of egress bytes
- Time stamp

Enabling PFC QoS Statistics Data Export for a Named Aggregate Policer

To enable PFC QoS statistics data export for a named aggregate policer, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos statistics-export aggregate-policer <i>aggregate_policer_name</i>	Enables PFC QoS statistics data export for a named aggregate policer.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable PFC QoS statistics data export for an aggregate policer named `aggr1M` and verify the configuration:

```
Router# configure terminal
Router(config)# platform qos statistics-export aggregate-policer aggr1M
Router(config)# end
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
Export Status : enabled
Export Interval : 300 seconds
Export Delimiter : |
Export Destination : Not configured

QoS Statistics Data Export is enabled on following ports:
-----
GigabitEthernet5/24

QoS Statistics Data export is enabled on following shared aggregate policers:
-----
aggr1M
Router#
```

When enabled for a named aggregate policer, PFC QoS statistics data export contains the following fields, separated by the delimiter character:

- Export type (“3” for an aggregate policer)
- Aggregate policer name
- Direction (“in”)
- PFC or DFC slot number
- Number of in-profile bytes
- Number of bytes that exceed the CIR
- Number of bytes that exceed the PIR
- Time stamp

Enabling PFC QoS Statistics Data Export for a Class Map

To enable PFC QoS statistics data export for a class map, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos statistics-export class-map <i>classmap_name</i>	Enables PFC QoS statistics data export for a class map.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to enable PFC QoS statistics data export for a class map named class3 and verify the configuration:

```
Router# configure terminal
Router(config)# platform qos statistics-export class-map class3
Router(config)# end
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
Export Status : enabled
Export Interval : 300 seconds
Export Delimiter : |
Export Destination : Not configured

QoS Statistics Data Export is enabled on following ports:
-----
GigabitEthernet5/24

QoS Statistics Data export is enabled on following shared aggregate policers:
-----
aggr1M

QoS Statistics Data Export is enabled on following class-maps:
-----
class3
Router#
```

When enabled for a class map, PFC QoS statistics data export contains the following fields, separated by the delimiter character:

- For data from a physical port:
 - Export type (“4” for a classmap and port)
 - Class map name
 - Direction (“in”)
 - Slot/port
 - Number of in-profile bytes
 - Number of bytes that exceed the CIR
 - Number of bytes that exceed the PIR
 - Time stamp
- For data from a VLAN interface:
 - Export type (“5” for a class map and VLAN)
 - Classmap name
 - Direction (“in”)

- PFC or DFC slot number
- VLAN ID
- Number of in-profile bytes
- Number of bytes that exceed the CIR
- Number of bytes that exceed the PIR
- Time stamp
- For data from a port channel interface:
 - Export type (“6” for a class map and port channel)
 - Class map name
 - Direction (“in”)
 - PFC or DFC slot number
 - Port channel ID
 - Number of in-profile bytes
 - Number of bytes that exceed the CIR
 - Number of bytes that exceed the PIR
 - Time stamp

Setting the PFC QoS Statistics Data Export Time Interval

To set the time interval for the PFC QoS statistics data export, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos statistics-export interval interval_in_seconds	Sets the time interval for the PFC QoS statistics data export. Note The interval needs to be short enough to avoid counter wraparound with the activity in your configuration, but because exporting PFC QoS statistic creates a significant load on the switch, be careful when decreasing the interval.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to set the PFC QoS statistics data export interval and verify the configuration:

```

Router# configure terminal
Router(config)# platform qos statistics-export interval 250
Router(config)# end
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
Export Status : enabled
Export Interval : 250 seconds
Export Delimiter : |
Export Destination : Not configured

QoS Statistics Data Export is enabled on following ports:
-----
GigabitEthernet5/24

```

QoS Statistics Data export is enabled on following shared aggregate policers:

aggr1M

QoS Statistics Data Export is enabled on following class-maps:

class3
Router#

Configuring PFC QoS Statistics Data Export Destination Host and UDP Port

To configure the PFC QoS statistics data export destination host and UDP port number, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos statistics-export destination {host_name host_ip_address} {port port_number syslog [facility facility_name] [severity severity_value]}	Configures the PFC QoS statistics data export destination host and UDP port number.
Step 2	Router(config)# end	Exits configuration mode.



Note

When the PFC QoS data export destination is a syslog server, the exported data is prefaced with a syslog header.

Table 30-1 Supported PFC QoS Data Export Facility Parameter Values

Name	Definition	Name	Definition
kern	kernel messages	cron	cron/at subsystem
user	random user-level messages	local0	reserved for local use
mail	mail system	local1	reserved for local use
daemon	system daemons	local2	reserved for local use
auth	security/authentication messages	local3	reserved for local use
syslog	internal syslogd messages	local4	reserved for local use
lpr	line printer subsystem	local5	reserved for local use
news	netnews subsystem	local6	reserved for local use
uucp	uucp subsystem	local7	reserved for local use

Table 30-2 Supported PFC QoS Data Export Severity Parameter Values

Severity Parameter		
Name	Number	Definition
emerg	0	system is unusable
alert	1	action must be taken immediately
crit	2	critical conditions

Table 30-2 Supported PFC QoS Data Export Severity Parameter Values (continued)

Severity Parameter		
Name	Number	Definition
err	3	error conditions
warning	4	warning conditions
notice	5	normal but significant condition
info	6	informational
debug	7	debug-level messages

This example shows how to configure 172.20.52.3 as the destination host and syslog as the UDP port number and verify the configuration:

```
Router# configure terminal
Router(config)# platform qos statistics-export destination 172.20.52.3 syslog
Router(config)# end
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
Export Status : enabled
Export Interval : 250 seconds
Export Delimiter : |
Export Destination : 172.20.52.3, UDP port 514 Facility local6, Severity debug

QoS Statistics Data Export is enabled on following ports:
-----
GigabitEthernet5/24

QoS Statistics Data export is enabled on following shared aggregate policers:
-----
aggr1M

QoS Statistics Data Export is enabled on following class-maps:
-----
class3
```

Setting the PFC QoS Statistics Data Export Field Delimiter

To set the PFC QoS statistics data export field delimiter, perform this task:

	Command	Purpose
Step 1	Router(config)# platform qos statistics-export delimiter delimiter_character	Sets the PFC QoS statistics data export field delimiter.
Step 2	Router(config)# end	Exits configuration mode.

This example shows how to set the PFC QoS statistics data export field delimiter and verify the configuration:

```
Router# configure terminal
Router(config)# platform qos statistics-export delimiter ,
Router(config)# end
Router# show platform qos statistics-export info
QoS Statistics Data Export Status and Configuration information
-----
```

```
Export Status : enabled
Export Interval : 250 seconds
Export Delimiter : ,
Export Destination : 172.20.52.3, UDP port 514 Facility local6, Severity debug
```

QoS Statistics Data Export is enabled on following ports:

GigabitEthernet5/24

QoS Statistics Data export is enabled on following shared aggregate policers:

aggr1M

QoS Statistics Data Export is enabled on following class-maps:

class3

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

31

Cisco IOS ACL Support

- [Restrictions for Cisco IOS ACLs, page 31-1](#)
- [Restrictions for Layer 4 Operators in ACLs, page 31-2](#)
- [Information About ACL Support, page 31-4](#)
- [Policy-Based ACLs \(PBACLs\), page 31-5](#)
- [IPv6 Object Group-Based ACLs \(IPv6 OG ACLs\), page 31-8](#)
- [MAC ACLs, page 31-11](#)
- [ARP ACLs, page 31-15](#)
- [Optimized ACL Logging, page 31-15](#)
- [Dry Run Support for ACLs, page 31-17](#)
- [Hardware ACL Statistics, page 31-19](#)

**Note**

- For complete information about configuring Cisco IOS ACLs, see this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/sec_data_acl/configuration/15-sy/sec-data-acl-15-sy-book.html
- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Restrictions for Cisco IOS ACLs

- You can apply Cisco IOS ACLs directly to Layer 3 ports and to VLAN interfaces.
- You can apply VLAN ACLs and port ACLs to Layer 2 interfaces and VLANs (see [Chapter 35, “Port ACLs \(PACLs\)”](#) and [Chapter 36, “VLAN ACLs \(VACLs\)”](#)).

- Each type of ACL (IP, IPX, and MAC) filters only traffic of the corresponding type. A Cisco IOS MAC ACL never matches IP or IPX traffic unless the **mac packet-classify** configuration command is enabled. By default, the **mac packet-classify** configuration command is disabled.
- When you enter the **mac packet-classify** configuration command, MAC ACLs will be applied to all protocol traffic.
- The PFC does not provide hardware support for Cisco IOS IPX ACLs. Cisco IOS IPX ACLs are supported in software on the route processor (RP).
- By default, the RP sends Internet Control Message Protocol (ICMP) unreachable messages when a packet is denied by an access group.

With the **ip unreachable** command enabled (which is the default), the switch drops most of the denied packets in hardware and sends only a small number of packets to the RP to be dropped in software, which generates ICMP-unreachable messages.

The **ip unreachable** command does not impact the hardware behavior for ACL drop packets and the leaking of ACL deny packets is enabled by default. You can enter the **no ip unreachable** interface configuration command to disable ICMP unreachable messages.

- ICMP unreachable messages are not sent if a packet is denied by a VACL or a PACL.
- Use named ACLs (instead of numbered ACLs) because this causes less CPU usage when creating or modifying ACL configurations and during system restarts. When you create ACL entries (or modify existing ACL entries), the software performs a CPU-intensive operation called an ACL merge to load the ACL configurations into the PFC hardware. An ACL merge also occurs when the startup configuration is applied during a system restart.

With named ACLs, the ACL merge is triggered only when the user exits the **named-acl** configuration mode. However, with numbered ACLs, the ACL merge is triggered for every ACE definition and results in a number of intermediate merges during ACL configuration.

- Global default results are used when the Hitless update does not succeed or when features are not configured on a given interface.

Restrictions for Layer 4 Operators in ACLs

- [Determining Layer 4 Operation Usage, page 31-2](#)
- [Determining Logical Operation Unit Usage, page 31-3](#)

Determining Layer 4 Operation Usage

You can specify these types of operations:

- gt (greater than)
- lt (less than)
- neq (not equal)
- eq (equal)
- range (inclusive range)

We recommend that you do not specify more than *ten different* operations on the same ACL. If you exceed this number, each new operation might cause the affected ACE to be translated into more than one ACE.

Use the following two guidelines to determine Layer 4 operation usage:

- Layer 4 operations are considered different if the operator or the operand differ. For example, in this ACL there are three different Layer 4 operations (“gt 10” and “gt 11” are considered two different Layer 4 operations):

```
... gt 10 permit
... lt 9 deny
... gt 11 deny
```



Note There is no limit to the use of “eq” operators as the “eq” operator does not use a logical operator unit (LOU) or a Layer 4 operation bit. See the [“Determining Logical Operation Unit Usage” section on page 31-3](#) for a description of LOUs.

- Layer 4 operations are considered different if the same operator/operand couple applies once to a source port and once to a destination port. For example, in this ACL there are two different Layer 4 operations because one ACE applies to the source port and one applies to the destination port.

```
... Src gt 10 ...
... Dst gt 10
```

Determining Logical Operation Unit Usage

Logical operation units (LOUs) are registers that store operator-operand couples. All ACLs use LOUs. There can be up to 104 LOUs and each LOU can store two different operator-operand couples, making the total number of LOU registers to be 208. LOU usage per Layer 4 operation is as follows:

- gt uses 1/2 LOU
- lt uses 1/2 LOU
- neq uses 1/2 LOU
- range uses 1 LOU
- eq does not require a LOU

For example, this ACL would use a single LOU to store two different operator-operand couples:

```
... Src gt 10 ...
... Dst gt 10
```

A more detailed example follows:

```
ACL1
... (dst port) gt 10 permit
... (dst port) lt 9 deny
... (dst port) gt 11 deny
... (dst port) neq 6 permit
... (src port) neq 6 deny
... (dst port) gt 10 deny

ACL2
... (dst port) gt 20 deny
... (src port) lt 9 deny
... (src port) range 11 13 deny
... (dst port) neq 6 permit
```

The Layer 4 operations and LOU usage is as follows:

- ACL1 Layer 4 operations: 5
- ACL2 Layer 4 operations: 4
- LOUs: 4

An explanation of the LOU usage follows:

- LOU 1 stores “gt 10” and “lt 9”
- LOU 2 stores “gt 11” and “neq 6”
- LOU 3 stores “gt 20” (with space for one more)
- LOU 4 stores “range 11 13” (range needs the entire LOU)

Information About ACL Support

ACLs can be processed in hardware by the Policy Feature Card (PFC), any Distributed Forwarding Cards (DFCs), or in software by the route processor (RP):

- ACL flows that match a “deny” statement in standard and extended ACLs (input and output) are dropped in hardware if “ip unreachable” is disabled.
- ACL flows that match a “permit” statement in standard and extended ACLs (input and output) are processed in hardware.
- VLAN ACL (VACL) and port ACL (PACL) flows are processed in hardware. If a field that is specified in a VACL or PACL is not supported by hardware processing, then that field is ignored (for example, the **log** keyword in an ACL), or the whole configuration is rejected (for example, a VACL containing IPX ACL parameters).
- IPv6 ACLs use 32 bit encoding.
- VACL logging is processed in software.
- VACL is not supported for IPX access lists.
- VACL supports only deny packet logging.
- Dynamic ACL flows are processed in hardware.
- Idle timeout is processed in software.

**Note**

Idle timeout is not configurable. Cisco IOS Release 15.4SY does not support the **access-enable host timeout** command.

- Except on MPLS interfaces, reflexive ACL flows are processed in hardware after the first packet in a session is processed in software on the RP.
- Reflexive ACL flows are not accelerated in hardware for traffic from IP to various tags and traffic from various tags to IP. Reflexive ACL flows are also not accelerated in hardware for any traffic coming in and going out of all tunnel interfaces.
- IP accounting for an ACL access violation on a given port is supported only for the ACL packets that are deny leaked, by forwarding all denied packets for that port to the RP for software processing without impacting other flows.
- MAC ACLs are supported in hardware on switch ports (MAC PACLs) or on VLANs as part of VACLs.

- The PFC does not provide hardware support for Cisco IOS IPX ACLs. Cisco IOS IPX ACLs are supported in software on the RP.
- Extended name-based MAC address ACLs are supported in hardware.
- The following ACL types are processed in software:
 - Internetwork Packet Exchange (IPX) access lists
 - Standard XNS access list
 - Extended XNS access list
 - DECnet access list
 - Protocol type-code access list

**Note**

IP packets with a header length of less than five will not be access controlled.

- Unless you configure optimized ACL logging (OAL), flows that require logging are processed in software without impacting nonlogged flow processing in hardware (see the [“Optimized ACL Logging” section on page 31-15](#)).
- The forwarding rate for software-processed flows is substantially less than for hardware-processed flows.
- With the hardware statistics feature enabled, when you enter the **show ip access-list** command, the match count displayed includes packets processed in hardware.
- When you enter the **ip unreachable config** command on the PFC interface, the hardware behavior remains unaltered.
- The Hitless TCAM update for IPv4 and IPv6 provides the capability to apply existing features to the incoming traffic while updating new features in the TCAM. The Hitless feature update is mandatory for IPv6 traffic where any changes in the IPv6 ACL on a given interface would trigger a reprogramming of all the IPv6 features on all the interfaces.
- Hitless update is enabled by default. To disable the Hitless update, enter the **no platform hardware acl update-mode hitless** command.

**Note**

See the [“Restrictions for eFSU” section on page 5-2](#) for information about some release-specific restrictions.

- With the Hitless update enabled, when the FM (Feature Manager) or the switch is making updates to the recently modified ACL, a copy of each TCAM entry will be programmed in hardware. If the ACLs are not modified, the TCAM space reserved for hitless update will equal the number of TCAM entries used by the largest ACL.

Policy-Based ACLs (PBACLs)

- [Restrictions for PBACLs, page 31-6](#)
- [Information about PBACLs, page 31-6](#)
- [How to Configure PBACLs, page 31-6](#)

Restrictions for PBACLs

- PBACLs are supported on Layer 3 interfaces (such as routed interfaces and VLAN interfaces).
- The PBACL feature only supports IPv4 ACEs.
- The PBACL feature supports only Cisco IOS ACLs. It is not supported with any other features. The keywords **reflexive** and **evaluate** are not supported.
- The PBACL feature supports only named Cisco IOS ACLs. It does not support numbered ACLs.
- Feature interactions for policy-based ACLs are the same as for Cisco IOS ACLs.

Information about PBACLs

PBACLs provide the capability to apply access control policies across object groups. An object group is a group of users or servers.

You define an object group as a group of IP addresses or as a group of protocol ports. You then create access control entries (ACEs) that apply a policy (such as permit or deny) to the object group. For example, a policy-based ACE can permit a group of users to access a group of servers.

An ACE defined using a group name is equivalent to multiple ACEs (one applied to each entry in the object group). The system expands the PBACL ACE into multiple Cisco IOS ACEs (one ACE for each entry in the group) and populates the ACEs in the TCAM. Therefore, the PBACL feature reduces the number of entries you need to configure but does not reduce TCAM usage.

If you make changes in group membership, or change the contents of an ACE that uses an access group, the system updates the ACEs in the TCAM. The following types of changes trigger the update:

- Adding a member to a group
- Deleting a member from a group
- Modifying the policy statements in an ACE that uses an access group

You configure a PBACL using extended Cisco IOS ACL configuration commands. As with regular ACEs, you can associate the same access policy with one or more interfaces.

When you configure an ACE, you can use an object group to define the source, the destination, or both.

How to Configure PBACLs

- [Configuring a PBACL IP Address Object Group, page 31-6](#)
- [Configuring a PBACL Protocol Port Object Group, page 31-7](#)
- [Configuring ACLs with PBACL Object Groups, page 31-8](#)
- [Configuring PBACL on an Interface, page 31-8](#)

Configuring a PBACL IP Address Object Group

To create or modify a PBACL IP address object group, perform this task:

	Command	Purpose
Step 1	Router(config)# object-group ip address <i>object_group_name</i>	Defines object group name and enters IP-address object-group configuration mode.
Step 2	Router(config-ipaddr-ogroup)# { <i>ip_address mask</i> } { host { <i>name</i> <i>ip_address</i> } }	Configures a member of the group. The member is either a network address plus mask or a host (identified by host name or IP address).
Step 3	Router(config-ipaddr-ogroup)# { end } { exit }	To leave the configuration mode, enter the end command. To leave the IP-address object-group configuration mode, enter the exit command.

The following example creates an object group with three hosts and a network address:

```
Router(config)# object-group ip address myAG
Router(config-ipaddr-pgroup)# host 10.20.20.1
Router(config-ipaddr-pgroup)# host 10.20.20.5
Router(config-ipaddr-pgroup)# 10.30.0.0 255.255.0.0
```

Configuring a PBACL Protocol Port Object Group

To create or modify a PBACL protocol port object group, perform this task:

Command	Purpose
Router(config)# object-group ip port <i>object_group_name</i>	Defines object group name and enters port object-group configuration mode.
Router(config-port-ogroup)# { eq <i>number</i> } { gt <i>number</i> } { lt <i>number</i> } { neq <i>number</i> } { range <i>number number</i> }	Configures a member of the group. The member is either equal to or not equal to a port number, less than or greater than a port number, or a range of port numbers.
Router(config-port-ogroup)# end exit	To leave the configuration mode, enter the end command. To leave the port object-group configuration mode, enter the exit command.

The following example creates a port object group that matches protocol port 100 and any port greater than 200, except 300:

```
Router(config)# object-group ip port myPG
Router(config-port-pgroup)# eq 100
Router(config-port-pgroup)# gt 200
Router(config-port-pgroup)# neq 300
```

Configuring ACLs with PBACL Object Groups

To configure an ACL to use a PBACL object group, perform this task:

	Command	Purpose
Step 1	Router(config)# ip access-list extended <i>acl_name</i>	Defines an extended ACL with the specified name and enters extended-ACL configuration mode.
Step 2	Router(config-ext-nacl)# permit tcp addrgroup <i>object_group_name</i> addrgroup <i>object_group_name</i>	Configures an ACE for TCP traffic using IP address object group as the source policy and an object group as the destination policy.
Step 3	Router(config-ext-nacl)# exit	Exits extended ACL configuration mode.

The following example creates an access list that permits packets from the users in myAG if the protocol ports match the ports specified in myPG:

```
Router(config)# ip access-list extended my-pbacl-policy
Router(config-ext-nacl)# permit tcp addrgroup myAG portgroup myPG any
Router(config-ext-nacl)# deny tcp any any
Router(config-ext-nacl)# exit
Router(config)# exit
Router# show ip access-list my-pbacl-policy
Extended IP access list my-pbacl-policy
10 permit tcp addrgroup AG portgroup PG any
20 permit tcp any any
Router# show ip access-list my-pbacl-policy expand
Extended IP access list my-pbacl-policy expanded
20 permit tcp host 10.20.20.1 eq 100 any
20 permit tcp host 10.20.20.1 gt 200 any
20 permit tcp host 10.20.20.1 neq 300 any
20 permit tcp host 10.20.20.5 eq 100 any
20 permit tcp host 10.20.20.5 gt 200 any
20 permit tcp host 10.20.20.5 neq 300 any
20 permit tcp 10.30.0.0 255.255.0.0 eq 100 any
20 permit tcp 10.30.0.0 255.255.0.0 gt 200 any
20 permit tcp 10.30.0.0 255.255.0.0 neq 300 any
```

Configuring PBACL on an Interface

To configure a PBACL on an interface, use the **ip access-group** command. The command syntax and usage is the same as for Cisco IOS ACLs. For additional information, see the [“Restrictions for Cisco IOS ACLs” section on page 31-1](#).

The following example shows how to associate access list my-pbacl-policy with VLAN 100:

```
Router(config)# int vlan 100
Router(config-if)# ip access-group mp-pbacl-policy in
```

IPv6 Object Group-Based ACLs (IPv6 OG ACLs)

- [Restrictions for IPv6 OG ACLs, page 31-9](#)
- [Information about IPv6 OG ACLs, page 31-9](#)
- [How to Configure IPv6 OG ACLs, page 31-9](#)

Restrictions for IPv6 OG ACLs

- IPv6 OG ACLs are supported on Layer 3 interfaces (such as routed interfaces and VLAN interfaces).
- The IPv6 OG ACLs feature supports only Cisco IOS ACLs. It is not supported with any other features. The keywords **reflexive** and **evaluate** are not supported.
- The IPv6 OG ACLs feature supports only named Cisco IOS ACLs. It does not support numbered ACLs.
- Feature interactions for IPv6 OG ACLs are the same as for Cisco IOS ACLs.
- The highest number of object group-based ACEs supported in an IPv6 OG ACL is 2048.

Information about IPv6 OG ACLs

IPv6 OG ACLs provide the capability to apply access control policies across object groups. An object group is a group of users or servers.

You define an object group as a group of IPv6 addresses or as a group of protocol ports. You then create access control entries (ACEs) that apply a policy (such as permit or deny) to the object group. For example, a policy-based ACE can permit a group of users to access a group of servers.

An ACE defined using a group name is equivalent to multiple ACEs (one applied to each entry in the object group). The system expands the IPv6 OG ACL ACE into multiple Cisco IOS ACEs (one ACE for each entry in the group) and populates the ACEs in the TCAM. Therefore, the IPv6 OG ACL feature reduces the number of entries you need to configure but does not reduce TCAM usage.

If you make changes in group membership, or change the contents of an ACE that uses an access group, the system updates the ACEs in the TCAM. The following types of changes trigger the update:

- Adding a member to a group
- Deleting a member from a group
- Modifying the policy statements in an ACE that uses an access group

You configure a IPv6 OG ACL using extended Cisco IOS ACL configuration commands. As with regular ACEs, you can associate the same access policy with one or more interfaces.

When you configure an ACE, you can use an object group to define the source, the destination, or both.

How to Configure IPv6 OG ACLs

- [Configuring a IPv6 Address Object Group, page 31-9](#)
- [Configuring a IPv6 Protocol Object Group, page 31-10](#)
- [Configuring ACLs with IPv6 Object Groups, page 31-11](#)
- [Configuring IPv6 OG ACL on an Interface, page 31-11](#)

Configuring a IPv6 Address Object Group

To create or modify a IPv6 OG ACL IPv6 address object group, perform this task:

	Command	Purpose
Step 1	Router(config)# object-group v6-network <i>object_group_name</i>	Defines object group name and enters IPv6-address object-group configuration mode.
Step 2	Router(config-v6network-group)# { <i>ip_address mask</i> } { host { <i>name</i> <i>ip_address</i> } }	Configures a member of the group. The member is either a network address plus mask or a host (identified by host name or IPv6 address).
Step 3	Router(config-v6network-group)# { end } { exit }	To leave the configuration mode, enter the end command. To leave the IPv6-address object-group configuration mode, enter the exit command.

The following example creates an object group with hosts and a network address:

```
Router(config)# object-group v6-network myAG
Router(config-v6network-group)# host 2001::1
Router(config-v6network-group)# 2002::1/64
Router(config-v6network-group)# exit
```

Configuring a IPv6 Protocol Object Group

To create or modify a IPv6 OG ACL protocol object group, perform this task:

Command	Purpose
Router(config)# object-group v6-service <i>object_group_name</i>	Defines object group name and enters protocol object-group configuration mode.
Router(config-v6service-group)# { ahp } { esp } { ipv6 } { pcp } { tcp }	Configures a member of the group.
Router(config-port-ogroup)# end exit	To leave the configuration mode, enter the end command. To leave the protocol object-group configuration mode, enter the exit command.

The following example creates a protocol object group:

```
Router(config)# object-group v6-service myPG
Router(config-v6service-group)# esp
Router(config-v6service-group)# hb
Router(config-v6service-group)# exit
```

Configuring ACLs with IPv6 Object Groups

To configure an ACL to use a IPv6 OG ACL object group, perform this task:

	Command	Purpose
Step 1	Router(config)# ipv6 access-list og-acl	Defines an object group ACL with the specified name and enters IPv6-ACL configuration mode.
Step 2	Router(config-ipv6-acl)# permit tcp object-group <object_group_name> any	Configures an ACE for TCP traffic using IPv6 address object group as the source policy and an object group as the destination policy.
Step 3	Router(config-ext-nacl)# exit	Exits IPv6-ACL configuration mode.

The following example creates an access list that permits packets from the users in myAG if the protocol ports match the ports specified in myPG:

```
Router(config)# ip access-list og-acl
Router(config-ipv6-acl)# permit tcp object-group
Router(config-ipv6-acl)# deny tcp any any
Router(config-ipv6-acl)# exit
Router(config)# exit
```

Configuring IPv6 OG ACL on an Interface

To configure a IPv6 OG ACL on an interface, use the **ipv6 traffic-filter access-list-name {in | out}** command in interface mode. The command syntax and usage is the same as for Cisco IOS ACLs. For additional information, see the [“Restrictions for Cisco IOS ACLs”](#) section on page 31-1.

MAC ACLs

- [How to Configure Protocol-Independent MAC ACL Filtering, page 31-11](#)
- [How to Enable VLAN-Based MAC QoS Filtering, page 31-13](#)
- [Configuring MAC ACLs, page 31-13](#)



Note

You can use MAC ACLs with VLAN ACLs (VACLs). For more information, see [Chapter 36, “VLAN ACLs \(VACLs\).”](#)

How to Configure Protocol-Independent MAC ACL Filtering

Protocol-independent MAC ACL filtering applies MAC ACLs to all ingress traffic types (for example, IPv4 traffic, IPv6 traffic, and MPLS traffic, in addition to MAC-layer traffic).

You can configure these interface types for protocol-independent MAC ACL filtering:

- VLAN interfaces
- Routed interfaces
- Physical LAN ports

- Logical LAN subinterfaces

Ingress traffic permitted or denied by a MAC ACL on an interface configured for protocol-independent MAC ACL filtering is processed by egress interfaces as MAC-layer traffic. You cannot apply egress IP ACLs to traffic that was permitted or denied by a MAC ACL on an interface configured for protocol-independent MAC ACL filtering.

To configure protocol-independent MAC ACL filtering, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port[.subinterface]} {port-channel number[.subinterface]}}	Selects the interface to configure.
Step 2	Router(config-if)# [no] mac packet-classify {input output use {ce_cos {input output} dscp {input output}}}	Enables protocol-independent MAC ACL filtering on the interface. By default, the mac packet-classify configuration command is disabled.

- Do not configure protocol-independent MAC ACL filtering on VLAN interfaces where you have configured an IP address.
- When the mac acl filtering is enabled, all other protocol features such as RACL, microflow policing will be ignored in the hardware.

This example shows how to configure VLAN interface 4018 for protocol-independent MAC ACL filtering and how to verify the configuration:

```
Router(config)# interface vlan 4018
Router(config-if)# mac packet-classify
Router(config-if)# end
Router# show running-config interface vlan 4018 | begin 4018
interface Vlan4018
mtu 9216
ipv6 enable
mac packet-classify
end
```

This example shows how to configure Gigabit Ethernet interface 6/1 for protocol-independent MAC ACL filtering and how to verify the configuration:

```
Router(config)# interface gigabitethernet 6/1
Router(config-if)# mac packet-classify
Router(config-if)# end
Router# show running-config interface gigabitethernet 6/1 | begin 6/1
interface GigabitEthernet6/1
mtu 9216
no ip address
mac packet-classify
mpls l2transport route 4.4.4.4 4094
end
```

This example shows how to configure Gigabit Ethernet interface 3/24, subinterface 4000, for protocol-independent MAC ACL filtering and how to verify the configuration:

```
Router(config)# interface gigabitethernet 3/24.4000
Router(config-if)# mac packet-classify
Router(config-if)# end
Router# show running-config interface gigabitethernet 3/24.4000 | begin 3/24.4000
interface GigabitEthernet3/24.4000
encapsulation dot1Q 4000
mac packet-classify
```

```
mpls l2transport route 4.4.4.4 4000
end
```

How to Enable VLAN-Based MAC QoS Filtering

You can globally enable or disable VLAN-based QoS filtering in MAC ACLs. VLAN-based QoS filtering in MAC ACLs is disabled by default.

To enable VLAN-based QoS filtering in MAC ACLs, perform this task:

Command	Purpose
Router(config)# mac packet-classify use outer-vlan	Enables VLAN-based QoS filtering in MAC ACLs. The VLAN field in MAC ACLs will be matched for outer-vlan tag. The options are in (apply to ingress MAC ACLs) and out (apply to egress MAC ACLs).

To disable VLAN-based QoS filtering in MAC ACLs, perform this task:

Command	Purpose
Router(config)# no mac packet-classify use outer-vlan	Disables VLAN-based QoS filtering in MAC ACLs.

Configuring MAC ACLs

You can configure named ACLs that filter IP, IPX, DECnet, AppleTalk, VINES, or XNS traffic based on MAC addresses.

You can configure MAC ACLs that do VLAN-based filtering or CoS-based filtering or both.

You can globally enable or disable VLAN-based QoS filtering in MAC ACLs (disabled by default).

To configure a MAC ACL, perform this task:

	Command	Purpose
Step 1	Router(config)# mac host <i>name mac_addr</i>	(Optional) Assigns a name to a MAC address.
Step 2	Router(config)# mac access-list extended <i>list_name</i>	Configures a MAC ACL.
Step 3	Router(config-ext-macl)# { permit deny } { <i>src_mac_mask</i> { host name <i>src_mac_name</i> } any } { <i>dest_mac_mask</i> { host name <i>dst_mac_name</i> } any } [{ <i>protocol_keyword</i> { <i>ethertype_number</i> <i>ethertype_mask</i> }] [vlan <i>vlan_ID</i>] [cos <i>cos_value</i>]	Configures an access control entry (ACE) in a MAC ACL. The source and destination MAC addresses can be specified by MAC address masks or by names created with the mac host command.

- Cisco IOS Release 15.4SY supports the **vlan** and **cos** keywords.
- The **vlan** keyword for VLAN-based QoS filtering in MAC ACLs can be globally enabled or disabled and is disabled by default.
- You can enter MAC addresses as three 2-byte values in dotted hexadecimal format. For example, 0030.9629.9f84.

- You can enter MAC address masks as three 2-byte values in dotted hexadecimal format. Use 1 bits as wildcards. For example, to match an address exactly, use 0000.0000.0000 (can be entered as 0.0.0).
- You can enter an EtherType and an EtherType mask as hexadecimal values.
- Entries without a protocol parameter match any protocol.
- ACL entries are scanned in the order you enter them. The first matching entry is used. To improve performance, place the most commonly used entries near the beginning of the ACL.
- An implicit **deny any any** entry exists at the end of an ACL unless you include an explicit **permit any any** entry at the end of the list.
- All new entries to an existing list are placed at the end of the list. You cannot add entries to the middle of a list.
- This list shows the EtherType values and their corresponding protocol keywords:
 - 0x0600—xns-idp—Xerox XNS IDP
 - 0x0BAD—vines-ip—Banyan VINES IP
 - 0x0baf—vines-echo—Banyan VINES Echo
 - 0x6000—etype-6000—DEC unassigned, experimental
 - 0x6001—mop-dump—DEC Maintenance Operation Protocol (MOP) Dump/Load Assistance
 - 0x6002—mop-console—DEC MOP Remote Console
 - 0x6003—decnet-iv—DEC DECnet Phase IV Route
 - 0x6004—lat—DEC Local Area Transport (LAT)
 - 0x6005—diagnostic—DEC DECnet Diagnostics
 - 0x6007—lavc-sca—DEC Local-Area VAX Cluster (LAVC), SCA
 - 0x6008—amber—DEC AMBER
 - 0x6009—mumps—DEC MUMPS
 - 0x0800—ip—Malformed, invalid, or deliberately corrupt IP frames
 - 0x8038—dec-spanning—DEC LANBridge Management
 - 0x8039—dsm—DEC DSM/DDP
 - 0x8040—netbios—DEC PATHWORKS DECnet NETBIOS Emulation
 - 0x8041—msdos—DEC Local Area System Transport
 - 0x8042—etype-8042—DEC unassigned
 - 0x809B—appletalk—Kinetics EtherTalk (AppleTalk over Ethernet)
 - 0x80F3—arp—Kinetics AppleTalk Address Resolution Protocol (AARP)

This example shows how to create a MAC-Layer ACL named `mac_layer` that denies dec-phase-iv traffic with source address 0000.4700.0001 and destination address 0000.4700.0009, but permits all other traffic:

```
Router(config)# mac access-list extended mac_layer
Router(config-ext-macl)# deny 0000.4700.0001 0.0.0 0000.4700.0009 0.0.0 dec-phase-iv
Router(config-ext-macl)# permit any any
```

ARP ACLs

This section describes how to configure ARP ACLs. You can configure named ACLs that filter ARP traffic (EtherType 0x0806). To configure an ARP ACL, perform this task:

	Command	Purpose
Step 1	Router(config)# arp access-list <i>list_name</i>	Configures an ARP ACL.
Step 2	Router(config-arp-nacl)# { permit deny } { ip { any host <i>sender_ip</i> <i>sender_ip</i> <i>sender_ip_wildcardmask</i> } mac any }	Configures an access control entry (ACE) in an ARP ACL.

- This publication describes the ARP ACL syntax that is supported in hardware by the PFC. Any other ARP ACL syntax displayed by the CLI help when you enter a question mark (“?”) is not supported and cannot be used to filter ARP traffic for QoS.
- ACLs entries are scanned in the order you enter them. The first matching entry is used. To improve performance, place the most commonly used entries near the beginning of the ACL.
- An implicit **deny ip any mac any** entry exists at the end of an ACL unless you include an explicit **permit ip any mac any** entry at the end of the list.
- All new entries to an existing list are placed at the end of the list. You cannot add entries to the middle of a list.
- The PFC does not apply IP ACLs to ARP traffic.
- You cannot apply microflow policing to ARP traffic.

This example shows how to create an ARP ACL named `arp_filtering` that only permits ARP traffic from IP address 1.1.1.1:

```
Router(config)# arp access-list arp_filtering
Router(config-arp-nacl)# permit ip host 1.1.1.1 mac any
```

Optimized ACL Logging

- [Restrictions for OAL, page 31-15](#)
- [Information about OAL, page 31-16](#)
- [How to Configure OAL, page 31-16](#)

Restrictions for OAL

- OAL and VACL capture are incompatible. Do not configure both features on the switch. With OAL configured, use SPAN to capture traffic.
- OAL checks for conflicts with other features using capture like VACL capture, Lawful Intercept (LI), and IPv6 learning.
- OAL supports only IPv4 unicast packets.
- OAL is not supported with port ACLs (PACLs).

- OAL does not provide hardware support for the following:
 - Reflexive ACLs
 - ACLs used to filter traffic for other features (for example, QoS)
 - ACLs for unicast reverse path forwarding (uRPF) check exceptions
 - Exception packets (for example, TTL failure and MTU failure)
 - Packets with IP options
 - Packets addressed at Layer 3 to the router
 - Packets sent to the RP to generate ICMP unreachable messages
 - Packets being processed by features not accelerated in hardware

Information about OAL

OAL provides hardware support for ACL logging. Unless you configure OAL, packets that require logging are processed completely in software on the RP. OAL permits or drops packets in hardware on the PFC or DFC and uses an optimized routine to send information to the RP to generate the logging messages.

How to Configure OAL

- [Configuring OAL Global Parameters, page 31-16](#)
- [Configuring OAL on an Interface, page 31-17](#)
- [Displaying OAL Information, page 31-17](#)
- [Clearing Cached OAL Entries, page 31-17](#)

Configuring OAL Global Parameters

To configure global OAL parameters, perform this task:

Command	Purpose
Router(config)# logging ip access-list cache {{ entries <i>number_of_entries</i> } { interval <i>seconds</i> } { rate-limit <i>number_of_packets</i> } { threshold <i>number_of_packets</i> }}	Sets OAL global parameters.

- **entries** *number_of_entries*
 - Sets the maximum number of entries cached.
 - Range: 0–1,048,576 (entered without commas).
 - Default: 8192.
- **interval** *seconds*
 - Sets the maximum time interval before an entry is sent to be logged. Also if the entry is inactive for this duration it is removed from the cache.
 - Range: 5–86,400 (1440 minutes or 24 hours, entered without commas).

- Default: 300 seconds (5 minutes).
- **rate-limit** *number_of_packets*
 - Sets the number of packets logged per second in software.
 - Range: 10–1,000,000 (entered without commas).
 - Default: 0 (rate limiting is off and all packets are logged).
- **threshold** *number_of_packets*
 - Sets the number of packet matches before an entry is logged.
 - Range: 1–1,000,000 (entered without commas).
 - Default: 0 (logging is not triggered by the number of packet matches).

Configuring OAL on an Interface

To configure OAL on an interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port}}	Specifies the interface to configure.
Step 2	Router(config-if)# logging ip access-list cache in	Enables OAL for ingress traffic on the interface.
Step 3	Router(config-if)# logging ip access-list cache out	Enables OAL for egress traffic on the interface.

Displaying OAL Information

To display OAL information, perform this task:

Command	Purpose
Router# show logging ip access-list cache	Displays OAL information.

Clearing Cached OAL Entries

To clear cached OAL entries, perform this task:

Command	Purpose
Router# clear logging ip access-list cache	Clears cached OAL entries.

Dry Run Support for ACLs

- [Restrictions for Dry Run Support, page 31-18](#)
- [Information About Dry Run Support, page 31-18](#)
- [How to Configure Dry Run Support for ACLs, page 31-18](#)

Restrictions for Dry Run Support

- Dry Run is supported only for IPv4 RACLs and can only be applied on interfaces.
- Dry Run is supported only with named ACLs (Standard or Extended) and not with numbered ACLs.
- Only one Dry Run session is allowed at a time with a single ACL or multiple ACLs in the Dry Run session.
- The Dry Run session ACL or ACLs are removed when an ACL or ACLs are changed under configuration mode.
- Exiting the Dry Run session does not clear the existing configuration. Clear the existing session before starting a new configuration.
- The validation process may abort if there are configuration or hardware changes made during the validation process.
- Dry Run mode does not support committing the changes to the running configuration.
- Dry Run is not supported for ACLs used in QoS policies.
- Dry Run is not supported for ACLs having hardware statistics enabled.
- You can access the switch using another Telnet session while a Dry Run session is in progress.

Information About Dry Run Support

In other releases, when a new feature is applied on an interface configured along with other features, and if the new feature does not fit in the TCAM, then existing features are also affected and removed from the TCAM. To incrementally update the feature and see whether the feature fits into the TCAM without installing it, the switch provides a Dry Run support, where applications can send regular requests to test whether the request can be programmed successfully or not. The switch receives the dry run request and calculates the total TCAM resources required for the request and compares those resources against the available free resources. If the request fits in successfully, then the switch returns a success, else the switch returns a failure. The Dry Run support helps applications make intelligent decisions.

How to Configure Dry Run Support for ACLs

To configure the Dry Run support, perform this task:

	Command	Purpose
Step 1	Router(config)# configure session <i>session_name</i>	Creates a configuration session and enters the dry run mode
Step 2	Router(dry-run-config)# {default exit ip no validate}	Choose the option to configure the dry run session
Step 3	Router(dry-run-config)# ip access-list {extended standard} <i>acl_name</i>	Choose the type of ACL

The following example configures the dry run support for a session with an existing ACL RACL10K:

```
Router(config)# configure session test
Router(dry-run-config)# ip access-list extended RACL10K
Router(dr-config-ext-nacl)# permit tcp host 10.20.0.1 host 11.20.0.1
Router(dr-config-ext-nacl)# permit tcp host 10.20.0.2 host 11.20.0.2
```

```

Router(dr-config-ext-nacl)# permit tcp host 10.20.0.3 host 11.20.0.3
Router(dr-config-ext-nacl)# permit tcp host 10.20.0.4 host 11.20.0.4
Router(dr-config-ext-nacl)# permit tcp host 10.20.0.5 host 11.20.0.5

Router(dr-config-ext-nacl)# exit

Router(dry-run-config)# validate

Router(dry-run-config)# exit
Router#
.Feb 23 2010 13:46:52.528: Validation is in progress !!
.Feb 23 2010 13:46:52.528: Please try again later.
.Feb 23 2010 13:46:53.136: %FM-6-SESSION_VALIDATION_RESULT_INFO: Session Validation Result
: "Validation Completed Successfully."
. Please use 'show config session test status' to get more details of the config
validation status

Router# show configuration session test status
=====
Status of last config validation:
Timestamp: 2010-02-23@13:46:51
=====
SLOT = [1]    Result = Configuration will fit in TCAM
SLOT = [2]    Result = Configuration will fit in TCAM
SLOT = [5]    Result = Configuration will fit in TCAM

Router# clear configuration session test

Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.

Router(config)# ip access-list extended RACL10K
Router(config-ext-nacl)# permit tcp host 10.20.0.1 host 11.20.0.1
Router(config-ext-nacl)# permit tcp host 10.20.0.2 host 11.20.0.2
Router(config-ext-nacl)# permit tcp host 10.20.0.3 host 11.20.0.3
Router(config-ext-nacl)# permit tcp host 10.20.0.4 host 11.20.0.4
Router(config-ext-nacl)# permit tcp host 10.20.0.5 host 11.20.0.5
Router(config-ext-nacl)# end

Router#

```

Hardware ACL Statistics

- [Restrictions for Hardware ACL Statistics, page 31-19](#)
- [Information About Hardware ACL Statistics, page 31-20](#)
- [How to Configure Hardware ACL Statistics, page 31-20](#)

Restrictions for Hardware ACL Statistics

- Hardware ACL statistics are supported for IPv4 and IPv6 RACLs in both ingress and egress directions.
- In IPv4, hardware ACL statistics are supported for both numbered and named ACLs.
- The hardware statistics is retrieved by polling the hardware once every 60 seconds.
- The hardware statistics is lost after SSO (Stateful Switchover).

- The hardware statistics is maintained on an ACL basis. If there are multiple interfaces using the same ACL, the statistics will be aggregated.
- Hardware statistics is disabled when ODM (Order Dependent Merge) optimizations are enabled.

Information About Hardware ACL Statistics

Using the hardware ACL statistics, the hardware counters for a given ACL are gathered, aggregated, and displayed in the IOS access-list output.

The ACE hit count is retrieved from the hardware and can be viewed using the following commands:

show ip access-list and **show ipv6 access-list**

Hardware statistics is disabled by default. To enable or disable hardware statistics, enter the command for hardware statistics.

How to Configure Hardware ACL Statistics

The following example enables hardware statistics for ACL `rac11`:

```
Router(config)# ip access-list extended rac11
Router(config-ext-nacl)# [no] hardware statistics
Router(config-ext-nacl)# permit ip host 1.1.1.1 host 2.2.2.2
Router(config-ext-nacl)# permit ip host 3.3.3.3 host 4.4.4.4
Router(config-ext-nacl)# deny ip any any
Router(config-ext-nacl)# end
```

The following example displays the hardware statistics for ACL `rac11`:

```
Router# show ip access-lists rac11
Extended IP access list rac11
    hardware statistics
    10 permit ip host 1.1.1.1 host 2.2.2.2
acl hw hit count 5
    20 permit ip host 3.3.3.3 host 4.4.4.4
acl hw hit count 20
    30 deny ip any any
```

The hardware statistics for each ACE is seen after the **acl hw hit count** string and indicates the number of packets switched in hardware.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

32

Cisco TrustSec (CTS)

Cisco TrustSec is an umbrella term for security improvements to Cisco network devices based on the capability to strongly identify users, hosts and network devices within a network. TrustSec provides topology independent and scalable access controls by uniquely classifying data traffic for a particular role. TrustSec ensures data confidentiality and integrity by establishing trust among authenticated peers and encrypting links with those peers.

The key component of Cisco TrustSec is the [Cisco Identity Services Engine](#). It is typical for the Cisco ISE to provision switches with TrustSec Identities and Security Group ACLs (SGACLs), though these may be configured manually.

To configure Cisco Trustsec on the switch, see the publication, “*Cisco TrustSec Switch Configuration Guide*” at the following URL:

<http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/trustsec.html>

Release Notes for Cisco TrustSec General Availability releases are at the following URL:

http://www.cisco.com/en/US/docs/switches/lan/trustsec/release/notes/rn_cts_crossplat.html

Additional information on the Cisco TrustSec Solution, including overviews, datasheets, and case studies, is available at:

<http://www.cisco.com/en/US/netsol/ns1051/index.html>

[Table 1](#) lists the TrustSec features to be eventually implemented on TrustSec-enabled Cisco switches. Successive general availability releases of TrustSec will expand the number of switches supported and the number of TrustSec features supported per switch.

See the section, “[Hardware Supported](#)” for information on TrustSec features supported on switching modules.

Table 1 Cisco TrustSec Key Features—TrustSec 1.0 General Availability 2010 Release

Cisco TrustSec Feature	Description
802.1AE Tagging (MACSec)	Protocol for IEEE 802.1AE-based wire-rate hop-to-hop layer 2 encryption. Between MACSec-capable devices, packets are encrypted on egress from the transmitting device, decrypted on ingress to the receiving device, and in the clear within the devices. This feature is only available between TrustSec hardware-capable devices.
Endpoint Admission Control (EAC)	EAC is an authentication process for an endpoint user or a device connecting to the TrustSec domain. Usually EAC takes place at the access level switch. Successful authentication and authorization in the EAC process results in Security Group Tag assignment for the user or device. Currently EAC can be 802.1X, MAC Authentication Bypass (MAB), and Web Authentication Proxy (WebAuth).
Network Device Admission Control (NDAC)	NDAC is an authentication process where each network device in the TrustSec domain can verify the credentials and trustworthiness of its peer device. NDAC utilizes an authentication framework based on IEEE 802.1X port-based authentication and uses EAP-FAST as its EAP method. Successful authentication and authorization in NDAC process results in Security Association Protocol negotiation for IEEE 802.1AE encryption.
Security Group Access Control List (SGACL)	A Security Group Access Control List (SGACL) associates a Security Group Tag with a policy. The policy is enforced upon SGT-tagged traffic egressing the TrustSec domain.
Security Association Protocol (SAP)	After NDAC authentication, the Security Association Protocol (SAP) automatically negotiates keys and the cipher suite for subsequent MACSec link encryption between TrustSec peers. SAP is defined in IEEE 802.11i.
Security Group Tag (SGT)	An SGT is a 16-bit single label indicating the security classification of a source in the TrustSec domain. It is appended to an Ethernet frame or an IP packet.
SGT Exchange Protocol (SXP)	Security Group Tag Exchange Protocol (SXP). Devices that are not TrustSec-hardware capable can, with SXP, receive from the Cisco ACS, SGT attributes for authenticated users or devices then forward the sourceIP-to-SGT binding to a TrustSec-hardware capable device for tagging and SGACL enforcement.

Flexible MACsec Replay Protection

Flexible MACsec replay protection feature provides flexible and configurable out-of-order mode replay protection window on CTS links. You can configure MACsec window size, as defined by the number of out-of-order frames that are accepted. This value is used while installing the security associations in the MACsec. A value of 0 means that frames are accepted only in the correct order.

**Note**

The MACsec replay protection window configuration change will come into affect during a shut/no-shut or link flap.

The *macsec replay-protection window-size* command would be visible on all CTS supported ports, however value range option would be restricted to <0-0> for non-supported cases.

To configure MACsec replay protection window size, perform this task:

<i>switch1(config-if)# macsec replay-protection window-size value</i>	Enables replay protection, and configure the window size in number of frames. The range is from 0 to 4294967295. The default window size is 0.

This example shows how to verify whether MACsec replay protection is supported on your switch:

```
switch1(config-if)# macsec replay-protection window-size ?
<\<0-4294967295> window is the number of frames (packets)
```

This example shows how to fis the MACsec replay protection window size:

```
switch1(config-if)# macsec replay-protection window-size 2000
Term1#sh running-config interface t4/1
Building configuration...
!
interface TenGigabitEthernet4/1
 ip address 4.1.1.1 255.255.255.0
 macsec replay-protection window-size 2000
 cts dot1x
end
```

Hardware Supported

Table 32-2 lists the level of Cisco TrustSec supported switching modules.

Table 32-2 Switching Module Support Levels for Cisco TrustSec

Cisco TrustSec Support Level	Description	Line Card
Cisco TrustSec Capable	Supports full Cisco TrustSec capabilities with hardware acceleration for Security Group Tag imposition and IEEE 802.1AE MACsec	Supervisor Engine 6T, and all 6900 Series line cards
Cisco TrustSec Aware	Does not support Security Group Tag imposition or IEEE 802.1AE MACsec. These line cards are capable of understanding forwarding decisions, which include the Security Group Tag information. This allows them to forward traffic to a Cisco TrustSec capable line card for egress.	- WS-X6816-10T-2T, WS-X6716-10T - WS-X6816-10G-2T, WS-X6716-10GE
Not Capable of Using Cisco TrustSec	Do not support Security Group Tag imposition or IEEE 802.1AE MACsec, nor can they interpret forwarding decisions with Security Group Tag information.	- WS-X6824-SFP-2T WS-X6724-SFP - WS-X6848-SFP-2T WS-X6748-SFP - WS-X6848-TX-2T WS-X6748-GE-TX - WS-X6704-10G - WS-X6148 series (all)

For all Cisco TrustSec hardware platform and feature support information, please see TrustSec Product Bulletins at the following URL:

<http://www.cisco.com/en/US/netsol/ns1051/index.html>

33

AutoSecure

- [Prerequisites for AutoSecure, page 33-1](#)
- [Restrictions for AutoSecure, page 33-1](#)
- [Information About AutoSecure, page 33-2](#)
- [How to Configure AutoSecure, page 33-6](#)
- [Examples for AutoSecure Configuration, page 33-8](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for AutoSecure

Be ready to answer these questions:

- Is the device going to be connected to the Internet?
- How many interfaces are connected to the Internet?
- What are the names of the interfaces connected to the Internet?
- What will be your local username and password?
- What will be the switch hostname and domain name?

Restrictions for AutoSecure

- Because there is no command to undo configuration changes made by AutoSecure, always save your running configuration before configuring AutoSecure.

- The AutoSecure configuration can be configured at run time or setup time. If any related configuration is modified after AutoSecure has been enabled, the AutoSecure configuration may not be fully effective.
- After AutoSecure has been enabled, tools that use SNMP to monitor or configure a device will be unable to communicate with the device using SNMP.
- If your device is managed by a network management (NM) application, securing the management plane could turn off some services such as HTTP server and disrupt the NM application support.
- If you are using Security Device Manager (SDM), you must manually enable the HTTP server using the **ip http server** command.
- NM applications that use CDP to discover network topology will not be able to perform discovery.

Information About AutoSecure

- [AutoSecure Overview, page 33-2](#)
- [Management Plane Security Enabled by AutoSecure, page 33-3](#)
- [Forwarding Plane Security Enabled by AutoSecure, page 33-6](#)



Caution

Although AutoSecure helps to secure a switch, it does not guarantee the complete security of the switch.

AutoSecure Overview

- [AutoSecure Benefits, page 33-2](#)
- [Simplified Switch Security Configuration, page 33-2](#)
- [Enhanced Password Security Enabled by AutoSecure, page 33-3](#)
- [System Logging Message Support, page 33-3](#)

AutoSecure Benefits

Use the AutoSecure feature to secure the switch without understanding all the security features. AutoSecure is a simple security configuration process that disables nonessential system services and enables a basic set of recommended security policies to ensure secure networking services.

Simplified Switch Security Configuration

AutoSecure automates a thorough configuration of security features of the switch. AutoSecure disables certain features that are enabled by default that could be exploited for security holes. You can execute AutoSecure in either of two modes, depending on your individual needs:

- Interactive mode—Prompts with options to enable and disable services and other security features, suggesting a default setting for each option.
- Noninteractive mode—Automatically executes the recommended Cisco default settings.

Enhanced Password Security Enabled by AutoSecure

- You can specify a required minimum password length, which can eliminate weak passwords that are prevalent on most networks, such as “lab” and “cisco.”

To configure a minimum password length, use the **security passwords min-length** command.

- You can cause a syslog message to be generated after the number of unsuccessful login attempts exceeds the configured threshold.

To configure the number of allowable unsuccessful login attempts (the threshold rate), use the **security authentication failure rate** command.

System Logging Message Support

System logging messages capture any subsequent changes to the AutoSecure configuration that are applied on the running configuration. As a result, a more detailed audit trail is provided when AutoSecure is executed.

Management Plane Security Enabled by AutoSecure

- [Management Plane Security Overview, page 33-3](#)
- [Global Services Disabled by AutoSecure, page 33-3](#)
- [Per-Interface Services Disabled by AutoSecure, page 33-4](#)
- [Global Services Enabled by AutoSecure, page 33-4](#)
- [Switch Access Secured by AutoSecure, page 33-5](#)
- [Logging Options Enabled by AutoSecure, page 33-5](#)



Caution

If your device is managed by a network management (NM) application, securing the management plane could turn off some services such as the HTTP server and disrupt the NM application support.

Management Plane Security Overview

AutoSecure provides protection for the switch management interfaces (the management plane) and the data routing and switching functions (the forwarding plane, described in the [“Forwarding Plane Security Enabled by AutoSecure”](#) section on page 33-6.) Securing the management plane is done by turning off certain global and interface services that can be potentially exploited for security attacks and turning on global services that help minimize the threat of attacks. Secure access and secure logging are also configured for the switch.

Global Services Disabled by AutoSecure

- Finger—Collects information about the system (reconnaissance) before an attack.
- PAD—Enables all packet assembler and disassembler (PAD) commands and connections between PAD devices and access servers.
- Small servers—Causes TCP and User Datagram Protocol (UDP) diagnostic port attacks: a sender transmits a volume of fake requests for UDP diagnostic services on the switch, consuming all CPU resources.

- Bootp server—Bootp is an insecure protocol that can be exploited for an attack.
- HTTP server—Without secure-HTTP or authentication embedded in the HTTP server with an associated ACL, the HTTP server is insecure and can be exploited for an attack. (If you must enable the HTTP server, you will be prompted for the proper authentication or access list.)



Note If you are using Security Device Manager (SDM), you must manually enable the HTTP server using the **ip http server** command.

- Identification service—An unsecure protocol (defined in RFC 1413) that allows an external host to query a TCP port for identification. An attacker can access private information about the user from the ID server.
- CDP—If a large number of Cisco Discovery Protocol (CDP) packets are sent to the switch, the available memory of the switch can be consumed, which causes the switch to crash.



Note NM applications that use CDP to discover network topology will not be able to perform discovery.

- NTP—Without authentication or access control, Network Time Protocol (NTP) is insecure and can be used by an attacker to send NTP packets to crash or overload the switch.

If you require NTP, you must configure NTP authentication using Message Digest 5 (MD5) and the **ntp access-group** command. If NTP is enabled globally, disable it on all interfaces on which it is not needed.

- Source routing—Source routing is provided only for debugging purposes, and should be disabled in all other cases. Otherwise, packets may avoid some of the access control mechanisms of the switch.

Per-Interface Services Disabled by AutoSecure

- ICMP redirects—Disabled on all interfaces. Does not add a useful functionality to a correctly configured network, but could be used by attackers to exploit security holes.
- ICMP unreachable—Disabled on all interfaces. Internet Control Management Protocol (ICMP) unreachable are a known method for some ICMP-based denial of service (DoS) attacks.
- ICMP mask reply messages—Disabled on all interfaces. ICMP mask reply messages can give an attacker the subnet mask for a particular subnetwork in the internetwork.
- Proxy-arp—Disabled on all interfaces. Proxy-arp requests are a known method for DoS attacks because the available bandwidth and resources of the switch can be consumed in an attempt to respond to the repeated requests sent by an attacker.
- Directed broadcast—Disabled on all interfaces. Potential cause of SMURF attacks for DoS.
- Maintenance Operations Protocol (MOP) service—Disabled on all interfaces.

Global Services Enabled by AutoSecure

- The **service password-encryption** command—Prevents passwords from being visible in the configuration.
- The **service tcp-keepalives-in** and **service tcp-keepalives-out** commands—Ensures that abnormally terminated TCP sessions are removed.

Switch Access Secured by AutoSecure



Caution

If your device is managed by an NM application, securing access to the switch could turn off vital services and may disrupt the NM application support.

- If a text banner does not exist, you will be prompted to add a banner. This feature provides the following sample banner:

Authorized access only

```
This system is the property of ABC Enterprise
Disconnect IMMEDIATELY if you are not an authorized user!
Contact abc@example.com +1 408 5551212 for help.
```

- The login and password (preferably a secret password, if supported) are configured on the console, AUX, vty, and tty lines. The **transport input** and **transport output** commands are also configured on all of these lines. (Telnet and secure shell (SSH) are the only valid transport methods.) The **exec-timeout** command is configured on the console and AUX as 10.
- When the image on the device is a crypto image, AutoSecure enables SSH and secure copy (SCP) for access and file transfer to and from the switch. The **timeout seconds** and **authentication-retries integer** options for the **ip ssh** command are configured to a minimum number. (Telnet and FTP are not affected by this operation and remain operational.)
- If the user specifies that the switch does not use Simple Network Management Protocol (SNMP), one of the following functionalities will occur:
 - In interactive mode, the user is asked whether to disable SNMP regardless of the values of the community strings, which act like passwords to regulate access to the agent on the switch.
 - In noninteractive mode, SNMP will be disabled if the community string is public or private.



Note

After AutoSecure has been enabled, tools that use SNMP to monitor or configure a device will be unable to communicate with the device using SNMP.

- If authentication, authorization, and accounting (AAA) is not configured, AutoSecure configures local AAA. AutoSecure will prompt the user to configure a local username and password on the switch.

Logging Options Enabled by AutoSecure

- Sequence numbers and time stamps for all debug and log messages. This option is useful when auditing logging messages.
- Logging messages for login-related events. For example, the message “Blocking Period when Login Attack Detected” will be displayed when a login attack is detected and the switch enters quiet mode. (Quiet mode means that the switch will not allow any login attempts using Telnet, HTTP, or SSH.)
- The **logging console critical** command, which sends system logging (syslog) messages to all available TTY lines and limits messages based on severity.
- The **logging buffered** command, which copies logging messages to an internal buffer and limits messages logged to the buffer based on severity.
- The **logging trap debugging** command, which allows all commands with a severity higher than debugging to be sent to the logging server.

Forwarding Plane Security Enabled by AutoSecure

- Strict Unicast Reverse Path Forwarding (uRPF) can be configured to help mitigate problems that are caused by the introduction of forged (spoofed) IP source addresses. uRPF discards IP packets that lack a verifiable IP source address.
- Hardware rate limiting—AutoSecure will enable hardware rate-limiting of the following types of traffic without prompting the user:
 - IP errors
 - RPF failures
 - ICMP no-route messages
 - ICMP acl-drop messages
 - IPv4 multicast FIB miss messages
 - IPv4 multicast partially switch flow messages

AutoSecure will provide the option for hardware rate-limiting of the following types of traffic:

- ICMP redirects
- TTL failures
- MTU failures
- IP unicast options
- IP multicast options
- Ingress and egress ACL bridged packets

**Note**

Rate-limiting of ingress and egress ACL bridged packets can interfere with ACL logging and can increase session setup rates for hardware accelerated features such as TCP intercept, NAT, and Layer 3 WCCP.

How to Configure AutoSecure

- [Configuring AutoSecure Parameters, page 33-6](#)
- [Configuring Additional Security, page 33-7](#)
- [Verifying AutoSecure, page 33-8](#)

Configuring AutoSecure Parameters

The **auto secure** command guides you through a semi-interactive session (also known as the AutoSecure session) to secure the management and forwarding planes. You can use this command to secure just the management plane or the forwarding plane; if neither option is selected in the command line, you can choose to configure one or both planes during the session.

This command also allows you to go through all noninteractive configuration portions of the session before the interactive portions. The noninteractive portions of the session can be enabled by selecting the optional **no-interact** keyword.

At any prompt you may enter a question mark (?) for help or Ctrl-C to abort the session.

In interactive mode, you will be asked at the end of the session whether to commit the generated configuration to the running configuration of the switch. In noninteractive mode, the changes will be automatically applied to the running configuration.

**Note**

There is no undo command for configuration changes made by AutoSecure. You should always save the running configuration before executing the **auto secure** command.

To execute the AutoSecure configuration process, beginning in privileged EXEC mode, perform this task:

Command	Purpose
Router# auto secure [management forwarding] [no-interact full]	Executes the AutoSecure session for configuring one or both planes of the switch. • management—Only the management plane will be secured. • forwarding—Only the forwarding plane will be secured. • no-interact—The user will not be prompted for any interactive configurations. • full—The user will be prompted for all interactive questions. This is the default.

For an example of the AutoSecure session, see the [“Examples for AutoSecure Configuration” section on page 33-8](#).

Configuring Additional Security

After completing the AutoSecure configuration, you can further enhance the security of management access to your switch by performing this task:

	Command or Action	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# security passwords min-length <i>length</i>	Ensures that all configured passwords are at least a specified length. • <i>length</i>—Minimum length of a configured password. The range is 0 to 16 characters.

	Command or Action	Purpose
Step 3	Router(config)# enable password {password [encryption-type] password}	Sets a local password to control access to various privilege levels. <i>encryption-type</i>—A value of 0 indicates that an unencrypted password follows. A value of 7 indicates that a hidden password follows. Note You usually will not enter an encryption type unless you enter a password that has already been encrypted by a Cisco router or switch.
Step 4	Router(config)# security authentication failure rate threshold-rate log	Configures the number of allowable unsuccessful login attempts. <i>threshold-rate</i>—Number of allowable unsuccessful login attempts. The range is 1 to 1024. log—Syslog authentication failures if the number of failures in one minute exceeds the threshold.

The following example shows how to configure the switch for a minimum password length of 10 characters and a threshold of 3 password failures in one minute. The example also shows how to set a hidden local password.

```
Router# configure terminal
Router(config)# security passwords min-length 10
Router(config)# security authentication failure rate 3
Router(config)# enable password 7 elephant123
```

Verifying AutoSecure

To verify that the AutoSecure feature has executed successfully, perform this task:

Command	Purpose
Router# show auto secure config	Displays all configuration commands that have been added as part of the AutoSecure configuration. The output is the same as the configuration generated output

Examples for AutoSecure Configuration

The following example is a sample AutoSecure session. After you execute the **auto secure** command, the feature will automatically prompt you with a similar response unless you enable the **no-interact** keyword. (For information on which features are disabled and which features are enabled, see the [“Management Plane Security Enabled by AutoSecure”](#) section on page 33-3 and the [“Forwarding Plane Security Enabled by AutoSecure”](#) section on page 33-6.)

```
Router# auto secure
--- AutoSecure Configuration ---

*** AutoSecure configuration enhances the security of
the router, but it will not make it absolutely resistant
to all security attacks ***
```


AutoSecure will modify the configuration of your device.

All the configuration changes will be shown. For a detailed explanation of how the configuration changes enhance security and any possible side effects, please refer to Cisco.com for AutoSecure documentation.

At any prompt you may enter '?' for help.

Use ctrl-c to abort this session at any prompt.

If this device is being managed by a network management station, AutoSecure configuration may block network management traffic.

Continue with AutoSecure? [no]: y

Gathering information about the router for AutoSecure

Is this router connected to internet? [no]: y

Enter the number of interfaces facing the internet [1]: 1

Interface	IP-Address	OK?	Method	Status	Protocol
Vlan1	unassigned	YES	NVRAM	administratively down	down
Vlan77	77.1.1.4	YES	NVRAM	down	down
GigabitEthernet6/1	unassigned	YES	NVRAM	administratively down	down
GigabitEthernet6/2	21.30.30.1	YES	NVRAM	up	up
Loopback0	3.3.3.3	YES	NVRAM	up	up
Tunnel1	unassigned	YES	NVRAM	up	up

Enter the interface name that is facing the internet: Vlan77

Securing Management plane services...

Disabling service finger

Disabling service pad

Disabling udp & tcp small servers

Enabling service password encryption

Enabling service tcp-keepalives-in

Enabling service tcp-keepalives-out

Disabling the cdp protocol

Disabling the bootp server

Disabling the http server

Disabling the finger service

Disabling source routing

Disabling gratuitous arp

Here is a sample Security Banner to be shown at every access to device. Modify it to suit your enterprise requirements.

Authorized Access only

This system is the property of <Name of Enterprise>.

UNAUTHORIZED ACCESS TO THIS DEVICE IS PROHIBITED.

You must have explicit permission to access this device. All activities performed on this device are logged. Any violations of access policy will result in disciplinary action.

Enter the security banner {Put the banner between k and k, where k is any character}:

k

banner

k

Enter the new enable secret:

Confirm the enable secret :

Enable password is not configured or its length is less than minimum no. of characters configured

Enter the new enable password:

Confirm the enable password:

Configuration of local user database

Enter the username: cisco

Enter the password:

Confirm the password:

Configuring AAA local authentication

Configuring Console, Aux and VTY lines for

local authentication, exec-timeout, and transport

Securing device against Login Attacks

Configure the following parameters

Blocking Period when Login Attack detected (in seconds): 5

Maximum Login failures with the device: 3

Maximum time period for crossing the failed login attempts (in seconds): ?

% A decimal number between 1 and 32767.

Maximum time period for crossing the failed login attempts (in seconds): 5

Configure SSH server? [yes]: no

Configuring interface specific AutoSecure services

Disabling mop on Ethernet interfaces

Securing Forwarding plane services...

Enabling unicast rpf on all interfaces connected
to internet

The following rate-limiters are enabled by default:

...

Would you like to enable the following rate-limiters also?

...

Enable the above rate-limiters also? [yes/no]: yes

Would you like to enable the rate-limiters for Ingress/EgressACL bridged packets also?

NOTE: Enabling the ACL in/out rate-limiters can affect ACL logging
and session setup rate for hardware accelerated features such
as NAT, Layer 3 WCCP and TCP Intercept

...

Enable the ACL in/out rate-limiters also? [yes/no]: no

This is the configuration generated:

```
no service finger
no service pad
no service udp-small-servers
no service tcp-small-servers
service password-encryption
service tcp-keepalives-in
service tcp-keepalives-out
no cdp run
no ip bootp server
no ip http server
no ip finger
no ip source-route
no ip gratuitous-arps
no ip identd
banner k
```

```

banner
k
security passwords min-length 6
security authentication failure rate 10 log
enable secret 5 $1$30kP$f.KDndYPz/Hv/.yTlJStN/
enable password 7 08204E4D0D48574446
username cisco password 7 08204E4D0D48574446
aaa new-model
aaa authentication login local_auth local
line console 0
  login authentication local_auth
  exec-timeout 5 0
  transport output telnet
line vty 0 15
  login authentication local_auth
  transport input telnet
login block-for 5 attempts 3 within 5
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
logging facility local2
logging trap debugging
service sequence-numbers
logging console critical
logging buffered
int Vlan1
  no ip redirects
  no ip proxy-arp
  no ip unreachable
  no ip directed-broadcast
  no ip mask-reply
  no mop enabled
int Vlan77
  no ip redirects
  no ip proxy-arp
  no ip unreachable
  no ip directed-broadcast
  no ip mask-reply
  no mop enabled
int GigabitEthernet6/1
  no ip redirects
  no ip proxy-arp
  no ip unreachable
  no ip directed-broadcast
  no ip mask-reply
  no mop enabled
int GigabitEthernet6/2
  no ip redirects
  no ip proxy-arp
  no ip unreachable
  no ip directed-broadcast
  no ip mask-reply
  no mop enabled
interface Vlan77
  ip verify unicast source reachable-via rx
...
!
end

```

Apply this configuration to running-config? [yes]: yes

Applying the config generated to running-config

Router#

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

34

MAC Address-Based Traffic Blocking



Note

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

To block all traffic to or from a MAC address in a specified VLAN, perform this task:

Command	Purpose
Router(config)# mac address-table static <i>mac_address</i> vlan <i>vlan_ID</i> drop	Blocks all traffic to or from the configured MAC address in the specified VLAN.

This example shows how to block all traffic to or from MAC address 0050.3e8d.6400 in VLAN 12:

```
Router# configure terminal
Router(config)# mac address-table static 0050.3e8d.6400 vlan 12 drop
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

35

Port ACLs (PACLs)

- [Prerequisites for PACls, page 35-1](#)
- [Restrictions for PACls, page 35-1](#)
- [Information About PACls, page 35-2](#)
- [Interface Template ACLs, page 35-7](#)
- [Interface Template ACLs, page 35-7](#)
- [How to Configure Interface Template ACLs, page 35-10](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- Port ACLs do not support the access-list keywords **log** or **reflexive**. These keywords in the access list are ignored. OAL does not support PACls.
- PACls are not supported on private VLANs.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for PACls

None.

Restrictions for PACls

- There can be at most one IP access list and one MAC access list applied to the same Layer 2 interface per direction.
- PACls are not applied to MPLS or ARP messages.

- An IP access list filters only IPv4 and IPv6 packets. For IP access lists, you can define a standard, extended, or named access-list.
- A MAC access list filters ingress packets that are of an unsupported type (not IP, ARP, or MPLS packets) based on the fields of the Ethernet datagram. A MAC access list is not applied to IP, MPLS, or ARP messages. You can define only named MAC access lists.
- The number of ACLs and ACEs that can be configured as part of a PACL are bounded by the hardware resources on the switch. Those hardware resources are shared by various ACL features (such as VACLs) that are configured on the system. If there are insufficient hardware resources to program a PACL in hardware, the PACL is not applied.
- PACL does not support the access-list **log** and **reflect/evaluate** keywords. These keywords are ignored if you add them to the access list for a PACL.
- OAL does not support PACLs.
- The access group mode can change the way PACLs interact with other ACLs. To maintain consistent behavior across Cisco platforms, use the default access group mode (merge mode).
- PACLs cannot filter Physical Link Protocols and Logical Link Protocols, such as CDP, VTP, DTP, PAGP, UDLD, and STP, because the protocols are redirected to the RP before the ACL takes effect. You can apply CoPP or QoS to Physical Link Protocol and Logical Link Protocol traffic.
- Layer 2 ports do not support Layer 3 ACLs.
- An Interface Template ACL can be applied on an interface only when there is no other ACL directly configured on the same interface. An ACL that is configured directly on an interface takes priority over an Interface Template ACL that is applied on the same interface. The Interface Template takes effect once the ACL directly applied on the interface has been removed.

Information About PACLs

- [PACL Overview, page 35-2](#)
- [EtherChannel and PACL Interactions, page 35-3](#)
- [Dynamic ACLs \(Applies to Merge Mode Only\), page 35-4](#)
- [Trunk Ports, page 35-4](#)
- [Layer 2 to Layer 3 Port Conversion, page 35-4](#)
- [Port-VLAN Association Changes, page 35-4](#)

PACL Overview

PACLs filter incoming traffic on Layer 2 interfaces, using Layer 3 information, Layer 4 header information, or non-IP Layer 2 information.

The PACL feature uses standard or extended IP ACLs or named MAC-extended ACLs that you want to apply to the port.

Port ACLs perform access control on all traffic entering the specified Layer 2 port.

PACLs and VACLs can provide access control based on the Layer 3 addresses (for IP protocols) or Layer 2 MAC addresses (for non-IP protocols).

The port ACL (PACL) feature provides the ability to perform access control on specific Layer 2 ports. A Layer 2 port is a physical LAN or trunk port that belongs to a VLAN. Port ACLs are applied only on the ingress traffic. The port ACL feature is supported only in hardware (port ACLs are not applied to any packets routed in software).

When you create a port ACL, an entry is created in the ACL TCAM. You can use the **show tcam counts** command to see how much TCAM space is available.

The PACL feature does not affect Layer 2 control packets received on the port.

You can use the **access-group mode** command to change the way that PACLs interact with other ACLs.

PACLs use the following modes:

- **Prefer port mode**—If a PACL is configured on a Layer 2 interface, the PACL takes effect and overwrites the effect of other ACLs (Cisco IOS ACL and VACL). If no PACL feature is configured on the Layer 2 interface, other features applicable to the interface are merged and are applied on the interface.
- **Merge mode**—In this mode, the PACL, VACL, and Cisco IOS ACLs are merged in the ingress direction following the logical serial model shown in [Figure 35-2](#). This is the default access group mode.

You configure the **access-group mode** command on each interface. The default is merge mode.

**Note**

A PACL can be configured on a trunk port only after prefer port mode has been selected. Trunk ports do not support merge mode.

To illustrate access group mode, assume a physical port belongs to VLAN100, and the following ACLs are configured:

- Cisco IOS ACL R1 is applied on routed interface VLAN100.
- VACL (VLAN filter) V1 is applied on VLAN100.
- PACL P1 is applied on the physical port.

In this situation, the following ACL interactions occur:

- In prefer port mode, Cisco IOS ACL R1 and VACL V1 are ignored.
- In merge mode, Cisco IOS ACL R1, VACL V1 and PACL P1 are merged and applied on the port.

**Note**

The CLI syntax for creating a PACL is identical to the syntax for creating a Cisco IOS ACL. An instance of an ACL that is mapped to a Layer 2 port is called a PACL. An instance of an ACL that is mapped to a Layer 3 interface is called a Cisco IOS ACL. The same ACL can be mapped to both a Layer 2 port and a Layer 3 interface.

The PACL feature supports MAC ACLs, IPv4, and IPv6 ACLs. The PACL feature does not support ACLs for ARP or Multiprotocol Label Switching (MPLS) traffic.

EtherChannel and PACL Interactions

This section describes the guidelines for the EtherChannel and PACL interactions:

- PACLs are supported on the main Layer 2 channel interface but not on the port members. A port that has a PACL configured on it may not be configured as an EtherChannel member port. The EtherChannel configuration commands are unavailable on ports that are configured with a PACL.

- Changing the configuration on the logical port affects all the ports in the channel. When an ACL is mapped to the logical port belonging to a channel, it is mapped to all ports in the channel.

Dynamic ACLs (Applies to Merge Mode Only)

Dynamic ACLs are VLAN-based and are used by two features: CBAC and GWIP. The merge mode *does not* support the merging of the dynamic ACLs with the PACLs. In merge mode, the following configurations are not allowed:

- Attempting to apply a PACL on a port where its corresponding VLAN has a dynamic ACL mapped. In this case, the PACL is not applied to traffic on the port.
- Configuring a dynamic ACL on a VLAN where one of its constituent ports has a PACL installed. In this case, the dynamic ACL is not applied.

Trunk Ports

To configure a PACL on a trunk port, you must first configure port prefer mode. The configuration commands to apply a PACL on a trunk or dynamic port will not be available until you configure the port in port prefer mode by entering the **access-group mode prefer port** interface command. Trunk ports do not support merge mode.

Layer 2 to Layer 3 Port Conversion

If you reconfigure a port from Layer 2 to Layer 3, any PACL configured on the port becomes inactive but remains in the configuration. If you subsequently configure the port as Layer 2, any PACL configured on the port becomes active again.

Port-VLAN Association Changes

You can enter port configuration commands that alter the port-VLAN association, which triggers an ACL remerge.

Unmapping and then mapping a PACL, VACL, or Cisco IOS ACL automatically triggers a remerge.

In merge mode, online insertion or removal of a switching module also triggers a remerge, if ports on the module have PACLs configured.

PACL and VACL Interactions

- [PACL Interaction with VACLs and Cisco IOS ACLs, page 35-5](#)
- [Bridged Packets, page 35-5](#)
- [Routed Packets, page 35-5](#)
- [Multicast Packets, page 35-6](#)

PACL Interaction with VACLs and Cisco IOS ACLs

This section describes the guidelines for the PACL interaction with the VACLs and Cisco IOS ACLs.

For an incoming packet on a physical port, the PACL is applied first. If the packet is permitted by the PACL, the VACL on the ingress VLAN is applied next. If the packet is Layer 3 forwarded and is permitted by the VACL, it is filtered by the Cisco IOS ACL on the same VLAN. The same process happens in reverse in the egress direction. However, there is currently no hardware support for output PACLs.

The PACLs override both the VACLs and Cisco IOS ACLs when the port is configured in prefer port mode. The one exception to this rule is when the packets are forwarded in the software by the route processor (RP). The RP applies the ingress Cisco IOS ACL regardless of the PACL mode. Two examples where the packets are forwarded in the software are as follows:


- Packets that are egress bridged (due to logging or features such as NAT)
- Packets with IP options

Bridged Packets

Figure 35-1 shows a PACL and a VACL applied to bridged packets. In merge mode, the ACLs are applied in the following order:

1. PACL for the ingress port
2. VACL for the ingress VLAN
3. VACL for the egress VLAN

Figure 35-1 Applying ACLs on Bridged Packets



In prefer port mode, only the PACL is applied to the ingress packets (the input VACL is not applied).


Routed Packets

Figure 35-2 shows how ACLs are applied on routed and Layer 3-switched packets. In merge mode, the ACLs are applied in the following order:

1. PACL for the ingress port
2. VACL for the ingress VLAN
3. Input Cisco IOS ACL
4. Output Cisco IOS ACL
5. VACL for the egress VLAN

In prefer port mode, only the PACL is applied to the ingress packets (the input VACL and Cisco IOS ACL are not applied).

Figure 35-2 Applying ACLs on Routed Packets




Multicast Packets

Figure 35-3 shows how ACLs are applied on packets that need multicast expansion. For packets that need multicast expansion, the ACLs are applied in the following order:

1. Packets that need multicast expansion:
 - a. PACL for the ingress port
 - b. VACL for the ingress VLAN
 - c. Input Cisco IOS ACL
2. Packets after multicast expansion:
 - a. Output Cisco IOS ACL
 - b. VACL for the egress VLAN
3. Packets originating from router:
 - a. Output Cisco IOS ACL
 - b. VACL for the egress VLAN

In prefer port mode, only the PACL is applied to the ingress packets (the input VACL and Cisco IOS ACL are not applied).

Figure 35-3 Applying ACLs on Multicast Packets

Interface Template ACLs

This section describes the Interface Template ACL feature on Catalyst 6500 series Switches. An interface template provides a mechanism to configure multiple commands at the same time and associate it with a target such as an interface. An interface template is a container of configurations or policies that can be applied to specific ports.

Interface Templates provide an efficient way to apply ACLs along with other commands on interfaces. Template ACLs can be applied on an interface by first configuring an ACL inside an interface template, and then applying the template to any number of desired interfaces. A single template having an ACL can be applied to any number of physical or virtual interfaces either statically or dynamically.

The support for configuring Interface Template ACLs was introduced on Catalyst 6500 series Switches starting with the 15.5(1)SY2 release.

How to Configure PACLs

- [Configuring IP and MAC ACLs on a Layer 2 Interface, page 35-8](#)
- [Configuring Access-group Mode on Layer 2 Interface, page 35-8](#)
- [Applying ACLs to a Layer 2 Interface, page 35-9](#)
- [Applying ACLs to a Port Channel, page 35-9](#)
- [Displaying an ACL Configuration on a Layer 2 Interface, page 35-9](#)

Configuring IP and MAC ACLs on a Layer 2 Interface

IP and MAC ACLs can be applied to Layer 2 physical interfaces. Standard (numbered, named) and Extended (numbered, named) IP ACLs, and Extended Named MAC ACLs are supported.

To apply IP or MAC ACLs on a Layer 2 interface, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface <i>interface</i>	Enters interface configuration mode for a Layer 2 port.
Step 3	Switch(config-if)# {ip mac} access-group {name number in out}	Applies a numbered or named ACL to the Layer 2 interface.
Step 4	Switch(config)# show running-config	Displays the access list configuration.

This example shows how to configure the Extended Named IP ACL `simple-ip-acl` to permit all TCP traffic and implicitly deny all other IP traffic:

```
Switch(config)# ip access-list extended simple-ip-acl
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# end
```

This example shows how to configure the Extended Named MAC ACL `simple-mac-acl` to permit source host 000.000.011 to any destination host:

```
Switch(config)# mac access-list extended simple-mac-acl
Switch(config-ext-macl)# permit host 000.000.011 any
Switch(config-ext-macl)# end
```

Configuring Access-group Mode on Layer 2 Interface

To configure the access mode on a Layer 2 interface, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface <i>interface</i>	Enters interface configuration mode for a Layer 2 port.
Step 3	Switch(config-if)# [no] access-group mode {prefer port merge}	Sets the mode for this Layer 2 interface. The no prefix sets the mode to the default value (which is merge).
Step 4	Switch(config)# show running-config	Displays the access list configuration.

This example shows how to configure an interface to use prefer port mode:

```
Switch# configure terminal
Switch(config)# interface gigabitEthernet 6/1
Switch(config-if)# access-group mode prefer port
```

This example shows how to configure an interface to use merge mode:

```
Switch# configure terminal
Switch(config)# interface gigabitEthernet 6/1
Switch(config-if)# access-group mode merge
```

Applying ACLs to a Layer 2 Interface

To apply IP and MAC ACLs to a Layer 2 interface, perform one of these tasks:

Command	Purpose
Switch(config-if)# ip access-group <i>ip-acl</i> in	Applies an IP ACL to the Layer 2 interface.
Switch(config-if)# mac access-group <i>mac-acl</i> in	Applies a MAC ACL to the Layer 2 interface.

This example applies the extended named IP ACL `simple-ip-acl` to interface GigabitEthernet 6/1 ingress traffic:

```
Switch# configure t
Switch(config)# interface gigabitEthernet 6/1
Switch(config-if)# ip access-group simple-ip-acl in
```

This example applies the extended named MAC ACL `simple-mac-acl` to interface GigabitEthernet 6/1 ingress traffic:

```
Switch# configure t
Switch(config)# interface gigabitEthernet 6/1
Switch(config-if)# mac access-group simple-mac-acl in
```

Applying ACLs to a Port Channel

To apply IP and MAC ACLs to a port channel logical interface, perform this task:

Command	Purpose
Switch(config-if)# interface port-channel <i>number</i>	Enters configuration mode for the port channel.
Switch(config-if)# ip access-group <i>ip-acl</i> { in out }	Applies an IP ACL to the port channel interface.
Switch(config-if)# mac access-group <i>mac-acl</i> { in out }	Applies a MAC ACL to the port channel interface.

This example applies the extended named IP ACL `simple-ip-acl` to port channel 3 ingress traffic:

```
Switch# configure t
Switch(config)# interface port-channel 3
Switch(config-if)# ip access-group simple-ip-acl in
```

Displaying an ACL Configuration on a Layer 2 Interface

To display information about an ACL configuration on Layer 2 interfaces, perform one of these tasks:

Command	Purpose
Switch# show ip access-lists [interface <i>interface-name</i>]	Shows the IP access group configuration on the interface.
Switch# show mac access-group [interface <i>interface-name</i>]	Shows the MAC access group configuration on the interface.
Switch# show access-group mode [interface <i>interface-name</i>]	Shows the access group mode configuration on the interface.

This example shows that the IP access group `simple-ip-acl` is configured on the inbound direction of interface `fa6/1`:

```
Switch# show ip interface gigabitethernet 6/1
GigabitEthernet6/1 is up, line protocol is up
  Inbound access list is simple-ip-acl
  Outgoing access list is not set
```

This example shows that MAC access group `simple-mac-acl` is configured on the inbound direction of interface Gigabit Ethernet 6/1:

```
Switch# show mac access-group interface gigabitethernet 6/1
Interface GigabitEthernet6/1:
  Inbound access-list is simple-mac-acl
  Outbound access-list is not set
```

This example shows that access group merge is configured on interface Gigabit Ethernet 6/1:

```
Switch# show access-group mode interface gigabitethernet 6/1
Interface GigabitEthernet6/1:
  Access group mode is: merge
```

How to Configure Interface Template ACLs

- [Configuring an IPv4 ACL by Statically Applying an Interface Template, page 35-10](#)
- [Configuring an IPv6 ACL by Statically Applying an Interface Template, page 35-12](#)
- [Configuration Examples for Template ACL Configuration, page 35-12](#)

Configuring an IPv4 ACL by Statically Applying an Interface Template

To configure an IPv4 ACL on an interface by statically applying a template, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# ip access-list [extended] <name <i>number</i> >	Defines an IPv4 access list by name or number.
Step 3	Switch(config-ext-nacl)# [<i>sequence-number</i>] permit <i>protocol source source-wildcard destination</i> <i>destination-wildcard</i> [option <i>option-name</i>] [precedence <i>precedence</i>] [tos <i>tos</i>] [ttl <i>operator</i> <i>value</i>] [log] [time-range <i>time-range-name</i>] [fragments]	Sets conditions to allow a packet to pass a named IP access list. • Every access list must have at least one permit statement. • Use the no sequence-number form of this command to delete an entry.
Step 4	Switch(config-ext-nacl)# [<i>sequence-number</i>] deny <i>protocol source source-wildcard destination</i> <i>destination-wildcard</i> [option <i>option-name</i>] [precedence <i>precedence</i>] [tos <i>tos</i>] [ttl <i>operator</i> <i>value</i>] [log] [time-range <i>time-range-name</i>] [fragments]	(Optional) Specifies a deny statement in named IP access list mode. • Use the no sequence-number form of this command to delete an entry.
Step 5	Repeat Step 3 and Step 4 as necessary.	Allows you to revise the access list.

	Command	Purpose
Step 6	<code>Switch(config-ext-nacl)# exit</code>	Exits access list configuration mode and returns to global configuration mode.
Step 7	<code>Switch(config)# template <i>template-name</i></code>	Creates a user template and enters template configuration mode.
Step 8	<code>Switch(config-template)# ip access-group {<i>access-list-name</i> <i>access-list-number</i>} {<i>in</i> <i>out</i>}</code>	Associates the specified access list with the interface template.
Step 9	<code>Switch(config-template)# exit</code>	Exits template configuration mode and returns to global configuration mode.
Step 10	<code>Switch(config)# interface <i>interface-name</i></code>	Enters the interface configuration mode for the specified interface.
Step 11	<code>Switch(config-if)# source template <i>template-name</i></code>	Statically applies an interface template to a target interface.
Step 12	<code>Switch(config-if)# end</code>	Exits interface configuration mode and returns to privileged EXEC mode.

Configuring an IPv6 ACL by Statically Applying an Interface Template

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# ipv6 access-list name	Defines an IPv6 access list by name and enters the IPv6 access list configuration mode.
Step 3	Switch(config-ipv6-acl)# permit protocol {source-ipv6-prefix/ prefix-length any host source-ipv6-address} [operator [port-number]] {destination-ipv6-prefix/ prefix-length any host destination-ipv6-address} [operator [port-number]] [dscp value] [fragments] [log] [log-input] [sequence value] [time-range name]	Sets conditions to allow a packet to pass a named IP access list. - Every access list must have at least one permit statement. - Use the no sequence-number form of this command to delete an entry.
Step 4	Switch(config-ipv6-acl)# deny protocol {source-ipv6-prefix/ prefix-length any host source-ipv6-address} [operator [port-number]] {destination-ipv6-prefix/ prefix-length any host destination-ipv6-address} [operator [port-number]] [dscp value] [fragments] [log] [log-input] [sequence value] [time-range name]	(Optional) Specifies a deny statement in named IP access list mode. Use the no sequence-number form of this command to delete an entry.
Step 5	Repeat Step 3 and Step 4 as necessary.	Allows you to revise the access list.
Step 6	Switch(config-ipv6-acl)# exit	Exits access list configuration mode and returns to global configuration mode.
Step 7	Switch(config)# template template-name	Creates a user template and enters template configuration mode.
Step 8	Switch(config-template)# ipv6 traffic-filter access-list-name {in out}	Associates the specified access list with the interface template.
Step 9	Switch(config-template)# exit	Exits template configuration mode and returns to global configuration mode.
Step 10	Switch(config)# interface interface-name	Enters the interface configuration mode for the specified interface.
Step 11	Switch(config-if)# source template template-name	Statically applies an interface template to a target interface.
Step 12	Switch(config-if)# end	Exits interface configuration mode and returns to privileged EXEC mode.

Configuration Examples for Template ACL Configuration

- [Configuration of IPv4 and IPv6 ACLs by Statically Applying an Interface Template, page 35-13](#)
- [Verifying Template Details, page 35-13](#)
- [Viewing Access List details, page 35-13](#)
- [Verifying ACL Configuration on an Interface with a Statically Applied Template, page 35-14](#)
- [Verifying ACL Configuration on an Interface with a Dynamically Applied Template, page 35-15](#)

Configuration of IPv4 and IPv6 ACLs by Statically Applying an Interface Template

The following example shows the configuration of IPv4 and IPv6 ACLs on an interface using an interface template:

```
Switch#configure terminal
Switch(config)#ip access-list extended IPV4_ACL
Switch(config-ext-nacl)#10 permit ip any host 10.1.1.10
Switch(config-ext-nacl)#20 deny ip any host 100.1.1.10
Switch(config-ext-nacl)#30 permit ip any host 12.1.1.1
Switch(config-ext-nacl)#40 deny ip any host 20.1.1.1
Switch(config-ext-nacl)#exit
Switch(config)#ipv6 access-list IPV6_ACL
Switch(config-ipv6-acl)#permit ipv6 any 2001:DB8:0101:0101::/32
Switch(config-ipv6-acl)#deny ipv6 any 2001:DB8:1001:0000::/32
Switch(config-ipv6-acl)#deny ipv6 any 2001:DB8:1001:0123::/32
Switch(config)#exit
Switch(config)#template ACL_test
Switch(config-template)#ip access-group IPV4_ACL in
Switch(config-template)#ipv6 traffic-filter IPV6_ACL out
Switch(config-template)#exit
Switch(config)#interface Te1/1/2
Switch(config-if)#source template ACL_test
Switch(config-if)# end
```

Verifying Template Details

The following example shows the sample output for **show template binding target interface-name** command for a statically applied Interface Template:

```
Switch# show template binding target Te1/1/2
```

```
Interface Templates
=====
Interface: Te1/1/2

Method          Source          Template-Name
-----
static          User            ACL_test
```

The following example shows the sample output for **show template binding target interface-name** command for a dynamically applied Interface Template:

```
Switch# show template binding target Gi160/3/0/13
```

```
Interface Templates
=====
Interface: Gi160/3/0/13

Method          Source          Template-Name
-----
dynamic         User            INTERFACE_TEMPLATE_ACL_IPV6
```

Viewing Access List details

The following examples shows the access list details for a given ACL:

```
Switch# show ip access-list IPV4_ACL
Extended IP access list IPV4_ACL
 10 permit ip any host 10.1.1.10
 20 deny ip any host 100.1.1.10
 30 permit ip any host 12.1.1.1
```

```

40 deny ip any host 20.1.1.1

Switch# show ipv6 access-list IPV6_ACL
IPv6 access list IPV6_ACL
  permit ipv6 any 2001:DB8::/32 sequence 10
  deny ipv6 any 2001:DB8::/32 sequence 20

```

The following examples show the access list details configured inside a given Template:

```

Switch# show template interface source user ACL_test
Building configuration...

Template Name      : ACL_test
Template Definition :
  ip access-group test_ACL in
  ipv6 traffic-filter test_ACL2 out
!
end

Switch# show template interface source user INTERFACE_TEMPLATE_ACL_IPV6
Building configuration...

Template Name      : INTERFACE_TEMPLATE_ACL_IPV6
Template Definition :
  ipv6 traffic-filter Block_ipv6 in
  load-interval 30
  description ipv6 acl temp cca
!
end

```

The following example shows the ACL entries programmed in the hardware using the sample output for **show platform hardware acl entry interface *interface-name*** command

```

show platform hardware acl entry interface TenGigabitEthernet1/1/2 securityin ip
mls_if_index:20004052 dir:0 feature:0 proto:0
pass#0 features
fno:0
tcam:B, bank:0, prot:0    Aces
Permit      ip host 1.1.1.3 host 4.4.4.4
Permit      ip host 1.1.1.2 host 3.3.3.3
Permit      ip host 1.1.1.1 host 2.2.2.2
L3_Deny     ip any any

```

Verifying ACL Configuration on an Interface with a Statically Applied Template

The following example shows the sample output for **show running-config interface *interface-name*** command:

```

Switch#show running-config int Te1/1/2
Building configuration...

Current configuration : 576 bytes
!
interface TenGigabitEthernet1/1/2
  switchport
  switchport mode access
  switchport trunk allowed vlan 20,30,40,50,60
  switchport access vlan 20
  switchport voice vlan 30
  switchport port-security
  ip access-group directACL in
  authentication periodic
  authentication timer reauthenticate server

```

```
access-session host-mode multi-domain
access-session closed
access-session port-control auto
dot1x pae authenticator
source template ACL_test
spanning-tree portfast edge
spanning-tree bpdupfilter enable
service-policy type control subscriber POLICY_Gi160/3/0/13_dot1x
end
```

The following example shows the sample output for **show derived-config interface *interface-name*** command for an interface with a statically applied Template ACL:

```
Switch# show derived-config int Te1/1/2
Building configuration...
```

```
Derived configuration : 550 bytes
!
interface TenGigabitEthernet1/1/2
 switchport
 switchport access vlan 20
 switchport trunk allowed vlan 20,30,40,50,60
 switchport mode access
 switchport voice vlan 30
 switchport port-security
 ip access-group directACL in
 authentication periodic
 authentication timer reauthenticate server
 access-session host-mode multi-domain
 access-session closed
 access-session port-control auto
 dot1x pae authenticator
 spanning-tree portfast edge
 spanning-tree bpdupfilter enable
 service-policy type control subscriber POLICY_Gi160/3/0/13_dot1x
end
```

Verifying ACL Configuration on an Interface with a Dynamically Applied Template

The following example shows the sample output for **show running-config interface *interface-name*** command:

```
Switch# show running-config int Gi160/3/0/13
Building configuration...
```

```
Current configuration : 669 bytes
!
interface GigabitEthernet160/3/0/13
 switchport
 switchport trunk allowed vlan 20,30,40,50,60
 switchport mode access
 switchport access vlan 20
 switchport voice vlan 30
 switchport port-security
 access-group mode prefer port
 ipv6 traffic-filter Block_ipv6 in
 authentication periodic
 authentication timer reauthenticate server
 access-session closed
 access-session port-control auto
 mab
 dot1x pae authenticator
 spanning-tree portfast edge
 spanning-tree bpdupfilter enable
```

```

service-policy type control subscriber POLICY_Gi160/3/0/13_violation
service-policy input AutoQoS-FEX-INGRESS-QoS-POLICY
service-policy type lan-queuing output AutoQoS-FEX-EGRESS-QoS-POLICY
end

```

The following example shows the sample output for **show derived-config interface *interface-name*** command for an interface with a dynamically applied Template ACL:

```

Switch# show derived-config int Gig160/3/0/13
Building configuration...

```

```

Derived configuration : 753 bytes
!
interface GigabitEthernet160/3/0/13
description ipv6 acl temp cca
switchport
switchport access vlan 20
switchport trunk allowed vlan 20,30,40,50,60
switchport mode access
switchport voice vlan 30
switchport port-security
access-group mode prefer port
ipv6 traffic-filter Block_ipv6 in
load-interval 30
authentication periodic
authentication timer reauthenticate server
access-session closed
access-session port-control auto
mab
dot1x pae authenticator
spanning-tree portfast edge
spanning-tree bpduguard enable
service-policy type control subscriber POLICY_Gi160/3/0/13_violation
service-policy input AutoQoS-FEX-INGRESS-QoS-POLICY
service-policy type lan-queuing output AutoQoS-FEX-EGRESS-QoS-POLICY
end

```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

36

VLAN ACLs (VACLs)

- [Prerequisites for VACLs, page 36-1](#)
- [Restrictions for VACLs, page 36-1](#)
- [Information About VACLs, page 36-2](#)
- [How to Configure VACLs, page 36-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- Optimized ACL logging (OAL) and VACL capture are incompatible. Do not configure both features on the switch. With OAL configured (see the [“Optimized ACL Logging” section on page 31-15](#)), use SPAN to capture traffic.
- Also see the [“PACL Interaction with VACLs and Cisco IOS ACLs” section on page 35-5](#).

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for VACLs

None.

Restrictions for VACLs

- VACLs use standard and extended Cisco IOS IP and MAC layer-named ACLs (see [“MAC ACLs” section on page 31-11](#)) and VLAN access maps.
- IGMP packets are not checked against VACLs.
- VLAN access maps can be applied to VLANs for VACL capture.

- Each VLAN access map can consist of one or more map sequences; each sequence has a match clause and an action clause. The match clause specifies IP or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match occurs. When a flow matches a permit ACL entry, the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.
- To apply access control to both bridged and routed traffic, you can use VACLs alone or a combination of VACLs and ACLs. You can define ACLs on the VLAN interfaces to apply access control to both the ingress and egress routed traffic. You can define a VACL to apply access control to the bridged traffic.
- The following caveats apply to ACLs when used with VACLs:
 - Packets that require logging on the outbound ACLs are not logged if they are denied by a VACL.
 - VACLs are applied on packets before NAT translation. If the translated flow is not subject to access control, the flow might be subject to access control after the translation because of the VACL configuration.
- VACLs check for conflicts with other features using capture like OAL, Lawful Intercept (LI), and IPv6 learning.
- When VACL capture is configured with Policy Based Routing (PBR) on the same interface, do not select BDD as the ACL merge algorithm.
- When VACL capture is configured on an egress interface together with another egress feature that requires software processing of the traffic, packets of the overlapping traffic may be captured twice.
- The action clause in a VACL can be forward, drop, capture, or redirect. Traffic can also be logged.

**Note**

- VACLs have an implicit deny at the end of the map; a packet is denied if it does not match any ACL entry, and at least one ACL is configured for the packet type.
- If an empty or undefined ACL is specified in a VACL, any packets will match the ACL, and the associated action is taken.

Information About VACLs

VLAN ACLs (VACLs) can provide access control for all packets that are bridged within a VLAN or that are routed into or out of a VLAN for VACL capture. Unlike Cisco IOS ACLs that are applied on routed packets only, VACLs apply to all packets and can be applied to any VLAN. VACLs are processed in the ACL TCAM hardware. VACLs ignore any Cisco IOS ACL fields that are not supported in hardware.

You can configure VACLs for IP and MAC-layer traffic.

If a VACL is configured for a packet type, and a packet of that type does not match the VACL, the default action is to deny the packet.

Packets can either enter the VLAN through a Layer 2 port or through a Layer 3 port after being routed. You can also use VACLs to filter traffic between devices in the same VLAN.

How to Configure VACLs

- [Defining a VLAN Access Map, page 36-3](#)
- [Configuring a Match Clause in a VLAN Access Map Sequence, page 36-3](#)
- [Configuring an Action Clause in a VLAN Access Map Sequence, page 36-4](#)
- [Applying a VLAN Access Map, page 36-4](#)
- [Verifying VLAN Access Map Configuration, page 36-5](#)
- [VLAN Access Map Configuration and Verification Examples, page 36-5](#)
- [Configuring a Capture Port, page 36-6](#)
- [Configuring VACL Logging, page 36-7](#)

Defining a VLAN Access Map

To define a VLAN access map, perform this task:

Command	Purpose
Router(config)# vlan access-map <i>map_name</i> [0-65535]	Defines the VLAN access map. Optionally, you can specify the VLAN access map sequence number.

- To insert or modify an entry, specify the map sequence number.
- If you do not specify the map sequence number, a number is automatically assigned.
- You can specify only one match clause and one action clause per map sequence.
- Use the **no** keyword with a sequence number to remove a map sequence.
- Use the **no** keyword without a sequence number to remove the map.

See the [“VLAN Access Map Configuration and Verification Examples” section on page 36-5](#).

Configuring a Match Clause in a VLAN Access Map Sequence

To configure a match clause in a VLAN access map sequence, perform this task:

Command	Purpose
Router(config-access-map)# match {[ip ipv6] address { 1-199 1300-2699 <i>acl_name</i> } { mac address <i>acl_name</i> }}	Configures the match clause in a VLAN access map sequence.

- Release 15.0(1)SY1 and later releases support IPv6 ACLs.
- You can select one or more ACLs.
- Use the **no** keyword to remove a match clause or specified ACLs in the clause.
- For information about named MAC-Layer ACLs, see [“MAC ACLs” section on page 31-11](#).
- For information about Cisco IOS ACLs, see [Chapter 31, “Cisco IOS ACL Support”](#) and the [“VLAN Access Map Configuration and Verification Examples” section on page 36-5](#).

Configuring an Action Clause in a VLAN Access Map Sequence

To configure an action clause in a VLAN access map sequence, perform this task:

Command	Purpose
Router(config-access-map)# action { drop [log]} { forward [capture vlan <i>vlan_ID</i>]} { redirect [{ fastethernet gigabitethernet tengigabitethernet } <i>slot/port</i>] { port-channel <i>channel_id</i> }}	Configures the action clause in a VLAN access map sequence.

- You can set the action to drop, forward, forward capture, or redirect packets.
- Forwarded packets are still subject to any configured Cisco IOS security ACLs.
- The **capture** action sets the capture bit for the forwarded packets so that ports with the capture function enabled can receive the packets. Only forwarded packets can be captured. For more information about the **capture** action, see the [“Configuring a Capture Port” section on page 36-6](#).
- The **forward vlan** action implements Policy-Based Forwarding (PBF), bridging between VLANs.
- When the **log** action is specified, dropped packets are logged in software. Only dropped IP packets can be logged.
- The **redirect** action allows you to specify up to five interfaces, which can be physical interfaces or EtherChannels. You cannot specify packets to be redirected to an EtherChannel member or a VLAN.
- The redirect interface must be in the VLAN for which the VACL access map is configured.
- If a VACL is redirecting traffic to an egress SPAN source port, SPAN does not copy the VACL-redirection traffic.
- SPAN and RSPAN destination ports transmit VACL-redirection traffic.
- Use the **no** keyword to remove an action clause or specified redirect interfaces.

See the [“VLAN Access Map Configuration and Verification Examples” section on page 36-5](#).

Applying a VLAN Access Map

To apply a VLAN access map, perform this task:

Command	Purpose
Router(config)# vlan filter <i>map_name</i> vlan-list	Applies the VLAN access map to the specified VLANs.

- You can apply the VLAN access map to one or more VLANs.
- The *vlan_list* parameter can be a single VLAN ID or a comma-separated list of VLAN IDs or VLAN ID ranges (*vlan_ID–vlan_ID*).
- You can apply only one VLAN access map to each VLAN.
- VACLs applied to VLANs are active only for VLANs with a Layer 3 VLAN interface configured. Applying a VLAN access map to a VLAN without a Layer 3 VLAN interface creates an administratively down Layer 3 VLAN interface to support the VLAN access map.
- VACLs applied to VLANs are inactive if the Layer 2 VLAN does not exist or is not operational.

- You cannot apply a VACL to a secondary private VLAN. VACLs applied to primary private VLANs also apply to secondary private VLANs.
- Use the **no** keyword to clear VLAN access maps from VLANs.

See the “[VLAN Access Map Configuration and Verification Examples](#)” section on page 36-5.

Verifying VLAN Access Map Configuration

To verify VLAN access map configuration, perform this task:

Command	Purpose
Router# show vlan access-map [map_name]	Verifies VLAN access map configuration by displaying the content of a VLAN access map.
Router# show vlan filter [access-map map_name vlan vlan_id]	Verifies VLAN access map configuration by displaying the mappings between VACLs and VLANs.

VLAN Access Map Configuration and Verification Examples

Assume IP-named ACL **net_10** and **any_host** are defined as follows:

```
Router# show ip access-lists net_10
Extended IP access list net_10
    permit ip 10.0.0.0 0.255.255.255 any

Router# show ip access-lists any_host
Standard IP access list any_host
    permit any
```

This example shows how to define and apply a VLAN access map to forward IP packets. In this example, IP traffic matching **net_10** is forwarded and all other IP packets are dropped due to the default drop action. The map is applied to VLAN 12 to 16.

```
Router(config)# vlan access-map thor 10
Router(config-access-map)# match ip address net_10
Router(config-access-map)# action forward
Router(config-access-map)# exit
Router(config)# vlan filter thor vlan-list 12-16
```

This example shows how to define and apply a VLAN access map to drop and log IP packets. In this example, IP traffic matching **net_10** is dropped and logged and all other IP packets are forwarded:

```
Router(config)# vlan access-map ganymede 10
Router(config-access-map)# match ip address net_10
Router(config-access-map)# action drop log
Router(config-access-map)# exit
Router(config)# vlan access-map ganymede 20
Router(config-access-map)# match ip address any_host
Router(config-access-map)# action forward
Router(config-access-map)# exit
Router(config)# vlan filter ganymede vlan-list 7-9
```

This example shows how to define and apply a VLAN access map to forward and capture IP packets. In this example, IP traffic matching **net_10** is forwarded and captured and all other IP packets are dropped:

```
Router(config)# vlan access-map mordred 10
Router(config-access-map)# match ip address net_10
```

```
Router(config-access-map)# action forward capture
Router(config-access-map)# exit
Router(config)# vlan filter mordred vlan-list 2, 4-6
```

Configuring a Capture Port



Note

- A port configured to capture VACL-filtered traffic is called a capture port.
- To apply IEEE 802.1Q tags to the captured traffic, configure the capture port to trunk unconditionally (see the [“Configuring the Layer 2 Switching Port as an 802.1Q Trunk”](#) section on page 21-8 and the [“Configuring the Layer 2 Trunk Not to Use DTP”](#) section on page 21-9).

To configure a capture port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port}}	Specifies the interface to configure.
Step 2	Router(config-if)# switchport capture allowed vlan {add all except remove} <i>vlan_list</i>	(Optional) Filters the captured traffic on a per-destination-VLAN basis. The default is all .
Step 3	Router(config-if)# switchport capture	Configures the port to capture VACL-filtered traffic.

- You can configure any port as a capture port.
- The *vlan_list* parameter can be a single VLAN ID or a comma-separated list of VLAN IDs or VLAN ID ranges (*vlan_ID–vlan_ID*).
- To encapsulate captured traffic, configure the capture port with the **switchport trunk encapsulation** command (see the [“Configuring a Layer 2 Switching Port as a Trunk”](#) section on page 21-8) before you enter the **switchport capture** command.
- For unencapsulated captured traffic, configure the capture port with the **switchport mode access** command (see the [“Configuring a LAN Interface as a Layer 2 Access Port”](#) section on page 21-14) before you enter the **switchport capture** command.
- The capture port supports only egress traffic. No traffic can enter the switch through a capture port.

This example shows how to configure a Gigabit Ethernet interface 5/1 as a capture port:

```
Router(config)# interface gigabitEthernet 5/1
Router(config-if)# switchport capture
Router(config-if)# end
```

This example shows how to display VLAN access map information:

```
Router# show vlan access-map mymap
Vlan access-map "mymap" 10
    match: ip address net_10
    action: forward capture
Router#
```

This example shows how to display mappings between VACLs and VLANs. For each VACL map, there is information about the VLANs that the map is configured on and the VLANs that the map is active on. A VACL is not active if the VLAN does not have an interface.

```
Router# show vlan filter
VLAN Map mordred:
```

```

Configured on VLANs: 2,4-6
Active on VLANs: 2,4-6
Router#

```

Configuring VACL Logging

When you configure VACL logging, IP packets that are denied generate log messages in these situations:

- When the first matching packet is received
- For any matching packets received during the last 5-minute interval
- If the threshold is reached before the 5-minute interval

Log messages are generated on a per-flow basis. A flow is defined as packets with the same IP addresses and Layer 4 (UDP or TCP) port numbers. When a log message is generated, the timer and packet count is reset.

These restrictions apply to VACL logging:

- Because of the rate-limiting function for redirected packets, VACL logging counters may not be accurate.
- Only denied IP packets are logged.

To configure VACL logging, use the **action drop log** command action in VLAN access map submode (see the [“Configuring an Action Clause in a VLAN Access Map Sequence”](#) section on page 36-4) and perform this task in global configuration mode to specify the global VACL logging parameters:

	Command	Purpose
Step 1	Router(config)# vlan access-log maxflow <i>max_number</i>	Sets the log table size. The content of the log table can be deleted by setting the maxflow number to 0. The default is 500 with a valid range of 0 to 2048. When the log table is full, logged packets from new flows are dropped by the software.
Step 2	Router(config)# vlan access-log ratelimit <i>pps</i>	Sets the maximum redirect VACL logging packet rate. The default packet rate is 2000 packets per second with a valid range of 0 to 5000. Packets exceeding the limit are dropped by the hardware.
Step 3	Router(config)# vlan access-log threshold <i>pkt_count</i>	Sets the logging threshold. A logging message is generated if the threshold for a flow is reached before the 5-minute interval. By default, no threshold is set.
Step 4	Router(config)# exit	Exits VLAN access map configuration mode.

This example shows how to configure global VACL logging in hardware:

```

Router(config)# vlan access-log maxflow 800
Router(config)# vlan access-log ratelimit 2200
Router(config)# vlan access-log threshold 4000

```

Displays the configured VACL logging properties.

```
Router# show vlan access-log config
```

Displays the content of the VACL log table.

```
Router# show vlan access-log flow protocol {{src_addr src_mask} | any | {host {hostname | host_ip}}} {{dst_addr dst_mask} | any | {host {hostname | host_ip}}}
```

```
[vlan vlan_id]
```

Displays packet and message counts and other statistics.

```
Router# show vlan access-log statistics
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

37

Policy-Based Forwarding (PBF)

- [Prerequisites for PBF, page 37-1](#)
- [Restrictions for PBF, page 37-1](#)
- [Information About PBF, page 37-2](#)
- [Default Settings for PBF, page 37-2](#)
- [How to Configure PBF, page 37-2](#)
- [Monitoring PBF, page 37-3](#)
- [Configuration Examples for PBF, page 37-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
- Optimized ACL logging (OAL) and VACL capture are incompatible. Do not configure both features on the switch. With OAL configured (see the [“Optimized ACL Logging” section on page 31-15](#)), use SPAN to capture traffic.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for PBF

None.

Restrictions for PBF

- PBF is performed in software, with optional rate limiters to control CPU usage.
- PBF is applied only to ingress traffic.

- To allow traffic in both directions between two VLANs, you must configure PBF in both VLANs.
- You can configure PBF between hosts in different switches.
- By default, PBF hosts in the same VLAN cannot communicate with each other. To allow local communication, use the **local** keyword.
- When configuring the **vlan filter** command, specify only one VLAN after the **vlan-list** keyword. If you specify more than one VLAN, PBF will ignore all but the last VLAN in the list.
- Layer 2 port ACLs (PACLs) take precedence over PBF.
- If the sending VLAN is shut down, PBF will still function. Shutting down a VLAN disables Layer 3 functionality, but PBF is a Layer 2 function.

Information About PBF

PBF is a MAC-address VACL that bridges packets between VLANs. PBF forwards packets based solely on the source and destination MAC addresses, ignoring any information above Layer 2.

Default Settings for PBF

None.

How to Configure PBF

To configure PBF, perform this task on each source VLAN:

	Command	Purpose
Step 1	Router(config)# mac host <i>my_host</i> <i>mac_addr</i>	(Optional) Assigns a name to the MAC address of the source host.
Step 2	Router(config)# mac access-list extended <i>macl_name</i>	Configures a MAC ACL.
Step 3	Router(config-ext-macl)# permit host <i>my_host</i> any	Configures an access control entry (ACE) to permit traffic from the named host to any other address. Hosts can be specified by a name or by a MAC address.
Step 4	Router(config-ext-macl)# permit host <i>my_host</i> host <i>other_host</i>	Configures an ACE to permit traffic from the named host to one other host.
Step 5	Router(config-ext-macl)# exit	Exits ACL configuration.
Step 6	Router(config)# vlan access-map <i>map_name</i>	Defines a VLAN access map.
Step 7	Router(config-access-map)# match mac address <i>macl_name</i>	Applies the MAC ACL to this VLAN access map.
Step 8	Router(config-access-map)# action forward vlan <i>other_vlan_ID</i> [local]	Forwards matching traffic to the other VLAN. Note By default, PBF-specified devices on the same VLAN cannot communicate with each other. To allow local communication by the host, use the local keyword.

	Command	Purpose
Step 9	Router(config-access-map)# exit	Exits access map configuration.
Step 10	Router(config)# vlan filter <i>map_name</i> vlan-list <i>my_vlan_ID</i>	Applies the VLAN access map to the specified VLAN.
Step 11	Router(config)# interface vlan <i>my_vlan_ID</i>	Enters interface configuration mode for the VLAN.
Step 12	Router(config-if)# mac packet-classify	Classifies incoming or outgoing Layer 3 packets on this VLAN as Layer 2 packets.
Step 13	Router(config-if)# exit	Exits interface configuration mode.
Step 14	Router(config)# exit	Exits global configuration mode.

Monitoring PBF

- The output of the **show vlan mac-pbf config** command displays the following fields for configured PBF paths:
 - Rcv Vlan — The number of the VLAN to which packets are forwarded by PBF.
 - Snd Vlan — The number of the VLAN which will forward packets by PBF.
 - DMAC — The MAC address of the destination host on the receiving VLAN.
 - SMAC — The MAC address of the source host on the sending VLAN.
 - (Local) — Displays 1 if the **local** keyword is configured in the **action forward vlan** command on the sending VLAN; displays 0 if the **local** keyword is not configured.
 - (Packet counter) — The number of packets that have been forwarded from the sending VLAN to the receiving VLAN. To clear this counter, enter the **clear vlan mac-pbf counters** command.
 - Pkts dropped — The number of packets that have been dropped by the sending VLAN. To clear this counter, enter the **clear vlan mac-pbf counters** command.

Configuration Examples for PBF

This example shows how to configure and display PBF to allow two hosts in separate VLANs (“red” VLAN 100 and “blue” VLAN 200) on the same switch to exchange packets:

```
Router(config)# mac host host_red3 0001.0002.0003
Router(config)# mac access-list extended macl_red
Router(config-ext-macl)# permit host host_red host host_blue
Router(config-ext-macl)# exit
Router(config)# vlan access-map red_to_blue
Router(config-access-map)# match mac address macl_red
Router(config-access-map)# action forward vlan 200 local
Router(config-access-map)# exit
Router(config)# vlan filter red_to_blue vlan-list 100
Router(config)# interface vlan 100
Router(config-if)# mac packet-classify
Router(config-if)# exit
Router(config)#
Router(config)# mac host host_blue5 0001.0002.0005
Router(config)# mac access-list extended macl_blue
Router(config-ext-macl)# permit host host_blue host host_red
Router(config-ext-macl)# exit
Router(config)# vlan access-map blue_to_red
```

```

Router(config-access-map)# match mac address mac1_blue
Router(config-access-map)# action forward vlan 100
Router(config-access-map)# exit
Router(config)# vlan filter blue_to_red vlan-list 200
Router(config)# interface vlan 200
Router(config-if)# mac packet-classify
Router(config-if)# exit
Router#
Router# show vlan mac-pbf config
  Rcv Vlan 100, Snd Vlan 200, DMAC 0001.0002.0003, SMAC 0001.0002.0005 1 15
  Rcv Vlan 200, Snd Vlan 100, DMAC 0001.0002.0005, SMAC 0001.0002.0003 0 23
  Pkts Dropped 0
Router#

```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

38

Denial of Service (DoS) Protection

- [Security ACLs and VACLs, page 38-1](#)
- [QoS Rate Limiting, page 38-2](#)
- [Global Protocol Packet Policing, page 38-3](#)
- [Unicast Reverse Path Forwarding \(uRPF\) Check, page 38-6](#)
- [Configuring Sticky ARP, page 38-9](#)
- [Monitoring Packet Drop Statistics, page 38-10](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - Also see:
 - [Chapter 34, “MAC Address-Based Traffic Blocking”](#)
 - [Chapter 43, “Traffic Storm Control”](#)
 - [Chapter 39, “Control Plane Policing \(CoPP\)”](#)
 - http://www.cisco.com/en/US/docs/ios-xml/ios/security/config_library/15-sy/secdata-15-sy-library.html
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Security ACLs and VACLs

If the network is under a DoS attack, ACLs can be an efficient method for dropping the DoS packets before they reach the intended target. Use security ACLs if an attack is detected from a particular host.

In this example, the host 10.1.1.10 and all traffic from that host is denied:

```
Router(config)# access-list 101 deny ip host 10.1.1.10 any
Router(config)# access-list 101 permit ip any any
```

Security ACLs also protect against the spoofing of addresses. For example, assume that a source address A is on the inside of a network and a switch interface that is pointing to the Internet. You can apply an inbound ACL on the switch Internet interface that denies all addresses with a source of A (the inside address). This action stops attacks where the attackers spoof inside source addresses. When the packet arrives at the switch interface, it matches on that ACL and drops the packet before it causes damage.

When the switch is used with a Cisco Intrusion Detection Module (CIDM), you can dynamically install the security ACL as a response to the detection of the attack by the sensing engine.

VACLs are a security enforcement tool based on Layer 2, Layer 3, and Layer 4 information. The result of a VACL lookup against a packet can be a permit, a deny, a permit and capture, or a redirect. When you associate a VACL with a particular VLAN, all traffic must be permitted by the VACL before the traffic is allowed into the VLAN. VACLs are enforced in hardware, so there is no performance penalty for applying VACLs to a VLAN.

See [Chapter 31, “Cisco IOS ACL Support,”](#) and [Chapter 36, “VLAN ACLs \(VACLs\).”](#)

QoS Rate Limiting

QoS ACLs limit the amount of a particular type of traffic that is processed by the RP. If a DoS attack is initiated against the RP, QoS ACLs can prevent the DoS traffic from reaching the RP data path and congesting it. The PFC and DFCs perform QoS in hardware, which offers an efficient means of limiting DoS traffic (once that traffic has been identified) to protect the switch from impacting the RP.

For example, if the network is experiencing ping-of-death or smurf attacks, the administrator should rate limit the ICMP traffic to counteract the DoS attack and still allow legitimate traffic through the processor, or allow it to be forwarded to the RP or host. This rate limiting configuration must be done for each flow that should be rate limited and the rate-limiting policy action should be applied to the interface.

In the following example, the access-list 101 permits and identifies ping (echo) ICMP messages from any source to any destination as traffic. Within the policy map, a policing rule defines a specified committed information rate (CIR) and burst value (96000 bps and 16000 bps) to rate limit the ping (ICMP) traffic through the chassis. The policy map then is applied to an interface or VLAN. If the ping traffic exceeds the specified rate on the VLAN or interface where the policy map is applied, it is dropped as specified in the markdown map (the markdown map for the normal burst configurations is not shown in the example).

```
Router(config)# access-list 101 permit icmp any any echo
Router(config)# class-map match-any icmp_class
Router(config-cmap)# match access-group 101
Router(config-cmap)# exit
Router(config)# policy-map icmp_policer
Router(config-pmap)# class icmp_class
Router(config-pmap-c)# police 96000 16000 conform-action transmit exceed-action
policed-dscp-transmit drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
```

See [Chapter 25, “Classification, Marking, and Policing.”](#)

Global Protocol Packet Policing

- [Prerequisites for Global Protocol Packet Policing, page 38-3](#)
- [Restrictions for Global Protocol Packet Policing, page 38-3](#)
- [Information About Global Protocol Packet Policing, page 38-5](#)
- [How to Configure Single-Command Global Protocol Packet Policing, page 38-5](#)
- [How to Configure Policy-Based Global Protocol Packet Policing, page 38-6](#)

Prerequisites for Global Protocol Packet Policing

None.

Restrictions for Global Protocol Packet Policing

- The minimum values supported by the **platform qos protocol arp police** command are too small for use in production networks.
- ARP packets are approximately 40 bytes long and ARP reply packets are approximately 60 bytes long. The policer rate value is in bits per second. The burst value is in bytes per second. Together, an ARP request and reply are approximately 800 bits.
- The configured rate limits are applied separately to the PFC and each DFC. The RP CPU will receive the configured value times the number of forwarding engines.
- Policy-based protocol packet policing is applied per-forwarding engine (PFC and any DFCs).
- With Supervisor Engine 6T, policy-based protocol packet policing supports distributed aggregate policing (see the [“Enabling Distributed Aggregate Policing” section on page 25-7](#)).
- The protocol packet policing mechanism effectively protects the RP CPU against attacks such as line-rate ARP attacks, but it polices both routing protocols and ARP packets to the switch and also polices traffic through the switch with less granularity than CoPP.
- The policing mechanism shares the root configuration with a policing-avoidance mechanism. The policing-avoidance mechanism lets the routing protocol and ARP packets flow through the network when they reach a QoS policer. This mechanism can be configured using the **platform qos protocol protocol_name pass-through** command.
- Policy-based protocol packet policing does not support microflow policers.
- Only ingress Policy-based protocol packet policing is supported.
- Policy-based protocol packet policing does not support Layer 4 ACL operators (see the [“Restrictions for Layer 4 Operators in ACLs” section on page 31-2](#)), which imposes these subsequent restrictions:
 - For IPv4 or IPv6 traffic, no support for UDP or TCP port range matching
 - For IPv6 traffic, no support for precedence or DSCP matching
- Protocol packet policing policies can share an aggregate policer with QoS policies.
- An aggregate policer cannot be applied to both ingress and egress traffic.
- With Supervisor Engine 6T, policy-based protocol packet policing supports a maximum of 1,000 global TCAM entries.

- Policy-based protocol packet policing supports the **class default** and **permit protocol_name any any** commands, but traffic flow might be affected significantly because the protocol packet policing policy processes all matched traffic.
- With Supervisor Engine 6T, policy-based protocol packet policing works with any configured port trust state.
- You can configure both single-command protocol packet policing and policy-based protocol packet policing. Single-command protocol packet policing is applied first and then policy-based protocol packet policing.

**Note**

The software does not detect or attempt to resolve any configuration conflicts between single-command protocol packet policing and policy-based protocol packet policing.

- You can configure both policy-based protocol packet policing and control plane policing (see [Chapter 39, “Control Plane Policing \(CoPP\)”](#)). Policy-based protocol packet policing is applied first and then CoPP.
- Single-command protocol packet policing programs the configured protocol-specific action for ingress traffic and automatically programs a corresponding egress-traffic pass-through action to preserve the ingress result egress traffic.
- Policy-based protocol packet policing does not automatically preserve the ingress policing result in egress traffic.
 - To preserve the ingress policing result in egress traffic with policy-based protocol packet policing, configure an appropriate output policy. To pass egress traffic through unchanged, duplicate each ingress class in the output policy and configure **trust dscp** as the class-map action.
 - Without an output policy-map, egress traffic is processed by any configured interface-based policy-map and ingress global policy result will be overwritten.
- The PFC and any DFCs supports a single **match** command in **class-map match-all** class maps, except that the **match protocol** command can be configured in a class map with the **match dscp** or **match precedence** command.
- The PFC and any DFCs supports multiple **match** commands in **class-map match-any** class maps.
- Class maps can use the **match** commands listed in [Table 38-1](#) to configure a traffic class that is based on the match criteria.

Table 38-1 Traffic Classification Class Map match Commands and Match Criteria

match Commands	Direction	Match Criteria
match access-group { <i>access_list_number</i> name <i>access_list_name</i> }	Ingress	Access control list (ACL). Note Use ACLs to match the following: —CoS value —VLAN ID —Packet length
match any	Ingress	Any match criteria.
match cos	Ingress	CoS value.
match discard-class	Ingress	Discard class value.

Table 38-1 Traffic Classification Class Map match Commands and Match Criteria (continued)

match Commands	Direction	Match Criteria
match dscp Note The match protocol command can be configured in a class map with the match dscp command.	Ingress	DSCP value.
match l2 miss	Ingress	Layer 2 traffic flooded in a VLAN because it is addressed to a currently unlearned MAC-Layer destination address.
match mpls experimental topmost	Ingress	MPLS EXP value in the topmost label.
match precedence Note The match protocol command can be configured in a class map with the match precedence command.	Ingress	IP precedence values.
match protocol {arp ip ipv6} Note The match protocol command can be configured in a class map with the match dscp or match precedence command.	Ingress	Protocol.
match qos-group	Ingress	QoS group ID.

The PFC and any DFCs supports these ACL types for use with the **match access group** command:

Protocol	Numbered ACLs	Extended ACLs	Named ACLs
IPv4	Yes: 1 to 99 1300 to 1999	Yes: 100 to 199 2000 to 2699	Yes
IPv6	Not applicable	Yes (named)	Yes
MAC Layer	Not applicable	Not applicable	Yes
ARP	Not applicable	Not applicable	Yes

Information About Global Protocol Packet Policing

Attackers may try to overwhelm the RP CPU with routing protocol control packets (for example, ARP packets). Protocol packet policing rate limits this traffic in hardware. Release 15.1(1)SY1 and later releases support policy-based global protocol packet policing, shown in Cisco Feature Navigator as the Global QoS Policy feature.

How to Configure Single-Command Global Protocol Packet Policing

Enter the **platform qos protocol ?** to display the supported routing protocols.

The **platform qos protocol arp police** command rate limits ARP packets. This example shows how to allow 200 ARP requests and replies per second:

```
Router(config)# platform qos protocol arp police 200000 6000
```

This example shows how to display the available protocols to use with protocol packet policing:

```
Router(config)# platform qos protocol ?
isis
eigrp
ldp
ospf
rip
bgp
ospfv3
bgpv2
ripng
neigh-discover
wlccp
arp
```

This example shows how to display the available keywords to use with the **platform qos protocol** command:

```
Router(config)# platform qos protocol protocol_name ?
pass-through  pass-through keyword
police         police keyword
precedence     change ip-precedence(used to map the dscp to cos value)
```

How to Configure Policy-Based Global Protocol Packet Policing

Use these QoS sections and the global protocol packet policing policy map configuration section:

- [Configuring a Class Map, page 25-8](#)
- [Configuring a Policy Map, page 25-9](#)
- [Configuring a Global Protocol Packet Policing Policy Map, page 38-6](#)

Configuring a Global Protocol Packet Policing Policy Map

To configure a global protocol packet policing policy map, perform this task:

Command	Purpose
Router(config)# platform qos service-policy input <i>policy_map_name</i>	Configures a global protocol packet policing policy map. Note You can configure one input policy.

Unicast Reverse Path Forwarding (uRPF) Check

- [Prerequisites for uRPF Check, page 38-7](#)
- [Restrictions for uRPF Check, page 38-7](#)
- [Information about uRPF Check, page 38-7](#)
- [Configuring the Unicast RPF Check Mode, page 38-8](#)
- [Enabling Self-Pinging, page 38-9](#)

Prerequisites for uRPF Check

None.

Restrictions for uRPF Check

- Unicast RPF does not provide complete protection against spoofing. Spoofed packets can enter a network through unicast RPF-enabled interfaces if an appropriate return route to the source IP address exists.
- You can configure a unicast RPF mode on each interface.
- The “allow default” options of the unicast RPF modes do not offer significant protection against spoofing.
 - Strict unicast RPF Check with Allow Default—Received IP traffic that is sourced from a prefix that exists in the routing table passes the unicast RPF check if the prefix is reachable through the input interface. If a default route is configured, any IP packet with a source prefix that is not in the routing table passes the unicast RPF check if the ingress interface is a reverse path for the default route.
 - Loose unicast RPF Check with Allow Default—If a default route is configured, any IP packet passes the unicast RPF check.
- Unicast RPF Strict Mode—The unicast RPF strict mode provides the greatest security against spoofed traffic. If, on all unicast RPF-check enabled interfaces, the switch receives valid IP traffic through interfaces that are reverse paths for the traffic, then strict mode is an option.
- Unicast RPF Loose Mode—The unicast RPF loose mode provides less protection than strict mode, but it is an option on switches that receive valid IP traffic on interfaces that are not reverse paths for the traffic. The unicast RPF loose mode verifies that received traffic is sourced from a prefix that exists in the routing table, regardless of the interface on which the traffic arrives.

Information about uRPF Check

The unicast RPF check verifies that the source address of received IP packets is reachable. The unicast RPF check discards IP packets that lack a verifiable IP source prefix (route), which helps mitigate problems that are caused by traffic with malformed or forged (spoofed) IP source addresses.

The PFC4 and DFC4s provide hardware support for the unicast RPF check on up to 16 paths, both with and without ACL filtering, for both IPv4 and IPv6 traffic.

To ensure that no more than 16 reverse-path interfaces exist in the routing table for each prefix, enter the **maximum-paths 16** command in config-router mode when configuring OSPF, EIGRP, or BGP.

How to Configure Unicast RPF Check

- [Configuring the Unicast RPF Check Mode, page 38-8](#)
- [Enabling Self-Pinging, page 38-9](#)



Note

The following commands exist in the CLI, but have no function:

- **platform ip cef rpf interface-group**
- **platform ip cef rpf multipath interface-group**
- **platform ip cef rpf multipath pass**
- **platform ip cef rpf multipath punt**

Configuring the Unicast RPF Check Mode

To configure unicast RPF check mode, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port} {port-channel number}}	Selects an interface to configure. Note Based on the input port, unicast RPF check verifies the best return path before forwarding the packet on to the next destination.
Step 2	Router(config-if)# ip verify unicast source reachable-via {rx any} [allow-default] [list]	Configures the IPv4 unicast RPF check mode.
Step 3	Router(config-if)# ipv6 verify unicast source reachable-via {rx any} [allow-default] [list]	Configures the IPv6 unicast RPF check mode.
Step 4	Router(config-if)# exit	Exits interface configuration mode.
Step 5	Router# show platform hardware cef ip rpf	Verifies the IPv4 configuration.
Step 6	Router# show platform hardware cef ipv6 rpf	Verifies the IPv6 configuration.

- Use the **rx** keyword to enable strict check mode.
- Use the **any** keyword to enable exist-only check mode.
- Use the **allow-default** keyword to allow use of the default route for RPF verification.
- Use the **list** option to identify an access list.
 - If the access list denies network access, denied packets are dropped at the port.
 - If the access list permits network access, packets are forwarded to the destination address. Forwarded packets are counted in the interface statistics.
 - If the access list includes the logging action, information about the packets is sent to the log server.

This example shows how to enable unicast RPF exist-only check mode on Gigabit Ethernet port 4/1:

```
Router(config)# interface gigabitethernet 4/1
Router(config-if)# ipv6 verify unicast source reachable-via any
Router(config-if)# ip verify unicast source reachable-via any
Router(config-if)# end
Router#
```

This example shows how to enable unicast RPF strict check mode on Gigabit Ethernet port 4/2:

```
Router(config)# interface gigabitethernet 4/2
Router(config-if)# ipv6 verify unicast source reachable-via rx
Router(config-if)# ip verify unicast source reachable-via rx
Router(config-if)# end
Router#
```

Enabling Self-Pinging

With unicast RPF check enabled, by default the switch cannot ping itself. To enable self-pinging, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{vlan vlan_ID} {type slot/port} {port-channel number}}	Selects the interface to configure.
Step 2	Router(config-if)# ip verify unicast source reachable-via any allow-self-ping	Enables the switch to ping itself or a secondary address.
Step 3	Router(config-if)# exit	Exits interface configuration mode.

This example shows how to enable self-pinging:

```
Router(config)# interface gigabitethernet 4/1
Router(config-if)# ip verify unicast source reachable-via any allow-self-ping
Router(config-if)# end
```

Configuring Sticky ARP

Sticky ARP prevents MAC address spoofing by ensuring that ARP entries (IP address, MAC address, and source VLAN) do not get overridden. The switch maintains ARP entries in order to forward traffic to end devices or other switches. ARP entries are usually updated periodically or modified when ARP broadcasts are received. During an attack, ARP broadcasts are sent using a spoofed MAC address (with a legitimate IP address) so that the switch learns the legitimate IP address with the spoofed MAC address and begins to forward traffic to that MAC address. With sticky ARP enabled, the switch learns the ARP entries and does not accept modifications received through ARP broadcasts. If you attempt to override the sticky ARP configuration, you will receive an error message.

To configure sticky ARP on a Layer 3 interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface type ¹ slot/port	Selects the interface on which sticky ARP is applied.
Step 2	Router(config-if)# ip sticky-arp	Enables sticky ARP.
Step 3	Router(config-if)# ip sticky-arp ignore	Disables sticky ARP.

1. *type* = fastethernet, gigabitethernet, or tengigabitethernet

This example shows how to enable sticky ARP on interface 5/1:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/1
Router(config-if)# ip sticky-arp
```

```
Router(config-if)# end
Router#
```

Monitoring Packet Drop Statistics

- [Prerequisites for Packet Drop Statistics, page 38-10](#)
- [Restrictions for Packet Drop Statistics, page 38-10](#)
- [Information About Packet Drop Statistics, page 38-10](#)
- [How to Monitor Dropped Packets, page 38-10](#)

Prerequisites for Packet Drop Statistics

None.

Restrictions for Packet Drop Statistics

- The incoming captured traffic is not filtered.
- The incoming captured traffic is not rate limited to the capture destination.

Information About Packet Drop Statistics

You can use show commands to display packet drop statistics. You can capture the traffic on an interface and send a copy of this traffic to a traffic analyzer connected to a port, which can aggregate packet drop statistics.

How to Monitor Dropped Packets

- [Using show Commands, page 38-10](#)
- [Using SPAN, page 38-11](#)
- [Using VACL Capture, page 38-12](#)

Using show Commands

The PFC and DFCs support ACL hit counters in hardware. You can use the **show platform hardware acl entry interface** command to display each entry in the ACL TCAM. You can also use the TTL and IP options counters to monitor the performance of the Layer 3 forwarding engine.

This example shows how to use the **show platform hardware acl entry interface** command to display packet statistics and errors associated with the Layer 3 forwarding engine:

```
Router# show platform hardware statistics
```

```
--- Hardware Statistics for Module 6 ---
```

```
L2 Forwarding Engine
```

```

Switched in L2 : 59624 @ 7 pps

L3 Forwarding Engine
  Processed in L3 : 59624 @ 7 pps
  Switched in L3 : 13 @ 0 pps

  Bridged : 4602
  FIB Switched
    IPv4 Ucast : 7
    IPv6 Ucast : 1
    EoMPLS : 1
    MPLS : 1
    (S , *) : 0
    IGMP MLD : 0
    IPv4 Mcast : 2
    IPv6 Mcast : 0
    Mcast Leak : 0
  ACL Routed
    Input : 1
    Output : 518
  Netflow Switched
    Input : 2
    Output : 0
  Exception Redirected
    Input : 0
    Output : 1
  Mcast Bridge Disable & No Redirect
    : 0
  Total packets with TOS Changed : 3
  Total packets with TC Changed : 0
  Total packets with COS Changed : 64
  Total packets with EXP Changed : 0
  Total packets with QOS Tunnel Encap Changed : 1
  Total packets with QOS Tunnel Decap Changed : 1
  Total packets dropped by ACL : 1
  Total packets dropped by Policing : 0

Errors
  MAC/IP length inconsistencies : 0
  Short IP packets received : 0
  IP header checksum errors : 0
  TTL failures : 0
  MTU failures : 0

Total packets L3 Processed by all Modules: 59624 @ 7 pps

```

Using SPAN

This example shows how to use the **monitor session** command to capture and forward traffic to an external interface:

```

Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# monitor session 1 source vlan 44 both
Router(config)# monitor session 1 destination interface g9/1
Router(config)# end
Router#

```

This example shows how to use the **show monitor session** command to display the destination port:

```

Router# show monitor session 1
Session 1
-----
Source Ports:

```

```

RX Only:      None
TX Only:      None
Both:         None
Source VLANs:
RX Only:      None
TX Only:      None
Both:         44
Destination Ports: Gi9/1
Filter VLANs:  None

```

For more information, see [Chapter 18, “Local SPAN, RSPAN, and ERSPAN.”](#)

Using VACL Capture

The VACL capture feature allows you to direct traffic to ports configured to forward captured traffic. The capture action sets the capture bit for the forwarded packets so that ports with the capture function enabled can receive the packets. Only forwarded packets can be captured.

You can use VACL capture to assign traffic from each VLAN to a different interface.

VACL capture does not allow you to send one type of traffic, such as HTTP, to one interface and another type of traffic, such as DNS, to another interface. Also, VACL capture granularity is only applicable to traffic switched locally; you cannot preserve the granularity if you direct traffic to a remote switch.

For more information, see [Chapter 36, “VLAN ACLs \(VACLs\).”](#)



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

39

Control Plane Policing (CoPP)

- [Prerequisites for CoPP, page 39-1](#)
- [Restrictions for CoPP, page 39-1](#)
- [Information About CoPP, page 39-3](#)
- [Default Settings for CoPP, page 39-3](#)
- [How to Configure CoPP, page 39-5](#)
- [Monitoring CoPP, page 39-9](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- For more information about CoPP, see this document:
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11-663623.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for CoPP

None.

Restrictions for CoPP

- The PFC and DFC provide hardware support for classes that match multicast traffic.
- CoPP is not supported in hardware for broadcast packets. The combination of ACLs, traffic storm control, and CoPP software protection provides protection against broadcast DoS attacks.

- CoPP does not support non-IP classes except for the default non-IP class. ACLs can be used instead of non-IP classes to drop non-IP traffic, and the default non-IP CoPP class can be used to limit non-IP traffic that reaches the RP CPU.
- Do not use the **log** keyword in CoPP policy ACLs.
- If you have a large QoS configuration, the system may run out of TCAM space. If this is the case, CoPP may be performed in software.
- When there is a large QoS configuration for other interfaces, you can run out of TCAM space. When this situation occurs, CoPP may be performed entirely in software and result in performance degradation and CPU cycle consumption.
- You must ensure that the CoPP policy does not filter critical traffic such as routing protocols or interactive access to the switches. Filtering this traffic could prevent remote access to the switch, requiring a console connection.
- The PFC and DFCs support built-in special-case rate limiters, which are useful for situations where an ACL cannot be used (for example, TTL, MTU, and IP options). When you enable the special-case rate limiters, you should be aware that the special-case rate limiters will override the CoPP policy for packets matching the rate-limiter criteria.
- Neither egress CoPP nor silent mode is supported. CoPP is only supported on ingress (service-policy output CoPP cannot be applied to the control plane interface).
- ACE hit counters in hardware are only for ACL logic. You can rely on software ACE hit counters and the **show access-list**, **show policy-map control-plane**, and **show platform ip qos** commands to troubleshoot evaluate CPU traffic.
- CoPP is performed on a per-forwarding-engine basis and software CoPP is performed on an aggregate basis.
- CoPP does not support ACEs with the **log** keyword.
- CoPP uses hardware QoS TCAM resources. Enter the **show tcam utilization** command to verify the TCAM utilization.
- To avoid matching the filtering and policing that are configured in a subsequent class, configure policing in each class. CoPP does not apply the filtering in a class that does not contain a police command. A class without a police command matches no traffic.
- The ACLs used for classification are QoS ACLs. The supported QoS ACLs are IP standard, extended, and named.
- These are the only match types supported:
 - **ip precedence**
 - **ip dscp**
 - **access-group**
- Only IP ACLs are supported in hardware.
- MAC-based matching is done in software only.
- You can enter one **match** command in a single class map only.
- When defining the service policy, the **police** policy-map action is the only supported action.
- When applying the service policy to the control plane, the **input** direction is only supported.

Information About CoPP

The traffic managed by the RP is divided into three functional components or *planes*:

- Data plane
- Management plane
- Control plane

The control plane policing (CoPP) feature increases security on the switch by protecting the RP from unnecessary or DoS traffic and giving priority to important control plane and management traffic. The PFC and DFCs provide hardware support for CoPP. CoPP works with the hardware rate limiters.

The PFC and DFCs support the built-in “special case” rate limiters that can be used when an ACL cannot classify particular scenarios, such as IP options cases, TTL and MTU failure cases, packets with errors, and multicast packets. When enabling the special-case rate limiters, the special-case rate limiters override the CoPP policy for packets matching the rate-limiter criteria.

The majority of traffic managed by the RP is handled by way of the control and management planes. You can use CoPP to protect the control and management planes, and ensure routing stability, reachability, and packet delivery. CoPP uses a dedicated control plane configuration through the modular QoS CLI (MQC) to provide filtering and rate-limiting capabilities for the control plane packets.

You must first identify the traffic to be classified by defining a class map. The class map defines packets for a particular traffic class. After you have classified the traffic, you can create policy maps to enforce policy actions for the identified traffic.

Default Settings for CoPP

CoPP is enabled by default. To disable the default CoPP configuration, enter the **no service-policy input policy-default-autocopp** control plane configuration mode command.

These are the default CoPP class maps:

```
class-map match-any class-copp-icmp-redirect-unreachable
class-map match-all class-copp-glean
class-map match-all class-copp-receive
class-map match-all class-copp-options
class-map match-all class-copp-broadcast
class-map match-all class-copp-mcast-acl-bridged
class-map match-all class-copp-slb
class-map match-all class-copp-mtu-fail
class-map match-all class-copp-ttl-fail
class-map match-all class-copp-arp-snooping
class-map match-any class-copp-mcast-copy
class-map match-any class-copp-ip-connected
class-map match-any class-copp-match-igmp
  match access-group name acl-copp-match-igmp
class-map match-all class-copp-unknown-protocol
class-map match-any class-copp-vacl-log
class-map match-all class-copp-mcast-ipv6-control
class-map match-any class-copp-match-pimv6-data
  match access-group name acl-copp-match-pimv6-data
class-map match-any class-copp-mcast-punt
class-map match-all class-copp-unsupp-rewrite
class-map match-all class-copp-ucast-egress-acl-bridged
class-map match-all class-copp-ip-admission
class-map match-all class-copp-service-insertion
class-map match-all class-copp-mac-pbf
```

```

class-map match-any class-copp-match-mld
  match access-group name acl-copp-match-mld
class-map match-all class-copp-ucast-ingress-acl-bridged
class-map match-all class-copp-dhcp-snooping
class-map match-all class-copp-wccp
class-map match-all class-copp-nd
class-map match-any class-copp-ipv6-connected
class-map match-all class-copp-mcast-rpf-fail
class-map match-any class-copp-ucast-rpf-fail
class-map match-all class-copp-mcast-ip-control
class-map match-any class-copp-match-pim-data
  match access-group name acl-copp-match-pim-data
class-map match-any class-copp-match-ndv6
  match access-group name acl-copp-match-ndv6
class-map match-any class-copp-mcast-v4-data-on-routedPort
class-map match-any class-copp-mcast-v6-data-on-routedPort

```

This is the default CoPP policy map:

```

policy-map policy-default-autocopp
  class class-copp-mcast-v4-data-on-routedPort
    police rate 10 pps burst 1 packets
    conform-action drop
    exceed-action drop
  class class-copp-mcast-v6-data-on-routedPort
    police rate 10 pps burst 1 packets
    conform-action drop
    exceed-action drop
  class class-copp-icmp-redirect-unreachable
    police rate 100 pps burst 10 packets
    conform-action transmit
    exceed-action drop
  class class-copp-ucast-rpf-fail
    police rate 100 pps burst 10 packets
    conform-action transmit
    exceed-action drop
  class class-copp-vacl-log
    police rate 2000 pps burst 1 packets
    conform-action transmit
    exceed-action drop
  class class-copp-mcast-punt
    police rate 1000 pps burst 256 packets
    conform-action transmit
    exceed-action drop
  class class-copp-mcast-copy
    police rate 1000 pps burst 256 packets
    conform-action transmit
    exceed-action drop
  class class-copp-ip-connected
    police rate 1000 pps burst 256 packets
    conform-action transmit
    exceed-action drop
  class class-copp-ipv6-connected
    police rate 1000 pps burst 256 packets
    conform-action transmit
    exceed-action drop
  class class-copp-match-pim-data
    police rate 1000 pps burst 1000 packets
    conform-action transmit
    exceed-action drop
  class class-copp-match-pimv6-data
    police rate 1000 pps burst 1000 packets
    conform-action transmit
    exceed-action drop

```

```

class class-copp-match-mld
  police rate 10000 pps burst 10000 packets
  conform-action set-discard-class-transmit 48
  exceed-action transmit
class class-copp-match-igmp
  police rate 10000 pps burst 10000 packets
  conform-action set-discard-class-transmit 48
  exceed-action transmit
class class-copp-match-ndv6
  police rate 1000 pps burst 1000 packets
  conform-action set-discard-class-transmit 48
  exceed-action drop

```

How to Configure CoPP

- [Configuring CoPP, page 39-5](#)
- [Defining CoPP Traffic Classification, page 39-6](#)

Configuring CoPP

To configure CoPP, perform this task:

	Command	Purpose
Step 1	Router(config)# ip access-list extended <i>access_list_name</i>	Creates an extended ACL. Note You must configure ACLs in most cases to identify the important or unimportant traffic.
Step 2	Router(config-ext-nacl)# { permit deny } <i>protocol</i> source <i>source_wildcard</i> destination <i>destination_wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [log log-input] [time-range <i>time_range_name</i>] [fragments]	Configures filtering in the ACL: • permit sets the conditions under which a packet passes a named IP access list. • deny sets the conditions under which a packet does not pass a named IP access list.
Step 3	Router(config)# class-map <i>traffic_class_name</i>	Creates a class map.
Step 4	Router(config-cmap)# match { ip precedence ip dscp <i>access_group</i> }	Configures matching in the class map.
Step 5	Router(config)# policy-map <i>service-policy-name</i>	Defines a service policy map.
Step 6	Router(config-pmap)# class <i>traffic_class_name</i>	Creates a policy map class.

	Command	Purpose
Step 7	<pre>Router(config-pmap-c)# police bits_per_second [normal_burst_bytes [maximum_burst_bytes]] [pir peak_rate_bps] [[conform-action selected_action] exceed-action selected_action] violate-action selected_action] Router(config-pmap-c)# police rate units bps [burst burst_bytes bytes] [peak-rate peak_rate_bps bps] [peak-burst peak_burst_bytes bytes] [conform-action selected_action] [exceed-action selected_action] [violate-action selected_action] Router(config-pmap-c)# police rate units pps [burst burst_packets packets] [peak-rate peak_rate_pps pps] [peak-burst peak_burst_packets packets] [conform-action selected_action] [exceed-action selected_action] [violate-action selected_action] Router(config-pmap-c)# police flow [mask {src-only dest-only full-flow}] bits_per_second normal_burst_bytes [[conform-action {drop set-dscp-transmit dscp_value set-prec-transmit ip_precedence_value transmit}] exceed-action {drop policed-dscp transmit}] violate-action {drop policed-dscp transmit}]</pre>	Configures policing in the service policy map. You can configure any of the following: • Byte-based policing. • Packet-based policing. • Flow-based policing. See the “Configuring Policy Map Class Policing” section on page 25-11.
Step 8	Router(config)# control-plane	Enters the control plane configuration mode.
Step 9	Router(config-cp)# service-policy input service-policy-name	Applies the QoS service policy to the control plane.

Defining CoPP Traffic Classification

- [Traffic Classification Overview, page 39-6](#)
- [Traffic Classification Restrictions, page 39-7](#)
- [Sample Basic ACLs for CoPP Traffic Classification, page 39-8](#)

Traffic Classification Overview

You can define any number of classes, but typically traffic is grouped into classes that are based on relative importance. The following provides a sample grouping:

- Border Gateway Protocol (BGP)—Traffic that is crucial to maintaining neighbor relationships for BGP routing protocol, for example, BGP keepalives and routing updates. Maintaining BGP routing protocol is crucial to maintaining connectivity within a network or to a service provider. Sites that do not run BGP do not need to use this class.

- **Interior Gateway Protocol (IGP)**—Traffic that is crucial to maintaining IGP routing protocols, for example, open shortest path first OSPF, enhanced interior gateway routing protocol (EIGRP), and routing information protocol (RIP). Maintaining IGP routing protocols is crucial to maintaining connectivity within a network.
- **Management**—Necessary, frequently used traffic that is required during day-to-day operations. For example, traffic used for remote network access, and Cisco IOS image upgrades and management, such as Telnet, secure shell (SSH), network time protocol (NTP), simple network management protocol (SNMP), terminal access controller access control system (TACACS), hypertext transfer protocol (HTTP), trivial file transfer protocol (TFTP), and file transfer protocol (FTP).
- **Reporting**—Traffic used for generating network performance statistics for the purpose of reporting. For example, using Cisco IOS IP service level agreements (SLAs) to generate ICMP with different DSCP settings in order to report on response times within different QoS data classes.
- **Monitoring**—Traffic used for monitoring a switch. Traffic should be permitted but should never be a risk to the switch; with CoPP, this traffic can be permitted but limited to a low rate. For example, ICMP echo request (ping) and traceroute.
- **Critical Applications**—Critical application traffic that is specific and crucial to a particular customer environment. Traffic included in this class should be tailored specifically to the required application requirements of the user (in other words, one customer may use multicast, while another uses IPsec or generic routing encapsulation (GRE). For example, GRE, hot standby router protocol (HSRP), virtual router redundancy protocol (VRRP), session initiation protocol (SIP), data link switching (DLSw), dynamic host configuration protocol (DHCP), multicast source discovery protocol (MSDP), Internet group management protocol (IGMP), protocol independent multicast (PIM), multicast traffic, and IPsec.
- **Layer 2 Protocols**—Traffic used for address resolution protocol (ARP). Excessive ARP packets can potentially monopolize RP resources, starving other important processes; CoPP can be used to rate limit ARP packets to prevent this situation. ARP is the only Layer 2 protocol that can be specifically classified using the match protocol classification criteria.
- **Undesirable**—Explicitly identifies bad or malicious traffic that should be unconditionally dropped and denied access to the RP. The undesirable classification is particularly useful when known traffic destined for the switch should always be denied and not placed into a default category. If you explicitly deny traffic, then you can enter **show** commands to collect approximate statistics on the denied traffic and estimate its rate.
- **Default**—All remaining traffic destined for the RP that has not been identified. MQC provides the default class, so the user can specify the treatment to be applied to traffic not explicitly identified in the other user-defined classes. This traffic has a highly reduced rate of access to the RP. With a default classification in place, statistics can be monitored to determine the rate of otherwise unidentified traffic destined for the control plane. After this traffic is identified, further analysis can be performed to classify it and, if needed, the other CoPP policy entries can be updated to accommodate this traffic.

After you have classified the traffic, the ACLs build the classes of traffic that are used to define the policies. For sample basic ACLs for CoPP classification, see the [“Sample Basic ACLs for CoPP Traffic Classification” section on page 39-8](#).

Traffic Classification Restrictions

- Before you develop the actual CoPP policy, you must identify and separate the required traffic into different classes. Traffic is grouped into nine classes that are based on relative importance. The actual number of classes needed might differ and should be selected based on your local requirements and security policies.

- You do not have to define policies that match bidirectionally. You only need to identify traffic unidirectionally (from the network to the RP) since the policy is applied on ingress only.

Sample Basic ACLs for CoPP Traffic Classification

This section shows sample basic ACLs for CoPP classification. In the samples, the commonly required traffic is identified with these ACLs:

- ACL 120—Critical traffic
- ACL 121—Important traffic
- ACL 122—Normal traffic
- ACL 123—Explicitly denies unwanted traffic
- ACL 124—All other traffic

This example shows how to define ACL 120 for critical traffic:

```
Router(config)# access-list 120 remark CoPP ACL for critical traffic
```

This example shows how to allow BGP from a known peer to this switch's BGP TCP port:

```
Router(config)# access-list 120 permit tcp host 47.1.1.1 host 10.9.9.9 eq bgp
```

This example shows how to allow BGP from a peer's BGP port to this switch:

```
Router(config)# access-list 120 permit tcp host 47.1.1.1 eq bgp host 10.9.9.9
Router(config)# access-list 120 permit tcp host 10.86.183.120 host 10.9.9.9 eq bgp
Router(config)# access-list 120 permit tcp host 10.86.183.120 eq bgp host 10.9.9.9
```

This example shows how to define ACL 121 for the important class:

```
Router(config)# access-list 121 remark CoPP Important traffic
```

This example shows how to permit return traffic from TACACS host:

```
Router(config)# access-list 121 permit tcp host 1.1.1.1 host 10.9.9.9 established
```

This example shows how to permit SSH access to the switch from a subnet:

```
Router(config)# access-list 121 permit tcp 10.0.0.0 0.0.0.255 host 10.9.9.9 eq 22
```

This example shows how to allow full access for Telnet to the switch from a host in a specific subnet and police the rest of the subnet:

```
Router(config)# access-list 121 deny tcp host 10.86.183.3 any eq telnet
Router(config)# access-list 121 permit tcp 10.86.183.0 0.0.0.255 any eq telnet
```

This example shows how to allow SNMP access from the NMS host to the switch:

```
Router(config)# access-list 121 permit udp host 1.1.1.2 host 10.9.9.9 eq snmp
```

This example shows how to allow the switch to receive NTP packets from a known clock source:

```
Router(config)# access-list 121 permit udp host 1.1.1.3 host 10.9.9.9 eq ntp
```

This example shows how to define ACL 122 for the normal traffic class:

```
Router(config)# access-list 122 remark CoPP normal traffic
```

This example shows how to permit switch-originated traceroute traffic:

```
Router(config)# access-list 122 permit icmp any any ttl-exceeded
Router(config)# access-list 122 permit icmp any any port-unreachable
```

This example shows how to permit receipt of responses to the switch that originated the pings:

```
Router(config)# access-list 122 permit icmp any any echo-reply
```

This example shows how to allow pings to the switch:

```
Router(config)# access-list 122 permit icmp any any echo
```

This example shows how to define ACL 123 for the undesirable class.

```
Router(config)# access-list 123 remark explicitly defined "undesirable" traffic
```



Note

In the following example, ACL 123 is a permit entry for classification and monitoring purposes, and traffic is dropped as a result of the CoPP policy.

This example shows how to permit all traffic destined to UDP 1434 for policing:

```
Router(config)# access-list 123 permit udp any any eq 1434
```

This example shows how to define ACL 124 for all other traffic:

```
Router(config)# access-list 124 remark rest of the IP traffic for CoPP
Router(config)# access-list 124 permit ip any any
```

Monitoring CoPP

You can enter the **show policy-map control-plane** command for developing site-specific policies, monitoring statistics for the control plane policy, and troubleshooting CoPP. This command displays dynamic information about the actual policy applied, including rate information and the number of bytes (and packets) that conformed or exceeded the configured policies both in hardware and in software.

The output of the **show policy-map control-plane** command is as follows:

```
Router# show policy-map control-plane
Control Plane Interface
  Service policy CoPP-normal
Hardware Counters:
class-map: CoPP-normal (match-all)
  Match: access-group 130
  police :
    96000 bps 3000 limit 3000 extended limit
  Earl in slot 3 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps
  Earl in slot 5 :
    0 bytes
    5 minute offered rate 0 bps
    aggregate-forwarded 0 bytes action: transmit
    exceeded 0 bytes action: drop
    aggregate-forward 0 bps exceed 0 bps

Software Counters:
Class-map: CoPP-normal (match-all) 0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: access-group 130
  police:
    96000 bps, 3125 limit, 3125 extended limit
```

```

conformed 0 packets, 0 bytes; action: transmit
exceeded 0 packets, 0 bytes; action: drop
conformed 0 bps, exceed 0 bps, violate 0 bps
Router#

```

To display the hardware counters for bytes dropped and forwarded by the policy, enter the **show platform qos ip** command:

```

Router# show platform qos ip
QoS Summary [IP]:      (* - shared aggregates, Mod - switch module)

Int Mod Dir  Class-map DSCP  Agg  Trust Fl  AgForward-By  AgPoliced-By
      Id      Id
-----
CPP  5  In CoPP-normal    0   1  dscp  0      505408      83822272
CPP  9  In CoPP-normal    0   4  dscp  0           0           0
Router#

```

To display the CoPP access list information, enter the **show access-lists coppacl-bgp** command:

```

Router# show access-lists coppacl-bgp
Extended IP access list coppacl-bgp
10 permit tcp host 47.1.1.1 host 10.9.9.9 eq bgp (4 matches)
20 permit tcp host 47.1.1.1 eq bgp host 10.9.9.9
30 permit tcp host 10.86.183.120 host 10.9.9.9 eq bgp (1 match)
40 permit tcp host 10.86.183.120 eq bgp host 10.9.9.9

```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

40

Dynamic Host Configuration Protocol (DHCP) Snooping

- [Prerequisites for DHCP Snooping, page 40-1](#)
- [Restrictions for DHCP Snooping, page 40-1](#)
- [Information About DHCP Snooping, page 40-3](#)
- [Default Configuration for DHCP Snooping, page 40-8](#)
- [How to Configure DHCP Snooping, page 40-9](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for DHCP Snooping

None.

Restrictions for DHCP Snooping

- [DHCP Snooping Configuration Restrictions, page 40-2](#)
- [DHCP Snooping Configuration Guidelines, page 40-2](#)
- [Minimum DHCP Snooping Configuration, page 40-2](#)

DHCP Snooping Configuration Restrictions

- The DHCP snooping database stores at least 12,000 bindings.
- When DHCP snooping is enabled, these Cisco IOS DHCP commands are not available on the switch:
 - **ip dhcp relay information check** global configuration command
 - **ip dhcp relay information policy** global configuration command
 - **ip dhcp relay information trust-all** global configuration command
 - **ip dhcp relay information option** global configuration command
 - **ip dhcp relay information trusted** interface configuration command

If you enter these commands, the switch returns an error message, and the configuration is not applied.

DHCP Snooping Configuration Guidelines

- DHCP snooping is not active until you enable the feature on at least one VLAN, and enable DHCP globally on the switch.
- 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.
- For DHCP server configuration information, see this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book.html
- If a Layer 2 LAN port is connected to a DHCP server, configure the port as trusted by entering the **ip dhcp snooping trust** interface configuration command.
- If a Layer 2 LAN port is connected to a DHCP client, configure the port as untrusted by entering the **no ip dhcp snooping trust** interface configuration command.
- You can enable DHCP snooping on private VLANs:
 - If DHCP snooping is enabled, any primary VLAN configuration is propagated to its associated secondary VLANs.
 - If DHCP snooping is configured on the primary VLAN and you configure DHCP snooping with different settings on an associated secondary VLAN, the configuration on the secondary VLAN does not take effect.
 - If DHCP snooping is not configured on the primary VLAN and you configure DHCP snooping on a secondary VLAN, the configuration takes effect only on the secondary VLAN.
 - When you manually configure DHCP snooping on a secondary VLAN, this message appears:
 DHCP Snooping configuration may not take effect on secondary vlan XXX
 - The **show ip dhcp snooping** command displays all VLANs (both primary and secondary) that have DHCP snooping enabled.

Minimum DHCP Snooping Configuration

1. Define and configure the DHCP server. See this publication:

http://www.cisco.com/en/US/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-sy/dhcp-15-sy-book.html

2. Enable DHCP snooping on at least one VLAN.

By default, DHCP snooping is inactive on all VLANs. See the “[Enabling DHCP Snooping on VLANs](#)” section on page 40-11

3. Ensure that DHCP server is connected through a trusted interface.

By default, the trust state of all interfaces is untrusted. See the “[Configuring the DHCP Trust State on Layer 2 LAN Interfaces](#)” section on page 40-12

4. Configure the DHCP snooping database agent.

This step ensures that database entries are restored after a restart or switchover. See the “[The DHCP Snooping Database Agent](#)” section on page 40-14

5. Enable DHCP snooping globally.

The feature is not active until you complete this step. See the “[Enabling DHCP Snooping Globally](#)” section on page 40-9

If you are configuring the switch for DHCP relay, the following additional steps are required:

1. Define and configure the DHCP relay agent IP address.

If the DHCP server is in a different subnet from the DHCP clients, configure the server IP address in the helper address field of the client side VLAN.

2. Configure DHCP option-82 on untrusted port.

See the “[Enabling the DHCP Option-82 on Untrusted Port Feature](#)” section on page 40-10

Information About DHCP Snooping

- [Overview of DHCP Snooping, page 40-3](#)
- [Trusted and Untrusted Sources, page 40-4](#)
- [DHCP Snooping Binding Database, page 40-4](#)
- [Packet Validation, page 40-5](#)
- [DHCP Snooping Option-82 Data Insertion, page 40-5](#)
- [Overview of the DHCP Snooping Database Agent, page 40-7](#)

Overview of DHCP Snooping

DHCP snooping is a security feature that acts like a firewall between untrusted hosts and trusted DHCP servers. The DHCP snooping feature performs the following activities:

- Validates DHCP messages received from untrusted sources and filters out invalid messages.
- Rate-limits DHCP traffic from trusted and untrusted sources.
- Builds and maintains the DHCP snooping binding database, which contains information about untrusted hosts with leased IP addresses.
- Utilizes the DHCP snooping binding database to validate subsequent requests from untrusted hosts.

Other security features, such as dynamic ARP inspection (DAI), also use information stored in the DHCP snooping binding database.

DHCP snooping is enabled on a per-VLAN basis. By default, the feature is inactive on all VLANs. You can enable the feature on a single VLAN or a range of VLANs.

The DHCP snooping feature is implemented in software on the route processor (RP). Therefore, all DHCP messages for enabled VLANs are intercepted in the PFC and directed to the RP for processing.

Trusted and Untrusted Sources

The DHCP snooping feature determines whether traffic sources are trusted or untrusted. An untrusted source may initiate traffic attacks or other hostile actions. To prevent such attacks, the DHCP snooping feature filters messages and rate-limits traffic from untrusted sources.

In an enterprise network, devices under your administrative control are trusted sources. These devices include the switches, routers, and servers in your network. Any device beyond the firewall or outside your network is an untrusted source. Host ports and unknown DHCP servers are generally treated as untrusted sources.

A DHCP server that is on your network without your knowledge on an untrusted port is called a *spurious DHCP server*. A spurious DHCP server is any piece of equipment that is loaded with DHCP server enabled. Some examples are desktop systems and laptop systems that are loaded with DHCP server enabled, or wireless access points honoring DHCP requests on the wired side of your network. If spurious DHCP servers remain undetected, you will have difficulties troubleshooting a network outage. You can detect spurious DHCP servers by sending dummy DHCPDISCOVER packets out to all of the DHCP servers so that a response is sent back to the switch.

In a service provider environment, any device that is not in the service provider network is an untrusted source (such as a customer switch). Host ports are untrusted sources.

In the switch, you indicate that a source is trusted by configuring the trust state of its connecting interface.

The default trust state of all interfaces is untrusted. You must configure DHCP server interfaces as trusted. You can also configure other interfaces as trusted if they connect to devices (such as switches or routers) inside your network. You usually do not configure host port interfaces as trusted.



Note

For DHCP snooping to function properly, all DHCP servers must be connected to the switch through trusted interfaces, as untrusted DHCP messages will be forwarded only to trusted interfaces.

DHCP Snooping Binding Database

The DHCP snooping binding database is also referred to as the DHCP snooping binding table.

The DHCP snooping feature dynamically builds and maintains the database using information extracted from intercepted DHCP messages. The database contains an entry for each untrusted host with a leased IP address if the host is associated with a VLAN that has DHCP snooping enabled. The database does not contain entries for hosts connected through trusted interfaces.

The DHCP snooping feature updates the database when the switch receives specific DHCP messages. For example, the feature adds an entry to the database when the switch receives a DHCPACK message from the server. The feature removes the entry in the database when the IP address lease expires or the switch receives a DHCPRELEASE message from the host.

Each entry in the DHCP snooping binding database includes the MAC address of the host, the leased IP address, the lease time, the binding type, and the VLAN number and interface information associated with the host.

Packet Validation

The switch validates DHCP packets received on the untrusted interfaces of VLANs with DHCP snooping enabled. The switch forwards the DHCP packet unless any of the following conditions occur (in which case the packet is dropped):

- The switch receives a packet (such as a DHCP OFFER, DHCP ACK, DHCP NAK, or DHCP REQUEST packet) from a DHCP server outside the network or firewall.
- The switch receives a packet on an untrusted interface, and the source MAC address and the DHCP client hardware address do not match. This check is performed only if the DHCP snooping MAC address verification option is turned on.
- The switch receives a DHCP RELEASE or DHCP DECLINE message from an untrusted host with an entry in the DHCP snooping binding table, and the interface information in the binding table does not match the interface on which the message was received.
- The switch receives a DHCP packet that includes a relay agent IP address that is not 0.0.0.0.

To support trusted edge switches that are connected to untrusted aggregation-switch ports, you can enable the DHCP option-82 on untrusted port feature, which enables untrusted aggregation-switch ports to accept DHCP packets that include option-82 information. Configure the port on the edge switch that connects to the aggregation switch as a trusted port.




Note

With the DHCP option-82 on untrusted port feature enabled, use dynamic ARP inspection on the aggregation switch to protect untrusted input interfaces.

DHCP Snooping 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 snooping 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.

Figure 40-1 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 is configured with a helper address to enable broadcast forwarding and to transfer DHCP messages between the clients and the server.

Figure 40-1 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 option-82 information contains the switch MAC address (the remote ID suboption) and the port identifier, vlan-mod-port, from which the packet is received (the circuit ID suboption).
- If IEEE 802.1X port-based authentication is enabled, the switch will also add the host's 802.1X authenticated user identity information (the RADIUS attributes suboption) to the packet. See the [“802.1X Authentication with DHCP Snooping”](#) section on page 45-15.
- If the IP address of the relay agent is configured, the switch adds the 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, or 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. The DHCP server then 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. When the client and server are on the same subnet, the server broadcasts the reply. 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 previously described sequence of events occurs, the values in these fields in [Figure 40-2](#) 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 circuit ID type


Figure 40-2 shows the packet formats for the remote ID suboption and the circuit ID suboption. The switch uses the packet formats when DHCP snooping is globally enabled and when the **ip dhcp snooping information option** global configuration command is entered. For the circuit ID suboption, the module field is the slot number of the module.

Figure 40-2 Suboption Packet Formats

Circuit ID Suboption Frame Format



Remote ID Suboption Frame Format



Overview of the DHCP Snooping Database Agent

To retain the bindings across reloads, you must use the DHCP snooping database agent. Without this agent, the bindings established by DHCP snooping are lost upon reload, and connectivity is lost as well.

The database agent stores the bindings in a file at a configured location. Upon reload, the switch reads the file to build the database for the bindings. The switch keeps the file current by writing to the file as the database changes.

The format of the file that contains the bindings is as follows:

```
<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 that is used to validate the entries whenever the file is read. The <initial-checksum> entry on the first line helps distinguish entries associated with the latest write from entries that are associated with a previous write.

This is a sample bindings file:

```
3ebe1518
TYPE DHCP-SNOOPING
VERSION 1
BEGIN
1.1.1.1 512 0001.0001.0005 3EBE2881 Gi1/1 e5e1e733
1.1.1.1 512 0001.0001.0002 3EBE2881 Gi1/1 4b3486ec
1.1.1.1 1536 0001.0001.0004 3EBE2881 Gi1/1 f0e02872
1.1.1.1 1024 0001.0001.0003 3EBE2881 Gi1/1 ac41adf9
1.1.1.1 1 0001.0001.0001 3EBE2881 Gi1/1 34b3273e
END
```

Each entry holds an IP address, VLAN, MAC address, lease time (in hex), and the interface associated with a binding. At the end of each entry is a checksum that is based on all the bytes from the start of the file through all the bytes associated with the entry. Each entry consists of 72 bytes of data, followed by a space, followed by a checksum.

Upon bootup, when the calculated checksum equals the stored checksum, the switch reads entries from the file and adds the bindings to the DHCP snooping database. If the calculated checksum does not equal the stored checksum, the entry read from the file is ignored and so are all the entries following the failed entry. The switch also ignores all those entries from the file whose lease time has expired. (This is possible because the lease time might indicate an expired time.) An entry from the file is also ignored if the interface referred to in the entry no longer exists on the system, or if it is a router port or a DHCP snooping-trusted interface.

When the switch learns of new bindings or when it loses some bindings, the switch writes the modified set of entries from the snooping database to the file. The writes are performed with a configurable delay to batch as many changes as possible before the actual write happens. Associated with each transfer is a timeout after which a transfer is aborted if it is not completed. These timers are referred to as the write delay and abort timeout.

Default Configuration for DHCP Snooping

Option	Default Value/State
DHCP snooping	Disabled
DHCP snooping information option	Enabled
DHCP option-82 on untrusted port feature	Disabled
DHCP snooping limit rate	None
DHCP snooping trust	Untrusted
DHCP snooping vlan	Disabled
DHCP snooping spurious server detection	Disabled
DHCP snooping detect spurious interval	30 minutes

How to Configure DHCP Snooping

- [Enabling DHCP Snooping Globally, page 40-9](#)
- [Enabling DHCP Option-82 Data Insertion, page 40-10](#)
- [Enabling the DHCP Option-82 on Untrusted Port Feature, page 40-10](#)
- [Enabling DHCP Snooping MAC Address Verification, page 40-11](#)
- [Enabling DHCP Snooping on VLANs, page 40-11](#)
- [Configuring the DHCP Trust State on Layer 2 LAN Interfaces, page 40-12](#)
- [Configuring Spurious DHCP Server Detection, page 40-13](#)
- [Configuring DHCP Snooping Rate Limiting on Layer 2 LAN Interfaces, page 40-14](#)
- [The DHCP Snooping Database Agent, page 40-14](#)
- [Configuration Examples for the Database Agent, page 40-15](#)
- [Displaying the DHCP Snooping Binding Table, page 40-18](#)

Enabling DHCP Snooping Globally



Note

Configure this command as the last configuration step (or enable the DHCP feature during a scheduled maintenance period) because after you enable DHCP snooping globally, the switch drops DHCP requests until you configure the ports.

To enable DHCP snooping globally, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping	Enables DHCP snooping globally.
Step 2	Router(config)# do show ip dhcp snooping include Switch	Verifies the configuration.

This example shows how to enable DHCP snooping globally:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip dhcp snooping
Router(config)# do show ip dhcp snooping | include Switch
Switch DHCP snooping is enabled
Router(config)#
```



Note

When DHCP snooping is disabled and DAI is enabled, the switch shuts down all the hosts because all ARP entries in the ARP table will be checked against a nonexistent DHCP database. When DHCP snooping is disabled or in non-DHCP environments, use ARP ACLs to permit or to deny ARP packets.

Enabling DHCP Option-82 Data Insertion

To enable DHCP option-82 data insertion, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping information option	Enables DHCP option-82 data insertion.
Step 2	Router(config)# do show ip dhcp snooping include 82	Verifies the configuration.

This example shows how to disable DHCP option-82 data insertion:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# no ip dhcp snooping information option
Router(config)# do show ip dhcp snooping | include 82
Insertion of option 82 is disabled
Router(config)#
```

This example shows how to enable DHCP option-82 data insertion:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip dhcp snooping information option
Router(config)# do show ip dhcp snooping | include 82
Insertion of option 82 is enabled
Router(config)#
```

Enabling the DHCP Option-82 on Untrusted Port Feature



Note

With the DHCP option-82 on untrusted port feature enabled, the switch does not drop DHCP packets that include option-82 information that are received on untrusted ports. Do not enter the **ip dhcp snooping information option allowed-untrusted** command on an aggregation switch to which any untrusted devices are connected.

To enable untrusted ports to accept DHCP packets that include option-82 information, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping information option allow-untrusted	(Optional) Enables untrusted ports to accept incoming DHCP packets with option-82 information. The default setting is disabled.
Step 2	Router(config)# do show ip dhcp snooping	Verifies the configuration.

This example shows how to enable the DHCP option-82 on untrusted port feature:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip dhcp snooping information option allow-untrusted
Router(config)#
```

Enabling DHCP Snooping MAC Address Verification

With DHCP snooping MAC address verification enabled, DHCP snooping verifies that the source MAC address and the client hardware address match in DHCP packets that are received on untrusted ports. The source MAC address is a Layer 2 field associated with the packet, and the client hardware address is a Layer 3 field in the DHCP packet.

To enable DHCP snooping MAC address verification, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping verify mac-address	Enables DHCP snooping MAC address verification.
Step 2	Router(config)# do show ip dhcp snooping include hwaddr	Verifies the configuration.

This example shows how to disable DHCP snooping MAC address verification:

```
Router(config)# no ip dhcp snooping verify mac-address
Router(config)# do show ip dhcp snooping | include hwaddr
Verification of hwaddr field is disabled
Router(config)#
```

This example shows how to enable DHCP snooping MAC address verification:

```
Router(config)# ip dhcp snooping verify mac-address
Router(config)# do show ip dhcp snooping | include hwaddr
Verification of hwaddr field is enabled
Router(config)#
```

Enabling DHCP Snooping on VLANs

By default, the DHCP snooping feature is inactive on all VLANs. You may enable the feature on a single VLAN or a range of VLANs.

When enabled on a VLAN, the DHCP snooping feature creates four entries in the VACL table in the MFC3. These entries cause the PFC or DFC to intercept all DHCP messages on this VLAN and send them to the RP. The DHCP snooping feature is implemented in software on the RP.

To enable DHCP snooping on VLANs, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping vlan {{vlan_ID [vlan_ID]} {vlan_range}}	Enables DHCP snooping on a VLAN or VLAN range.
Step 2	Router(config)# do show ip dhcp snooping	Verifies the configuration.

You can configure DHCP snooping for a single VLAN or a range of VLANs:

- To configure a single VLAN, enter a single VLAN number.
- To configure a range of VLANs, enter a beginning and an ending VLAN number or a dash-separated pair of VLAN numbers.
- You can enter a comma-separated list of VLAN numbers and dash-separated pairs of VLAN numbers.

This example shows how to enable DHCP snooping on VLANs 10 through 12:

```
Router# configure terminal
Router(config)# ip dhcp snooping vlan 10 12
Router(config)#
```

This example shows another way to enable DHCP snooping on VLANs 10 through 12:

```
Router# configure terminal
Router(config)# ip dhcp snooping vlan 10-12
```

This example shows another way to enable DHCP snooping on VLANs 10 through 12:

```
Router# configure terminal
Router(config)# ip dhcp snooping vlan 10,11,12
```

This example shows how to enable DHCP snooping on VLANs 10 through 12 and VLAN 15:

```
Router# configure terminal
Router(config)# ip dhcp snooping vlan 10-12,15
```

This example shows how to verify the configuration:

```
Router(config)# do show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
10-12,15
DHCP snooping is operational on following VLANs:
none
DHCP snooping is configured on the following Interfaces:

Insertion of option 82 is enabled
Verification of hwaddr field is enabled
Interface                Trusted      Rate limit (pps)
-----
Router#
```

Configuring the DHCP Trust State on Layer 2 LAN Interfaces

To configure DHCP trust state on a Layer 2 LAN interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel number}	Selects the interface to configure. Note Select only LAN ports configured with the switchport command or Layer 2 port-channel interfaces.
Step 2	Router(config-if)# ip dhcp snooping trust	Configures the interface as trusted.
Step 3	Router(config-if)# do show ip dhcp snooping begin pps	Verifies the configuration.

This example shows how to configure Gigabit Ethernet port 5/12 as trusted:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/12
Router(config-if)# ip dhcp snooping trust
Router(config-if)# do show ip dhcp snooping | begin pps
Interface                Trusted      Rate limit (pps)
-----
GigabitEthernet5/12      yes        unlimited
Router#
```

This example shows how to configure Gigabit Ethernet port 5/12 as untrusted:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/12
Router(config-if)# no ip dhcp snooping trust
Router(config-if)# do show ip dhcp snooping | begin pps
Interface                Trusted      Rate limit (pps)
-----
GigabitEthernet5/12      no        unlimited
Router#
```

Configuring Spurious DHCP Server Detection

To detect and locate spurious DHCP servers, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping detect spurious vlan range	Enables detection of spurious DHCP servers on a specified VLAN range.
Step 2	Router(config)# ip dhcp snooping detect spurious interval time	Sets the interval time, the default is 30 minutes.
Step 3	Router# show ip dhcp snooping detect spurious	Verifies spurious DHCP server detection.

This example shows how to configure DHCP spurious server detection on VLANs 20 to 25 and set the interval to 50 minutes:

```
Router# configure terminal
Router(config)# ip dhcp snooping detect spurious vlan 20-25
Router(config)# ip dhcp snooping detect spurious interval 50
Router# do show ip dhcp snooping detect spurious
Spurious DHCP server detection is enabled.

Detection VLAN list : 20-25
Detection interval : 50 minutes
Router#
```

Configuring DHCP Snooping Rate Limiting on Layer 2 LAN Interfaces

To configure DHCP snooping rate limiting on a Layer 2 LAN interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel number}	Selects the interface to configure. Note Select only LAN ports configured with the switchport command or Layer 2 port-channel interfaces.
Step 2	Router(config-if)# ip dhcp snooping limit rate rate	Configures DHCP packet rate limiting.
Step 3	Router(config-if)# do show ip dhcp snooping begin pps	Verifies the configuration.

When configuring DHCP snooping rate limiting on a Layer 2 LAN interface, note the following information:

- We recommend an untrusted rate limit of not more than 100 packets per second (pps).
- If you configure rate limiting for trusted interfaces, you might need to increase the rate limit on trunk ports carrying more than one VLAN on which DHCP snooping is enabled.
- DHCP snooping puts ports where the rate limit is exceeded into the error-disabled state.

This example shows how to configure DHCP packet rate limiting to 100 pps on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/12
Router(config-if)# ip dhcp snooping limit rate 100
Router(config-if)# do show ip dhcp snooping | begin pps
Interface                Trusted      Rate limit (pps)
-----
GigabitEthernet5/12      no          100
Router#
```

The DHCP Snooping Database Agent

- [Prerequisites for the DHCP Snooping Database Agent, page 40-14](#)
- [Restrictions for the DHCP Snooping Database Agent, page 40-14](#)
- [Default Settings for the DHCP Snooping Database Agent, page 40-15](#)
- [How to Configure the DHCP Snooping Database Agent, page 40-15](#)
- [Configuration Examples for the Database Agent, page 40-15](#)

Prerequisites for the DHCP Snooping Database Agent

None.

Restrictions for the DHCP Snooping Database Agent

- The DHCP snooping database stores at least 8,000 bindings.
- Store the file on a TFTP server to avoid consuming storage space on the switch storage devices.

- When a switchover occurs, if the file is stored in a remote location accessible through TFTP, the newly active supervisor engine can use the binding list.
- Network-based URLs (such as TFTP and FTP) require that you create an empty file at the configured URL before the switch can write the set of bindings for the first time.

Default Settings for the DHCP Snooping Database Agent

None.

How to Configure the DHCP Snooping Database Agent

To configure the DHCP snooping database agent, perform one or more of the following tasks:

Command	Purpose
Router(config)# ip dhcp snooping database { _url write-delay seconds timeout seconds }	Configures a URL for the database agent (or file) and the related timeout values.
Router# show ip dhcp snooping database [detail]	Displays the current operating state of the database agent and statistics associated with the transfers.
Router# clear ip dhcp snooping database statistics	Clears the statistics associated with the database agent.
Router# renew ip dhcp snooping database [validation none] [url]	Requests the read entries from a file at the given URL.
Router# ip dhcp snooping binding mac_address vlan vlan_ID ip_address interface ifname expiry lease_in_seconds	Adds bindings to the snooping database.

Configuration Examples for the Database Agent

- [Example 1: Enabling the Database Agent, page 40-15](#)
- [Example 2: Reading Binding Entries from a TFTP File, page 40-17](#)
- [Example 3: Adding Information to the DHCP Snooping Database, page 40-18](#)

Example 1: Enabling the Database Agent

The following example shows how to configure the DHCP snooping database agent to store the bindings at a given location and to view the configuration and operating state:

```
Router# configure terminal
Router(config)# ip dhcp snooping database tftp://10.1.1.1/directory/file
Router(config)# end
Router# show ip dhcp snooping database detail
Agent URL : tftp://10.1.1.1/directory/file
Write delay Timer : 300 seconds
Abort Timer : 300 seconds

Agent Running : No
Delay Timer Expiry : 7 (00:00:07)
Abort Timer Expiry : Not Running

Last Succeeded Time : None
Last Failed Time : 17:14:25 UTC Sat Jul 7 2001
Last Failed Reason : Unable to access URL.
```

```

Total Attempts      :      21   Startup Failures :      0
Successful Transfers :      0   Failed Transfers  :     21
Successful Reads    :      0   Failed Reads    :      0
Successful Writes   :      0   Failed Writes   :     21
Media Failures      :      0

First successful access: Read

Last ignored bindings counters :
Binding Collisions   :      0   Expired leases   :      0
Invalid interfaces   :      0   Unsupported vlans :      0
Parse failures       :      0
Last Ignored Time : None

Total ignored bindings counters:
Binding Collisions   :      0   Expired leases   :      0
Invalid interfaces   :      0   Unsupported vlans :      0
Parse failures       :      0

Router#

```

The first three lines of output show the configured URL and related timer-configuration values. The next three lines show the operating state and the amount of time left for expiry of write delay and abort timers.

Among the statistics shown in the output, startup failures indicate the number of attempts to read or create the file that failed on bootup.



Note

Create a temporary file on the TFTP server with the **touch** command in the TFTP server daemon directory. With some UNIX implementations, the file should have full read and write access permissions (777).

DHCP snooping bindings are keyed on the MAC address and VLAN combination. If an entry in the remote file has an entry for a given MAC address and VLAN set for which the switch already has a binding, the entry from the remote file is ignored when the file is read. This condition is referred to as the *binding collision*.

An entry in a file may no longer be valid because the lease indicated by the entry may have expired by the time it is read. The expired leases counter indicates the number of bindings that are ignored because of this condition. The Invalid interfaces counter refers to the number of bindings that have been ignored when the interface referred by the entry either does not exist on the system or is a router or DHCP snooping trusted interface (if it exists) when the read happened. Unsupported VLANs refers to the number of entries that have been ignored because the indicated VLAN is not supported on the system. The Parse failures counter provides the number of entries that have been ignored when the switch is unable to interpret the meaning of the entries from the file.

The switch maintains two sets of counters for these ignored bindings. One provides the counters for a read that has at least one binding ignored by at least one of these conditions. These counters are shown as the “Last ignored bindings counters.” The total ignored bindings counters provides a sum of the number of bindings that have been ignored because of all the reads since the switch bootup. These two sets of counters are cleared by the **clear** command. The total counter set may indicate the number of bindings that have been ignored since the last clear.

Example 2: Reading Binding Entries from a TFTP File

To manually read the entries from a TFTP file, perform this task:

	Command	Purpose
Step 1	Router# show ip dhcp snooping database	Displays the DHCP snooping database agent statistics.
Step 2	Router# renew ip dhcp snoop data url	Directs the switch to read the file from the URL.
Step 3	Router# show ip dhcp snoop data	Displays the read status.
Step 4	Router# show ip dhcp snoop bind	Verifies whether the bindings were read successfully.

This is an example of how to manually read entries from the tftp://10.1.1.1/directory/file:

```

Router# show ip dhcp snooping database
Agent URL :
Write delay Timer : 300 seconds
Abort Timer : 300 seconds

Agent Running : No
Delay Timer Expiry : Not Running
Abort Timer Expiry : Not Running

Last Succeeded Time : None
Last Failed Time : None
Last Failed Reason : No failure recorded.

Total Attempts      :          0   Startup Failures :          0
Successful Transfers :          0   Failed Transfers :          0
Successful Reads     :          0   Failed Reads    :          0
Successful Writes    :          0   Failed Writes   :          0
Media Failures       :          0

Router# renew ip dhcp snoop data tftp://10.1.1.1/directory/file
Loading directory/file from 10.1.1.1 (via GigabitEthernet1/1): !
[OK - 457 bytes]
Database downloaded successfully.

Router#
00:01:29: %DHCP_SNOOPING-6-AGENT_OPERATION_SUCCEEDED: DHCP snooping database Read
succeeded.
Router# show ip dhcp snoop data
Agent URL :
Write delay Timer : 300 seconds
Abort Timer : 300 seconds

Agent Running : No
Delay Timer Expiry : Not Running
Abort Timer Expiry : Not Running

Last Succeeded Time : 15:24:34 UTC Sun Jul 8 2001
Last Failed Time : None
Last Failed Reason : No failure recorded.

Total Attempts      :          1   Startup Failures :          0
Successful Transfers :          1   Failed Transfers :          0
Successful Reads     :          1   Failed Reads    :          0
Successful Writes    :          0   Failed Writes   :          0
Media Failures       :          0

Router#
Router# show ip dhcp snoop bind
  
```

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface

```

-----
00:01:00:01:00:05  1.1.1.1      49810    dhcp-snooping  512  GigabitEthernet1/1
00:01:00:01:00:02  1.1.1.1      49810    dhcp-snooping  512  GigabitEthernet1/1
00:01:00:01:00:04  1.1.1.1      49810    dhcp-snooping  1536 GigabitEthernet1/1
00:01:00:01:00:03  1.1.1.1      49810    dhcp-snooping  1024 GigabitEthernet1/1
00:01:00:01:00:01  1.1.1.1      49810    dhcp-snooping  1    GigabitEthernet1/1
Router# clear ip dhcp snoop bind
Router# show ip dhcp snoop bind
-----
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
-----
Router#

```

Example 3: Adding Information to the DHCP Snooping Database

To manually add a binding to the DHCP snooping database, perform this task:

	Command	Purpose
Step 1	Router# show ip dhcp snooping binding	Views the DHCP snooping database.
Step 2	Router# ip dhcp snooping binding <i>binding_id</i> vlan <i>vlan_id</i> interface <i>interface</i> expiry <i>lease_time</i>	Adds the binding using the ip dhcp snooping exec command.
Step 3	Router# show ip dhcp snooping binding	Checks the DHCP snooping database.

This example shows how to manually add a binding to the DHCP snooping database:

```

Router# show ip dhcp snooping binding
-----
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
-----
Router#
Router# ip dhcp snooping binding 1.1.1 vlan 1 1.1.1.1 interface g11/1 expiry 1000

Router# show ip dhcp snooping binding
-----
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
-----
00:01:00:01:00:01  1.1.1.1      992        dhcp-snooping  1    GigabitEthernet1/1
Router#

```

Displaying the DHCP Snooping Binding Table

The DHCP snooping binding table for each switch contains binding entries that correspond to untrusted ports. The table does not contain information about hosts interconnected with a trusted port because each interconnected switch will have its own DHCP snooping binding table.

This example shows how to display the DHCP snooping binding information for a switch:

```

Router# show ip dhcp snooping binding
-----
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
-----
00:02:B3:3F:3B:99  55.5.5.2      6943       dhcp-snooping  10    GigabitEthernet6/10

```

Table 40-1 describes the fields in the **show ip dhcp snooping binding** command output.

Table 40-1 *show ip dhcp snooping binding Command Output*

Field	Description
MAC Address	Client hardware MAC address
IP Address	Client IP address assigned from the DHCP server
Lease (seconds)	IP address lease time
Type	Binding type: dynamic binding learned by DHCP snooping or statically-configured binding
VLAN	VLAN number of the client interface
Interface	Interface that connects to the DHCP client host

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

1

Configuring IGMP Proxy

- [Prerequisites for IGMP Proxy, page 1-1](#)
- [Information About IGMP Proxy, page 1-2](#)
- [How to Configure IGMP Proxy, page 1-4](#)
- [Configuration Examples for IGMP Proxy, page 1-6](#)
- [Feature Information for IGMP Proxy, page 1-7](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11845/prod_command_reference_list.html
 - Cisco IOS Release 15.0SY supports only Ethernet interfaces. Cisco IOS Release 15.0SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IGMP Proxy

Before you configure VPLS, ensure that the network is configured as follows:

- All devices on the IGMP UDL have the same subnet address. If all devices on the UDL cannot have the same subnet address, the upstream device must be configured with secondary addresses to match all of the subnets to which the downstream devices are attached.
- IP multicast is enabled and the PIM interfaces are configured.

**Note**

Use the following guidelines when configuring PIM interfaces for IGMP proxy:

- Use PIM sparse mode (PIM-SM) when the interface is operating in a sparse-mode region and you are running static RP, bootstrap (BSR), or Auto-RP with the Auto-RP listener capability.
- Use PIM sparse-dense mode when the interface is running in a sparse-dense mode region and you are running Auto-RP without the Auto-RP listener capability.

- Use PIM dense mode (PIM-DM) when the interface is operating in dense mode and is, thus, participating in a dense-mode region
- Use PIM-DM with the proxy-register capability when the interface is receiving source traffic from a dense-mode region that needs to reach receivers that are in a sparse-mode region

Information About IGMP Proxy

- [IGMP Proxy, page 1-2](#)

IGMP Proxy

An IGMP proxy enables hosts in a unidirectional link routing (UDLR) environment that are not directly connected to a downstream router to join a multicast group sourced from an upstream network.

The figure below illustrates a sample topology that shows two UDLR scenarios:


- Traditional UDL routing scenario--A UDL device with directly connected receivers.
- IGMP proxy scenario--UDL device without directly connected receivers.

**Note**

IGMP UDLs are needed on the upstream and downstream devices.

**Note**

Although the following illustration and example uses routers in the configuration, any device (router or switch) can be used.



Scenario 1--Traditional UDLR Scenario (UDL Device with Directly Connected Receivers)

For scenario 1, no IGMP proxy mechanism is needed. In this scenario, the following sequence of events occurs:

- User 2 sends an IGMP membership report requesting interest in group G.
- Router B receives the IGMP membership report, adds a forwarding entry for group G on LAN B, and proxies the IGMP report to Router A, which is the UDLR upstream device.
- The IGMP report is then proxied across the Internet link.
- Router A receives the IGMP proxy and maintains a forwarding entry on the unidirectional link.

Scenario 2--IGMP Proxy Scenario (UDL Device without Directly Connected Receivers)

For scenario 2, the IGMP proxy mechanism is needed to enable hosts that are not directly connected to a downstream device to join a multicast group sourced from an upstream network. In this scenario, the following sequence of events occurs:

- User 1 sends an IGMP membership report requesting interest in group G.
- Router C sends a PIM Join message hop-by-hop to the RP (Router B).
- Router B receives the PIM Join message and adds a forwarding entry for group G on LAN B.

- Router B periodically checks its mroute table and proxies the IGMP membership report to its upstream UDL device across the Internet link.
- Router A creates and maintains a forwarding entry on the unidirectional link (UDL).

In an enterprise network, it is desirable to be able to receive IP multicast traffic via satellite and forward the traffic throughout the network. With unidirectional link routing (UDLR) alone, scenario 2 would not be possible because receiving hosts must be directly connected to the downstream device, Router B. The IGM proxy mechanism overcomes this limitation by creating an IGMP report for (*, G) entries in the multicast forwarding table. To make this scenario functional, therefore, you must enable IGMP report forwarding of proxied (*, G) multicast static route (mroute) entries (using the `ip igmp mroute-proxy` command) and enable the mroute proxy service (using the `ip igmp proxy-service` command) on interfaces leading to PIM-enabled networks with potential members.

**Note**

Because PIM messages are not forwarded upstream, each downstream network and the upstream network have a separate domain.

How to Configure IGMP Proxy

- [Configuring the Upstream UDL Device for IGMP UDLR, page 1-4](#)
- [Configuring the Downstream UDL Device for IGMP UDLR with IGMP Proxy Support, page 1-5](#)

Configuring the Upstream UDL Device for IGMP UDLR

Perform this task to configure the upstream UDL device for IGMP UDLR.

	Command or Action	Purpose
Step 1	Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	Device# configure	Enters global configuration mode.
Step 3	Device (config)# interface <i>interface-type number</i>	Enters the interface configuration mode. • For the type and number arguments, specify the interface to be used as the UDL on the upstream device.
Step 4	Device (config-if)# ip igmp unidirectional-link	Configures IGMP on the interface to be unidirectional for IGMP UDLR
Step 5	Device (config-if)# end	Ends the current configuration session and returns to privileged EXEC mode.

Configuring the Downstream UDL Device for IGMP UDLR with IGMP Proxy Support

Perform this task to configure the downstream UDL device for IGMP UDLR with IGMP proxy support

	Command or Action	Purpose
Step 1	Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	Device# configure	Enters global configuration mode.
Step 3	Device (config)# interface <i>interface-type number</i>	Enters the interface configuration mode. For the type and number arguments, specify the interface to be used as the UDL on the downstream device for IGMP UDLR.
Step 4	Device (config-if)# ip igmp unidirectional-link	Configures IGMP on the interface to be unidirectional for IGMP UDLR
Step 5	Device (config-if)# exit	Exits interface configuration mode and returns to global configuration mode.
Step 6	Device (config)# interface <i>interface-type number</i>	Enters the interface configuration mode. For the type and number arguments, select an interface that is facing the nondirectly connected hosts.
Step 7	Device (config-if)# ip igmp mroute-proxy <i>type number</i>	Enables IGMP report forwarding of proxied (*, G) multicast static route (mroute) entries. - This step is performed to enable the forwarding of IGMP reports to a proxy service interface for all (*, G) forwarding entries in the multicast forwarding table. - In this example, the ip igmp mroute-proxy command is configured on Gigabit Ethernet interface 1/0/0 to request that IGMP reports be sent to loopback interface 0 for all groups in the mroute table that are forwarded to Gigabit Ethernet interface 1/0/0.
Step 8	Device (config-if) exit	Exits interface configuration mode and returns to global configuration mode.
Step 9	Device (config)# interface <i>interface-type number</i>	Enters interface configuration mode for the specified interface. In this example, loopback interface 0 is specified.

Step 10	Device (config-if)# ip igmp helper-address udl <i>interface-type interface-number</i>	Configures IGMP helpering for UDL. • This step allows the downstream device to helper IGMP reports received from hosts to an upstream device connected to a UDL associated with the interface specified for the interface-type and interface-number arguments. • In the example topology, IGMP helpering is configured over loopback interface 0 on the downstream device. Loopback interface 0, thus, is configured to helper IGMP reports from hosts to an upstream device connected to Gigabit Ethernet interface 0/0/0.
Step 11	Device (config-if)# ip igmp proxy-service	Enables the mroute proxy service. • When the mroute proxy service is enabled, the device periodically checks the static mroute table for (*, G) forwarding entries that match interfaces configured with the ip igmp mroute-proxy command (see Step 7) based on the IGMP query interval. Where there is a match, one IGMP report is created and received on this interface.  Note The ip igmp proxy-service command is intended to be used with the ip igmp helper-address (UDL) command. • In this example, the ip igmp proxy-service command is configured on loopback interface 0 to enable the forwarding of IGMP reports out the interface for all groups on interfaces registered through the ip igmp mroute-proxy command (see Step 7).
Step 12	Device (config-if)# end	Ends the current configuration session and returns to privileged EXEC mode.
Step 13	Device# show ip igmp interface	(Optional) Displays multicast-related information about an interface.
Step 14	Device# show ip igmp udlr	(Optional) Displays UDLR information for directly connected multicast groups on interfaces that have a UDL helper address configured.

Configuration Examples for IGMP Proxy

[Example: IGMP Proxy Configuration, page 1-7](#)

Example: IGMP Proxy Configuration

The following example shows how to configure the upstream UDL device for IGMP UDLR and the downstream UDL device for IGMP UDLR with IGMP proxy support.

Upstream Device Configuration

```
interface gigabitethernet 0/0/0
ip address 10.1.1.1 255.255.255.0
ip pim dense-mode
!
interface gigabitethernet 1/0/0
ip address 10.2.1.1 255.255.255.0
ip pim dense-mode
ip igmp unidirectional-link
!
interface gigabitethernet 2/0/0
ip address 10.3.1.1 255.255.255.0
```

Downstream Device Configuration

```
ip pim rp-address 10.5.1.1 5
access-list 5 permit 239.0.0.0 0.255.255.255
!
interface loopback 0
ip address 10.7.1.1 255.255.255.0
ip pim dense-mode
ip igmp helper-address udl ethernet 0
ip igmp proxy-service
!
interface gigabitethernet 0/0/0
ip address 10.2.1.2 255.255.255.0
ip pim dense-mode
ip igmp unidirectional-link
!
interface gigabitethernet 1/0/0
ip address 10.5.1.1 255.255.255.0
ip pim sparse-mode
ip igmp mroute-proxy loopback 0
!
interface gigabitethernet 2/0/0
ip address 10.6.1.1 255.255.255.0
```

Feature Information for IGMP Proxy

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Table 1-1 lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Table 1-1 Feature Information for Power over Ethernet

Feature Name	Releases	Feature Information
PIM Join to IGMP Proxy Support	Cisco IOS Release 15.5(1)SY3	This feature was introduced.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

41

IP Source Guard

- [Prerequisites for IP Source Guard, page 41-1](#)
- [Restrictions for IP Source Guard, page 41-1](#)
- [Information About IP Source Guard, page 41-2](#)
- [Default Settings for IP Source Guard, page 41-3](#)
- [How to Configure IP Source Guard, page 41-3](#)
- [Displaying IP Source Guard PACL Information, page 41-4](#)
- [Displaying IP Source Binding Information, page 41-6](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for IP Source Guard

None.

Restrictions for IP Source Guard

Because the IP source guard feature is supported only in hardware, IP source guard is not applied if there are insufficient hardware resources available. These hardware resources are shared by various other ACL features that are configured on the system. The following restrictions apply to IP source guard:

- Supported only on ingress Layer 2 ports.
- Supported only in hardware; not applied to any traffic that is processed in software.

- Does not support filtering of traffic based on MAC address.
- Is not supported on private VLANs.

Information About IP Source Guard

- [Overview of IP Source Guard, page 41-2](#)
- [IP Source Guard Interaction with VLAN-Based Features, page 41-2](#)
- [Channel Ports, page 41-2](#)
- [Layer 2 and Layer 3 Port Conversion, page 41-3](#)
- [IP Source Guard and Voice VLAN, page 41-3](#)
- [IP Source Guard and Web-Based Authentication, page 41-3](#)

Overview of IP Source Guard

IP source guard provides source IP address filtering on a Layer 2 port to prevent a malicious host from impersonating a legitimate host by assuming the legitimate host's IP address. The feature uses dynamic DHCP snooping and static IP source binding to match IP addresses to hosts on untrusted Layer 2 access ports.

Initially, all IP traffic on the protected port is blocked except for DHCP packets. After a client receives an IP address from the DHCP server, or after static IP source binding is configured by the administrator, all traffic with that IP source address is permitted from that client. Traffic from other hosts is denied. This filtering limits a host's ability to attack the network by claiming a neighbor host's IP address. IP source guard is a port-based feature that automatically creates an implicit port access control list (PACL).

IP Source Guard Interaction with VLAN-Based Features

Use the **access-group mode** command to specify how IP source guard interacts with VLAN-based features (such as VACL and Cisco IOS ACL and RACL).

In prefer port mode, if IP source guard is configured on an interface, IP source guard overrides other VLAN-based features. If IP source guard is not configured on the interface, other VLAN-based features are merged in the ingress direction and applied on the interface.

In merge mode, IP source guard and VLAN-based features are merged in the ingress direction and applied on the interface. This is the default access-group mode.

Channel Ports

IP source guard is supported on Layer 2 port-channel interfaces but not on the port members. When IP source guard is applied to a Layer 2 port-channel channel interface, it is applied to all the member ports in the EtherChannel.

Layer 2 and Layer 3 Port Conversion

When an IP source guard policy is configured on a Layer 2 port, if the port is reconfigured as a Layer 3 port, the IP source guard policy no longer functions but is still present in the configuration. If the port is reconfigured as a Layer 2 port, the IP source guard policy becomes effective again.

IP Source Guard and Voice VLAN

IP source guard is supported on a Layer 2 port that belongs to a voice VLAN. For IP source guard to be active on the voice VLAN, DHCP snooping must be enabled on the voice VLAN. In merge mode, the IP source guard feature is merged with VACL and Cisco IOS ACL configured on the access VLAN.

IP Source Guard and Web-Based Authentication

You can configure IP source guard and web-based authentication (see [Chapter 46, “Web-Based Authentication”](#)) on the same interface. Other VLAN-based features are not supported when IP Source Guard and web-based authentication are combined.

Default Settings for IP Source Guard

None.

How to Configure IP Source Guard

To enable IP source guard, perform this task:

	Command	Purpose
Step 1	Router(config)# ip dhcp snooping	Enables DHCP snooping globally.
Step 2	Router(config)# ip dhcp snooping vlan <i>number</i> [<i>number</i>]	Enables DHCP snooping on your VLANs.
Step 3	Router(config)# interface <i>interface-name</i>	Selects the interface to be configured.
Step 4	Router(config-if)# no ip dhcp snooping trust	Use the no keyword to configure the interface as untrusted.
Step 5	Router(config-if)# ip verify source vlan dhcp-snooping [port-security]	Enables IP source guard, source IP address filtering on the port. The following are the command parameters: vlan applies the feature to only specific VLANs on the interface. The dhcp-snooping option applies the feature to all VLANs on the interface that have DHCP snooping enabled. port-security enables MAC address filtering. This feature is currently not supported.
Step 6	Router(config-if)# exit	Returns to global configuration mode.

	Command	Purpose
Step 7	Router(config)# ip source binding <i>mac_address</i> vlan <i>vlan-id</i> ip-address interface <i>interface_name</i>	(Optional) Configures a static IP binding on the port.
Step 8	Router(config)# end	Exits configuration mode.
Step 9	Router# show ip verify source [interface <i>interface_name</i>]	Verifies the configuration.

**Note**

The static IP source binding can only be configured on a Layer 2 port. If you enter the **ip source binding vlan interface** command on a Layer 3 port, you receive this error message:

Static IP source binding can only be configured on switch port.

The **no** keyword deletes the corresponding IP source binding entry. This command requires an exact match of all the required parameters in order for the deletion to be successful.

This example shows how to enable per-Layer 2 port IP source guard on VLANs 10 through 20:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip dhcp snooping
Router(config)# ip dhcp snooping vlan 10 20
Router(config)# interface gigabitEthernet 6/1
Router(config-if)# switchport mode access
Router(config-if)# switchport access vlan 10
Router(config-if)# no ip dhcp snooping trust
Router(config-if)# ip verify source vlan dhcp-snooping
Router(config-if)# end
Router# show ip verify source interface gigabitEthernet 6/1
Interface  Filter-type  Filter-mode  IP-address      Mac-address      Vlan
-----
Gi6/1      ip           active       10.0.0.1        10
Gi6/1      ip           active       deny-all        11-20
Router#
```

The output shows that there is one valid DHCP binding to VLAN 10.

This example shows how to configure an interface to use prefer port mode:

```
Router# configure terminal
Router(config)# interface gigabitEthernet 6/1
Router(config-if)# access-group mode prefer port
```

This example shows how to configure an interface to use merge mode:

```
Router# configure terminal
Router(config)# interface gigabitEthernet 6/1
Router(config-if)# access-group mode merge
```

Displaying IP Source Guard PACL Information

To display IP source guard PACL information for all interfaces on a switch, perform this task:

Command	Purpose
Router# show ip verify source [interface interface-name]	Displays IP source guard PACL information for all interfaces on a switch or for a specified interface.

This example shows that DHCP snooping is enabled on VLAN 10 through 20, interface fa6/1 is configured for IP filtering, and there is an existing IP address binding 10.0.0.1 on VLAN 10:

```
Router# show ip verify source interface fa6/1
Interface  Filter-type  Filter-mode  IP-address  Mac-address  Vlan
-----
fa6/1      ip           active       10.0.0.1    -----
fa6/1      ip           active       deny-all   -----  10
fa6/1      ip           active       deny-all   -----  11-20
```

**Note**

The second entry shows that a default PACL (deny all IP traffic) is installed on the port for those snooping-enabled VLANs that do not have a valid IP source binding.

This example shows the displayed PACL information for a trusted port:

```
Interface  Filter-type  Filter-mode  IP-address  Mac-address  Vlan
-----
fa6/2      ip           inactive-trust-port
```

This example shows the displayed PACL information for a port in a VLAN not configured for DHCP snooping:

```
Interface  Filter-type  Filter-mode  IP-address  Mac-address  Vlan
-----
fa6/3      ip           inactive-no-snooping-vlan
```

This example shows the displayed PACL information for a port with multiple bindings configured for an IP/MAC filtering:

```
Interface  Filter-type  Filter-mode  IP-address  Mac-address  Vlan
-----
fa6/4      ip           active       10.0.0.2    aaaa.bbbb.cccc  10
fa6/4      ip           active       11.0.0.1    aaaa.bbbb.cccd  11
fa6/4      ip           active       deny-all    deny-all      12-20
```

This example shows the displayed PACL information for a port configured for IP/MAC filtering but not for port security:

```
Interface  Filter-type  Filter-mode  IP-address  Mac-address  Vlan
-----
fa6/5      ip           active       10.0.0.3    permit-all    10
fa6/5      ip           active       deny-all    permit-all    11-20
```

**Note**

The MAC address filter shows permit-all because port security is not enabled, so the MAC filter cannot apply to the port/VLAN and is effectively disabled. Always enable port security first.

This example shows an error message when you enter the **show ip verify source** command on a port that does not have an IP source filter mode configured:

```
Router# show ip verify source interface fa6/6
IP Source Guard is not configured on the interface fa6/6.
```

This example shows how to display all interfaces on the switch that have IP source guard enabled:

```
Router# show ip verify source
Interface  Filter-type  Filter-mode  IP-address      Mac-address      Vlan
-----
fa6/1      ip           active       10.0.0.1        -----
fa6/1      ip           active       deny-all       11-20
fa6/2      ip           inactive-trust-port
fa6/3      ip           inactive-no-snooping-vlan
fa6/4      ip           active       10.0.0.2        aaaa.bbbb.cccc  10
fa6/4      ip           active       11.0.0.1        aaaa.bbbb.cccd  11
fa6/4      ip           active       deny-all       deny-all       12-20
fa6/5      ip           active       10.0.0.3        permit-all     10
fa6/5      ip           active       deny-all       permit-all     11-20
```

Displaying IP Source Binding Information

To display all IP source bindings configured on all interfaces on a switch, perform this task:

Command	Purpose
Router# show ip source binding [<i>ip_address</i>] [<i>mac_address</i>] [dhcp-snooping static] [vlan <i>vlan_id</i>] [interface <i>interface_name</i>]	Displays IP source bindings using the optional specified display filters. The dhcp-snooping filter displays all VLANs on the interface that have DHCP snooping enabled.

This example shows how to display all IP source bindings configured on all interfaces on the switch.

```
Router# show ip source binding
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
-----
00:02:B3:3F:3B:99  55.5.5.2      6522       dhcp-snooping  10    GigabitEthernet6/10
00:00:00:0A:00:0B  11.0.0.1      infinite    static         10    GigabitEthernet6/10
Router#
```

[Table 41-1](#) describes the fields in the **show ip source binding** command output.

Table 41-1 *show ip source binding Command Output*

Field	Description
MAC Address	Client hardware MAC address
IP Address	Client IP address assigned from the DHCP server
Lease (seconds)	IP address lease time
Type	Binding type; static bindings configured from CLI to dynamic binding learned from DHCP snooping
VLAN	VLAN number of the client interface
Interface	Interface that connects to the DHCP client host

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

42

Dynamic ARP Inspection (DAI)

- [Prerequisites for DAI, page 42-1](#)
- [Restrictions for DAI, page 42-1](#)
- [Information About DAI, page 42-2](#)
- [Default Settings for DAI, page 42-6](#)
- [How to Configure DAI, page 42-7](#)
- [Configuration Examples for DAI, page 42-16](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
 - The PFC and any DFCs support DAI in hardware.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for DAI

None.

Restrictions for DAI

- Hardware-accelerated DAI is enabled by default.
- When DAI is hardware-accelerated, you can configure CoPP to rate limit ARP traffic that would be processed by the RP (for example, packets with a broadcast destination MAC address or the MAC address of the RP; see [Chapter 39, “Control Plane Policing \(CoPP\)”](#)).

- [DAI logging](#), including both ACL logging and DHCP logging, is not compatible with DAI hardware acceleration. When DAI is hardware-accelerated, DAI logging is disabled.

**Note**

Regardless of the enable state of DAI hardware acceleration, DAI configured to use an ARP ACL with the [acl-match matchlog](#) keywords is processed in software and supports logging.

- Because DAI is an ingress security feature, it does not perform any egress checking.
- DAI is not effective for hosts connected to switches that do not support DAI 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 DAI checks from the one with no checking. This action secures the ARP caches of hosts in the domain enabled for DAI.
- DAI 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 40, “Dynamic Host Configuration Protocol \(DHCP\) Snooping.”](#)
- When DHCP snooping is disabled or in non-DHCP environments, use ARP ACLs to permit or to deny packets.
- DAI is supported on access ports, trunk ports, EtherChannel ports, and private VLAN ports.
- A physical port can join an EtherChannel port channel only when the trust state of the physical port and the channel port match. 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 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 DAI-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.

Information About DAI

- [Information about ARP, page 42-3](#)
- [ARP Spoofing Attacks, page 42-3](#)
- [DAI and ARP Spoofing Attacks, page 42-4](#)
- [Interface Trust States and Network Security, page 42-4](#)

- [Rate Limiting of ARP Packets, page 42-5](#)
- [Relative Priority of ARP ACLs and DHCP Snooping Entries, page 42-6](#)
- [Logging of Dropped Packets, page 42-6](#)

Information about ARP


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.

ARP Spoofing Attacks

ARP spoofing attacks and ARP cache poisoning can occur because ARP allows a gratuitous reply from a host even if an ARP request was not received. After the attack, all traffic from the device under attack flows through the attacker's computer and then to the router, switch, or host.

An ARP spoofing attack can target 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 42-1](#) shows an example of ARP cache poisoning.

Figure 42-1 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 for 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, which is the topology of the classic *man-in-the middle* attack.

DAI and ARP Spoofing Attacks

The PFC and any DFCs provide hardware support for DAI. DAI is a security feature that validates ARP packets in a network. DAI intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings. This capability protects the network from some man-in-the-middle attacks.

DAI 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

DAI 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.

DAI can validate ARP packets against user-configured ARP access control lists (ACLs) for hosts with statically configured IP addresses (see [“Applying ARP ACLs for DAI Filtering” section on page 42-9](#)). The switch logs dropped packets (see the [“Logging of Dropped Packets” section on page 42-6](#)).

You can configure DAI 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 (see the [“Enabling Additional Validation” section on page 42-11](#)).

Interface Trust States and Network Security

DAI associates a trust state with each interface on the switch. Packets arriving on trusted interfaces bypass all DAI validation checks, and those arriving on untrusted interfaces undergo the DAI 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 42-2](#), assume that both Switch A and Switch B are running DAI 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 42-2 ARP Packet Validation on a VLAN Enabled for DAI

Configuring interfaces to be trusted when they are actually untrusted leaves a security hole in the network. If Switch A is not running DAI, 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 DAI.

DAI ensures that hosts (on untrusted interfaces) connected to a switch running DAI do not poison the ARP caches of other hosts in the network. However, DAI does not prevent hosts in other portions of the network from poisoning the caches of the hosts that are connected to a switch running DAI.

In cases in which some switches in a VLAN run DAI and other switches do not, configure the interfaces connecting such switches as untrusted. However, to validate the bindings of packets from switches where DAI is not configured, configure ARP ACLs on the switch running DAI. When you cannot determine such bindings, isolate switches running DAI at Layer 3 from switches not running DAI. For configuration information, see the [“One Switch Supports DAI”](#) section on page 42-21.

**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 performs DAI validation checks, which rate limits incoming ARP packets 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.

For configuration information, see the [“Configuring ARP Packet Rate Limiting”](#) section on page 42-10.

Relative Priority of ARP ACLs and DHCP Snooping Entries

DAI 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** 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 DAI Logging” section on page 42-13](#).

Default Settings for DAI

Feature	Default Setting
DAI	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 Layer 2-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.
Log buffer	When DAI is enabled, all denied or dropped ARP packets are logged. The number of entries in the log is 32. The number of system messages is limited to 5 per second. The logging-rate interval is 1 second.
Per-VLAN logging	All denied or dropped ARP packets are logged.

How to Configure DAI

- [Enabling DAI on VLANs, page 42-7](#)
- [Configuring DAI Hardware Acceleration, page 42-8](#)
- [Configuring the DAI Interface Trust State, page 42-8](#)
- [Applying ARP ACLs for DAI Filtering, page 42-9](#)
- [Configuring ARP Packet Rate Limiting, page 42-10](#)
- [Enabling DAI Error-Disabled Recovery, page 42-11](#)
- [Enabling Additional Validation, page 42-11](#)
- [Configuring DAI Logging, page 42-13](#)
- [Displaying DAI Information, page 42-15](#)

Enabling DAI on VLANs

To enable DAI on VLANs, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip arp inspection vlan {vlan_ID vlan_range}	Enables DAI on VLANs.
Step 3	Router(config-if)# do show ip arp inspection vlan {vlan_ID vlan_range} begin Vlan	Verifies the configuration.

You can enable DAI on a single VLAN or a range of VLANs:

- To enable a single VLAN, enter a single VLAN number.
- To enable a range of VLANs, enter a dash-separated pair of VLAN numbers.
- You can enter a comma-separated list of VLAN numbers and dash-separated pairs of VLAN numbers.

This example shows how to enable DAI on VLANs 10 through 12:

```
Router# configure terminal
Router(config)# ip arp inspection vlan 10-12
```

This example shows another way to enable DAI on VLANs 10 through 12:

```
Router# configure terminal
Router(config)# ip arp inspection vlan 10,11,12
```

This example shows how to enable DAI on VLANs 10 through 12 and VLAN 15:

```
Router# configure terminal
Router(config)# ip arp inspection vlan 10-12,15
```

This example shows how to verify the configuration:

```
Router(config)# do show ip arp inspection vlan 10-12,15 | begin Vlan
Vlan      Configuration      Operation      ACL Match      Static ACL
----      -
10        Enabled             Inactive
```

11	Enabled	Inactive
12	Enabled	Inactive
15	Enabled	Inactive
Vlan	ACL Logging	DHCP Logging
----	-----	-----
10	Deny	Deny
11	Deny	Deny
12	Deny	Deny
15	Deny	Deny

Configuring DAI Hardware Acceleration

When DAI is enabled, by default DAI hardware acceleration is also enabled. To configure the DAI hardware acceleration state, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip arp inspection accelerate	Enables DAI hardware acceleration.
	Router(config)# no ip arp inspection accelerate	Disables DAI hardware acceleration.
Step 3	Router(config)# do show ip arp inspection include Acceleration	Verifies the configuration.

This example shows how to reenable DAI hardware acceleration:

```
Router# configure terminal
Router(config)# ip arp inspection accelerate
Router(config)# do show ip arp inspection | include Acceleration
Hardware Acceleration Mode : Enabled
Router(config)#
```

Configuring the DAI Interface Trust State

The switch forwards ARP packets that it receives on a trusted interface, but does not check them.

On 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 DAI Logging” section on page 42-13](#).

To configure the DAI interface trust state, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface {type slot/port port-channel number}	Specifies the interface connected to another switch, and enter interface configuration mode.

	Command	Purpose
Step 3	Router(config-if)# ip arp inspection trust	Configures the connection between switches as trusted.
Step 4	Router(config-if)# do show ip arp inspection interfaces	Verifies the DAI configuration.

This example shows how to configure Gigabit Ethernet port 5/12 as trusted:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# ip arp inspection trust
Router(config-if)# do show ip arp inspection interfaces | include Int|--|5/12
Interface           Trust State      Rate (pps)      Burst Interval
-----
Gi5/12              Trusted          None            N/A
```

Applying ARP ACLs for DAI Filtering

To apply an ARP ACL, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router# ip arp inspection filter <i>arp_acl_name</i> vlan { <i>vlan_ID</i> <i>vlan_range</i> } [static]	Applies the ARP ACL to a VLAN.
Step 3	Router(config)# do show ip arp inspection vlan { <i>vlan_ID</i> <i>vlan_range</i> }	Verifies your entries.

- See the command reference for information about the **arp access-list** command.
- For *vlan_range*, you can specify a single VLAN or a range of VLANs:
 - To specify a single VLAN, enter a single VLAN number.
 - To specify a range of VLANs, enter a dash-separated pair of VLAN numbers.
 - You can enter a comma-separated list of VLAN numbers and dash-separated pairs of VLAN numbers.

- (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.

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.

- ARP packets containing only IP-to-MAC address bindings are compared against the ACL. Packets are permitted only if the access list permits them.

This example shows how to apply an ARP ACL named *example_arp_acl* to VLANs 10 through 12 and VLAN 15:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection filter example_arp_acl vlan 10-12,15
Router(config)# do show ip arp inspection vlan 10-12,15 | begin Vlan
Vlan      Configuration      Operation      ACL Match      Static ACL
-----
10        Enabled            Inactive      example_arp_acl No
```

11	Enabled	Inactive	example_arp_acl	No
12	Enabled	Inactive	example_arp_acl	No
15	Enabled	Inactive	example_arp_acl	No
Vlan	ACL Logging	DHCP Logging		
----	-----	-----		
10	Deny	Deny		
11	Deny	Deny		
12	Deny	Deny		
15	Deny	Deny		

Configuring ARP Packet Rate Limiting



Note

When DAI is hardware-accelerated, you can configure CoPP to rate limit ARP traffic that would be processed by the RP (for example, packets with a broadcast destination MAC address or the MAC address of the RP; see [Chapter 39, “Control Plane Policing \(CoPP\)”](#)).

When nonaccelerated DAI is enabled, the switch performs ARP packet validation checks, which makes the switch vulnerable to an ARP-packet denial-of-service attack. ARP packet rate limiting can prevent an ARP-packet denial-of-service attack.

To configure ARP packet rate limiting on a port, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface {type slot/port port-channel number}	Selects the interface to be configured.
Step 3	Router(config-if)# ip arp inspection limit {rate pps [burst interval seconds] none}	(Optional) Configures ARP packet rate limiting.
Step 4	Router(config-if)# do show ip arp inspection interfaces	Verifies the configuration.

- The default rate is 15 pps on untrusted interfaces and unlimited on trusted interfaces.
- For **rate pps**, specify an upper limit for the number of incoming packets processed per second. The range is 0 to 2048 pps.
- The **rate none** keywords specify that there is no upper limit for the rate of incoming ARP packets that can be processed.
- (Optional) For **burst interval seconds** (default is 1), 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.
- 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 the error-disabled state until you enable error-disabled recovery, which allows the port to emerge from the error-disabled state after a specified timeout period.
- Unless you configure a rate-limiting value on an interface, changing the trust state of the interface also changes its rate-limiting value to the default value for the configured trust state. After you configure the rate-limiting value, the interface retains the rate-limiting value even when you change its trust state. If you enter the **no ip arp inspection limit** interface configuration command, the interface reverts to its default rate-limiting value.

- For configuration guidelines about limiting the rate of incoming ARP packets on trunk ports and EtherChannel ports, see the [“Restrictions for DAI” section on page 42-1](#).

This example shows how to configure ARP packet rate limiting on Gigabit Ethernet port 5/14:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/14
Router(config-if)# ip arp inspection limit rate 20 burst interval 2
Router(config-if)# do show ip arp inspection interfaces | include Int|--|5/14
Interface          Trust State      Rate (pps)      Burst Interval
-----
Gi5/14             Untrusted              20              2
```

Enabling DAI Error-Disabled Recovery

To enable DAI error-disabled recovery, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# errdisable recovery cause arp-inspection	(Optional) Enables DAI error-disabled recovery.
Step 3	Router(config)# do show errdisable recovery include Reason --- arp-	Verifies the configuration.

This example shows how to enable DAI error disabled recovery:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# errdisable recovery cause arp-inspection
Router(config)# do show errdisable recovery | include Reason|---|arp-
ErrDisable Reason    Timer Status
-----
arp-inspection       Enabled
```

Enabling Additional Validation

DAI intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings. You can enable additional validation on the destination MAC address, the sender and target IP addresses, and the source MAC address.

To enable additional validation, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip arp inspection validate {[dst-mac] [ip] [src-mac]}	(Optional) Enables additional validation.
Step 3	Router(config)# do show ip arp inspection include abled\$	Verifies the configuration.

The additional validations do the following:

- **dst-mac**—Checks 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.
- **ip**—Checks 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.
- **src-mac**—Checks the source MAC address in the Ethernet header against the sender MAC address in the ARP body. This check is performed on both ARP requests and responses. When enabled, packets with different MAC addresses are classified as invalid and are dropped.

When enabling additional validation, note the following information:

- You must specify at least one of the keywords.
- Each **ip arp inspection validate** command overrides the configuration from any previous commands. If an **ip arp inspection validate** command enables **src-mac** and **dst-mac** validations, and a second **ip arp inspection validate** command enables IP validation only, the **src-mac** and **dst-mac** validations are disabled as a result of the second command.

This example shows how to enable **src-mac** additional validation:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection validate src-mac
Router(config)# do show ip arp inspection | include abled$
Source Mac Validation      : Enabled
Destination Mac Validation : Disabled
IP Address Validation      : Disabled
```

This example shows how to enable **dst-mac** additional validation:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection validate dst-mac
Router(config)# do show ip arp inspection | include abled$
Source Mac Validation      : Disabled
Destination Mac Validation : Enabled
IP Address Validation      : Disabled
```

This example shows how to enable **ip** additional validation:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection validate ip
Router(config)# do show ip arp inspection | include abled$
Source Mac Validation      : Disabled
Destination Mac Validation : Disabled
IP Address Validation      : Enabled
```

This example shows how to enable **src-mac** and **dst-mac** additional validation:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection validate src-mac dst-mac
Router(config)# do show ip arp inspection | include abled$
Source Mac Validation      : Enabled
Destination Mac Validation : Enabled
IP Address Validation      : Disabled
```


This example shows how to enable **src-mac**, **dst-mac**, and **ip** additional validation:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection validate src-mac dst-mac ip
Router(config)# do show ip arp inspection | include abled$
Source Mac Validation      : Enabled
Destination Mac Validation : Enabled
IP Address Validation      : Enabled
```

Configuring DAI Logging

- [DAI Logging Overview, page 42-13](#)
- [DAI Logging Restrictions, page 42-13](#)
- [Configuring the DAI Logging Buffer Size, page 42-13](#)
- [Configuring the DAI Logging System Messages, page 42-14](#)
- [Configuring DAI Log Filtering, page 42-15](#)

DAI Logging Overview

When DAI 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, DAI 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, DAI 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. Two dashes (“--”) appear instead of data except for 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.

DAI Logging Restrictions

DAI logging, including both ACL logging and DHCP logging, is not compatible with DAI hardware acceleration. When DAI is hardware-accelerated, DAI logging is disabled. Regardless of the enable state of DAI hardware acceleration, DAI configured to use an ARP ACL with the [acl-match matchlog](#) keywords is processed in software and supports logging.

Configuring the DAI Logging Buffer Size

To configure the DAI logging buffer size, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.

	Command	Purpose
Step 2	Router(config)# ip arp inspection log-buffer entries <i>number</i>	Configures the DAI logging buffer size (range is 0 to 1024).
Step 3	Router(config)# do show ip arp inspection log include Size	Verifies the configuration.

This example shows how to configure the DAI logging buffer for 64 messages:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection log-buffer entries 64
Router(config)# do show ip arp inspection log | include Size
Total Log Buffer Size : 64
```

Configuring the DAI Logging System Messages

To configure the DAI logging system messages, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip arp inspection log-buffer logs <i>number_of_messages interval length_in_seconds</i>	Configures the DAI logging buffer.
Step 3	Router(config)# do show ip arp inspection log	Verifies the configuration.

- For **logs** *number_of_messages* (default is 5), 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.
- For **interval** *length_in_seconds* (default is 1), 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). An interval setting of 0 overrides a log setting of 0.
- System messages are sent at the rate of *number_of_messages* per *length_in_seconds*.

This example shows how to configure DAI logging to send 12 messages every 2 seconds:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection log-buffer logs 12 interval 2
Router(config)# do show ip arp inspection log | include Syslog
Syslog rate : 12 entries per 2 seconds.
```

This example shows how to configure DAI logging to send 20 messages every 60 seconds.

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection log-buffer logs 20 interval 60
Router(config)# do show ip arp inspection log | include Syslog
Syslog rate : 20 entries per 60 seconds.
```

Configuring DAI Log Filtering

To configure DAI log filtering, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# ip arp inspection vlan <i>vlan_range</i> logging { acl-match { matchlog none } dhcp-bindings { all none permit }}	Configures log filtering for each VLAN.
Step 3	Router(config)# do show running-config include ip arp inspection vlan <i>vlan_range</i>	Verifies the configuration.

- By default, all denied packets are logged.
- For *vlan_range*, you can specify a single VLAN or a range of VLANs:
 - To specify a single VLAN, enter a single VLAN number.
 - To specify a range of VLANs, enter a dash-separated pair of VLAN numbers.
 - You can enter a comma-separated list of VLAN numbers and dash-separated pairs of VLAN numbers.
- **acl-match matchlog**—Logs packets based on the DAI ACL 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.
- **acl-match none**—Does not log packets that match ACLs.
- **dhcp-bindings all**—Logs all packets that match DHCP bindings.
- **dhcp-bindings none**—Does not log packets that match DHCP bindings.
- **dhcp-bindings permit**—Logs DHCP-binding permitted packets.

This example shows how to configure the DAI log filtering for VLAN 100 not to log packets that match ACLs:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip arp inspection vlan 100 logging acl-match none
Router(config)# do show running-config | include ip arp inspection vlan 100
ip arp inspection vlan 100 logging acl-match none
```

Displaying DAI 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 DAI for the specified VLAN. If no VLANs are specified or if a range is specified, displays information only for VLANs with DAI enabled (active).

Command	Description
<code>show ip arp inspection statistics [vlan <i>vlan_range</i>]</code>	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 DAI enabled (active). The switch increments the number of forwarded packets for each ARP request and response packet on a trusted DAI port. The switch increments the number of ACL-permitted 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.
<code>show ip arp inspection log</code>	Displays the configuration and contents of the DAI log buffer.

Configuration Examples for DAI

- [Two Switches Support DAI, page 42-16](#)
- [One Switch Supports DAI, page 42-21](#)

Two Switches Support DAI

- [Overview, page 42-16](#)
- [Configuring Switch A, page 42-17](#)
- [Configuring Switch B, page 42-18](#)

Overview

This procedure shows how to configure DAI 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 42-2 on page 42-5](#). Both switches are running DAI 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. Switch A has the bindings for Host 1 and Host 2, and Switch B has the binding for Host 2. Switch A Gigabit Ethernet port 6/3 is connected to the Switch B Gigabit Ethernet port 3/3.



Note

- DAI 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 40, “Dynamic Host Configuration Protocol \(DHCP\) Snooping.”](#)
- This configuration does not work if the DHCP server is moved from Switch A to a different location.
- To ensure that this configuration does not compromise security, configure Gigabit Ethernet port 6/3 on Switch A and Gigabit Ethernet port 3/3 on Switch B as trusted.

Configuring Switch A

To enable DAI and configure Gigabit Ethernet port 6/3 on Switch A as trusted, follow these steps:

Step 1 Verify the connection between switches Switch A and Switch B:

```
SwitchA# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
```

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
SwitchB	Fas 6/3	177	R S I	WS-C6506	Fas 3/3

```
SwitchA#
```

Step 2 Enable DAI on VLAN 1 and verify the configuration:

```
SwitchA# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
SwitchA(config)# ip arp inspection vlan 1
```

```
SwitchA(config)# end
```

```
SwitchA# show ip arp inspection vlan 1
```

```
Source Mac Validation      : Disabled
Destination Mac Validation : Disabled
IP Address Validation      : Disabled
```

Vlan	Configuration	Operation	ACL Match	Static ACL
1	Enabled	Active		

Vlan	ACL Logging	DHCP Logging
1	Deny	Deny

```
SwitchA#
```

Step 3 Configure Gigabit Ethernet port 6/3 as trusted:

```
SwitchA# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
SwitchA(config)# interface gigabitethernet 6/3
```

```
SwitchA(config-if)# ip arp inspection trust
```

```
SwitchA(config-if)# end
```

```
SwitchA# show ip arp inspection interfaces gigabitethernet 6/3
```

Interface	Trust State	Rate (pps)
Gi6/3	Trusted	None

```
SwitchA#
```

Step 4 Verify the bindings:

```
SwitchA# show ip dhcp snooping binding
```

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
00:02:00:02:00:02	1.1.1.2	4993	dhcp-snooping	1	GigabitEthernet6/4

```
SwitchA#
```

Step 5 Check the statistics before and after DAI processes any packets:

```
SwitchA# show ip arp inspection statistics vlan 1
```

Vlan	Forwarded	Dropped	DHCP Drops	ACL Drops
1	0	0	0	0

```

Vlan    DHCP Permits    ACL Permits    Source MAC Failures
----    -
1              0              0              0

Vlan    Dest MAC Failures    IP Validation Failures
----    -
1              0              0
SwitchA#

```

If Host 1 then sends out two ARP requests with an IP address of 1.1.1.2 and a MAC address of 0002.0002.0002, both requests are permitted, as reflected in the following statistics:

```
SwitchA# show ip arp inspection statistics vlan 1
```

```

Vlan    Forwarded    Dropped    DHCP Drops    ACL Drops
----    -
1              2              0              0              0

Vlan    DHCP Permits    ACL Permits    Source MAC Failures
----    -
1              2              0              0

Vlan    Dest MAC Failures    IP Validation Failures
----    -
1              0              0
SwitchA#

```

If Host 1 then tries to send an ARP request with an IP address of 1.1.1.3, the packet is dropped and an error message is logged:

```

00:12:08: %SW_DAI-4-DHCP_SNOOPING_DENY: 2 Invalid ARPs (Req) on Gi6/4, vlan
1. ([0002.0002.0002/1.1.1.3/0000.0000.0000/0.0.0.0/02:42:35 UTC Tue Jul 10 2001])
SwitchA# show ip arp inspection statistics vlan 1
SwitchA#

```

The statistics will display as follows:

```

Vlan    Forwarded    Dropped    DHCP Drops    ACL Drops
----    -
1              2              2              2              0

Vlan    DHCP Permits    ACL Permits    Source MAC Failures
----    -
1              2              0              0

Vlan    Dest MAC Failures    IP Validation Failures
----    -
1              0              0
SwitchA#

```

Configuring Switch B

To enable DAI and configure Gigabit Ethernet port 3/3 on Switch B as trusted, follow these steps:

Step 1 Verify the connectivity:

```
SwitchA# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
```

```

Device ID           Local Intrfce    Holdtme    Capability   Platform   Port ID
SwitchB             Fas 3/3         120        R S I       WS-C6506   Fas 6/3
SwitchB#

```

Step 2 Enable DAI on VLAN 1, and verify the configuration:

```

SwitchB# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SwitchB(config)# ip arp inspection vlan 1
SwitchB(config)# end
SwitchB# show ip arp inspection vlan 1

```

```

Source Mac Validation      : Disabled
Destination Mac Validation : Disabled
IP Address Validation      : Disabled

```

Vlan	Configuration	Operation	ACL Match	Static ACL
1	Enabled	Active		

Vlan	ACL Logging	DHCP Logging
1	Deny	Deny

```
SwitchB#
```

Step 3 Configure Gigabit Ethernet port 3/3 as trusted:

```

SwitchB# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SwitchB(config)# interface gigabitethernet 3/3
SwitchB(config-if)# ip arp inspection trust
SwitchB(config-if)# end
SwitchB# show ip arp inspection interfaces

```

Interface	Trust State	Rate (pps)
Gi1/1	Untrusted	15
Gi1/2	Untrusted	15
Gi3/1	Untrusted	15
Gi3/2	Untrusted	15
Gi3/3	Trusted	None
Gi3/4	Untrusted	15
Gi3/5	Untrusted	15
Gi3/6	Untrusted	15
Gi3/7	Untrusted	15

<output truncated>

```
SwitchB#
```

Step 4 Verify the list of DHCP snooping bindings:

```

SwitchB# show ip dhcp snooping binding

```

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
00:01:00:01:00:01	1.1.1.1	4995	dhcp-snooping	1	GigabitEthernet3/4

```

SwitchB#

```

Step 5 Check the statistics before and after DAI processes any packets:

```
SwitchB# show ip arp inspection statistics vlan 1
```

Vlan	Forwarded	Dropped	DHCP Drops	ACL Drops
1	0	0	0	0

```

      1              0              0              0              0
Vlan  DHCP Permits    ACL Permits    Source MAC Failures
----  -
1      0              0              0

Vlan  Dest MAC Failures    IP Validation Failures
----  -
1      0              0
SwitchB#

```

If Host 2 then sends out an ARP request with the IP address 1.1.1.1 and the MAC address 0001.0001.0001, the packet is forwarded and the statistics are updated appropriately:

```

SwitchB# show ip arp inspection statistics vlan 1

Vlan    Forwarded    Dropped    DHCP Drops    ACL Drops
----    -
1        1            0          0            0

Vlan  DHCP Permits    ACL Permits    Source MAC Failures
----  -
1      1            0              0

Vlan  Dest MAC Failures    IP Validation Failures
----  -
1      0              0
SwitchB#

```

If Host 2 attempts to send an ARP request with the IP address 1.1.1.2, DAI drops the request and logs a system message:

```

00:18:08: %SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Gi3/4, vlan
1.([0001.0001.0001/1.1.1.2/0000.0000.0000/0.0.0.0/01:53:21 UTC Fri May 23 2003])
SwitchB#

```

The statistics display as follows:

```

SwitchB# show ip arp inspection statistics vlan 1

Vlan    Forwarded    Dropped    DHCP Drops    ACL Drops
----    -
1        1            1          1            0

Vlan  DHCP Permits    ACL Permits    Source MAC Failures
----  -
1      1            0              0

Vlan  Dest MAC Failures    IP Validation Failures
----  -
1      0              0
SwitchB#

```


One Switch Supports DAI

This procedure shows how to configure DAI when Switch B shown in [Figure 42-2 on page 42-5](#) does not support DAI or DHCP snooping.

If switch Switch B does not support DAI or DHCP snooping, configuring Gigabit Ethernet port 6/3 on Switch A as trusted creates a security hole because both Switch A and Host 1 could be attacked by either Switch B or Host 2.

To prevent this possibility, you must configure Gigabit Ethernet port 6/3 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, which would make it 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.

To set up an ARP ACL on switch Switch A, follow these steps:

- Step 1** Configure the access list to permit the IP address 1.1.1.1 and the MAC address 0001.0001.0001, and verify the configuration:

```
SwitchA# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SwitchA(config)# arp access-list H2
SwitchA(config-arp-nacl)# permit ip host 1.1.1.1 mac host 1.1.1
SwitchA(config-arp-nacl)# end
SwitchA# show arp access-list
ARP access list H2
    permit ip host 1.1.1.1 mac host 0001.0001.0001
```

- Step 2** Apply the ACL to VLAN 1, and verify the configuration:

```
SwitchA# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SwitchA(config)# ip arp inspection filter H2 vlan 1
SwitchA(config)# end
SwitchA#
```

```
SwitchA# show ip arp inspection vlan 1
```

```
Source Mac Validation      : Disabled
Destination Mac Validation : Disabled
IP Address Validation      : Disabled
```

Vlan	Configuration	Operation	ACL Match	Static ACL
1	Enabled	Active	H2	No

Vlan	ACL Logging	DHCP Logging
1	Deny	Deny

```
SwitchA#
```

- Step 3** Configure Gigabit Ethernet port 6/3 as untrusted, and verify the configuration:

```
SwitchA# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SwitchA(config)# interface gigabitethernet 6/3
SwitchA(config-if)# no ip arp inspection trust
SwitchA(config-if)# end
Switch# show ip arp inspection interfaces gigabitethernet 6/3
```

Interface	Trust State	Rate (pps)
-----	-----	-----

```
Gi6/3          Untrusted          15
```

```
Switch#
```

When Host 2 sends 5 ARP requests through Gigabit Ethernet port 6/3 on Switch A and a “get” is permitted by Switch A, the statistics are updated appropriately:

```
Switch# show ip arp inspection statistics vlan 1
Vlan      Forwarded      Dropped      DHCP Drops      ACL Drops
-----
1          5                0            0                0
Vlan      DHCP Permits      ACL Permits      Source MAC Failures
-----
1          0                5                0
Vlan      Dest MAC Failures      IP Validation Failures
-----
1          0                0
Switch#
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

43

Traffic Storm Control

- [Prerequisites for Traffic Storm Control, page 43-1](#)
- [Restrictions for Traffic Storm Control, page 43-1](#)
- [Information About Traffic Storm Control, page 43-2](#)
- [Default Setting for Traffic Storm Control, page 43-3](#)
- [How to Enable Traffic Storm Control, page 43-4](#)
- [Displaying Traffic Storm Control Settings, page 43-5](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Traffic Storm Control

None.

Restrictions for Traffic Storm Control

- The following LAN switching modules do not support traffic storm control:
 - WS-X6148A-GE-45AF
 - WS-X6148A-GE-TX
- The switch supports multicast and unicast traffic storm control on WS-X6148A-RJ-45, Gigabit Ethernet, and 10-Gigabit Ethernet LAN ports.

- The switch supports broadcast traffic storm control on all LAN ports except on those modules previously noted.
- Except for BPDUs, traffic storm control does not differentiate between control traffic and data traffic.
- When multicast suppression is enabled, traffic storm control suppresses BPDUs when the multicast suppression threshold is exceeded on these modules:
 - WS-X6848-SFP-2T, WS-X6748-SFP
 - WS-X6824-SFP-2T, WS-X6724-SFP
 - WS-X6848-TX-2T, WS-X6748-GE-TX
 - WS-X6704-10GE

When multicast suppression is enabled on the listed modules, do not configure traffic storm control on STP-protected ports that need to receive BPDUs.

Except on the listed modules, traffic storm control does not suppress BPDUs.

Information About Traffic Storm Control

A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. The traffic storm control feature prevents LAN ports from being disrupted by a broadcast, multicast, or unicast traffic storm on physical interfaces.


Traffic storm control (also called traffic suppression) monitors incoming traffic levels over a 1-second traffic storm control interval and, during the interval, compares the traffic level with the traffic storm control level that you configure. The traffic storm control level is a percentage of the total available bandwidth of the port. Each port has a single traffic storm control level that is used for all types of traffic (broadcast, multicast, and unicast).

Traffic storm control monitors the level of each traffic type for which you enable traffic storm control in 1-second traffic storm control intervals. Within an interval, when the ingress traffic for which traffic storm control is enabled reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the traffic storm control interval ends.

By default, within an interval, when the ingress traffic for which traffic storm control is enabled reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the traffic storm control interval ends. These are the configurable traffic storm control optional actions:

- Shutdown—When a traffic storm occurs, traffic storm control puts the port into the error-disabled state. To reenab ports, use the error-disable detection and recovery feature or the **shutdown** and **no shutdown** commands.
- Trap—When a traffic storm occurs, traffic storm control generates an SNMP trap.

Figure 43-1 shows the broadcast traffic patterns on a LAN interface over a given interval. In this example, traffic storm control occurs between times T1 and T2 and between T4 and T5. During those intervals, the amount of broadcast traffic exceeded the configured threshold.

Figure 43-1 Broadcast Suppression

The traffic storm control threshold numbers and the time interval combination make the traffic storm control algorithm work with different levels of granularity. A higher threshold allows more packets to pass through.

Traffic storm control is implemented in hardware. The traffic storm control circuitry monitors packets passing from a LAN interface to the switching bus. Using the Individual/Group bit in the packet destination address, the traffic storm control circuitry determines if the packet is unicast or broadcast, keeps track of the current count of packets within the 1-second interval, and when a threshold is reached, filters out subsequent packets.

Because hardware traffic storm control uses a bandwidth-based method to measure traffic, the most significant implementation factor is setting the percentage of total available bandwidth that can be used by controlled traffic. Because packets do not arrive at uniform intervals, the 1-second interval during which controlled traffic activity is measured can affect the behavior of traffic storm control.

The following are examples of traffic storm control behavior:

- If you enable broadcast traffic storm control, and broadcast traffic exceeds the level within a 1-second traffic storm control interval, traffic storm control drops all broadcast traffic until the end of the traffic storm control interval.
- If you enable broadcast and multicast traffic storm control, and the combined broadcast and multicast traffic exceeds the level within a 1-second traffic storm control interval, traffic storm control drops all broadcast and multicast traffic until the end of the traffic storm control interval.
- If you enable broadcast and multicast traffic storm control, and broadcast traffic exceeds the level within a 1-second traffic storm control interval, traffic storm control drops all broadcast and multicast traffic until the end of the traffic storm control interval.
- If you enable broadcast and multicast traffic storm control, and multicast traffic exceeds the level within a 1-second traffic storm control interval, traffic storm control drops all broadcast and multicast traffic until the end of the traffic storm control interval.

Default Setting for Traffic Storm Control

None.

How to Enable Traffic Storm Control

To enable traffic storm control, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {{type slot/port} {port-channel number}}	Selects an interface to configure.
Step 2	Router(config-if)# storm-control broadcast level level[.level]	Enables broadcast traffic storm control on the interface, configures the traffic storm control level, and applies the traffic storm control level to all traffic storm control modes enabled on the interface.
Step 3	Router(config-if)# storm-control multicast level level[.level] Note The storm-control multicast command is supported only on Gigabit Ethernet interfaces.	Enables multicast traffic storm control on the interface, configures the traffic storm control level, and applies the traffic storm control level to all traffic storm control modes enabled on the interface.
Step 4	Router(config-if)# storm-control unicast level level[.level] Note The storm-control unicast command is supported only on Gigabit Ethernet interfaces.	Enables unicast traffic storm control on the interface, configures the traffic storm control level, and applies the traffic storm control level to all traffic storm control modes enabled on the interface.
Step 5	Router(config-if)# end	Exits configuration mode.

- You can configure traffic storm control on port channel interfaces.
- Do not configure traffic storm control on ports that are members of an EtherChannel. Configuring traffic storm control on ports that are configured as members of an EtherChannel puts the ports into a suspended state.
- Specify the level as a percentage of the total interface bandwidth:
 - The level can be from 0 to 100.
 - The optional fraction of a level can be from 0 to 99.
 - 100 percent means no traffic storm control.
 - 0.0 percent suppresses all traffic.
- On these modules, these levels suppress all traffic:
 - WS-X6704-10GE: 0.33 percent or less
 - WS-X6824-SFP-2T, WS-X6724-SFP 10Mbps ports: 0.33 percent or less
 - WS-X6848-SFP-2T, WS-X6748-SFP 100Mbps ports: 0.03 percent or less
 - WS-X6848-TX-2T, WS-X6748-GE-TX 100Mbps ports: 0.03 percent or less
 - WS-X6816-10G-2T, WS-X6716-10G
Oversubscription Mode: 0.29 percent or less

Because of hardware limitations and the method by which packets of different sizes are counted, the level percentage is an approximation. Depending on the sizes of the frames making up the incoming traffic, the actual enforced level might differ from the configured level by several percentage points.

This example shows how to enable multicast traffic storm control on Gigabit Ethernet interface 3/16 and how to configure the traffic storm control level at 70.5 percent:

```
Router# configure terminal
Router(config)# interface gigabitethernet 3/16
Router(config-if)# storm-control multicast level 70.5
Router(config-if)# end
```

This example shows how the traffic storm control level configured for one mode affects all other modes that are already configured on the Gigabit Ethernet interface 4/10:

```
Router# show run inter gig4/10
Building configuration...

Current configuration : 176 bytes
!
Router# interface GigabitEthernet4/10
Router# switchport
Router# switchport mode access
Router# storm-control broadcast level 70.00
Router# storm-control multicast level 70.00
Router# spanning-tree portfast edge
Router# end

Router# configure terminal
Router(config)# interface gigabitethernet 4/10
Router(config-if)# storm-control unicast level 20
Router(config-if)# end

Router# show interfaces gigabitethernet 4/10 counters storm-control

Port      UcastSupp %    McastSupp %    BcastSupp %    TotalSuppDiscards
Gi4/10    20.00          20.00          20.00          0

Router#
```

Displaying Traffic Storm Control Settings

Command	Purpose
Router# show interfaces [{type slot/port} {port-channel number}] switchport	Displays the administrative and operational status of all Layer 2 LAN ports or the specified Layer 2 LAN port.
Router# show interfaces [{type slot/port} {port-channel number}] counters storm-control	Displays the total number of packets discarded for all three traffic storm control modes, on all interfaces or on the specified interface.
Router# show interfaces counters storm-control [module slot_number]	



Note

The **show interfaces** [{interface_type slot/port} | {port-channel number}] **counters** command does not display the discard count. You must enter the **storm-control** keyword to display the discard count.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

44

Unknown Unicast and Multicast Flood Control

- [Prerequisites for Unknown Traffic Flood Control, page 44-1](#)
- [Restrictions for Unknown Traffic Flood Control, page 44-1](#)
- [Information About Unknown Traffic Flood Control, page 44-2](#)
- [Default Settings for Unknown Traffic Flood Control, page 44-2](#)
- [How to Configure Unknown Traffic Flood Control, page 44-2](#)
- [Configuration Examples for Unknown Traffic Flood Control, page 44-3](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Unknown Traffic Flood Control

None.

Restrictions for Unknown Traffic Flood Control

- Entering the **switchport block multicast** command on nonreceiver (router) ports of the VLAN could disrupt routing protocols. This command could also disrupt ARP functionality and other protocols, such as Network Time Protocol (NTP), that make use of local subnetwork multicast control groups in the 224.0.0.0/24 range.

- When unknown unicast flood rate-limiting (UUFRL) is enabled, per-VLAN learning must be enabled on all the Layer 3 routed ports, otherwise, any unicast flooded packet coming into a routed port will also be rate-limited by UUFRL.

Information About Unknown Traffic Flood Control

By default, unknown unicast and multicast traffic is flooded to all Layer 2 ports in a VLAN. You can use the unknown unicast flood blocking (UUFB), unknown multicast flood blocking (UMFB), and unknown unicast flood rate-limiting (UUFRL) features to prevent or limit this traffic.

The UUFB and UMFB features block unknown unicast and multicast traffic flooding at a specific port, only permitting egress traffic with MAC addresses that are known to exist on the port. The UUFB and UMFB features are supported on all ports that are configured with the **switchport** command, including private VLAN (PVLAN) ports.

The UUFRL feature globally rate limits unknown unicast traffic on all VLANs.

Default Settings for Unknown Traffic Flood Control

None.

How to Configure Unknown Traffic Flood Control

- [How to Configure UUFB or UMFB, page 44-2](#)
- [How to Configure UUFRL, page 44-3](#)

How to Configure UUFB or UMFB

To configure UUFB or UFMB, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface {{type slot/port} {port-channel number}}	Selects the interface to configure.
Step 3	Router(config-if)# switchport	Configures the port for Layer 2 switching.
Step 4	Router(config-if)# switchport block {unicast multicast}	Enables unknown unicast or multicast flood blocking on the port.

How to Configure UUFRL

To configure UUFRL, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# platform rate-limit layer2 unknown rate-in-pps [burst-size]	Enables UUFRL and sets the maximum packet rate. (Optional) Specify a burst size limit.
Step 3	Router(config)# exit	Exits configuration mode.

When you configure UUFRL, note the following information:

- For the *rate-in-pps* value:
 - The range is 10 through 1,000,000 (entered as 1000000).
 - There is no default value.
 - Values lower than 1,000 (entered as 1000) should offer sufficient protection.
- For the *burst-size* value:
 - The range is 1 through 255.
 - The default is 10.
 - The default value should provide sufficient protection.

Configuration Examples for Unknown Traffic Flood Control

This example shows how to configure UUFB on Gigabit Ethernet port 5/12 and how to verify the configuration:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport
Router(config-if)# switchport block unicast
Router(config-if)# do show interface gigabitethernet 5/12 switchport | include Unknown
Unknown unicast blocked: enabled
```

This example shows how to configure UUFRL with a rate limit of 1000 pps with a burst of 20 packets:

```
Router# configure terminal
Router(config)# platform rate-limit layer2 unknown 1000 20
Router(config)# exit
```



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

45

IEEE 802.1X Port-Based Authentication

- [Prerequisites for 802.1X Authentication, page 45-1](#)
- [Restrictions for 802.1X Authentication, page 45-2](#)
- [Information About 802.1X Port-Based Authentication, page 45-6](#)
- [Default Settings for 802.1X Port-Based Authentication, page 45-31](#)
- [How to Configure 802.1X Port-Based Authentication, page 45-32](#)
- [Displaying Authentication Status and Information, page 45-57](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for 802.1X Authentication

None.

Restrictions for 802.1X Authentication

- [802.1X Authentication, page 45-2](#)
- [VLAN Assignment, Guest VLAN, Restricted VLAN, and Inaccessible Authentication Bypass, page 45-3](#)
- [MAC Authentication Bypass, page 45-5](#)
- [Web-Based Authentication, page 45-5](#)
- [Network Edge Access Topology \(NEAT\) and Client Information Signalling Protocol \(CISP\), page 45-5](#)

802.1X Authentication

- When 802.1X authentication is enabled, ports are authenticated before any other Layer 2 or Layer 3 features are enabled.
- If you try to change the mode of an 802.1X-enabled port (for example, from access to trunk), an error message appears, and the port mode is not changed.
- Although not recommended, you can configure port security and 802.1X port-based authentication on the same port.



Note

802.1X authentication is not compatible with port security with sticky MAC addresses or static secure MAC addresses. With Release 15.1(1)SY1 and later, you cannot configure 802.1X authentication with port security with sticky MAC addresses or static secure MAC addresses.

- 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 reauthentication.
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 802.1X protocol is supported on Layer 2 static-access ports, voice VLAN ports, and Layer 3 routed ports, but it is not supported on these port types:
 - Trunk port—If you try to enable 802.1X authentication on a trunk 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 trunk, an error message appears, and the port mode is not changed.
 - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X 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.
- Because all traffic from unauthenticated hosts is forwarded to the switch processor, we recommend that you apply rate limiting to this traffic.

802.1X Host Mode

- In most cases when the host mode is changed on a port, any existing 802.1X authentications on that port are deleted. Exceptions are when changing from the single-host mode to any other mode, and when changing from multidomain mode to multiauth mode. In these two cases, existing 802.1X authentications are retained.
- If you enter the **authentication open** interface configuration command, any new MAC address detected on the port will be allowed unrestricted Layer 2 access to the network even before any authentication has succeeded. If you use this command, you should use static default ACLs to restrict Layer 3 traffic. For additional details, see the [“Pre-Authentication Open Access” section on page 45-15](#).
- When configuring multiple-hosts mode, if the multiple-hosts port becomes unauthorized (reauthentication fails or an EAPOL-logoff message is received), all attached clients are denied access to the network.
- When configuring MDA host mode, A third-party IP phone’s MAC address will initially be assigned to the data VLAN. When tagged voice packets are observed, the device will be removed from the data VLAN and placed on the voice VLAN.
- When configuring multiauth host mode, note the following guidelines:
 - If one client on a multiauth port becomes unauthorized (reauthentication fails or an EAPOL-logoff message is received from that client), the authorization status of the other attached clients is not changed.
 - RADIUS-assigned VLANs are not supported on multiauth ports, which can have only one data VLAN. If the authentication server sends VLAN-related attributes, the authentication will succeed but the VLAN assignment will be ignored.
 - Although multiple hosts are allowed on the data VLAN, only one host is allowed on the voice VLAN. When one IP phone has been authenticated, further IP phones on the same port will be denied authentication.
 - A multiauth port does not support a guest VLAN, authentication-fail VLAN, or a critical VLAN.

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, a private primary PVLAN, or a voice VLAN as an 802.1X guest VLAN. The guest VLAN feature is not supported on internal VLANs (routed ports) or trunk ports; it is supported only on access ports.
- 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 (**dot1x timeout quiet-period** and **dot1x timeout tx-period** interface configuration commands). The amount to decrease the settings depends on the connected 802.1X client type.
- When configuring the 802.1X VLAN user distribution feature, follow these guidelines:
 - A maximum of 100 VLAN groups can be configured, and a maximum of 4094 VLANs can be mapped to a VLAN group.
 - A VLAN can be mapped to more than one VLAN group.
 - A guest VLAN, a critical VLAN, or a restricted VLAN can be mapped to a VLAN group.
 - A VLAN group name cannot be specified as a guest VLAN, a critical VLAN, or a restricted VLAN.
 - You can modify a VLAN group by adding or removing a VLAN, but at least one VLAN must be mapped to the VLAN group. If you remove the last VLAN from the VLAN group, the VLAN group is deleted.
 - Removing an existing VLAN from the VLAN group name does not revoke the authentication status of the ports in the VLAN, but the mappings are removed from the existing VLAN group.
 - Deleting an existing VLAN group name does not revoke the authentication status of the ports in any VLAN within the group, but the VLAN mappings to the VLAN group are removed.
- When configuring the inaccessible authentication bypass feature, follow these guidelines:
 - The inaccessible authentication bypass feature is supported on 802.1X ports in single-host mode, multiple-hosts mode, and MDA 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 reinitiate the DHCP configuration process.
 - You can configure the inaccessible authentication bypass feature and the critical VLAN on an 802.1X port. If the switch tries to reauthenticate a critical port in a critical VLAN and all the RADIUS servers are unavailable, the switch changes the port state to the critical authentication state and the port remains in the critical VLAN.
 - You can configure the inaccessible bypass feature and port security on the same port.
- 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 internal VLANs (routed ports) or 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 45-2](#).
- If you disable MAC authentication bypass from a port after the port has been authorized with its MAC address, the session will be removed.
- When MAC authentication bypass with EAP has been enabled on an interface, it is not disabled by a subsequent **default interface** command on the interface.
- 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 reauthorize the port.
- If the port is in the authorized state, the port remains in this state until reauthorization occurs.
- To use MAC authentication bypass on a routed port, make sure that MAC address learning is enabled on the port.
- You can optionally configure a timeout period for hosts that are connected by MAC authentication bypass but are inactive. The range is 1 to 65535 seconds, but should be set to a value less than the reauthentication timeout. You must enable port security before configuring a timeout value. For more information, see the [“How to Configure Port Security” section on page 47-4](#).

Web-Based Authentication

- Fallback to web-based authentication is configured on switch ports in access mode. Ports in trunk mode are not supported.
- Fallback to web-based authentication is not supported on EtherChannels or EtherChannel members.
- Although fallback to web-based authentication is an interface-specific configuration, the web-based authentication fallback behavior is defined in a global fallback profile. If the global fallback configuration changes, the new profile will not be used until the next instance of authentication fallback.

For detailed information on configuring web-based authentication, see [Chapter 46, “Web-Based Authentication.”](#)

Network Edge Access Topology (NEAT) and Client Information Signalling Protocol (CISP)

- 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
- Starting Release 15.2(1)SY1 NEAT support is available on FEX ports.

Information About 802.1X Port-Based Authentication

- [802.1X Overview, page 45-6](#)
- [802.1X Device Roles, page 45-7](#)
- [Port-based Authentication Process, page 45-8](#)
- [Authentication Initiation and Message Exchange, page 45-10](#)
- [Ports in Authorized and Unauthorized States, page 45-12](#)
- [802.1X Host Modes, page 45-13](#)
- [802.1X Authentication with DHCP Snooping, page 45-15](#)
- [802.1X Accounting, page 45-16](#)
- [802.1X Authentication with VLAN Assignment, page 45-17](#)
- [Multiple VLANs and VLAN User Distribution with VLAN Assignment, page 45-18](#)
- [802.1X Authentication with Guest VLAN, page 45-19](#)
- [802.1X Authentication with Restricted VLAN, page 45-20](#)
- [802.1X Authentication with Inaccessible Authentication Bypass, page 45-21](#)
- [802.1X Authentication with Voice VLAN Ports, page 45-22](#)
- [802.1X Authentication with Port Security, page 45-23](#)
- [802.1X Authentication with ACL Assignments and Redirect URLs, page 45-23](#)
- [802.1X Authentication with Port Descriptors, page 45-26](#)
- [802.1X Authentication with MAC Authentication Bypass, page 45-26](#)
- [Network Admission Control Layer 2 IEEE 802.1X Validation, page 45-27](#)
- [802.1X Authentication with Wake-on-LAN, page 45-28](#)

802.1X Overview

This section describes the role of 802.1X port-based authentication as a part of a system of authentication, authorization, and accounting (AAA). The IEEE 802.1X standard defines a client and server-based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through publicly accessible ports. The authentication server authenticates each client connected to a switch port and assigns the port to a VLAN before making available any services offered by the switch or the LAN.

Until the client is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the client is connected. After authentication is successful, normal traffic can pass through the port.

802.1X Device Roles

With 802.1X port-based authentication, the devices in the network have specific roles as shown in Figure 45-1.

Figure 45-1 802.1X Device Roles

The specific roles shown in Figure 45-1 are as follows:

- *Client*—The device (workstation) that requests access to the LAN and switch services and responds to requests from the switch. The workstation must be running 802.1X-compliant client software such as that offered in the Microsoft Windows XP operating system. (The client is the *supplicant* in the IEEE 802.1X specification.)



Note To resolve Windows XP network connectivity and 802.1X port-based authentication issues, read the Microsoft Knowledge Base article at this URL:
<http://support.microsoft.com/kb/q303597/>

- *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. The Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server; it is available in Cisco Secure Access Control Server (ACS), version 3.0. RADIUS uses a client-server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.
- *Switch* (also called the *authenticator* and *back-end authenticator*)—Controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch includes the RADIUS client, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped and the remaining EAP frame is reencapsulated in the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

Port-based Authentication Process

When 802.1X port-based authentication is enabled, these events occur:

- If the client supports 802.1X-compliant client software and the client's identity is valid, the 802.1X authentication succeeds and the switch grants the client access to the network.
- If 802.1X authentication times out while waiting for an EAPOL message exchange, the switch can use a fallback authentication method, such as MAC authentication bypass (MAB) or web-based authentication (webauth), if either or both are enabled:
 - If MAC authentication bypass is enabled, the switch relays the client's MAC address to the AAA server for authorization. If the client's MAC address is valid, the authorization succeeds and the switch grants the client access to the network.
 - If web-based authentication is enabled, the switch sends an HTTP login page to the client. The switch relays the client's username and password to the AAA server for authorization. If the login succeeds, the switch grants the client access to the network.


**Note**

The default order for authentication methods is 802.1X, and then MAB, then web-based authentication. You can change the order, and you can disable any of these methods.

- If fallback authentication methods are not enabled or are not successful, and if a guest VLAN is configured, the switch assigns the client to a guest VLAN that provides limited services.
- If the switch receives 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 user-specified critical VLAN.

**Note**

Inaccessible authentication bypass is also referred to as critical authentication or the AAA fail policy.

Figure 45-2 Authentication Flowchart

1 = This occurs if the switch does not detect EAPOL packets from the client.

The switch reauthenticates a client when one of these situations occurs:

- Periodic reauthentication is enabled, and the reauthentication timer expires.

You can configure the reauthentication 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 reauthentication occurs.

The Termination-Action RADIUS attribute (Attribute [29]) specifies the action to take during reauthentication. 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 reauthentication. When the ReAuthenticate action is set (the attribute value is RADIUS-Request), the session is not affected during reauthentication.

- You manually reauthenticate the client by entering the **dot1x re-authenticate interface type slot/port** privileged EXEC command.

Authentication Initiation and Message Exchange

The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x pae authenticator** and **authentication port-control auto** interface configuration commands, the switch must initiate authentication when it determines that the port link state transitions from down to up. The switch then sends an EAP-request/identity frame to the client to request its identity (typically, the switch sends an initial identity/request frame followed by one or more requests for authentication information). When the client receives the frame, it responds with an EAP-response/identity frame.

If the client does not receive an EAP-request/identity frame from the switch during bootup, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.




Note

If 802.1X is not enabled or supported on the network access device, any EAPOL frames from the client are dropped. If the client does not receive an EAP-request/identity frame after three attempts to start authentication, the client transmits 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 45-12](#).


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 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 45-12](#).

The specific exchange of EAP frames depends on the authentication method being used. [Figure 45-3](#) shows a message exchange initiated by the client using the One-Time-Password (OTP) authentication method with a RADIUS server.

Figure 45-3 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 MAB 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 starts 802.1X authentication.

Figure 45-4 Message Exchange During MAC Authentication Bypass

Ports in Authorized and Unauthorized States

The switch port state determines whether or not the client is granted access to the network. The port starts in the *unauthorized* state. While in this state, the port disallows all ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authenticated, the port transitions to the *authorized* state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1X 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 protocol, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **authentication port-control** interface configuration command and these keywords:

- **force-authorized**—Disables 802.1X port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized**—Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto**—Enables 802.1X port-based authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the client's MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

802.1X Host Modes

- [Host Mode Overview, page 45-13](#)
- [Single-Host Mode, page 45-13](#)
- [Multiple-Hosts Mode, page 45-13](#)
- [Multidomain Authentication Mode, page 45-14](#)
- [Multi-Auth VLAN Assignment, page 45-14](#)
- [Multiauthentication Mode, page 45-15](#)
- [Pre-Authentication Open Access, page 45-15](#)

Host Mode Overview

The 802.1X port's host mode determines whether more than one client can be authenticated on the port and how authentication will be enforced. You can configure an 802.1X port to use any of the four host modes described in the following sections. In addition, each mode may be modified to allow pre-authentication open access.


Single-Host Mode

In single-host mode (see [Figure 45-1 on page 45-7](#)), only one client can be connected to the 802.1X-enabled 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.

Multiple-Hosts Mode

In multiple-hosts mode, you can attach multiple hosts to a single 802.1X-enabled port. [Figure 45-5](#) 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 (reauthentication 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.


With the multiple-hosts mode enabled, you can use 802.1X authentication to authenticate the port and you can use port security to manage network access for all MAC addresses, including the client's MAC address.

Figure 45-5 Multiple Host Mode Example

Multidomain Authentication Mode

Multidomain authentication (MDA) mode allows an IP phone (Cisco or third-party) and a single host behind the IP phone to authenticate independently, using 802.1X, MAC authentication bypass (MAB), or (for the host only) web-based authentication. In this application, multidomain refers to two domains — data and voice — and only two MAC addresses are allowed per port. The switch can place the host in the data VLAN and the IP phone in the voice VLAN, though they appear on the same switch port. The data VLAN assignment can be obtained from the vendor-specific attributes (VSAs) received from the authentication, authorization, and accounting (AAA) server during authentication.

Figure 45-6 shows a typical MDA application with a single host behind an IP phone connected to the 802.1X-enabled port. Because the client is not directly connected to the switch, the switch cannot detect a loss of port link if the client is disconnected. To prevent the possibility of another device using the established authentication of the disconnected client, later Cisco IP phones send a Cisco Discovery Protocol (CDP) host presence type length value (TLV) to notify the switch of changes in the attached client's port link state.

Figure 45-6 Multidomain Authentication Mode Example

Multi-Auth VLAN Assignment

Multi-auth VLAN assignment uses existing commands to support the assignment of a RADIUS server-supplied VLAN in multiauth mode when these conditions occur:

- 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.
- 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.
- The behavior of the critical-auth VLAN is not changed for multiauth mode. When a host tries to authenticate and the server is not reachable, all authorized hosts are reinitialized in the configured VLAN.

**Note**

- Only one voice VLAN is supported on a multiauth port.
- You cannot configure a guest VLAN or an auth-fail VLAN in multiauth mode.

Multiauthentication Mode

Multiauthentication (multiauth) mode allows one 802.1X/MAB client on the voice VLAN and multiple authenticated 802.1X/MAB/webauth clients on the data VLAN. When a hub or access point is connected to an 802.1X port, as in [Figure 45-5](#), multiauth mode provides enhanced security over multiple-hosts mode by requiring authentication of each connected client. For non-802.1X devices, MAB or web-based authentication can be used as the fallback method for individual host authentications, allowing different hosts to be authenticated through different methods on a single port.

Multiauth also supports MDA functionality on the voice VLAN by assigning authenticated devices to either a data or voice VLAN depending on the VSAs received from the authentication server.

Pre-Authentication Open Access

Any of the four host modes may be additionally configured to allow a device to gain network access before authentication. This pre-authentication open access is useful in an application such as the Pre-boot eXecution Environment (PXE), where a device must access the network to download a bootable image containing an authentication client.

Pre-authentication open access is enabled by entering the **authentication open** command after host mode configuration, and acts as an extension to the configured host mode. For example, if pre-authentication open access is enabled with single-host mode, then the port will allow only one MAC address. When pre-authentication open access is enabled, initial traffic on the port is restricted only by whatever other access restriction, independent of 802.1X, is configured on the port. If no access restriction other than 802.1X is configured on the port, then a client device will have full access on the configured VLAN.

802.1X Authentication with DHCP Snooping

When the Dynamic Host Configuration Protocol (DHCP) snooping option-82 with data insertion feature is enabled, the switch can insert a client's 802.1X authenticated user identity information into the DHCP discovery process, allowing the DHCP server to assign IP addresses from different IP address pools to different classes of end users. This feature allows you to secure the IP addresses given to the end users for accounting purposes and to allow services based on Layer 3 criteria.

After a successful 802.1X authentication between a supplicant and the RADIUS server, the switch puts the port in the forwarding state and stores the attributes that it receives from the RADIUS server. While performing DHCP snooping, the switch acts as a DHCP relay agent, receiving DHCP messages and regenerating those messages for transmission on another interface. When a client, after 802.1X authentication, sends a DHCP discovery message, the switch receives the packet. The switch adds to the packet a RADIUS attributes suboption section containing the stored RADIUS attributes of the client. The switch then submits the discovery broadcast again. The DHCP server receives the modified DHCP discovery packet and can, if configured to do so, use the authenticated user identity information when creating the IP address lease. The mapping of user-to-IP address can be on a one-to-one, one-to-many, or many-to-many basis. The one-to-many mapping allows the same user to authenticate through the 802.1X hosts on multiple ports.

The switch will automatically insert the authenticated user identity information when 802.1X authentication and DHCP snooping option-82 with data insertion features are enabled. To configure DHCP snooping option-82 with data insertion, see the [“DHCP Snooping Option-82 Data Insertion” section on page 40-5](#).

For information about the data inserted in the RADIUS attributes suboption, see RFC 4014, “Remote Authentication Dial-In User Service (RADIUS) Attributes Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Information Option.”

802.1X Accounting

The IEEE 802.1X standard defines how users are authorized and authenticated for network access but does not keep track of network usage. IEEE 802.1X accounting is disabled by default. You can enable 802.1X accounting to monitor the following activities on 802.1X-enabled ports:

- User successfully authenticates.
- User logs off.
- Link-down occurs.
- Reauthentication successfully occurs.
- Reauthentication fails.

The switch does not log IEEE 802.1X accounting information. Instead, it sends this information to the RADIUS server, which must be configured to log accounting messages.

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 45-1 lists the AV pairs and indicates when they are sent by the switch.

Table 45-1 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	Sometimes
			Note The Framed-IP-Address AV pair is sent only if a valid DHCP binding exists for the host in the DHCP snooping bindings table.	
Attribute[25]	Class	Always	Always	Always
Attribute[26]	Vendor-Specific	—	—	—
	Note Vendor-specific attributes (VSAs) are used by other 802.1X features.			
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	Never	Always
Attribute[43]	Acct-Output-Octets	Never	Never	Always
Attribute[44]	Acct-Session-ID	Always	Always	Always
Attribute[45]	Acct-Authentic	Always	Always	Always
Attribute[46]	Acct-Session-Time	Never	Never	Always
Attribute[49]	Acct-Terminate-Cause	Never	Never	Always
Attribute[61]	NAS-Port-Type	Always	Always	Always

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 AV pairs, see RFC 3580, “IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines.”

802.1X Authentication with VLAN Assignment

After successful 802.1X authentication of a port, the RADIUS server sends the VLAN assignment to configure the port. The RADIUS server database maintains the username-to-VLAN mappings, assigning the VLAN based on the username of the client connected to the port. You can use this feature to limit network access for certain users.

When configured on the switch and the RADIUS server, 802.1X authentication with VLAN assignment has these characteristics:

- If 802.1X authentication is enabled on a port, and if all information from the RADIUS server is valid, the port is placed in the RADIUS server-assigned VLAN after authentication.
- If the multiple-hosts mode is enabled on an 802.1X port, all hosts on the port are placed in the same RADIUS server-assigned VLAN as the first authenticated host.

- If the multiauth mode is enabled on an 802.1X port, the VLAN assignment will be ignored.
- If no VLAN number is supplied by the RADIUS server, the port is configured in its access VLAN after successful authentication. 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, the port returns to the unauthorized state and remains in the configured access VLAN. This prevents ports from appearing unexpectedly in an inappropriate VLAN because of a configuration error.
Configuration errors could include specifying a VLAN for a routed port, a malformed VLAN ID, a nonexistent or internal (routed port) VLAN ID, or an attempted assignment to a voice VLAN ID.
- If 802.1X authentication is disabled on the port, the port is returned to the configured access VLAN.
- The Radius VLAN assignment is supported for both DATA clients and VOICE clients.

When the port is in the force-authorized, force-unauthorized, unauthorized, or shutdown state, the port is put into the configured access VLAN.

If an 802.1X port is authenticated and put in the RADIUS server-assigned VLAN, any change to the port access VLAN configuration does not take effect.

The 802.1X 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, perform this task:

-
Step 1	Enable AAA authorization by using the network keyword to allow interface configuration from the RADIUS server.
Step 2	Enable 802.1X authentication.
Step 3	The VLAN assignment feature is automatically enabled when you configure 802.1X authentication on an access port.
Step 4	Assign vendor-specific tunnel attributes in the RADIUS server. The RADIUS server must return these attributes to the switch: • [64] Tunnel-Type = VLAN • [65] Tunnel-Medium-Type = 802 • [81] Tunnel-Private-Group-ID = VLAN name or VLAN ID

Attribute [64] must contain the value *VLAN* (type 13). Attribute [65] must contain the value *802* (type 6). Attribute [81] specifies the *VLAN name* or *VLAN ID* assigned to the 802.1X-authenticated user.

Multiple VLANs and VLAN User Distribution with VLAN Assignment

The RADIUS-supplied VLAN assignment can provide load balancing by distributing 802.1X-authenticated users among multiple VLANs.

In earlier releases, the RADIUS server can supply a single VLAN name or ID for the assignment of an authenticating user. The RADIUS server can supply multiple VLAN names and IDs or the name of a VLAN group that contains multiple VLANs. Use either of the following two methods to load balance the users between the different VLANs:

- Configure the RADIUS server to send more than one VLAN ID or VLAN name as part of the response to the authenticating user. The 802.1X VLAN user group feature tracks the users in a particular VLAN and achieves load balancing by placing newly authenticated users in the least populated VLAN of the RADIUS-supplied VLAN IDs.

Perform the steps shown in the [“802.1X Authentication with VLAN Assignment”](#) section on page 45-17 with the following exception:

Attribute [81] Tunnel-Private-Group-ID specifies multiple VLAN names or VLAN IDs

- Define a VLAN group that contains multiple VLANs. Configure the RADIUS server to supply the VLAN group name instead of a VLAN ID as part of the response to the authenticating user. If the supplied VLAN group name is found among the VLAN group names that you have defined, the newly authenticated user is placed in the least populated VLAN within the VLAN group.

Perform the steps shown in the [“802.1X Authentication with VLAN Assignment”](#) section on page 45-17 with the following exception:

Attribute [81] Tunnel-Private-Group-ID specifies a defined VLAN group name

For more information, see the [“Configuring VLAN User Distribution”](#) section on page 45-44.

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 non-802.1X-compliant clients, such as for downloading the 802.1X client software. 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 and no fallback authentication methods are enabled.

In addition, 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 will not change to the guest VLAN state. The EAPOL packet history is cleared if the interface link status goes down.

Use the **dot1x guest-vlan supplicant** global configuration command to allow an interface to change to the guest VLAN state regardless of the EAPOL packet history. That is, a host that is not 802.1X-capable will be assigned to the guest VLAN even if a previous host on that interface was 802.1X-capable.



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 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.

When operating as an 802.1X guest VLAN, a port functions in multiple-hosts mode regardless of the configured host mode of the port.

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

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 45-26 and the [“Configuring a Guest VLAN”](#) section on page 45-45.

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 to provide limited services to clients that failed authentication and 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 port remains in the spanning-tree blocking state. With this feature, you can configure the port to be in the restricted VLAN after a specified number of authentication attempts.

The authenticator counts the failed authentication attempts for the client. The failed attempt count increments when the RADIUS server replies with either an Access-Reject EAP failure or an empty response without an EAP packet. When this count exceeds the configured maximum number of authentication attempts, the port moves to the restricted VLAN, the failed attempt counter resets, and subsequent EAPOL-start messages from the failed client are ignored.

Users who fail authentication remain in the restricted VLAN until the next switch-initiated reauthentication attempt. A port in the restricted VLAN tries to reauthenticate at configured intervals (the default is 60 seconds). If reauthentication fails, the port remains in the restricted VLAN. If reauthentication is successful, the port moves either to the configured VLAN or to a VLAN sent by the RADIUS server. You can disable reauthentication. 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 reauthentication 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.

When operating as an 802.1X restricted VLAN, a port functions in single-host mode regardless of the configured host mode of the port. Only the client that failed authentication is allowed access on the port. An exception is that a port configured in MDA mode can still authenticate a voice supplicant from the restricted VLAN.

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 routed or trunk ports; it is supported only on access ports.

This feature works with port security. As soon as the port is authorized, a MAC address is provided to port security. If port security does not permit the MAC address or if the maximum secure address count is reached, the port becomes unauthorized and error disabled.

Other port 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 45-46](#).

802.1X Authentication with Inaccessible Authentication Bypass

When the switch cannot reach the configured RADIUS servers and hosts cannot be authenticated, you can configure the switch to allow network access to the hosts connected to critical ports. A critical port is enabled for the inaccessible authentication bypass feature, also referred to as critical authentication or the AAA fail policy.

When this feature is enabled, the switch checks the status of the configured RADIUS servers whenever the switch tries to authenticate a host connected to a critical port. 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.

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 user-specified critical 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 exchanges times out, and the switch puts the critical port in the critical-authentication state during the next authentication attempt.

When a RADIUS server that can authenticate the host is available, all critical ports in the critical-authentication state are automatically reauthenticated.

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 user-specified critical 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.

802.1X Authentication with Voice VLAN Ports

A Multi-VLAN Access Port (MVAP) is a port that belongs to two VLANs. A voice VLAN port is an MVAP that allows separating a port's voice traffic and data traffic on different VLANs. A voice VLAN port is associated with two VLAN identifiers:

- **Voice VLAN identifier (VVID)** to carry voice traffic to and from the IP phone. The VVID is used to configure the IP phone connected to the port.
- **Port VLAN identifier (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.

In order to recognize an IP phone, the switch will allow CDP traffic on a port regardless of the authorization state of the port. 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.



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 voice VLAN configuration information, see [Chapter 19, “Cisco IP Phone Support.”](#)

802.1X Authentication with Port Security

You can configure an 802.1X port with port security in either single-host or multiple-hosts mode. (You also must configure port security on the port by using the **switchport port-security** interface configuration command.) When you enable port security and 802.1X authentication on a port, 802.1X 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 an 802.1X port.

These are some examples of the interaction between 802.1X authentication and port security on the switch:

- When a client is authenticated, and the port security table is not full, the client MAC address is added to the port security list of secure hosts. The port then proceeds to come up normally.

When a client is authenticated and manually configured for port security, it is guaranteed an entry in the secure host table.

A security violation occurs if the client is authenticated, but the port security table is full. This can happen if the maximum number of secure hosts has been statically configured or if the client ages out of the secure host table. If the client address is aged, its place in the secure host table can be taken by another host.

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

The port security violation modes determine the action for security violations. For more information, see the [“Configuring the Port Security Violation Mode on a Port”](#) section on page 47-6.

- When you manually remove an 802.1X client address from the port security table by using the **no switchport port-security mac-address mac_address** interface configuration command, you should reauthenticate the 802.1X client by using the **dot1x re-authenticate interface type slot/port** privileged EXEC command.
- When an 802.1X client logs off, the port changes to an unauthenticated state, and all dynamic entries in the secure host table are cleared, including the entry for the client. Normal authentication then takes place.
- If the port is administratively shut down, the port becomes unauthenticated, and all dynamic entries are removed from the secure host table.
- Port security and a voice VLAN can be configured simultaneously on an 802.1X port that is in either single-host or multiple-hosts mode. Port security applies to both the voice VLAN identifier (VVID) and the port VLAN identifier (PVID).

For more information about enabling port security on your switch, see the [“How to Configure Port Security”](#) section on page 47-4.

802.1X Authentication with ACL Assignments and Redirect URLs

- [Overview, page 45-24](#)
- [Downloadable ACLs Using the Cisco Secure ACS, page 45-24](#)
- [Filter-ID ACLs, page 45-25](#)
- [Redirect URLs, page 45-25](#)
- [Static Sharing of ACLs, page 45-25](#)

Overview

Per-host policies such as ACLs and redirect URLs can be downloaded to the switch from the authentication server (AS) in a RADIUS Access-Accept packet at the end of an 802.1X, MAB, or web-based authentication exchange.

Per-host policies are activated during authentication as follows:

- Downloadable ACLs (DACLS) are defined in the Cisco Secure ACS and downloaded from the ACS to the switch using VSAs.
- Filter-ID ACLs are defined on the switch, and only the ACL name is downloaded from the AS to the switch using the RADIUS Filter-ID attribute.
- A redirection URL and an ACL name are downloaded from the ACS to the switch using VSAs. The redirection ACL is defined on the switch.

For information about configuring per-host policies, see the [“Configuring the Switch for DACLS or Redirect URLs” section on page 45-51](#).

Downloadable ACLs Using the Cisco Secure ACS

Following a successful host authentication, the Cisco Secure ACS can use a VSA to download an ACL to the switch. The switch combines the DACL with the default ACL on the port to which the host has connected. Because the DACL definition resides on the authentication server, this feature allows for centralized policy management.

Two methods are provided in the Cisco Secure ACS for configuring DACLS:

- Downloadable IP ACL

Downloading of the DACL is enabled by selecting Assign IP ACL in the ACS configuration, and the DACL is defined in the Downloadable IP ACL Content menu of the ACS. There is no restriction on the size of the DACL.

- Per-user ACL

The ACS can use the CiscoSecure-Defined-ACL [009\001 cisco-av-pair] VSAs to deliver the DACL. Because the entire DACL is delivered in a single RADIUS packet, the maximum size is limited by the 4096-byte maximum size for a RADIUS packet. The DACL must be defined on the ACS using the following format:

protocol:inacl#sequence_number=ace

as shown in this example:

```
ip:inacl#10=permit ip any 67.2.2.0 0.0.0.255
```

These guidelines apply when using DACLS:

- The source address for all ACEs must be defined as ANY.
- When the 802.1X host mode of the port is MDA or multiauth, the DACL will be modified to use the authenticated host's IP address as the source address. When the host mode is either single-host or multiple-host, the source address will be configured as ANY, and the downloaded ACLs or redirects will apply to all devices on the port.
- If no DACLS are provided during the authentication of a host, the static default ACL configured on the port will be applied to the host. On a voice VLAN port, only the static default ACL of the port will be applied to the phone.

Filter-ID ACLs

Following a successful host authentication, the authentication server can use the RADIUS Filter-ID attribute (Attribute[11]) rather than a VSA to deliver only the name of an extended ACL to the switch in the following format:

acl_name.in

The suffix “.in” indicates that the ACL should be applied in the inbound direction.

In this method, the ACL must be already defined on the switch. The switch matches the Filter-ID attribute value to a locally configured ACL that has the same name or number as the Filter-ID (for example, Filter-ID=101.in will match the extended numbered ACL 101, and Filter-ID= guest.in will match the extended named ACL “guest”). The specified ACL is then applied to the port. Because the ACL definition resides on the switch, this feature allows for local variation in a policy.

These guidelines apply when using Filter-ID ACLs:

- The guidelines for using DACLs also apply to Filter-ID ACLs.
- The Filter-ID attribute may be a number (100 to 199, or 2000 to 2699) or a name.

Redirect URLs

Following a successful host authentication, the Cisco Secure ACS can use a VSA to download information to the switch for intercepting and redirecting HTTP or HTTPS requests from the authenticated host. The ACS downloads a redirection ACL and URL. When an HTTP or HTTPS request from the host matches the downloaded ACL, the host’s web browser is redirected to the downloaded redirection URL.

The ACS uses these cisco-av-pair VSAs to configure the redirection:

- url-redirect-acl

This AV pair contains the name or number of an ACL that specifies the HTTP or HTTPS traffic to be redirected. The ACL must be defined on the switch, and the source address must be defined as ANY. Traffic that matches a permit entry in the redirect ACL will be redirected.

- url-redirect

This AV pair contains the HTTP or HTTPS URL to which the web browser will be redirected.

Static Sharing of ACLs

When a number of interfaces have the same PACL and VLAN-based features, the static sharing feature stores one copy of the PACL and inherited VLAN-based feature ACLs in the TCAM for all ports that use the same ACL set, freeing TCAM space for more ACLs. The switch automatically evaluates all configured or enabled interfaces for static sharing when any of these events occur:

- When an interface is configured.
- When a state change occurs on an interface.

Consider the following guidelines and restrictions:

- Static sharing is not supported for interfaces configured to support IPv6.
- Static sharing is supported only on switch ports in access mode with NAC or 802.1X DACL features configured.
- Static sharing is not supported on switch ports enabled with QoS, with the exception of VLAN-based QoS.

- When 802.1X is used with DACL, we recommend entering the **platform hardware acl dynamic setup static** command to avoid triggering a static sharing evaluation when the port is dynamically configured by the authentication server response. The static sharing evaluation may adversely affect the port/host linkup time.
- 802.1X interfaces with fallback authentication as active cannot form a static sharing group with interfaces on which fallback is not enabled or is not active.

802.1X Authentication with Port Descriptors

You can associate descriptive text with an 802.1X client's authentication information by configuring the Cisco vendor-specific attribute (VSA) **aaa:supplicant-name** on the RADIUS server. During a successful 802.1X authentication of the client on the port, the switch will receive the descriptive information from the RADIUS server as part of the Access-Accept packet and will display the information when the **show interface users** command is entered for the port. If the port is in a mode supporting multiple authenticated hosts, identity information for all the authenticated hosts will be displayed with the port description.

802.1X Authentication with MAC Authentication Bypass

You can configure the switch to authorize clients based on the client MAC address (see [Figure 45-4 on page 45-12](#)) 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 reauthentication occurs, the switch uses 802.1X authentication as the preferred reauthentication 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 reauthenticated. The reauthentication process is the same as that for clients that were authenticated with 802.1X. During reauthentication, the port remains in the previously assigned VLAN. If reauthentication is successful, the switch keeps the port in the same VLAN. If reauthentication fails, the switch assigns the port to the guest VLAN, if one is configured.

If reauthentication 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 reauthentication. If MAC authentication bypass is enabled

and the 802.1X authentication times out, the switch uses the MAC authentication bypass feature to initiate reauthorization. For more information about these AV pairs, see RFC 3580, “IEEE 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 45-23.
- Voice VLAN—See the “[802.1X Authentication with Voice VLAN Ports](#)” section on page 45-22.
- VLAN Membership Policy Server (VMPS)—802.1X and VMPS are mutually exclusive.
- Private VLAN—You can assign a client to a private VLAN.
- Network admission control (NAC) Layer 2 IP validation—This feature takes effect after an 802.1X port is authenticated with MAC authentication bypass, including hosts in the exception list.

Network Admission Control Layer 2 IEEE 802.1X Validation

Network Admission Control (NAC) Layer 2 IEEE 802.1X validation checks the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access. NAC Layer 2 IEEE 802.1X validation performs policy enforcement by assigning the authenticated port into a specified VLAN, which provides segmentation and quarantine of poorly postured hosts at Layer 2.

Configuring NAC Layer 2 IEEE 802.1X validation is similar to configuring 802.1X port-based authentication except that you must configure a posture token on the RADIUS server. You can view the NAC posture token, which shows the posture of the client, by using the **show dot1x** privileged EXEC command. For information about configuring NAC Layer 2 IEEE 802.1X validation, see the “[Configuring NAC Layer 2 IEEE 802.1X Validation](#)” section on page 45-50.

For more information about NAC, see the *Network Admission Control Software Configuration Guide*.

NAC Agentless Audit Support

MAB support is added for the Cisco NAC Audit Architecture, which uses an external audit server to check the antivirus posture of clients that do not run a Cisco Trust Agent (CTA) and cannot respond to NAC queries. To audit and report an agentless client’s antivirus posture, the NAC audit server must possess the client’s IP address and a unique session identifier for the client’s connection to the switch. To support the NAC audit architecture for agentless clients, the switch must snoop the client’s IP address, create and assign a unique session identifier for the agentless client, and pass this information to the RADIUS server for sharing with the NAC audit server.

Because MAB operates at Layer 2, the MAB authenticator does not normally know the IP address of the supplicant, and the supplicant might not have an IP address when it first contacts the authenticator. A supplicant that requires a DHCP-assigned IP address must be allowed access to a DHCP server before authentication. You must enable ARP and DHCP snooping on the switch to allow the MAB authenticator to learn the IP address of the supplicant. To allow the IP address and unique session identifier information to be shared with the NAC audit server, you must enable the sending of certain RADIUS attributes. See the “[Configuring NAC Agentless Audit Support](#)” section on page 45-51.

The client IP address and unique session identifier are shared in RADIUS Access-Requests and Access-Accepts using the following RADIUS *cisco-av-pair* vendor-specific attributes (VSAs):

- Cisco-AVPair="identity-request=*ip-address*"
ip-address is the client IP address obtained by the switch through ARP or DHCP snooping.
- Cisco-AVPair="audit-session-id=*audit session id string*"
audit session id string is a UTF-8 encoding of a unique 96-bit identifier derived by the switch from the network access server (NAS) IP address, a session count, and the session start timestamp.

802.1X Authentication with Wake-on-LAN

The 802.1X authentication with wake-on-LAN (WoL) feature allows dormant PCs to be powered up 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.

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.

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. Port security behavior remains the same when you configure MAC move.

**Note**

- MAC move is supported in all host modes. (The authenticated host can move to any port on the switch, no matter which host mode is enabled on the port.)
- MAC move is supported with port security.
- The MAC move feature applies to both voice and data hosts.
- 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 45-53](#).

MAC Replace

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**

- The Mac replace feature is not supported on ports in multiauth mode, because violations are not triggered in that mode.
- The Mac replace feature is not supported on 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 45-53](#).

802.1x Supplicant and Authenticator Switches with Network Edge Access Topology (NEAT)

NEAT 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. On the authenticator switch interface, multihost mode is not supported and in MDA mode voice client is not supported.

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 45-7.
- 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 45-7 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		

For more information, see the [“Configuring NEAT Authenticator and Supplicant Switches”](#) section on page 45-54.

Default Settings for 802.1X Port-Based Authentication

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 - IP address - UDP authentication port - Key	- None specified. 1812. - None specified.
Host mode	Single-host mode.
Control direction	Bidirectional control.
Periodic reauthentication	Disabled.
Number of seconds between reauthentication attempts	3600 seconds.
Reauthentication 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 retransmitting 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 retransmitting 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 retransmitting the response to the server).
Inactivity timeout	Disabled.
Guest VLAN	None specified.
Inaccessible authentication bypass	Disabled.
Restricted VLAN	None specified.

Feature	Default Setting
Authenticator (switch) mode	None specified.
MAC authentication bypass	Disabled. Note When MAC authentication bypass with EAP has been enabled on an interface, it is not disabled by a subsequent default interface command on the interface.

How to Configure 802.1X Port-Based Authentication

- [Enabling 802.1X Authentication, page 45-33](#)
- [Configuring Switch-to-RADIUS-Server Communication, page 45-34](#)
- [Configuring 802.1X Authenticator Host Mode, page 45-35](#)
- [Enabling Fallback Authentication, page 45-36](#)
- [Enabling Periodic Reauthentication, page 45-38](#)
- [Manually Reauthenticating the Client Connected to a Port, page 45-39](#)
- [Initializing Authentication for the Client Connected to a Port, page 45-39](#)
- [Removing 802.1X Client Information Globally, page 45-40](#)
- [Removing 802.1X Client Information from an Interface, page 45-40](#)
- [Clearing Authentication Sessions, page 45-40](#)
- [Changing 802.1X Timeouts, page 45-40](#)
- [Setting the Switch-to-Client Frame Retransmission Number, page 45-42](#)
- [Setting the Reauthentication Number, page 45-43](#)
- [Configuring IEEE 802.1X Accounting, page 45-43](#)
- [Configuring VLAN User Distribution, page 45-44](#)
- [Configuring a Guest VLAN, page 45-45](#)
- [Configuring a Restricted VLAN, page 45-46](#)
- [Configuring the Inaccessible Authentication Bypass Feature, page 45-47](#)
- [Configuring MAC Authentication Bypass, page 45-49](#)
- [Configuring NAC Layer 2 IEEE 802.1X Validation, page 45-50](#)
- [Configuring NAC Agentless Audit Support, page 45-51](#)
- [Configuring the Switch for DACLs or Redirect URLs, page 45-51](#)
- [Configuring 802.1X Authentication with WoL, page 45-52](#)
- [Enabling MAC Move, page 45-53](#)
- [Enabling MAC Replace, page 45-53](#)
- [Configuring NEAT Authenticator and Supplicant Switches, page 45-54](#)
- [Disabling 802.1X Authentication on the Port, page 45-56](#)
- [Resetting the 802.1X Configuration to the Default Values, page 45-57](#)

Enabling 802.1X Authentication

To enable 802.1X port-based authentication, you must enable AAA and specify the authentication method list.

A method list describes the sequence and authentication methods to be queried to authenticate a user. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle, the authentication process stops, and no other authentication methods are attempted.

To allow VLAN assignment, you must enable AAA authorization to configure the switch for all network-related service requests.

The 802.1X AAA process is as follows:

1. A user connects to a port on the switch.
2. Authentication is performed.
3. VLAN assignment is enabled, as appropriate, based on the RADIUS server configuration.
4. The switch sends a start message to an accounting server.
5. Reauthentication is performed, as necessary.
6. The switch sends an interim accounting update to the accounting server that is based on the result of reauthentication.
7. The user disconnects from the port.
8. The switch sends a stop message to the accounting server.

To configure 802.1X port-based authentication, perform this task:

	Command	Purpose
Step 1	Router(config)# aaa new-model	Enables AAA.
Step 2	Router(config)# aaa authentication dot1x { default } <i>method1</i> [<i>method2...</i>]	Creates an 802.1X port-based authentication method list. To create a default list that is used when a named list is <i>not</i> specified in the aaa 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. For <i>method1</i> , enter the group radius keywords to use the list of all RADIUS servers for authentication. Though other keywords are visible in the command-line help string, only the group radius keywords are supported.
Step 3	Router(config)# dot1x system-auth-control	Globally enables 802.1X port-based authentication.
Step 4	Router(config)# aaa authorization network { default } group radius	(Optional) Configures the switch to use user-RADIUS authorization for all network-related service requests such as VLAN assignment.
Step 5	Router(config)# radius-server host <i>ip-address</i>	Specifies the IP address of the RADIUS server.

	Command	Purpose
Step 6	Router(config)# radius-server key <i>string</i>	Specifies the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.
Step 7	Router(config)# interface <i>type slot/port</i>	Enters interface configuration mode and specifies the interface to be enabled for 802.1X authentication.
Step 8	Router(config-if)# switchport mode access	Sets the port to access mode only if you configured the RADIUS server in previous steps.
Step 9	Router(config-if)# authentication port-control auto	Enables port-based authentication on the interface. The no form of the command disables port-based authentication on the interface.
Step 10	Router(config-if)# dot1x pae authenticator	Enables 802.1X authentication on the interface.
Step 11	Router(config)# end	Returns to privileged EXEC mode.

This example shows how to enable AAA and 802.1X on Gigabit Ethernet port 5/1:

```
Router(config)# aaa new-model
Router(config)# aaa authentication dot1x default group radius
Router(config)# dot1x system-auth-control
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication port-control auto
Router(config-if)# dot1x pae authenticator
Router(config-if)# end
```

This example shows how to verify the configuration:

```
Router# show dot1x all

Sysauthcontrol           Enabled
Dot1x Protocol Version   2

Dot1x Info for GigabitEthernet1/7
-----
PAE                       = AUTHENTICATOR
PortControl               = AUTO
ControlDirection         = Both
HostMode                  = SINGLE_HOST
QuietPeriod               = 60
ServerTimeout             = 30
SuppTimeout               = 30
ReAuthMax                 = 2
MaxReq                    = 2
TxPeriod                  = 30
```

Configuring Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by any of the following:

- 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, 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 failover backup to the first one. The RADIUS host entries are tried in the order that they were configured.

To configure the RADIUS server parameters, perform this task:

	Command	Purpose
Step 1	Router(config)# ip radius source-interface <i>interface_name</i>	Specifies that the RADIUS packets have the IP address of the indicated interface.
Step 2	Router(config)# radius-server host { <i>hostname</i> <i>ip_address</i> }	Configures the RADIUS server host name or IP address on the switch. If you want to use multiple RADIUS servers, reenter this command.
Step 3	Router(config)# radius-server key <i>string</i>	Configures the authorization and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.

- For *hostname* or *ip_address*, specify the host name or IP address of the remote RADIUS server.
- 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**, 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.
- 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.



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.

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

```
Router(config)# ip radius source-interface Vlan80
Router(config)# radius-server host 172.120.39.46
Router(config)# radius-server key rad123
```

Configuring 802.1X Authenticator Host Mode

An 802.1X-enabled port can grant access to a single client or multiple clients as described in the [“802.1X Host Modes” section on page 45-13](#).

To configure the host mode of an 802.1X-authorized port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# authentication port-control auto	Enables port-based authentication on the interface.
Step 3	Router(config-if)# dot1x pae authenticator	Enables 802.1X authentication on the interface.
Step 4	Router(config-if)# authentication host-mode <i>host_mode</i>	Configures the host mode. The values for <i>host_mode</i> are: • single-host—Allows a single authenticated host (client) on an authorized port. • multi-host—Allows multiple clients on an authorized port when one client is authenticated. • multi-domain—Allows a single IP phone and a single data client to independently authenticate on an authorized port. • multi-auth—Allows a single IP phone and multiple data clients to independently authenticate on an authorized port.
Step 5	Router(config-if)# authentication open	(Optional) Enables pre-authentication open access.
Step 6	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to enable 802.1X on Gigabit Ethernet interface 5/1 and to allow multiple hosts:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication port-control auto
Router(config-if)# dot1x pae authenticator
Router(config-if)# authentication host-mode multi-host
```

Enabling Fallback Authentication

On a port in multiauth mode, either or both of MAB and web-based authentication can be configured as fallback authentication methods for non-802.1X hosts (those that do not respond to EAPOL). You can configure the order and priority of the authentication methods.

For detailed configuration information for MAB, see the [“Configuring MAC Authentication Bypass” section on page 45-49](#).

For detailed configuration information for web-based authentication, see [Chapter 46, “Web-Based Authentication.”](#)

To enable fallback authentication, perform this task:

	Command	Purpose
Step 1	Router(config)# ip admission name <i>rule-name</i> proxy http	Configures an authentication rule for web-based authentication.
Step 2	Router(config)# fallback profile <i>profile-name</i>	Creates a fallback profile for web-based authentication.

	Command	Purpose
Step 3	Router(config-fallback-profile)# ip access-group <i>rule-name</i> in	Specifies the default ACL to apply to network traffic before web-based authentication.
Step 4	Router(config-fallback-profile)# ip admission <i>name</i> <i>rule-name</i>	Associates an IP admission rule with the profile, and specifies that a client connecting by web-based authentication uses this rule.
Step 5	Router(config-fallback-profile)# exit	Returns to global configuration mode.
Step 6	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 7	Router(config-if)# authentication port-control auto	Enables authentication on the port.
Step 8	Router(config-if)# authentication order <i>method1</i> [<i>method2</i>] [<i>method3</i>]	(Optional) Specifies the fallback order of authentication methods to be used. The three values of <i>method</i> , in the default order, are dot1x , mab , and webauth . The specified order also determines the relative priority of the methods for reauthentication, from highest to lowest.
Step 9	Router(config-if)# authentication priority <i>method1</i> [<i>method2</i>] [<i>method3</i>]	(Optional) Overrides the relative priority of authentication methods to be used. The three values of <i>method</i> , in the default order of priority, are dot1x , mab , and webauth .
Step 10	Router(config-if)# authentication event fail action <i>next-method</i>	Specifies that the next configured authentication method will be used if authentication fails.
Step 11	Router(config-if)# mab [eap]	Enables MAC authentication bypass. The optional eap keyword specifies that the EAP extension is used during RADIUS authentication.
Step 12	Router(config-if)# authentication fallback <i>profile-name</i>	Enables web-based authentication using the specified profile.
Step 13	Router(config-if)# authentication violation [shutdown restrict]	(Optional) Configures the disposition of the port if a security violation occurs. The default action is to shut down the port. If the restrict keyword is configured, the port will not be shutdown, but trap entries will be installed for the violating MAC address, and traffic from that MAC address will be dropped.
Step 14	Router(config-if)# authentication timer inactivity { <i>seconds</i> server }	(Optional) Configures the inactivity timeout value for MAB and 802.1X. By default, inactivity aging is disabled for a port. <i>seconds</i>—Specifies inactivity timeout period. The range is from 1 to 65535 seconds. server—Specifies that the inactivity timeout period value will be obtained from the authentication server.
Step 15	Router(config-if)# authentication timer restart <i>seconds</i>	(Optional) Specifies a period after which the authentication process will restart in an attempt to authenticate an unauthorized port. <i>seconds</i>—Specifies the restart period. The range is from 1 to 65535 seconds.
Step 16	Router(config-if)# exit	Returns to global configuration mode.

	Command	Purpose
Step 17	Router(config)# ip device tracking	Enables the IP device tracking table, which is required for web-based authentication.
Step 18	Router(config)# exit	Returns to privileged EXEC mode.

This example shows how to enable 802.1X fallback to MAB, and then to enable web-based authentication, on an 802.1X-enabled port:

```
Router(config)# ip admission name rule1 proxy http
Router(config)# fallback profile fallback1
Router(config-fallback-profile)# ip access-group default-policy in
Router(config-fallback-profile)# ip admission rule1
Router(config-fallback-profile)# exit
Router(config)# interface gigabit1/1
Router(config-if)# switchport mode access
Router(config-if)# authentication port-control auto
Router(config-if)# dot1x pae authenticator
Router(config-if)# authentication order dot1x mab webauth
Router(config-if)# mab eap
Router(config-if)# authentication fallback fallback1
Router(config-if)# exit
Router(config)# ip device tracking
Router(config)# exit
```

Enabling Periodic Reauthentication

You can enable periodic 802.1X client reauthentication and specify how often it occurs. You can specify the reauthentication period manually or you can use the session-timeout period specified by the RADIUS server. If you enable reauthentication without specifying a time period, the number of seconds between reauthentication attempts is 3600.

To enable periodic reauthentication of the client and to configure the number of seconds between reauthentication attempts, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# authentication periodic	Enables periodic reauthentication of the client, which is disabled by default.
Step 3	Router(config-if)# authentication timer reauthenticate [<i>seconds</i> server]	Specifies the number of seconds between reauthentication attempts using these keywords: seconds—Sets the number of seconds from 1 to 65535; the default is 3600 seconds. server—Sets the number of seconds based on the value of the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute [29]). This command affects the operation of the switch only if periodic reauthentication is enabled.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to enable periodic reauthentication and set the number of seconds between reauthentication attempts to 4000:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication periodic
Router(config-if)# authentication timer reauthenticate 4000
```

Verifies your entries.

```
Router# show dot1x interface type slot/port
```

Manually Reauthenticating the Client Connected to a Port

To manually reauthenticate the client connected to a port, perform this task:

Command	Purpose
Router# dot1x re-authenticate interface <i>type slot/port</i>	Manually reauthenticates the client connected to a port.
	Note Reauthentication does not disturb the status of an already authorized port.

This example shows how to manually reauthenticate the client connected to Gigabit Ethernet port 5/1:

```
Router# dot1x re-authenticate interface gigabitethernet 5/1
```

Verifies your entries.

```
Router# show dot1x all
```

Initializing Authentication for the Client Connected to a Port

To initialize the authentication for the client connected to a port, perform this task:

Command	Purpose
Router# dot1x initialize interface <i>type slot/port</i>	Initializes the authentication for the client connected to a port.
	Note Initializing authentication disables any existing authentication before authenticating the client connected to the port.

This example shows how to initialize the authentication for the client connected to Gigabit Ethernet port 5/1:

```
Router# dot1x initialize interface gigabitethernet 5/1
```

Verifies your entries.

```
Router# show dot1x all
```

Removing 802.1X Client Information Globally

To completely delete all existing supplicants from all the interfaces on the switch, perform this task:

Command	Purpose
Router# clear dot1x all	Removes 802.1X client information for all clients connected to all ports.

This example shows how to remove 802.1X client information for all clients connected to all ports:

```
Router# clear dot1x all
```

Removing 802.1X Client Information from an Interface

To completely delete all existing supplicants from an interface or from all the interfaces on the switch, perform this task:

Command	Purpose
Router# clear dot1x interface <i>type slot/port</i>	Removes 802.1X client information for the client connected to the specified port.

This example shows how to remove 802.1X client information for the client connected to Gigabit Ethernet port 5/1:

```
Router# clear dot1x interface gigabitethernet 5/1
```

Clearing Authentication Sessions

To clear all or selected authentication sessions, perform this task:

Command	Purpose
Router# clear authentication sessions [handle <i>handle</i>] [interface <i>interface</i>] [mac <i>mac</i>] [method <i>method</i>]	Clears current authentication sessions. With no options specified, all current active sessions will be cleared. The keywords can be added and combined to clear specific sessions or subset of sessions.

This example shows how to clear all MAB authentication sessions connected to Gigabit Ethernet port 5/1:

```
Router# clear authentication sessions interface gigabitethernet 5/1 method mab
```

Changing 802.1X Timeouts

You can change several 802.1X timeout attributes using the **dot1x timeout {attribute} seconds** command form in the interface configuration mode. This section shows in detail how to change the quiet period timeout, followed by descriptions of how to change other 802.1X timeouts using the same command form.

Setting the Quiet Period

When the switch cannot authenticate the client, the switch remains idle for a set period of time and then tries again. The **dot1x timeout quiet-period** interface configuration command controls the idle period. A failed authentication of the client might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a number smaller than the default.

To change the quiet period, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# dot1x timeout quiet-period <i>seconds</i>	Sets the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client. The range is 0 to 65535 seconds; the default is 60.
Step 3	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to set the quiet period on the switch to 30 seconds:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# dot1x timeout quiet-period 30
```

Setting 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 retransmits the frame.



Note

You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific operational problems with certain clients and authentication servers.

To change the amount of time that the switch waits for a response to an EAP-request/identity frame from the client before retransmitting the request, use the **dot1x timeout tx-period** *seconds* command in the interface configuration mode. The range is 1 to 65535 seconds; the default is 30. To return to the default retransmission time, use the **no dot1x timeout tx-period** 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 retransmitting the request:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# dot1x timeout tx-period 60
```

Setting the Switch-to-Client Retransmission Time for EAP-Request Frames

The client notifies the switch that it received the EAP-request frame. If the switch does not receive this notification, the switch waits a set period of time, and then retransmits the frame.

To set the amount of time that the switch waits for notification, use the **dot1x timeout supp-timeout** *seconds* command in the interface configuration mode. The range is 1 to 65535 seconds; the default is 30. To return to the default retransmission time, use the **no dot1x supp-timeout** command.

This example shows how to set the switch-to-client retransmission time for the EAP-request frame to 25 seconds:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# dot1x timeout supp-timeout 25
```

Setting the Switch-to-Authentication-Server Retransmission Time for Layer 4 Packets

The authentication server notifies the switch each time it receives a Layer 4 packet. If the switch does not receive a notification after sending a packet, the switch waits a set period of time and then retransmits the packet.

To set the value for the retransmission of Layer 4 packets from the switch to the authentication server, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# dot1x timeout server-timeout <i>seconds</i>	Sets the value for the retransmission of Layer 4 packets from the switch to the authentication server. The range for <i>seconds</i> is 1 to 65535 seconds; the default is 30.
Step 3	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to set the switch-to-authentication-server retransmission time for Layer 4 packets to 25 seconds:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# dot1x timeout server-timeout 25
```

Setting the Switch-to-Client Frame Retransmission Number

In addition to changing the switch-to-client retransmission time, you can change the number of times that the switch sends an EAP-request/identity frame (assuming no response is received) to the client before restarting the authentication process.



Note

You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific operational problems with certain clients and authentication servers.

To set the switch-to-client frame retransmission number, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# dot1x max-req <i>count</i>	Sets 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 3	Router(config-if)# end	Returns to privileged EXEC mode.

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:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# dot1x max-req 5
```

Setting the Reauthentication 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 operational problems with certain clients and authentication servers.

To set the reauthentication number, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# dot1x max-reauth-req <i>count</i>	Sets 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 3	Router(config-if)# end	Returns to privileged EXEC mode.

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:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# dot1x max-reauth-req 4
```

Configuring IEEE 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 then can determine 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-3-NOACCOUNTINGRESPONSE: Accounting message Start for session
172.20.50.145 sam 11/06/03 07:01:16 11000002 failed to receive Accounting Response.
```

**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.

To configure 802.1X accounting after AAA is enabled on your switch, perform this task:

	Command	Purpose
Step 1	Router(config)# aaa accounting dot1x default start-stop group radius	Enables 802.1X accounting using the list of all RADIUS servers.
Step 2	Router(config)# 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 3	Router(config)# end	Returns to privileged EXEC mode.
Step 4	Router# show running-config	Verifies your entries.

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:

```
Router(config)# radius-server host 172.120.39.46 auth-port 1812 acct-port 1813 key rad123
Router(config)# aaa accounting dot1x default start-stop group radius
Router(config)# aaa accounting system default start-stop group radius
```

Configuring VLAN User Distribution

You can define a VLAN group that contains multiple VLANs. For VLAN load balancing, you can then configure the RADIUS server to supply a VLAN group name as part of the response to a user during 802.1X authentication. If the supplied VLAN group name is found among the VLAN group names that you have defined, the newly authenticated user is placed in the least populated VLAN within the VLAN group.

To configure a VLAN group, perform this task:

Command	Purpose
Router(config)# vlan group <i>group-name</i> vlan-list <i>vlan-list</i>	Creates a VLAN group or adds VLANs to an existing VLAN group. <i>group-name</i>—The name of the VLAN group. The name may contain up to 32 characters and must begin with a letter. vlan-list <i>vlan-list</i>—The VLANs that belong to the VLAN group. Group members can be specified as a single VLAN ID, a list of VLAN IDs, or a VLAN ID range. Multiple entries are separated by a hyphen (-) or a comma (,).

This example shows how to map VLANs 7 through 9 and 11 to a VLAN group:

```
Router(config)# vlan group ganymede vlan-list 7-9,11
```

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. When operating as a guest VLAN, a port functions in multiple-hosts mode regardless of the configured host mode of the port.

To configure a guest VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# switchport mode { access private-vlan host }	Sets the port to access mode or as a private VLAN host port. Routed and trunk ports do not support a guest VLAN.
Step 3	Router(config-if)# authentication port-control auto	Enables authentication on the port.
Step 4	Router(config-if)# authentication event no-response action authorize vlan <i>vlan-id</i>	Specifies an active VLAN as a guest VLAN. The range is 1 to 4094. You can configure any active VLAN except an internal VLAN (routed port), an RSPAN VLAN, a private primary PVLAN, or a voice VLAN as a guest VLAN.
Step 5	Router(config-if)# { dot1x pae authenticator mab }	Specifies whether the port authentication method is 802.1X or MAC address bypass.
Step 6	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to enable VLAN 2 as an 802.1X guest VLAN:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication port-control auto
Router(config-if)# authentication event no-response action authorize vlan 2
Router(config-if)# dot1x pae authenticator
```

This example shows how to set 3 seconds as the client notification timeout 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 resending the request, and to enable VLAN 2 as an 802.1X guest VLAN when an 802.1X port is connected to a DHCP client:

```
Router(config-if)# dot1x timeout supp-timeout 3
Router(config-if)# dot1x timeout tx-period 15
Router(config-if)# authentication event no-response action authorize vlan 2
Router(config-if)# dot1x pae authenticator
```

Configuring a Restricted VLAN

When you configure a restricted VLAN on 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. When operating as a restricted VLAN, a port functions in single-host mode regardless of the configured host mode of the port.

To configure a restricted VLAN, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# switchport mode { access private-vlan host }	Sets the port to access mode or as a private VLAN host port. Routed and trunk ports do not support a guest VLAN.
Step 3	Router(config-if)# authentication port-control auto	Enables port-based authentication on the port.
Step 4	Router(config-if)# authentication event fail [retry <i>retries</i>] action authorize vlan <i>vlan-id</i>	Specifies an active VLAN as a restricted VLAN. The range for <i>vlan-id</i> is 1 to 4094. (Optional) The retry keyword specifies a number of authentication attempts to allow before a port moves to the restricted VLAN.
Step 5	Router(config-if)# end	Returns to privileged EXEC mode.

To disable and remove the restricted VLAN, use the **no** form of the **authentication event fail** command or the **dot1x auth-fail** command. The port returns to the unauthorized state.

You can configure the maximum number of authentication attempts allowed before a user is assigned to the restricted VLAN. You can set the number of attempts by using the **retry** keyword in the **authentication event fail** [**retry** *retries*] **action authorize vlan** command. The range of *retries* (allowable authentication attempts) is 1 to 5. The default is 2 attempts.

This example shows how to enable VLAN 2 as a restricted VLAN, with assignment of a host after 3 failed attempts:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication port-control auto
Router(config-if)# authentication event fail retry 3 action authorize vlan 2
Router(config-if)# dot1x pae authenticator
```

Configuring the Inaccessible Authentication Bypass Feature

You can configure the inaccessible bypass feature, also referred to as critical authentication or the AAA fail policy.

To configure the port as a critical port and enable the inaccessible authentication bypass feature, perform this task:

	Command	Purpose
Step 1	Router(config)# radius-server dead-criteria <i>time</i> <i>tries</i> <i>tries</i>	(Optional) Sets the conditions that are used to decide when a RADIUS server is considered unavailable or <i>dead</i> . 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. 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.
Step 2	Router(config)# radius-server deadtime <i>minutes</i>	(Optional) Sets 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.

	Command	Purpose
Step 3	<pre>Router(config)# radius-server host ip-address [acct-port udp-port] [auth-port udp-port] [key string] [test username name [idle-time time] [ignore-acct-port] [ignore-auth-port]]</pre>	(Optional) Configures the RADIUS server parameters by using these keywords: • acct-port <i>udp-port</i>—Specifies 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>—Specifies the UDP port for the RADIUS authentication server. The range for the UDP port number is from 0 to 65536. The default is 1645. Note You should configure the UDP port for the RADIUS accounting server and the UDP port for the RADIUS authentication server to nondefault values. • key <i>string</i>—Specifies the authentication and encryption key for all RADIUS communication between the switch and the RADIUS daemon. Note You can also configure the authentication and encryption key by using the radius-server key {<i>0 string</i> <i>7 string</i> <i>string</i>} global configuration command. • test username <i>name</i>—Enables automated testing of the RADIUS server status, and specify the username to be used. • idle-time <i>time</i>—Sets 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—Disables testing on the RADIUS server accounting port. • ignore-auth-port—Disables testing on the RADIUS server authentication port.
Step 4	<pre>Router(config)# dot1x critical eapol</pre>	(Optional) Specifies that the switch sends an EAPOL-Success message when the switch successfully authenticates the critical port.
Step 5	<pre>Router(config-if)# authentication critical recovery delay milliseconds</pre>	(Optional) Sets the recovery delay period during which the switch waits to reinitialize 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 reinitialized every second).
Step 6	<pre>Router(config)# interface type slot/port</pre>	Specifies the port to be configured, and enters interface configuration mode.

	Command	Purpose
Step 7	Router(config-if)# authentication event server dead action authorize [vlan <i>vlan-id</i>]	Enables the inaccessible authentication bypass feature, authorizing ports on the specified VLAN when the AAA server is unreachable. If no VLAN is specified, the access VLAN will be used. Note The vlan keyword is only available on a switch port.
Step 8	Router(config-if)# authentication event server alive action reinitialize	Configures the inaccessible authentication bypass recovery feature, specifying that the recovery action is to authenticate the port when an authentication server becomes available.
Step 9	Router(config-if)# end	Returns to privileged EXEC mode.

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 dot1x critical eapol** global configuration command. To disable inaccessible authentication bypass, use the **no authentication event server dead action authorize** (or **no dot1x critical**) interface configuration command.

This example shows how to configure the inaccessible authentication bypass feature:

```
Router(config)# radius-server dead-criteria time 30 tries 20
Router(config)# radius-server deadtime 60
Router(config)# radius-server host 1.1.1.2 acct-port 1550 auth-port 1560 key abc1234 test
username user1 idle-time 30
Router(config)# dot1x critical eapol
Router(config)# authentication critical recovery delay 2000
Router(config)# interface gigabitethernet 0/1
Router(config-if)# authentication event server dead action authorize vlan 123
Router(config-if)# authentication event server alive action reinitialize
```

Configuring MAC Authentication Bypass

To configure MAC authentication bypass on an interface, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# authentication port-control { auto force-authorized force-unauthorized }	Enables 802.1X authentication on the port. The keywords have these meanings: • auto—Allows only EAPOL traffic until successful authentication. • force-authorized—Allows all traffic, requires no authentication. • force-unauthorized—Allows no traffic.

	Command	Purpose
Step 3	Router(config-if)# mab [eap]	Enables MAC authentication bypass on the interface. (Optional) Use the eap keyword to configure the switch to use EAP for authorization. - When MAC authentication bypass with EAP has been enabled on an interface, it is not disabled by a subsequent default interface command on the interface. - To use MAC authentication bypass on a routed port, make sure that MAC address learning is enabled on the port.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to enable MAC authentication bypass on a port:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication port-control auto
Router(config-if)# mab
```

Configuring NAC Layer 2 IEEE 802.1X Validation

You can configure NAC Layer 2 IEEE 802.1X validation, which is also referred to as 802.1X authentication with a RADIUS server. NAC Layer 2 IEEE 802.1X configuration is the same as 802.1X configuration with the additional step of configuring the RADIUS server with a posture token and VLAN assignment.

To configure NAC Layer 2 IEEE 802.1X validation, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# authentication port-control auto	Enables port-based authentication on the interface.
Step 3	Router(config-if)# authentication periodic	Enables periodic reauthentication of the client, which is disabled by default.
Step 4	Router(config-if)# authentication timer reauthenticate [<i>seconds</i> server]	Specifies the number of seconds between reauthentication attempts using these keywords: <i>seconds</i>—Sets the number of seconds from 1 to 65535; the default is 3600 seconds. server—Sets the number of seconds based on the value of the Session-Timeout RADIUS attribute (Attribute[27]) and the Termination-Action RADIUS attribute (Attribute [29]). This command affects the operation of the switch only if periodic reauthentication is enabled.
Step 5	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to configure NAC Layer 2 IEEE 802.1X validation:

```
Router(config)# interface gigabitethernet 5/1
Router(config)# authentication port-control auto
Router(config-if)# authentication periodic
Router(config-if)# authentication timer reauthenticate server
```

Configuring NAC Agentless Audit Support

To support the NAC audit architecture for agentless clients, the switch must snoop an authenticating 802.1X client's IP address, create and assign a unique session identifier for the agentless client, and pass this information to the RADIUS server for sharing with the NAC audit server. To allow the switch to obtain and share this information, you must enable ARP and DHCP snooping on the switch and you must enable the sending of certain RADIUS attributes.

To configure the RADIUS and tracking settings to support NAC agentless audit, perform this task:

	Command	Purpose
Step 1	Router(config)# radius-server attribute 8 include-in-access-req	Configures the switch to send the Framed-IP-Address RADIUS attribute (Attribute[8]) in access-request or accounting-request packets.
Step 2	Router(config)# radius-server vsa send authentication	Configures the network access server to recognize and use vendor-specific attributes (VSAs) (specifically audit-session-id) in RADIUS Access-Requests generated by the switch during the authentication phase.
Step 3	Router(config)# radius-server vsa send accounting	Allows VSAs to be included in subsequent RADIUS Accounting-Requests.
Step 4	Router(config)# ip device tracking	Enables the IP device tracking table.

Configuring the Switch for DACLs or Redirect URLs

To configure switch ports to accept DACLs or redirect URLs from the RADIUS server during authentication of an attached host, perform this task:

	Command	Purpose
Step 1	Router# config terminal	Enters global configuration mode.
Step 2	Router(config)# radius-server vsa send authentication	Configures the network access server to recognize and use vendor-specific attributes (VSAs) in RADIUS Access-Requests generated by the switch during the authentication phase. Note This step is necessary only with redirect URLs or when DACLs are downloaded using VSAs rather than the Filter-ID attribute.
Step 3	Router(config)# ip device tracking	Enables the IP device tracking table.

	Command	Purpose
Step 4	Router(config)# ip access-list extended <i>dac1-name</i>	Configures an ACL that will be referenced by the VSA or Filter-ID attribute. Note This step is not necessary for DACLs defined on the RADIUS server and downloaded using VSAs.
Step 5	Router(config-std-nacl)# { permit deny } ...	Defines the ACL. Note The source address must be ANY.
Step 6	Router(config-std-nacl)# exit	Returns to global configuration mode.
Step 7	Router(config)# ip access-list extended <i>acl-name</i>	Configures a default ACL for the ports.
Step 8	Router(config-std-nacl)# { permit deny } ...	Defines the ACL.
Step 9	Router(config-std-nacl)# exit	Returns to global configuration mode.
Step 10	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 11	Router(config-if)# exit	Returns to global configuration mode.

This example shows how to configure a switch for a downloadable policy:

```
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# radius-server vsa send authentication
Router(config)# ip device tracking
Router(config)# ip access-list extended my_dac1
Router(config-ext-nacl)# permit tcp any host 10.2.3.4
Router(config-ext-nacl)# exit
Router(config)# ip access-list extended default_acl
Router(config-ext-nacl)# permit ip any any
Router(config-ext-nacl)# exit
Router(config)# interface fastEthernet 2/13
Router(config-if)# ip access-group default_acl in
Router(config-if)# exit
```

Configuring 802.1X Authentication with WoL

To enable 802.1X authentication with wake-on-LAN (WoL), perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.

	Command	Purpose
Step 2	Router(config-if)# authentication control-direction {both in}	Enables 802.1X authentication with WoL on the port, and uses these keywords to configure the port as bidirectional or unidirectional. 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 3	Router(config-if)# end	Returns to privileged EXEC mode.

To disable 802.1X authentication with WoL, use the **no authentication control-direction** (or the **no dot1x control-direction**) interface configuration command.

This example shows how to enable 802.1X authentication with WoL and set the port as bidirectional:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# authentication control-direction both
```

Enabling MAC Move

To globally enable MAC move on the switch, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# authentication mac-move permit	Enables MAC move on the switch.
Step 3	Router(config)# end	Returns to privileged EXEC mode.
Step 4	Router# show running-config	(Optional) Verifies your entries.
Step 5	Router# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to globally enable MAC move on a switch:

```
Router(config)# authentication mac-move permit
```

Enabling MAC Replace

To enable MAC replace on an interface, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface interface-id	Specifies the port to be configured, and enter interface configuration mode.

	Command	Purpose
Step 3	Router(config-if)# authentication violation { protect replace restrict shutdown }	Uses 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: • 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	Router(config)# end	Returns to privileged EXEC mode.
Step 5	Router# show running-config	Verifies your entries.
Step 6	Router# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to enable MAC replace on an interface:

```
Router(config)# interface gigabitethernet2/2
Router(config-if)# authentication violation replace
```

Configuring NEAT Authenticator and Supplicant Switches

- [NEAT Authenticator Configuration, page 45-54](#)
- [NEAT Supplicant Configuration, page 45-55](#)



Note

- NEAT requires one switch to be configured as a supplicant and to be 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 45-30.
- The cisco-av-pairs value must be configured as “device-traffic-class=switch” on the ACS, which sets the interface as a trunk after the supplicant is successfully authenticated.

NEAT Authenticator Configuration

To configure a switch as an authenticator, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# cisp enable	Enables CISP.
Step 3	Router(config)# interface interface-id	Specifies the port to be configured, and enters interface configuration mode.

	Command	Purpose
Step 4	Router(config-if)# switchport mode access	Sets the port mode to access .
Step 5	Router(config-if)# authentication port-control auto	Sets the port-authentication mode to auto.
Step 6	Router(config-if)# dot1x pae authenticator	Configures the interface as a port access entity (PAE) authenticator.
Step 7	Router(config-if)# spanning-tree portfast	Enables PortFast on an access port connected to a single workstation or server.
Step 8	Router(config-if)# end	Returns to privileged EXEC mode.
Step 9	Router# show running-config interface interface-id	Verifies your configuration.
Step 10	Router# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to configure a switch as an 802.1x authenticator:

```
Router# configure terminal
Router(config)# cisp enable
Router(config)# interface gigabitethernet1/1
Router(config-if)# switchport mode access
Router(config-if)# authentication port-control auto
Router(config-if)# dot1x pae authenticator
Router(config-if)# spanning-tree portfast trunk
```

NEAT Supplicant Configuration

To configure a switch as a supplicant, perform this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# cisp enable	Enables CISP.
Step 3	Router(config)# dot1x credentials profile	Creates 802.1x credentials profile. This must be attached to the port that is configured as supplicant.
Step 4	Router(config)# username suppswitch	Creates a username.
Step 5	Router(config)# password password	Creates a password for the new username.
Step 6	Router(config)# dot1x supplicant force-multicast	Forces the switch to send <i>only</i> multicast EAPOL packets when it receives either unicast or multicast packets, which allows NEAT to work on the supplicant switch in all host modes.
Step 7	Router(config)# interface interface-id	Specifies the port to be configured, and enters interface configuration mode.
Step 8	Router(config-if)# switchport trunk encapsulation dot1q	Sets the port to trunk mode.
Step 9	Router(config-if)# switchport mode trunk	Configures the interface as a VLAN trunk port.
Step 10	Router(config-if)# dot1x pae supplicant	Configures the interface as a port access entity (PAE) supplicant.
Step 11	Router(config-if)# dot1x credentials profile-name	Attaches the 802.1x credentials profile to the interface.
Step 12	Router(config-if)# end	Returns to privileged EXEC mode.

	Command	Purpose
Step 13	Router# show running-config interface <i>interface-id</i>	Verifies your configuration.
Step 14	Router# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to configure a switch as a supplicant:

```
Router# configure terminal
Router(config)# cisp enable
Router(config)# dot1x credentials test
Router(config)# username suppswitch
Router(config)# password myswitch
Router(config)# dot1x supplicant force-multicast
Router(config)# interface gigabitethernet1/1
Router(config-if)# switchport trunk encapsulation dot1q
Router(config-if)# switchport mode trunk
Router(config-if)# dot1x pae supplicant
Router(config-if)# dot1x credentials test
Router(config-if)# end
```

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.

To disable 802.1X authentication on the port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# no dot1x pae authenticator	Disables 802.1X authentication on the port.
Step 3	Router(config-if)# end	Returns to privileged EXEC mode.

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:

```
Router(config)# interface gigabitethernet 5/1
Router(config-if)# no dot1x pae authenticator
```

Resetting the 802.1X Configuration to the Default Values

To reset the 802.1X configuration to the default values, perform this task:

	Command	Purpose
Step 1	Router(config)# interface <i>type slot/port</i>	Specifies the port to be configured, and enters interface configuration mode.
Step 2	Router(config-if)# dot1x default	Resets the configurable 802.1X parameters to the default values.
Step 3	Router(config-if)# end	Returns to privileged EXEC mode.

This example shows how to reset a port's 802.1X authentication settings to the default values:

```
Router(config)# interface gigabitethernet 3/27
Router(config-if)# dot1x default
```

Displaying Authentication Status and Information

- [Displaying 802.1X Status, page 45-57](#)
- [Displaying Authentication Methods and Status, page 45-58](#)
- [Displaying MAC Authentication Bypass Status, page 45-61](#)

Displaying 802.1X Status

To display the global 802.1X administrative and operational status for the switch or the 802.1X settings for individual ports, perform this task:

Command	Purpose
Router# show dot1x [all interface <i>type slot/port</i>]	Displays the global 802.1X administrative and operational status for the switch. (Optional) Use the all keyword to display the global 802.1X status and the 802.1X settings for all interfaces using 802.1X authentication. (Optional) Use the interface keyword to display the 802.1X settings for a specific interface.

This example shows how to view only the global 802.1X status:

```
Router# show dot1x
Sysauthcontrol          Disabled
Dot1x Protocol Version      2
Critical Recovery Delay    100
Critical EAPOL            Disabled

Router#
```

This example shows how to view the global 802.1X status and the 802.1X settings for all interfaces using 802.1X authentication:

```
Router# show dot1x all
Sysauthcontrol          Disabled
Dot1x Protocol Version      2
Critical Recovery Delay    100
Critical EAPOL           Disabled

Dot1x Info for GigabitEthernet3/27
-----
PAE                      = AUTHENTICATOR
PortControl              = FORCE_AUTHORIZED
ControlDirection        = Both
HostMode                 = SINGLE_HOST
ReAuthentication         = Disabled
QuietPeriod              = 60
ServerTimeout            = 30
SuppTimeout              = 30
ReAuthPeriod             = 3600 (Locally configured)
ReAuthMax                = 2
MaxReq                   = 2
TxPeriod                 = 30
RateLimitPeriod          = 0

Router#
```

Displaying Authentication Methods and Status

To display the authentication methods and status, perform any of these tasks:

Command	Purpose
Router# show authentication registrations	Displays details of all registered methods.
Router# show authentication interface <i>interface</i>	Displays authentication information for a specific interface
Router# show authentication method <i>method</i>	Lists current authentication sessions that were authorized using the specified method.
Router# show authentication sessions [handle <i>handle</i>] [interface <i>interface</i>] [mac <i>mac</i>] [method <i>method</i>] [session-id <i>session-id</i>]	Displays information about current authentication sessions. With no options specified, all current active sessions will be listed. The keywords can be added and combined to display detailed information about specific sessions or subset of sessions.

Table 45-2 Authentication Session States

State	Description
Idle	The session has been initialized and no methods have run yet.
Running	A method is running for this session.
No methods	No method has provided a result for this session.
Authc Success	A method has provided a successful authentication result for the session.

Table 45-2 Authentication Session States (continued)

State	Description
Authc Failed	A method has provided a failed authentication result for the session.
Authz Success	All features have been successfully applied for this session.
Authz Failed	A feature has failed to be applied for this session.

Table 45-3 Authentication Method States

State	Description
Not run	The method has not run for this session
Running	The method is running for this session.
Failed over	The method has failed and the next method is expected to provide a result.
Success	The method has provided a successful authentication result for the session.
Authc Failed	The method has provided a failed authentication result for the session.

This example shows how to display the registered authentication methods:

```
Router# show authentication registrations
Auth Methods registered with the Auth Manager:
  Handle  Priority  Name
    3         0    dot1x
    2         1    mab
    1         2    webauth
```

This example shows how to display the authentication details for a given interface:

```
Router# show authentication interface gigabitethernet 1/23
Client list:
  MAC Address      Domain   Status           Handle           Interface
  0123.4567.abcd   DATA   Authz Success    0xE0000000      GigabitEthernet1/23

Available methods list:
  Handle  Priority  Name
    3         0    dot1x
    2         1    mab
Runnable methods list:
  Handle  Priority  Name
    2         0    mab
    3         1    dot1x
```

This example shows how to display all authentication sessions on the switch:

```
Router# show authentication sessions

Interface  MAC Address      Method  Domain   Status           Session ID
Gi1/48     0015.63b0.f676   dot1x   DATA   Authz Success    0A3462B10000000102983C05C
Gi1/5      000f.23c4.a401   mab     DATA   Authz Success    0A3462B100000000D24F80B58
Gi1/5      0014.bf5d.d26d   dot1x   DATA   Authz Success    0A3462B100000000E29811B94
```

This example shows how to display sessions authorized using a specified authentication method:

```
Router# show authentication method dot1x
Interface  MAC Address      Method  Domain  Status      Session ID
Gi1/48     0015.63b0.f676  dot1x   DATA   Authz Success  0A3462B1000000102983C05C
Gi1/5      0014.bf5d.d26d  dot1x   DATA   Authz Success  0A3462B10000000E29811B94
```

This example shows how to display all authentication sessions on an interface:

```
Router# show authentication sessions interface gigabitethernet 1/47
```

```
Interface: GigabitEthernet1/47
      MAC Address: Unknown
      IP Address: Unknown
      Status: Authz Success
      Domain: DATA
      Oper host mode: multi-host
      Oper control dir: both
      Authorized By: Guest Vlan
      Vlan Policy: 20
      Session timeout: N/A
      Idle timeout: N/A
      Common Session ID: 0A3462C8000000000002763C
      Acct Session ID: 0x00000002
      Handle: 0x25000000
```

Runnable methods list:

```
Method  State
mab      Failed over
dot1x    Failed over
```

```
Interface: GigabitEthernet1/47
      MAC Address: 0005.5e7c.da05
      IP Address: Unknown
      User-Name: 00055e7cda05
      Status: Authz Success
      Domain: VOICE
      Oper host mode: multi-domain
      Oper control dir: both
      Authorized By: Authentication Server
      Session timeout: N/A
      Idle timeout: N/A
      Common Session ID: 0A3462C80000000010002A238
      Acct Session ID: 0x00000003
      Handle: 0x91000001
```

Runnable methods list:

```
Method  State
mab      Authc Success
dot1x    Not run
```

This example shows how to display the authentication session for a specified session ID:

```
Router# show authentication sessions session-id 0B0101C70000004F2ED55218
```

```
Interface: GigabitEthernet9/2
      MAC Address: 0000.0000.0011
      IP Address: 20.0.0.7
      Username: johndoe
      Status: Authz Success
      Domain: DATA
      Oper host mode: multi-host
      Oper control dir: both
```



```

Authorized By: Critical Auth
Vlan policy: N/A
Session timeout: N/A
Idle timeout: N/A
Common Session ID: 0B0101C70000004F2ED55218
Acct Session ID: 0x00000003
Handle: 0x91000001

```

```

Runnable methods list:
Method      State
mab         Authc Success
dot1x       Not run

```

This example shows how to display all clients authorized by the specified authentication method:

```
Router# show authentication sessions method mab
```

No Auth Manager contexts match supplied criteria

```
Router# show authentication sessions method dot1x
```

```

Interface  MAC Address      Domain   Status      Session ID
Gi9/2      0000.0000.0011  DATA   Authz Success 0B0101C70000004F2ED55218

```

Displaying MAC Authentication Bypass Status

To display the MAB status, perform this task:

Command	Purpose
Router# show mab {all interface type slot/port} [detail]	Displays MAB authentication details for all interfaces or for a specific interface.

Table 45-4 MAB Authentication States

State	Description
INITIALIZE	The authorization session is initialized.
ACQUIRING	The session is obtaining the client MAC address.
AUTHORIZING	The session is waiting for MAC-based authorization.
TERMINATE	The authorization session result has been obtained.

This example shows how to display the brief MAB status for a single interface:

```
Router# show mab interface fa1/1
```

```

MAB details for GigabitEthernet1/1
-----
Mac-Auth-Bypass           = Enabled
Inactivity Timeout        = None

```

This example shows how to display the detailed MAB status for a single interface:

```
Router# show mab interface fa1/1 detail
```

```

MAB details for GigabitEthernet1/1
-----

```

```
Mac-Auth-Bypass           = Enabled
Inactivity Timeout        = None

MAB Client List
-----
Client MAC                 = 000f.23c4.a401
MAB SM state               = TERMINATE
Auth Status                = AUTHORIZED
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

46

Web-Based Authentication

- [Prerequisites for Web-based Authentication, page 46-1](#)
- [Restrictions for Web-based Authentication, page 46-1](#)
- [Information About Web-Based Authentication, page 46-2](#)
- [Default Web-Based Authentication Configuration, page 46-7](#)
- [How to Configure Web-Based Authentication, page 46-7](#)
- [Displaying Web-Based Authentication Status, page 46-15](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Web-based Authentication

None.

Restrictions for Web-based Authentication

- 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.

- On Layer 2 interfaces, you cannot authenticate hosts 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 HTTP server on the switch. 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 may experience traffic disruption if an STP topology change results in the host traffic arriving on a different port. This is because ARP and DHCP updates may not be sent after a Layer 2 (STP) topology change.
- Web-based authentication does not support VLAN assignment as a downloadable host policy.
- Cisco IOS Release 15.4SY support downloadable ACLs (DACLS) from the RADIUS server.
- Web-based authentication is not supported for IPv6 traffic.

Information About Web-Based Authentication

- [Web-Based Authentication Overview, page 46-2](#)
- [Device Roles, page 46-3](#)
- [Host Detection, page 46-3](#)
- [Session Creation, page 46-4](#)
- [Authentication Process, page 46-4](#)
- [AAA Fail Policy, page 46-5](#)
- [Customization of the Authentication Proxy Web Pages, page 46-5](#)
- [Web-based Authentication Interactions with Other Features, page 46-5](#)

Web-Based Authentication Overview

The web-based authentication feature implements web-based authentication (also known as Web Authentication Proxy), which can function as part of the authentication, authorization, and accounting (AAA) system.

You can use the web-based authentication feature to authenticate end users on host systems that do not run the IEEE 802.1X supplicant. You can configure the web-based authentication feature on Layer 2 and Layer 3 interfaces.

When a user initiates an HTTP session, the web-based authentication feature intercepts ingress HTTP packets from the host and sends an HTML login page to the user. The user keys in their credentials, which the web-based authentication feature sends to the AAA server for authentication. If the 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 the authentication fails, web-based authentication feature sends a Login-Fail HTML page to the user, which prompts the user to retry the login attempt. If the user exceeds the maximum number of failed login attempts, web-based authentication sends a Login-Expired HTML page to the host and the user is placed on a watch list for a waiting period.

Device Roles

With web-based authentication, the devices in the network have specific roles as shown in [Figure 46-1](#).

Figure 46-1 *Web-based Authentication Device Roles*

The specific roles shown in [Figure 46-1](#) are as follows:

- *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 an HTML browser with Java Script enabled.
- *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.
- *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.

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 3 interfaces, web-based authentication sets an HTTP intercept ACL when the feature is configured on the interface (or when the interface is put in service).

For Layer 2 interfaces, web-based authentication detects IP hosts using the following mechanisms:

- ARP based trigger—ARP redirect ACL allows web-based authentication to detect hosts with static IP address or dynamically acquired 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:

- Checks 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 considered to be established.
- Checks for Auth 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 considered to be 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 web-based authentication is enabled, the following 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 on the login page, and the switch sends the entries to the authentication server.
- If the client identity is valid and 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, but if the maximum number of attempts fail, the switch sends the login expired page and the host is placed in a watch list. After a watch list timeout, 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 will apply the failure access policy to the host. The login success page is sent to the user.

The switch reauthenticates a client when the host does not respond to an ARP probe on a Layer 2 interface, or 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.

AAA Fail Policy

The AAA fail policy is a method for allowing a user to connect or to remain connected to the network if the AAA server is not available. If the AAA server cannot be reached when web-based authentication of a client is needed, instead of rejecting the user (that is, not providing the access to the network), an administrator can configure a default AAA fail policy that can be applied to the user.

This policy is advantageous for the following reasons:

- While AAA is unavailable, the user will still have connectivity to the network, although access may be restricted.
- When the AAA server is again available, a user can be revalidated, and the user's normal access policies can be downloaded from the AAA server.

**Note**

When the AAA server is down, the AAA fail policy is applied only if there is no existing policy associated with the user. Typically, if the AAA server is unavailable when a user session requires reauthentication, the policies currently in effect for the user are retained.

While the AAA fail policy is in effect, the session state is maintained as AAA Down.

Customization of the Authentication Proxy Web Pages

The switch's internal HTTP server hosts four HTML pages for delivery to an authenticating client during the web-based authentication process. The four pages allow the server to notify the user of the following four states of the authentication process:

- Login—The user's credentials are requested
- Success—The login was successful
- Fail—The login has failed
- Expire—The login session has expired due to excessive login failures

You can substitute your custom HTML pages for the four default internal HTML pages, or you can specify a URL to which the user will be redirected upon successful authentication, effectively replacing the internal Success page.

Web-based Authentication Interactions with Other Features

- [Port Security, page 46-6](#)
- [Gateway IP, page 46-6](#)
- [ACLs, page 46-6](#)
- [IP Source Guard, page 46-6](#)
- [EtherChannel, page 46-6](#)
- [Switchover, page 46-6](#)

Port Security

You can configure web-based authentication and port security on the same port. (You configure port security on the port by using the **switchport port-security** interface configuration command.) When you enable port security and web-based authentication on a 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 [“How to Configure Port Security” section on page 47-4](#).

Gateway IP

You cannot configure Gateway IP 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 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.

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 has VACL capture configured.

IP Source Guard

Configuring IP Source Guard and web-based authentication on the same interface is not supported.

You can configure IP Source Guard and web-based authentication on the same interface. If DHCP snooping is also enabled on the access VLAN, you must enter the **platform acl tcam override dynamic dhcp-snooping** command in global configuration mode to avoid conflict between the two features. Other VLAN-based features are not supported when IP Source Guard and web-based authentication are combined.

EtherChannel

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

Switchover

In RPR redundancy mode, information about currently authenticated hosts is maintained during a switchover. Users will not need to reauthenticate.

Default Web-Based Authentication Configuration

Feature	Default Setting
AAA	Disabled
RADIUS server	
• IP address	• None specified
• UDP authentication port	• 1812
• Key	• None specified
Default value of inactivity timeout	3600 seconds
Inactivity timeout	Enabled

How to Configure Web-Based Authentication

- [Default Web-Based Authentication Configuration, page 46-7](#)
- [Web-based Authentication Configuration Task List, page 46-8](#)
- [Configuring the Authentication Rule and Interfaces, page 46-8](#)
- [Configuring AAA Authentication, page 46-9](#)
- [Configuring Switch-to-RADIUS-Server Communication, page 46-9](#)
- [Configuring the HTTP Server, page 46-11](#)
- [Configuring the Web-based Authentication Parameters, page 46-14](#)
- [Removing Web-based Authentication Cache Entries, page 46-14](#)

Web-based Authentication Configuration Task List

- [Configuring the Authentication Rule and Interfaces, page 46-8](#)
- [Configuring AAA Authentication, page 46-9](#)
- [Configuring Switch-to-RADIUS-Server Communication, page 46-9](#)
- [Configuring the HTTP Server, page 46-11](#)
- [Configuring an AAA Fail Policy, page 46-13](#)
- [Configuring the Web-based Authentication Parameters, page 46-14](#)
- [Removing Web-based Authentication Cache Entries, page 46-14](#)

Configuring the Authentication Rule and Interfaces

To configure web-based authentication, perform this task:

	Command	Purpose
Step 1	Router(config)# ip admission name <i>name</i> proxy http	Configures an authentication rule for web-based authorization.
Step 2	Router(config)# interface <i>type slot/port</i>	Enters interface configuration mode and specifies the ingress Layer 2 or Layer 3 interface to be enabled for web-based authentication.
Step 3	Router(config-if)# ip access-group <i>name</i>	Applies the default ACL.
Step 4	Router(config-if)# ip admission <i>name</i>	Configures web-based authentication on the specified interface.
Step 5	Router(config-if)# authentication order <i>method1</i> [<i>method2</i>] [<i>method3</i>]	(Optional) Specifies the fallback order of authentication methods to be used. The three values of <i>method</i> , in the default order, are dot1x , mab , and webauth . Omitting a method disables that method on the interface.
Step 6	Router(config-if)# exit	Returns to configuration mode.
Step 7	Router(config)# ip device tracking	Enables the IP device tracking table.
Step 8	Router(config)# end	Returns to privileged EXEC mode.

This example shows how to enable web-based authentication, while disabling 802.1X or MAB authentication, on port 5/1:

```
Router(config)# ip admission name webauth1 proxy http
Router(config)# interface gigabitethernet 5/1
Router(config-if)# ip admission webauth1
Router(config-if)# authentication order webauth
Router(config-if)# exit
Router(config)# ip device tracking
```

This example shows how to verify the configuration:

```
Router# 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

To enable web-based authentication, you must enable AAA and specify the authentication method. perform this task:

	Command	Purpose
Step 1	Router(config)# aaa new-model	Enables AAA functionality.
Step 2	Router(config)# aaa authentication login default group {tacacs+ radius}	Defines the list of authentication methods at login.
Step 3	Router(config)# aaa authorization auth-proxy default group {tacacs+ radius}	Creates an authorization method list for web-based authorization.
Step 4	Router(config)# tacacs-server host {hostname ip_address}	Specifies an AAA server. For Radius servers, see the section “Configuring Switch-to-RADIUS-Server Communication” section on page 46-9.
Step 5	Router(config)# tacacs-server key {key-data}	Configures the authorization and encryption key used between the switch and the TACACS server.

This example shows how to enable AAA:

```

Router(config)# aaa new-model
Router(config)# aaa authentication login default group tacacs+
Router(config)# aaa authorization auth-proxy default group tacacs+

```

Configuring Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by any of the following:

- 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, 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 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	Router(config)# ip radius source-interface <i>interface_name</i>	Specifies that the RADIUS packets have the IP address of the indicated interface.
Step 2	Router(config)# radius-server host { <i>hostname</i> <i>ip-address</i> } test username <i>username</i>	Specifies 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 be used between the switch and the RADIUS server. To use multiple RADIUS servers, reenter this command.
Step 3	Router(config)# radius-server key <i>string</i>	Configures the authorization and encryption key used between the switch and the RADIUS daemon running on the RADIUS server.
Step 4	Router(config)# radius-server vsa send authentication	Enables downloading of an ACL from the RADIUS server.
Step 5	Router(config)# radius-server dead-criteria tries <i>num-tries</i>	Specifies the number of unanswered transmits to a RADIUS server before considering the server to be dead. The range of <i>num-tries</i> is 1 to 100.

- 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*, 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.
- 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.



Note

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch, the key string to be shared by both the server and the switch, and the downloadable ACL. For more information, see the RADIUS server documentation.

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

```
Router(config)# ip radius source-interface Vlan80
Router(config)# radius-server host 172.120.39.46 test username user1
Router(config)# radius-server key rad123
Router(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. To enable the server, perform one of these tasks in global configuration mode:

Command	Purpose
Router(config)# ip http server	Enables the HTTP server. The web-based authentication feature uses the HTTP server to communicate with the hosts for user authentication.
Router(config)# ip http secure-server	Enables HTTPS.

You can optionally configure custom authentication proxy web pages or specify a redirection URL for successful login, as described in the following sections:

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

Customizing the Authentication Proxy Web Pages

You have the option to provide four substitute HTML pages to be displayed to the user in place of the switch's internal 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's internal disk or flash memory, then perform this task in global configuration mode:

	Command	Purpose
Step 1	Router(config)# ip admission proxy http login page file <i>device:login-filename</i>	Specifies the location in the switch memory file system of the custom HTML file to be used in place of the default login page. The <i>device:</i> is either disk or flash memory, such as disk0:.
Step 2	Router(config)# ip admission proxy http success page file <i>device:success-filename</i>	Specifies the location of the custom HTML file to be used in place of the default login success page.
Step 3	Router(config)# ip admission proxy http failure page file <i>device:fail-filename</i>	Specifies the location of the custom HTML file to be used in place of the default login failure page.
Step 4	Router(config)# ip admission proxy http login expired page file <i>device:expired-filename</i>	Specifies the location of the custom HTML file to be used in place of the default login expired page.

- To enable the custom web pages feature, you must specify all four custom HTML files. If fewer than four files are specified, the internal default HTML pages will be used.
- The four custom HTML files must be present on the disk or flash of the switch.
- An image file has a size limit of 256 KB.
- All image files must have a filename that begins with “web_auth_” (like “web_auth_logo.jpg” instead of “logo.jpg”).
- All image file names must be less than 33 characters.

- Any images on the custom pages must be located on an accessible HTTP server. An intercept ACL must be configured within the admission rule to allow access to the HTTP server.
- Any external link from a custom page will require configuration of an intercept ACL within the admission rule.
- Any name resolution required for external links or images will require configuration of an intercept ACL within the admission rule to access a valid DNS server.
- If the custom web pages feature is enabled, a configured auth-proxy-banner will not be used.
- If the custom web pages feature is enabled, the redirection URL for successful login feature will not be 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 the following guidelines for this page:

- The login form must accept user input for the username and password and must POST the data 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.

The following example shows how to configure custom authentication proxy web pages:

```
Router(config)# ip admission proxy http login page file disk1:login.htm
Router(config)# ip admission proxy http success page file disk1:success.htm
Router(config)# ip admission proxy http fail page file disk1:fail.htm
Router(config)# ip admission proxy http login expired page file disk1:expired.htm
```

The following example shows how to verify the configuration of custom authentication proxy web pages:

```
Router# show ip admission configuration
```

```
Authentication proxy webpage
Login page           : disk1:login.htm
Success page         : disk1:success.htm
Fail Page            : disk1:fail.htm
Login expired Page   : disk1: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 have the option to specify a URL to which the user will be redirected upon successful authentication, effectively replacing the internal Success HTML page.

To specify a redirection URL for successful login, perform this task in global configuration mode:

Command	Purpose
Router(config)# ip admission proxy http success redirect url-string	Specifies a URL for redirection of the user in place of the default login success page.

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

- If the custom authentication proxy web pages feature is enabled, the redirection URL feature is disabled and will not be 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 will not be used.
- To remove the specification of a redirection URL, use the **no** form of the command.

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

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

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

```
Router# 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 an AAA Fail Policy

To configure an AAA fail policy, perform this task in global configuration mode:

	Command	Purpose
Step 1	Router(config)# ip admission name rule-name proxy http event timeout aaa policy identity identity_policy_name	Creates an AAA fail rule and associates an identity policy to be applied to sessions when the AAA server is unreachable. To remove the rule on the switch, use the no ip admission name rule-name proxy http event timeout aaa policy identity global configuration command.
Step 2	Router(config)# ip admission ratelimit aaa-down number_of_sessions	(Optional) To avoid flooding the AAA server when it returns to service, you can rate limit the authentication attempts from hosts in the AAA Down state.

The following example shows how to apply an AAA fail policy:

```
Router(config)# ip admission name AAA_FAIL_POLICY proxy http event timeout aaa policy identity GLOBAL_POLICY1
```

The following example shows how to determine whether any hosts are connected in the AAA Down state:

```
Router# show ip admission cache
Authentication Proxy Cache
Client IP 209.165.201.11 Port 0, timeout 60, state ESTAB (AAA Down)
```

The following example shows how to view detailed information about a particular session based on the host IP address:

```
Router# show ip admission cache 209.165.201.11
Address          : 209.165.201.11
MAC Address      : 0000.0000.0000
Interface        : Vlan333
Port             : 3999
Timeout          : 60
Age              : 1
State            : AAA Down
AAA Down policy  : AAA_FAIL_POLICY
```

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.

To configure the web-based authentication parameters, perform this task:

	Command	Purpose
Step 1	Router(config)# ip admission max-login-attempts <i>number</i>	Sets the maximum number of failed login attempts. The range is 1 to 2147483647 attempts; the default is 5.
Step 2	Router(config)# end	Returns to privileged EXEC mode.

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

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

Removing Web-based Authentication Cache Entries

To delete existing session entries, perform either of these tasks:

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

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


Displaying Web-Based Authentication Status

To display the web-based authentication settings for all interfaces or for specific ports, perform this task:

Command	Purpose
Router# show fm ip-admission l2http [all interface <i>type slot/port</i>]	Displays the web-based authentication settings. (Optional) Use the all keyword to display the settings for all interfaces using web-based authentication. (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:

```
Router# show fm ip-admission l2http all
```

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

```
Router# show fm ip-admission l2http interface gigabitethernet 3/27
```

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

47

Port Security

- [Prerequisites for Port Security, page 47-1](#)
- [Restrictions for Port Security, page 47-1](#)
- [Information About Port Security, page 47-2](#)
- [Default Port Security Configuration, page 47-4](#)
- [How to Configure Port Security, page 47-4](#)
- [Verifying the Port Security Configuration, page 47-11](#)

**Note**

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.4SY supports only Ethernet interfaces. Cisco IOS Release 15.4SY does not support any WAN features or commands.

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Port Security

None.

Restrictions for Port Security

- With the default port security configuration, to bring all secure ports out of the error-disabled state, enter the **errdisable recovery cause psecure-violation** global configuration command, or manually reenables the port by entering the **shutdown** and **no shut down** interface configuration commands.
- Enter the **clear port-security dynamic** global configuration command to clear all dynamically learned secure addresses.

- Port security learns unauthorized MAC addresses with a bit set that causes traffic to them or from them to be dropped. The **show mac address-table** command displays the unauthorized MAC addresses, but does not display the state of the bit. (CSCeb76844)
- To preserve dynamically learned sticky MAC addresses and configure them on a port following a bootup or a reload and after the dynamically learned sticky MAC addresses have been learned, you must enter a **write memory** or **copy running-config startup-config** command to save them in the startup-config file.
- Port security supports private VLAN (PVLAN) ports.
- Port security supports IEEE 802.1Q tunnel ports.
- Port security does not support Switch Port Analyzer (SPAN) destination ports.
- Port security supports access and trunking EtherChannel port-channel interfaces.
- You can configure port security and 802.1X port-based authentication on the same port.
- Port security supports nonnegotiating trunks.
 - Port security only supports trunks configured with these commands:
switchport
switchport trunk encapsulation
switchport mode trunk
switchport nonegotiate
 - If you reconfigure a secure access port as a trunk, port security converts all the sticky and static secure addresses on that port that were dynamically learned in the access VLAN to sticky or static secure addresses on the native VLAN of the trunk. Port security removes all secure addresses on the voice VLAN of the access port.
 - If you reconfigure a secure trunk as an access port, port security converts all sticky and static addresses learned on the native VLAN to addresses learned on the access VLAN of the access port. Port security removes all addresses learned on VLANs other than the native VLAN.



Note Port security uses the VLAN ID configured with the **switchport trunk native vlan** command.

- Take care when you enable port security on the ports connected to the adjacent switches when there are redundant links running between the switches because port security might error-disable the ports due to port security violations.
- In Cisco IOS Release 15.1(1)SY2 and later releases, the options starting with the number 6 for the **radius-server attribute** command changed from 6, 61, 69 to 6, 61, 66, 67, and 69.

Information About Port Security

- [Port Security with Dynamically Learned and Static MAC Addresses, page 47-3](#)
- [Port Security with Sticky MAC Addresses, page 47-3](#)
- [Port Security with IP Phones, page 47-4](#)

Port Security with Dynamically Learned and Static MAC Addresses

You can use port security with dynamically learned and static MAC addresses to restrict a port's ingress traffic by limiting the MAC addresses that are allowed to send traffic into the port. When you assign secure MAC addresses to a secure port, the port does not forward ingress traffic that has 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 device attached to that port has the full bandwidth of the port.

A security violation occurs in these situations:

- When the maximum number of secure MAC addresses is reached on a secure port and the source MAC address of the ingress traffic is different from any of the identified secure MAC addresses, port security applies the configured violation mode.
- If traffic with a secure MAC address that is configured or learned on one secure port attempts to access another secure port in the same VLAN, applies the configured violation mode.

**Note**

After a secure MAC address is configured or learned on one secure port, the sequence of events that occurs when port security detects that secure MAC address on a different port in the same VLAN is known as a MAC move violation.

- Running diagnostic tests with port security enabled.

See the [“Configuring the Port Security Violation Mode on a Port” section on page 47-6](#) for more information about the violation modes.

After you have set the maximum number of secure MAC addresses on a port, port security includes the secure addresses in the address table in one of these ways:

- You can statically configure all secure MAC addresses by using the **switchport port-security mac-address *mac_address*** interface configuration command.
- You can allow the port to dynamically configure secure MAC addresses with the MAC addresses of connected devices.
- You can statically configure a number of addresses and allow the rest to be dynamically configured.

If the port has a link-down condition, all dynamically learned addresses are removed.

Following bootup, a reload, or a link-down condition, port security does not populate the address table with dynamically learned MAC addresses until the port receives ingress traffic.

A security violation occurs if the maximum number of secure MAC addresses have been added to the address table and the port receives traffic from a MAC address that is not in the address table.

You can configure the port for one of three violation modes: protect, restrict, or shutdown. See the [“How to Configure Port Security” section on page 47-4](#).

To ensure that an attached device has the full bandwidth of the port, set the maximum number of addresses to one and configure the MAC address of the attached device.


Port Security with Sticky MAC Addresses

Port security with sticky MAC addresses provides many of the same benefits as port security with static MAC addresses, but sticky MAC addresses can be learned dynamically. Port security with sticky MAC addresses retains dynamically learned MAC addresses during a link-down condition.

If you enter a **write memory** or **copy running-config startup-config** command, then port security with sticky MAC addresses saves dynamically learned MAC addresses in the startup-config file and the port does not have to learn addresses from ingress traffic after bootup or a restart.

Port Security with IP Phones

Figure 47-1 Device Connected Through IP Phone



Because the device is not directly connected to the switch, the switch cannot physically detect a loss of port link if the device is disconnected. Later Cisco IP phones send a Cisco Discovery Protocol (CDP) host presence type length value (TLV) to notify the switch of changes in the attached device's port link state. The switch recognizes the host presence TLV. Upon receiving a host presence TLV notification of a link down on the IP phone's data port, port security removes from the address table all static, sticky, and dynamically learned MAC addresses. The removed addresses are added again only when the addresses are learned dynamically or configured.

Default Port Security Configuration

Feature	Default Setting
Port security	Disabled.
Maximum number of secure MAC addresses	1.
Violation mode	Shutdown. The port shuts down when the maximum number of secure MAC addresses is exceeded, and an SNMP trap notification is sent.

How to Configure Port Security

- [Enabling Port Security, page 47-5](#)
- [Configuring the Port Security Violation Mode on a Port, page 47-6](#)
- [Configuring the Maximum Number of Secure MAC Addresses on a Port, page 47-7](#)
- [Enabling Port Security with Sticky MAC Addresses on a Port, page 47-8](#)
- [Configuring a Static Secure MAC Address on a Port, page 47-9](#)
- [Configuring Secure MAC Address Aging on a Port, page 47-10](#)

Enabling Port Security

- [Enabling Port Security on a Trunk, page 47-5](#)
- [Enabling Port Security on an Access Port, page 47-6](#)

Enabling Port Security on a Trunk

Port security supports nonnegotiating trunks.



Caution

Because the default number of secure addresses is one and the default violation action is to shut down the port, configure the maximum number of secure MAC addresses on the port before you enable port security on a trunk (see [“Configuring the Maximum Number of Secure MAC Addresses on a Port” section on page 47-7](#)).

To enable port security on a trunk, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the interface to configure.
Step 2	Router(config-if)# switchport	Configures the port as a Layer 2 port.
Step 3	Router(config-if)# switchport trunk encapsulation {isl dot1q}	Configures the encapsulation as 802.1Q.
Step 4	Router(config-if)# switchport mode trunk	Configures the port to trunk unconditionally.
Step 5	Router(config-if)# switchport nonegotiate	Configures the trunk not to use DTP.
Step 6	Router(config-if)# switchport port-security	Enables port security on the trunk.
Step 7	Router(config-if)# do show port-security interface type slot/port include Port Security	Verifies the configuration.

This example shows how to configure Gigabit Ethernet port 5/36 as a nonnegotiating trunk and enable port security:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/36
Router(config-if)# switchport
Router(config-if)# switchport mode trunk
Router(config-if)# switchport nonegotiate
Router(config-if)# switchport port-security
Router(config-if)# do show port-security interface gigabitethernet 5/36 | include Port Security
Port Security                               : Enabled
```

Enabling Port Security on an Access Port

To enable port security on an access port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the interface to configure. Note The port can be a tunnel port or a PVLAN port.
Step 2	Router(config-if)# switchport	Configures the port as a Layer 2 port.
Step 3	Router(config-if)# switchport mode access	Configures the port as a Layer 2 access port. Note A port in the default mode (dynamic desirable) cannot be configured as a secure port.
Step 4	Router(config-if)# switchport port-security	Enables port security on the port.
Step 5	Router(config-if)# do show port-security interface type slot/port include Port Security	Verifies the configuration.

This example shows how to enable port security on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport
Router(config-if)# switchport mode access
Router(config-if)# switchport port-security
Router(config-if)# do show port-security interface gigabitethernet 5/12 | include Port Security
Port Security : Enabled
```

Configuring the Port Security Violation Mode on a Port

To configure the port security violation mode on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport port-security violation {protect restrict shutdown}	(Optional) Sets the violation mode and the action to be taken when a security violation is detected.
Step 3	Router(config-if)# do show port-security interface type slot/port include violation_mode	Verifies the configuration. The values for <i>violation_mode</i> are protect , restrict , or shutdown .

- **protect**—The PFC drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value.
- **restrict**—The PFC drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value and causes the security violation counter to increment.
- **shutdown**—Puts the interface into the error-disabled state immediately and sends an SNMP trap notification.

**Note**

To bring a secure port out of the error-disabled state, enter the **errdisable recovery cause violation_mode** global configuration command, or you can manually reenab it by entering the **shutdown** and **no shut down** interface configuration commands.

This example shows how to configure the protect security violation mode on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport port-security violation protect
Router(config-if)# do show port-security interface gigabitethernet 5/12 | include Protect
Violation Mode                : Protect
```

This example shows how to configure the restrict security violation mode on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport port-security violation restrict
Router(config-if)# do show port-security interface gigabitethernet 5/12 | include Restrict
Violation Mode                : Restrict
```

Configuring the Maximum Number of Secure MAC Addresses on a Port

To configure the maximum number of secure MAC addresses on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the interface to configure.
Step 2	Router(config-if)# switchport port-security maximum number_of_addresses vlan {vlan_ID vlan_range}	Sets the maximum number of secure MAC addresses for the port (default is 1). Note Per-VLAN configuration is supported only on trunks.

- The range for *number_of_addresses* is 1 to 4,097.
- Port security supports trunks.
 - On a trunk, you can configure the maximum number of secure MAC addresses both on the trunk and for all the VLANs on the trunk.
 - You can configure the maximum number of secure MAC addresses on a single VLAN or a range of VLANs.
 - For a range of VLANs, enter a dash-separated pair of VLAN numbers.

- You can enter a comma-separated list of VLAN numbers and dash-separated pairs of VLAN numbers.

This example shows how to configure a maximum of 64 secure MAC addresses on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport port-security maximum 64
Router(config-if)# do show port-security interface gigabitethernet 5/12 | include Maximum
Maximum MAC Addresses      : 64
```

Enabling Port Security with Sticky MAC Addresses on a Port

To enable port security with sticky MAC addresses on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the interface to configure.
Step 2	Router(config-if)# switchport port-security mac-address sticky	Enables port security with sticky MAC addresses on a port.

- When you enter the **switchport port-security mac-address sticky** command:
 - All dynamically learned secure MAC addresses on the port are converted to sticky secure MAC addresses.
 - Static secure MAC addresses are not converted to sticky MAC addresses.
 - Secure MAC addresses dynamically learned in a voice VLAN are not converted to sticky MAC addresses.
 - New dynamically learned secure MAC addresses are sticky.
- When you enter the **no switchport port-security mac-address sticky** command, all sticky secure MAC addresses on the port are converted to dynamic secure MAC addresses.
- To preserve dynamically learned sticky MAC addresses and configure them on a port following a bootup or a reload, after the dynamically learned sticky MAC addresses have been learned, you must enter a **write memory** or **copy running-config startup-config** command to save them in the startup-config file.

This example shows how to enable port security with sticky MAC addresses on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport port-security mac-address sticky
```

Configuring a Static Secure MAC Address on a Port

To configure a static secure MAC address on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport port-security mac-address sticky mac_address [vlan vlan_ID]	Configures a static MAC address as secure on the port. Note Per-VLAN configuration is supported only on trunks.
Step 3	Router(config-if)# end	Exits configuration mode.

- You can configure sticky secure MAC addresses if port security with sticky MAC addresses is enabled (see the [“Enabling Port Security with Sticky MAC Addresses on a Port”](#) section on page 47-8).
- The maximum number of secure MAC addresses on the port, configured with the **switchport port-security maximum** command, defines how many secure MAC addresses you can configure.
- If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are learned dynamically.
- Port security is supported on trunks.
 - On a trunk, you can configure a static secure MAC address in a VLAN.
 - On a trunk, if you do not configure a VLAN for a static secure MAC address, it is secure in the VLAN configured with the **switchport trunk native vlan** command.

This example shows how to configure a MAC address 1000.2000.3000 as secure on Gigabit Ethernet port 5/12 and verify the configuration:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport port-security mac-address 1000.2000.3000
Router(config-if)# end
Router# show port-security address
Secure Mac Address Table
```

```
-----
Vlan    Mac Address      Type              Ports
----    -
1       1000.2000.3000   SecureConfigured  Gi5/12
```

Configuring Secure MAC Address Aging on a Port

- [Configuring the Secure MAC Address Aging Type on a Port, page 47-10](#)
- [Configuring Secure MAC Address Aging Time on a Port, page 47-10](#)



Note

- Static secure MAC addresses and sticky secure MAC addresses do not age out.
- When the aging type is configured with the **absolute** keyword, all the dynamically learned secure addresses age out when the aging time expires. When the aging type is configured with the **inactivity** keyword, the aging time defines the period of inactivity after which all the dynamically learned secure addresses age out.

Configuring the Secure MAC Address Aging Type on a Port

You can configure the secure MAC address aging type on a port. To configure the secure MAC address aging type on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the LAN port to configure.
Step 2	Router(config-if)# switchport port-security aging type {absolute inactivity}	Configures the secure MAC address aging type on the port (default is absolute).

This example shows how to set the aging type to inactivity on Gigabit Ethernet port 5/12:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/12
Router(config-if)# switchport port-security aging type inactivity
Router(config-if)# do show port-security interface gigabitethernet 5/12 | include Type
Aging Type : Inactivity
```

Configuring Secure MAC Address Aging Time on a Port

To configure the secure MAC address aging time on a port, perform this task:

	Command	Purpose
Step 1	Router(config)# interface {type slot/port port-channel channel_number}	Selects the interface to configure.
Step 2	Router(config-if)# switchport port-security aging time aging_time	Configures the secure MAC address aging time on the port. The <i>aging_time</i> range is 1 to 1440 minutes (default is 0).

This example shows how to configure 2 hours (120 minutes) as the secure MAC address aging time on Gigabit Ethernet port 5/1:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitethernet 5/1
Router(config-if)# switchport port-security aging time 120
```

```
Router(config-if)# do show port-security interface gigabitethernet 5/12 | include Time
Aging Time : 120 mins
```

Verifying the Port Security Configuration

To display port security settings, enter this command:

Command	Purpose
Router# show port-security [interface {{ vlan <i>vlan_ID</i> { <i>type slot/port</i> }}] [address]	Displays port security settings for the switch or for the specified interface.

- Port security supports the **vlan** keyword only on trunks.
- Enter the **address** keyword to display secure MAC addresses, with aging information for each address, globally for the switch or per interface.
- The display includes these values:
 - 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
 - The violation mode

This example displays output from the **show port-security** command when you do not enter an interface:

```
Router# show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
              (Count)          (Count)          (Count)
-----
Gi5/1             11              11              0              Shutdown
Gi5/5             15              5               0              Restrict
Gi5/11            5               4               0              Protect
-----

Total Addresses in System: 21
Max Addresses limit in System: 128
```

This example displays output from the **show port-security** command for a specified interface:

```
Router# show port-security interface gigabitethernet 5/1
Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 11
Total MAC Addresses: 11
Configured MAC Addresses: 3
Aging time: 20 mins
Aging type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

This example displays the output from the **show port-security address** privileged EXEC command:

```
Router# show port-security address
Secure Mac Address Table
-----
Vlan    Mac Address      Type                Ports  Remaining Age
      (mins)
```

```

-----
1      0001.0001.0001      SecureDynamic      Gi5/1      15 (I)
1      0001.0001.0002      SecureDynamic      Gi5/1      15 (I)
1      0001.0001.1111      SecureConfigured   Gi5/1      16 (I)
1      0001.0001.1112      SecureConfigured   Gi5/1      -
1      0001.0001.1113      SecureConfigured   Gi5/1      -
1      0005.0005.0001      SecureConfigured   Gi5/5      23
1      0005.0005.0002      SecureConfigured   Gi5/5      23
1      0005.0005.0003      SecureConfigured   Gi5/5      23
1      0011.0011.0001      SecureConfigured   Gi5/11     25 (I)
1      0011.0011.0002      SecureConfigured   Gi5/11     25 (I)
-----

```

Total Addresses in System: 10

Max Addresses limit in System: 128



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

48

Lawful Intercept

- [Prerequisites for Lawful Intercept, page 48-1](#)
- [Restrictions for Lawful Intercept, page 48-1](#)
- [Information About Lawful Intercept, page 48-3](#)
- [How to Configure Lawful Intercept Support, page 48-9](#)

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Prerequisites for Lawful Intercept

- You must be running images that support secure shell (SSH), for example, s72033-adventerprise9-mz. Lawful intercept is not supported on images that do not support SSH.
- You must be logged in to the switch with the highest access level (level 15). To log in with level-15 access, enter the **enable** command and specify the highest-level password defined for the switch.
- You must issue commands in global configuration mode at the command-line interface (CLI). You can configure lawful intercept globally on all interfaces or on a specific interface.
- The time of day on the switches and the mediation device must be synchronized; use Network Time Protocol (NTP) on both the switches and the mediation device.
- (Optional) It might be helpful to use a loopback interface for the interface through which the switch communicates with the mediation device. If you do not use a loopback interface, you must configure the mediation device with multiple physical interfaces on the switch to handle network failures.

Restrictions for Lawful Intercept

- [General Configuration Restrictions, page 48-1](#)
- [MIB Guidelines, page 48-3](#)

General Configuration Restrictions

- VSS mode does not support lawful intercept.

- If the network administrator expects lawful intercept to be deployed at a node, do not configure optimized ACL logging (OAL), VLAN access control list (VACL) capture, or Intrusion Detection System (IDS) at the node. Deploying lawful intercept at the node causes unpredictable behavior in OAL, VACL capture, and IDS.
- To maintain switch performance, lawful intercept is limited to no more than 0.2% of active calls. For example, if the switch is handling 4000 calls, 8 of those calls can be intercepted.
- The CISCO-IP-TAP-MIB does not support the virtual routing and forwarding (VRF) OID `citapStreamVRF`.
- Captured traffic is rate limited to protect the CPU usage at the route processor. The rate limit is 8500 pps.
- The interface index is used during provisioning to select the index to enable lawful intercept on only; when set to 0, lawful intercept is applied to all interfaces.
- (Optional) The domain name for both the switch and the mediation device may be registered in the Domain Name System (DNS).
- The mediation device must have an access function (AF).
- You must add the mediation device to the SNMP user group that has access to the CISCO-TAP2-MIB view. Specify the username of the mediation device as the user to add to the group.

When you add the mediation device as a CISCO-TAP2-MIB user, you must include the mediation device's authorization password. The password must be at least eight characters in length.

- Dedicate an interface for lawful intercept processing. For example, you should not configure the interface to perform processor-intensive tasks such as QoS or routing.
- Supported for IPv4 unicast traffic only. In addition, for traffic to be intercepted, the traffic must be IPv4 on both the ingress and egress interfaces. For example, lawful intercept cannot intercept traffic if the egress side is MPLS and the ingress side is IPv4.
- IPv4 multicast, IPv6 unicast, and IPv6 multicast flows are not supported.
- Not supported on Layer 2 interfaces. However, lawful intercept can intercept traffic on VLANs that run over the Layer 2 interface.
- Not supported for packets that are encapsulated within other packets (for example, tunneled packets or Q-in-Q packets).
- Not supported for Q-in-Q packets. There is no support for Layer 2 taps for lawful intercept.
- Not supported for packets that are subject to Layer 3 or Layer 4 rewrite (for example, Network Address Translation [NAT] or TCP reflexive).
- In the ingress direction, the switch intercepts and replicates packets even if the packets are later dropped (for example, due to rate limiting or an access control list [ACL] **deny** statement). In the egress direction, packets are not replicated if they are dropped (for example, by ACL).
- Lawful intercept ACLs are applied internally to both the ingress and the egress directions of an interface.
- To intercept traffic from a specific user, a typical configuration consists of two flows, one for each direction.
- Packets that are subject to hardware rate limiting are processed by lawful intercept as follows:
 - Packets that are dropped by the rate limiter are not intercepted or processed.
 - Packets that are passed by the rate limiter are intercepted and processed.

- If multiple LEAs are using a single mediation device and each is executing a wiretap on the same target, the switch sends a single packet to the mediation device. It is up to the mediation device to duplicate the packet for each LEA.
- Lawful intercept can intercept IPv4 packets with values that match a combination of one or more of the following fields:
 - Destination IP address and mask
 - Destination port range
 - Source IP address and mask
 - Source port range
 - Protocol ID

MIB Guidelines

The following Cisco MIBs are used for lawful intercept processing. Include these MIBs in the SNMP view of lawful intercept MIBs to enable the mediation device to configure and execute wiretaps on traffic that flows through the switch.

- CISCO-TAP2-MIB—Required for both types of lawful intercepts: regular and broadband.
- CISCO-IP-TAP-MIB—Required for wiretaps on Layer 3 (IPv4) streams. Supported for regular and broadband lawful intercept.
- The CISCO-IP-TAB-MIB imposes limitations on the following features:
 - If one or all of the following features are configured and functioning and lawful intercept is enabled, lawful intercept takes precedence, and the feature behaves as follows:
 - Optimized ACL logging (OAL)—Does not function.
 - VLAN access control list (VACL) capturing—Does not function properly.
 - Intrusion detection system (IDS)—Does not function properly.The feature starts to function after you disable or unconfigure lawful intercept.
 - IDS cannot capture traffic on its own, but captures traffic that has been intercepted by lawful intercept only.

Information About Lawful Intercept

- [Lawful Intercept Overview, page 48-4](#)
- [Benefits of Lawful Intercept, page 48-4](#)
- [CALEA for Voice, page 48-5](#)
- [Network Components Used for Lawful Intercept, page 48-5](#)
- [Lawful Intercept Processing, page 48-7](#)
- [Lawful Intercept MIBs, page 48-7](#)

**Caution**

This guide does not address legal obligations for the implementation of lawful intercept. As a service provider, you are responsible to ensure that your network complies with applicable lawful intercept statutes and regulations. We recommend that you seek legal advice to determine your obligations.

Lawful Intercept Overview

Lawful intercept is a process that enables a Law Enforcement Agency (LEA) to perform electronic surveillance on an individual (a target) as authorized by a judicial or administrative order. To facilitate the lawful intercept process, certain legislation and regulations require service providers (SPs) and Internet service providers (ISPs) to implement their networks to explicitly support authorized electronic surveillance.

The surveillance is performed through the use of wiretaps on traditional telecommunications and Internet services in voice, data, and multiservice networks. The LEA delivers a request for a wiretap to the target's service provider, who is responsible for intercepting data communication to and from the individual. The service provider uses the target's IP address to determine which of its edge switches handles the target's traffic (data communication). The service provider then intercepts the target's traffic as it passes through the switch, and sends a copy of the intercepted traffic to the LEA without the target's knowledge.

The Lawful Intercept feature supports the Communications Assistance for Law Enforcement Act (CALEA), which describes how service providers in the United States must support lawful intercept. Currently, lawful intercept is defined by the following standards:

- Telephone Industry Association (TIA) specification J-STD-025
- Packet Cable Electronic Surveillance Specification (PKT-SP-ESP-101-991229)

For information about the Cisco lawful intercept solution, contact your Cisco account representative.

**Note**

The Lawful Intercept feature supports the interception of IPv4 protocol as defined by the object `citapStreamprotocol` in the `CISCO-IP-TAB-MIB` that includes voice and data interception.

Benefits of Lawful Intercept

- Allows multiple LEAs to run a lawful intercept on the same target without each other's knowledge.
- Does not affect subscriber services on the switch.
- Supports wiretaps in both the input and output direction.
- Supports wiretaps of Layer 1 and Layer 3 traffic. Layer 2 traffic is supported as IP traffic over VLANs.
- Supports wiretaps of individual subscribers that share a single physical interface.
- Cannot be detected by the target. Neither the network administrator nor the calling parties is aware that packets are being copied or that the call is being tapped.
- Uses Simple Network Management Protocol Version 3 (SNMPv3) and security features such as the View-based Access Control Model (SNMP-VACM-MIB) and User-based Security Model (SNMP-USM-MIB) to restrict access to lawful intercept information and components.

- Hides information about lawful intercepts from all but the most privileged users. An administrator must set up access rights to enable privileged users to access lawful intercept information.
- Provides two secure interfaces for performing an intercept: one for setting up the wiretap and one for sending the intercepted traffic to the LEA.

CALEA for Voice

The Communications Assistance for Law Enforcement Act (CALEA) for Voice feature allows the lawful interception of voice conversations that are running on Voice over IP (VoIP). Although the switches are not voice gateway devices, VoIP packets traverse the switches at the edge of the service provider network.


When an approved government agency determines that a telephone conversation is interesting, CALEA for Voice copies the IP packets comprising the conversation and sends the duplicate packets to the appropriate monitoring device for further analysis.

Network Components Used for Lawful Intercept

- [Mediation Device](#)
- [Lawful Intercept Administration](#)
- [Intercept Access Point](#)
- [Content Intercept Access Point](#)

For information about lawful intercept processing, see the [“Lawful Intercept Processing”](#) section on [page 48-7](#).

Figure 48-1 **Lawful Intercept Overview**



Mediation Device

A mediation device (supplied by a third-party vendor) handles most of the processing for the lawful intercept. The mediation device:

- Provides the interface used to set up and provision the lawful intercept.
- Generates requests to other network devices to set up and run the lawful intercept.
- Converts the intercepted traffic into the format required by the LEA (which can vary from country to country) and sends a copy of the intercepted traffic to the LEA without the target's knowledge.

**Note**

If multiple LEAs are performing intercepts on the same target, the mediation device must make a copy of the intercepted traffic for each LEA. The mediation device is also responsible for restarting any lawful intercepts that are disrupted due to a failure.

Lawful Intercept Administration

Lawful intercept administration (LIA) provides the authentication interface for lawful intercept or wiretap requests and administration.

Intercept Access Point

An intercept access point (IAP) is a device that provides information for the lawful intercept. There are two types of IAPs:

- Identification (ID) IAP—A device, such as an authentication, authorization, and accounting (AAA) server, that provides intercept-related information (IRI) for the intercept (for example, the target's username and system IP address) or call agents for voice over IP. The IRI helps the service provider determine which content IAP (switch) the target's traffic passes through.
- Content IAP—A device, such as the switch, that the target's traffic passes through. The content IAP:
 - Intercepts traffic to and from the target for the length of time specified in the court order. The switch continues to forward traffic to its destination to ensure that the wiretap is undetected.
 - Creates a copy of the intercepted traffic, encapsulates it in User Datagram Protocol (UDP) packets, and forwards the packets to the mediation device without the target's knowledge. IP option header is not supported.

**Note**

The content IAP sends a single copy of intercepted traffic to the mediation device. If multiple LEAs are performing intercepts on the same target, the mediation device must make a copy of the intercepted traffic for each LEA.

Content Intercept Access Point

Content IAP intercepts the interested data stream, duplicates the content, and sends the duplicated content to the mediation device. The mediation device receives the data from the ID IAP and Content IAP, converts the information to the required format depending on country specific requirement and forwards it to law enforcement agency (LEA).

Lawful Intercept Processing

After acquiring a court order or warrant to perform surveillance, the LEA delivers a surveillance request to the target's service provider. Service provider personnel use an administration function that runs on the mediation device to configure a lawful intercept to monitor the target's electronic traffic for a specific period of time (as defined in the court order).

After the intercept is configured, user intervention is no longer required. The administration function communicates with other network devices to set up and execute the lawful intercept. The following sequence of events occurs during a lawful intercept:

1. The administration function contacts the ID IAP for intercept-related information (IRI), such as the target's username and the IP address of the system, to determine which content IAP (switch) the target's traffic passes through.
2. After identifying the switch that handles the target's traffic, the administration function sends SNMPv3 **get** and **set** requests to the switch's Management Information Base (MIB) to set up and activate the lawful intercept. The CISCO-TAP2-MIB is the supported lawful intercept MIB to provide per-subscriber intercepts.
3. During the lawful intercept, the switch:
 - a. Examines incoming and outgoing traffic and intercepts any traffic that matches the specifications of the lawful intercept request.
 - b. Creates a copy of the intercepted traffic and forwards the original traffic to its destination so the target does not suspect anything.
 - c. Encapsulates the intercepted traffic in UDP packets and forwards the packets to the mediation device without the target's knowledge.



Note The process of intercepting and duplicating the target's traffic adds no detectable latency in the traffic stream.

4. The mediation device converts the intercepted traffic into the required format and sends it to a collection function running at the LEA. Here, the intercepted traffic is stored and processed.



Note If the switch intercepts traffic that is not allowed by the judicial order, the mediation device filters out the excess traffic and sends the LEA only the traffic allowed by the judicial order.

5. When the lawful intercept expires, the switch stops intercepting the target's traffic.

Lawful Intercept MIBs

- **CISCO-TAP2-MIB**—Used for lawful intercept processing.
- **CISCO-IP-TAP-MIB**—Used for intercepting Layer 3 (IPv4) traffic.

CISCO-TAP2-MIB

The CISCO-TAP2-MIB contains SNMP management objects that control lawful intercepts. The mediation device uses the MIB to configure and run lawful intercepts on targets whose traffic passes through the switch.

The CISCO-TAP2-MIB contains several tables that provide information for lawful intercepts that are running on the switch:

- **cTap2MediationTable**—Contains information about each mediation device that is currently running a lawful intercept on the switch. Each table entry provides information that the switch uses to communicate with the mediation device (for example, the device's address, the interfaces to send intercepted traffic over, and the protocol to use to transmit the intercepted traffic).
- **cTap2StreamTable**—Contains information used to identify the traffic to intercept. Each table entry contains a pointer to a filter that is used to identify the traffic stream associated with the target of a lawful intercept. Traffic that matches the filter is intercepted, copied, and sent to the corresponding mediation device application (cTap2MediationContentId).

The cTap2StreamTable table also contains counts of the number of packets that were intercepted, and counts of dropped packets that should have been intercepted, but were not.

- **cTap2DebugTable**—Contains debug information for troubleshooting lawful intercept errors.

The CISCO-TAP2-MIB also contains several SNMP notifications for lawful intercept events. For detailed descriptions of MIB objects, see the MIB itself.

CISCO-TAP2-MIB Processing

The administration function (running on the mediation device) issues SNMPv3 **set** and **get** requests to the switch's CISCO-TAP2-MIB to set up and initiate a lawful intercept. To do this, the administration function performs the following actions:

1. Creates a cTap2MediationTable entry to define how the switch is to communicate with the mediation device executing the intercept.



Note The cTap2MediationNewIndex object provides a unique index for the mediation table entry.

2. Creates an entry in the cTap2StreamTable to identify the traffic stream to intercept.
3. Sets cTap2StreamInterceptEnable to true(1) to start the intercept. The switch intercepts traffic in the stream until the intercept expires (cTap2MediationTimeout).

CISCO-IP-TAP-MIB

The CISCO-IP-TAP-MIB contains the SNMP management objects to configure and execute lawful intercepts on IPv4 traffic streams that flow through the switch. This MIB is an extension to the CISCO-TAP2-MIB.

You can use the CISCO-IP-TAP-MIB to configure lawful intercept on the switch to intercept IPv4 packets with values that match a combination of one or more of the following fields:

- Destination IP address and mask
- Destination port range
- Source IP address and mask
- Source port range
- Protocol ID

CISCO-IP-TAP-MIB Processing

When data is intercepted, two streams are created. One stream is for packets that originate from the target IP address to any other IP address using any port. The second stream is created for packets that are routed to the target IP address from any other address using any port. For VoIP, two streams are created, one for RTP packets from target and the second stream is for the RTP packets to target using the specific source and destination IP addresses and ports specified in SDP information used to setup RTP stream.

How to Configure Lawful Intercept Support

- [Security Considerations, page 48-9](#)
- [Accessing the Lawful Intercept MIBs, page 48-9](#)
- [Configuring SNMPv3, page 48-10](#)
- [Creating a Restricted SNMP View of Lawful Intercept MIBs, page 48-10](#)
- [Enabling SNMP Notifications for Lawful Intercept, page 48-12](#)

Security Considerations

- SNMP notifications for lawful intercept must be sent to UDP port 161 on the mediation device, not port 162 (which is the SNMP default). See the [“Enabling SNMP Notifications for Lawful Intercept” section on page 48-12](#) for instructions.
- The only users who should be allowed to access the lawful intercept MIBs are the mediation device and system administrators who need to know about lawful intercepts on the switch. In addition, these users must have authPriv or authNoPriv access rights to access the lawful intercept MIBs. Users with NoAuthNoPriv access cannot access the lawful intercept MIBs.
- You cannot use the SNMP-VACM-MIB to create a view that includes the lawful intercept MIBs.
- The default SNMP view excludes the following MIBs:
 - CISCO-TAP2-MIB
 - CISCO-IP-TAP-MIB
 - SNMP-COMMUNITY-MIB
 - SNMP-USM-MIB
 - SNMP-VACM-MIB

Also see the [“Restrictions for Lawful Intercept” section on page 48-1](#) and the [“Prerequisites for Lawful Intercept” section on page 48-1](#).

Accessing the Lawful Intercept MIBs

Due to its sensitive nature, the Cisco lawful intercept MIBs are only available in software images that support the lawful intercept feature. These MIBs are not accessible through the Network Management Software MIBs Support page (<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>).

Restricting Access to the Lawful Intercept MIBs

Only the mediation device and users who need to know about lawful intercepts should be allowed to access the lawful intercept MIBs. To restrict access to these MIBs, you must:

1. Create a view that includes the Cisco lawful intercept MIBs.
2. Create an SNMP user group that has read-and-write access to the view. Only users assigned to this user group can access information in the MIBs.
3. Add users to the Cisco lawful intercept user groups to define who can access the MIBs and any information related to lawful intercepts. Be sure to add the mediation device as a user in this group; otherwise, the switch cannot perform lawful intercepts.



Note

Access to the Cisco lawful intercept MIB view should be restricted to the mediation device and to system administrators who need to be aware of lawful intercepts on the switch. To access the MIB, users must have level-15 access rights on the switch.

Configuring SNMPv3

To perform the following procedures, SNMPv3 must be configured on the switch. See this publication:

<http://www.cisco.com/en/US/docs/ios-xml/ios/snmp/configuration/15-sy/snmp-15-sy-book.html>

Creating a Restricted SNMP View of Lawful Intercept MIBs

To create and assign users to an SNMP view that includes the Cisco lawful intercept MIBs, perform the following procedure at the CLI, in global configuration mode with level-15 access rights. For command examples, see the “Configuration Example” section on page 48-11.



Note

The command syntax in the following steps includes only those keywords required to perform each task. For details on command syntax, see the documents listed in the previous section (“Configuring SNMPv3”).

- Step 1** Make sure that SNMPv3 is configured on the switch. For instructions, see the documents listed in the “Configuring SNMPv3” section on page 48-10.
- Step 2** Create an SNMP view that includes the CISCO-TAP2-MIB (where *view_name* is the name of the view to create for the MIB). This MIB is required for both regular and broadband lawful intercept.

```
Router(config)# snmp-server view view_name ciscoTap2MIB included
```
- Step 3** Add one or both of the following MIBs to the SNMP view to configure support for wiretaps on IPv4 streams (where *view_name* is the name of the view you created in Step 2).

```
Router(config)# snmp-server view view_name ciscoIpTapMIB included
```
- Step 4** Create an SNMP user group (*groupname*) that has access to the lawful intercept MIB view and define the group’s access rights to the view.

```
Router(config)# snmp-server group groupname v3 noauth read view_name write view_name
```


- Step 5** Add users to the user group you just created (where *username* is the user, *groupname* is the user group, and *auth_password* is the authentication password):

```
Router(config)# snmp-server user username groupname v3 auth md5 auth_password
```



Note Be sure to add the mediation device to the SNMP user group; otherwise, the switch cannot perform lawful intercepts. Access to the lawful intercept MIB view should be restricted to the mediation device and to system administrators who need to know about lawful intercepts on the switch.

The mediation device is now able to access the lawful intercept MIBs, and issue SNMP **set** and **get** requests to configure and run lawful intercepts on the switch.

For instructions on how to configure the switch to send SNMP notifications to the mediation device, see the [“Enabling SNMP Notifications for Lawful Intercept”](#) section on page 48-12.

Configuration Example

The following commands show an example of how to enable the mediation device to access the lawful intercept MIBs.

```
Router(config)# snmp-server view tapV ciscoTap2MIB included
Router(config)# snmp-server view tapV ciscoIpTapMIB included
```

1. Create a view (tapV) that includes the appropriate lawful intercept MIBs (CISCO-TAP2-MIB and the CISCO-IP-TAP-MIB).
2. Create a user group (tapGrp) that has read, write, and notify access to MIBs in the tapV view.
3. Add the mediation device (ss8user) to the user group, and specify MD5 authentication with a password (ss8passwd).
4. (Optional) Assign a 24-character SNMP engine ID (for example, 123400000000000000000000) to the switch for administration purposes. If you do not specify an engine ID, one is automatically generated. Note that you can omit the trailing zeros from the engine ID, as shown in the last line of the example above.



Note Changing an engine ID has consequences for SNMP user passwords and community strings.

Enabling SNMP Notifications for Lawful Intercept

SNMP automatically generates notifications for lawful intercept events (see Table 48-1). This is because the default value of the `cTap2MediationNotificationEnable` object is `true(1)`.

To configure the switch to send lawful intercept notifications to the mediation device, issue the following CLI commands in global-configuration mode with level-15 access rights (where *MD-ip-address* is the IP address of the mediation device and *community-string* is the password-like community string to send with the notification request):

```
Router(config)# snmp-server host MD-ip-address community-string udp-port 161 snmp
Router(config)# snmp-server enable traps snmp authentication linkup linkdown coldstart
warmstart
```

- For lawful intercept, **udp-port** must be 161 and not 162 (the SNMP default).
- The second command configures the switch to send RFC 1157 notifications to the mediation device. These notifications indicate authentication failures, link status (up or down), and system restarts.

Table 48-1 SNMP Notifications for Lawful Intercept Events

Notification	Meaning
cTap2MIBActive	The switch is ready to intercept packets for a traffic stream configured in the CISCO-TAP2-MIB.
cTap2MediationTimedOut	A lawful intercept was terminated (for example, because cTap2MediationTimeout expired).
cTap2MediationDebug	Intervention is required for events related to cTap2MediationTable entries.
cTap2StreamDebug	Intervention is required for events related to cTap2StreamTable entries.

Disabling SNMP Notifications

You can disable SNMP notifications by entering the **no snmp-server enable traps** command.

To disable lawful intercept notifications, use SNMPv3 to set the CISCO-TAP2-MIB object `cTap2MediationNotificationEnable` to `false(2)`. To reenable lawful intercept notifications through SNMPv3, reset the object to `true(1)`.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

1

Online Diagnostic Tests

- [Global Health-Monitoring Tests, page 1-2](#)
- [Per-Port Tests, page 1-8](#)
- [PFC Layer 2 Tests, page 1-15](#)
- [DFC Layer 2 Tests, page 1-17](#)
- [PFC Layer 3 Tests, page 1-22](#)
- [DFC Layer 3 Tests, page 1-28](#)
- [Replication Engine Tests, page 1-34](#)
- [Fabric Tests, page 1-35s](#)
- [Exhaustive Memory Tests, page 1-39](#)
- [Service Module Tests, page 1-41](#)
- [Stress Tests, page 1-42](#)
- [General Tests, page 1-44](#)
- [Critical Recovery Tests, page 1-46](#)
- [ViSN Tests, page 1-47](#)

**Note**

- For information about configuring online diagnostic tests see [Chapter 16, “Online Diagnostics.”](#)
 - Before you enable any onlines diagnostics tests, enable console logging to see all warning messages.
 - We recommend that when you are running disruptive tests that you only run the tests when connected through console. When disruptive tests are complete a warning message on the console recommends that you reload the system to return to normal operation: strictly follow this warning.
 - While tests are running, all ports are shut down as a stress test is being performed with looping ports internally and external traffic might affect the test results. The switch must be rebooted to bring the switch to normal operation. When you issue the command to reload the switch, the system will ask you if the configuration should be saved.
 - Do not save the configuration.
 - If you are running the tests on other modules, after the test is initiated and complete, you must reset the module.
-

**Tip**

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

Global Health-Monitoring Tests

- [TestAsicSync, page 1-2](#)
- [TestEARLInternalTables, page 1-3](#)
- [TestErrorCounterMonitor, page 1-3](#)
- [TestFexModeLoopback, page 1-4](#)
- [TestIntPortLoopback, page 1-4](#)
- [TestL3TcamMonitoring, page 1-4](#)
- [TestLtlFpoeMemoryConsistency, page 1-5](#)
- [TestMacNotification, page 1-6](#)
- [TestPortTxMonitoring, page 1-6](#)
- [TestScratchRegister, page 1-7](#)
- [TestSnrMonitoring, page 1-7](#)
- [TestUnusedPortLoopback, page 1-8](#)

TestAsicSync

This test periodically verifies the status of bus and port synchronization ASICs.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Reset the module. After the module has ten consecutive failures or three consecutive resets, it powers down.
Hardware support:	All modules.

TestEARLInternalTables

This test detects most PFC and DFC hardware table problems by running consistency checks on the PFC and DFC hardware tables. The test runs every 5 minutes.

A failure of the test for the PFC results in one of these actions:

- Failover to the redundant supervisor engine.
- If a redundant supervisor engine is not installed, shutdown of the supervisor engine.

A failure of the test for the DFC results in one of these actions:

- Up to two resets of the DFC-equipped switching module.
- Shutdown following a third failure.

A CallHome message is generated if CallHome is configured on the system.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Reset the affected module.
Hardware support:	PFC and DFCs.

TestErrorCounterMonitor

This test monitors the errors and interrupts that occur on each module in the system by periodically polling for the error counters maintained in the module. If the errors exceed a threshold value, a syslog message is displayed with detailed information including the error-counter identifier, port number, total failures, consecutive failures, and the severity of the error counter.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable. This test is automatically disabled during CPU-usage spikes to maintain accuracy.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Display a syslog message indicating the error-counters detected on that port.
Hardware support:	All modules including supervisor engines.

TestFexModeLoopback

This test verifies the data path between the inband port and ports that can support an IA client link. 802.1q packets are sent to the port being tested. After VNTag encapsulation and deencapsulation, the test packet is looped back on the test port to the inband port. This test skips ports with active IA parent-client links. The test typically disrupts traffic for less than one second.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	
Default:	Off. This test is run automatically on bootup. On-demand test also supported.
Initial Release:	15.1(2)SY
Corrective action:	Displays the “%CONST_DIAG-3-BOOTUP_TEST_FAIL: Module <i>number</i> : TesFexModeLoopback failed on port(s) <i>range</i> ” message.
Hardware support:	WS-X6904-40G-2T

TestIntPortLoopback

This test uses the switching module internal port to run a non-disruptive loopback test. It can be used to detect fabric channel failure and also port ASIC failure. This test is similar to TestFabricCh0Health. The test runs every 15 seconds.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not turn this test off. Use as a health-monitoring test.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	The module resets after 10 consecutive failures. Three consecutive resets powers down the module.
Hardware support:	WS-X6148E-GE-45AT, WS-X6148A-GE-TX, WS-X6148A-GE-45AF, WS-X6148-FE-SFP, WS-X6148A-RJ-45, WS-X6148A-45AF.

TestL3TcamMonitoring

This test verifies Layer 3 packet switching and monitors the health of both FIB and CL TCAM using the diagnostic lookup key. This test is nondisruptive and runs periodically every 15 seconds. Ten consecutive failures are treated as fatal, which cause the module to reload during runtime.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	After the module has ten consecutive failures, the module reloads during runtime.
Hardware support:	Supervisor engines.

TestLtlFpoeMemoryConsistency

This test verifies that the LTL and FPOE memories are working properly. The test runs every 15 seconds. Self-correction is applied if an error is detected. If self-correction fails, corrective action is triggered to reset the module. The module is powered-down on the third consecutive module reset. If self-correction passes, no action is taken. If too many self-corrections occur within a short period of time (more than three self-corrections in less than 300 seconds), the module is reset.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Failure of this test causes the module to reset and power down after two resets.
Hardware support:	All modules including supervisor engines.

TestMacNotification

This test verifies the data and control path between DFC-equipped modules and supervisor engines. This test also ensures Layer 2 MAC address consistency across Layer 2 MAC address tables. The test runs every six seconds. Ten consecutive failures causes the module to reset during bootup or runtime (default). After three consecutive resets, the module powers down. This test runs every 15 seconds.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Reset the module. After the module has ten consecutive failures or three consecutive resets, it powers down.
Hardware support:	DFC-equipped modules.

TestPortTxMonitoring

This test periodically polls the transmit counters on each port. The test displays a syslog message and error disables the port if no activity is seen for the configured time interval and failure threshold. You configure the time interval and threshold by entering the **diagnostic monitor interval** and **diagnostic monitor threshold** commands. The test does not source any packets, but leverages the CDP protocol that transmits packets periodically. If the CDP protocol is disabled, the polling for that port is not performed. The test runs every 75 seconds, and the failure threshold is set to five by default.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Display a syslog message indicating the port(s) that failed. Error disable the port(s) that failed.
Hardware support:	All modules including supervisor engines.

TestScratchRegister

This test monitors the health of application-specific integrated circuits (ASICs) by writing values into registers and reading back the values from these registers. The test runs every 30 seconds. Five consecutive failures causes a supervisor engine to switchover (or reset), if you are testing the supervisor engine, or in the module powering down when testing a module.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Reset the malfunctioning supervisor engine or power down the module.
Hardware support:	Active and standby supervisor engine, DFC-equipped modules, WS-X6148A-GE-TX, WS-X6148A-GE-45AF, WS-X6148-FE-SFP, WS-X6148A-RJ-45, WS-X6148A-45AF.

TestSnrMonitoring

This test monitors the SNR (signal-to-noise ratio) margin for a port, which varies between -12.7 dB to +12.7 dB. The test uses the following two threshold levels to compare SNR:

- Minor threshold at +1.0 dB
- Major threshold at 0.0 dB

When the SNR value drops below the minor threshold, the test logs a minor warning message. When the SNR value drops below the major threshold, the test logs a major warning message. Similarly, recovery messages are logged when SNR recovers the two threshold levels. The default interval for the test is 30 seconds and can be configured to as low as 10 seconds for faster monitoring. The TestSnrMonitoring is not a bootup test and cannot be run on demand.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	None.
Hardware support:	WS-X6816-10T-2T, WS-X6716-10T.

TestUnusedPortLoopback

This test periodically verifies the data path between the supervisor engine and the network ports of a module in the runtime. In this test, a Layer 2 packet is flooded onto the VLAN associated with the test port and the inband port of the supervisor engine. The packet loops back in the test port and returns to the supervisor engine on the same VLAN. This test is similar to TestLoopback but only runs on unused (admin down) network ports and on only one unused port per port ASIC. This test substitutes the lack of a nondisruptive loopback test in current ASICs. This test runs every 60 seconds.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable. This test is automatically disabled during CPU-usage spikes to maintain accuracy.
Default:	On.
Initial Release:	
Corrective action:	Display a syslog message indicating the port(s) that failed. For modules other than the supervisor engines, if all port groups fail (for example, at least one port per port ASIC fails more than the failure threshold for all port ASICs), the default action is to reset the module and power down the module after two resets.
Hardware support:	All modules including the supervisor engines.

Per-Port Tests

- [TestActiveToStandbyLoopback](#), page 1-9
- [TestCCPLoopback](#), page 1-9
- [TestDataPortLoopback](#), page 1-10
- [TestDCPLoopback](#), page 1-10
- [TestL2CTSLoopback](#), page 1-11
- [TestL3CTSLoopback](#), page 1-11
- [TestLoopback](#), page 1-12
- [TestMediaLoopback](#), page 1-12
- [TestMgmtPortsLoopback](#), page 1-13
- [TestNetflowInlineRewrite](#), page 1-13
- [TestNonDisruptiveLoopback](#), page 1-14
- [TestNPLoopback](#), page 1-14
- [TestTransceiverIntegrity](#), page 1-15

TestActiveToStandbyLoopback

This test verifies the data path between the active supervisor engine and the network ports of the standby supervisor engine. In this test, a Layer 2 packet is flooded onto a VLAN that consists of only the test port and the active supervisor engine's inband port. The test packets are looped back in the targeted port and are flooded back onto the bus with only the active supervisor engine's inband port listening in on the flooded VLAN.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the loopback port (for example, Spanning Tree Protocol).
Recommendation:	Schedule during downtime.
Default:	Runs at bootup or after OIR.
Initial Release:	15.0(1)SY.
Corrective action:	Error disable a port if the loopback test fails on the port. Reset the standby supervisor engine if all of the ports fail.
Hardware support:	Standby supervisor engine only.

TestCCPLoopback

This test checks the control plane data path. This test sends an online diagnostics packet from the supervisor engine to service or high availability port on the Wireless Services Module (WiSM2). The TestCCPLoopback checks whether the test packet loops back. If the test fails, a syslog message is displayed to indicate the error. This test also can be run as health monitoring, on-demand, and scheduled tests.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY1.
Corrective action:	A syslog message is displayed after five consecutive failures.
Hardware support:	WS-SVC-WISM2-K9.

TestDataPortLoopback

This test sends a packet from the inband port of the supervisor to the data port on the Firewall or NAM service module to verify the data packet path. The packet is looped back to the supervisor in hardware. If the packet does not return from the supervisor, hardware counters are polled to isolate the faulty path. This test runs every 45 seconds.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable. If the test fails for 10 consecutive times, the module is reset. If the test fails persistently, the module is powered down.
Default:	On.
Initial Release:	15.0(1)SY1.
Corrective action:	None.
Hardware support:	WS-SVC-ASA-SM1-K9 and WS-SVC-NAM3-6G-K9.

TestDCPLoopback

This test checks the data plane data path. This test sends an online diagnostics packet from the supervisor engine to data ports on the Wireless Services Module (WiSM2). This test checks whether the test packet loops back. If the test fails, a syslog message is displayed to indicate the error. This test also can be run as health monitoring, on-demand, and scheduled tests.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY1.
Corrective action:	A syslog message is displayed after five consecutive failures.
Hardware support:	WS-SVC-WISM2-K9.

TestL2CTSLoopback

This test provides encapsulation for Layer 2 Ethernet packets sent from the supervisor engine inband port to each port inside the Ganita ASIC. The test sends back the Layer 2 Ethernet packet to the supervisor engine inband port after decapsulation with its original content.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test is run automatically on bootup. On-demand test also supported.
Default:	Off. This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines.

TestL3CTSLoopback

This test provides encapsulation for Layer 3 IPv4 packets sent from the supervisor engine inband port to each port inside the Ganita ASIC and sends back the Layer 3 IPv4 packet to the supervisor engine inband port after decapsulation with its original content.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test is run automatically on bootup. On-demand test also supported.
Default:	Off. This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines.

TestLoopback

This test verifies the data path between the supervisor engine and the network ports of a module. In this test, a Layer 2 packet is flooded onto a VLAN that consists of only the test port and the supervisor engine's inband port. The packet loops back in the port and returns to the supervisor engine on that same VLAN.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of looped-back port (for example, Spanning Tree Protocol).
Recommendation:	Schedule during downtime.
Default:	Runs at bootup or after online insertion and removal (OIR).
Initial Release:	15.0(1)SY.
Corrective action:	Error disable a port if the loopback test fails on the port. Reset the module if all of the ports fail.
Hardware support:	All modules including supervisor engines.

TestMediaLoopback

This test verifies the data path of MediaNet-like traffic. Index direct UDP packets are sent out to the MediaNet interface under test. The packets are looped back and forwarded to the inband port of the module.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Do not disable.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	WS-X6816-10T-2T, WS-X6716-10T, WS-X6816-10G-2T, WS-X6716-10GE, WS-X6848-SFP-2T, WS-X6748-SFP, WS-X6824-SFP-2T, WS-X6724-SFP, WS-X6848-TX-2T, WS-X6748-GE-TX.

TestMgmtPortsLoopback

This test sends a packet from the inband port of the supervisor to the Firewall or NAM service module to verify the health of the backplane ports. The packet is looped back to the supervisor in hardware. If the packet does not return from the supervisor, the service application is queried for the status of the packet and depending on the action suggested by the service module, a syslog message is displayed and the module is reset. This test runs every 30 seconds.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable. If the failure is isolated to the firewall module, then a syslog is printed indicating which port failed the test. If the test fails due to any other datapath issue for 10 consecutive times, the module is reset. If the test fails persistently, the module is powered down.
Default:	On.
Initial Release:	15.0(1)SY1.
Corrective action:	None.
Hardware support:	WS-SVC-ASA-SM1-K9 and WS-SVC-NAM3-6G-K9.

TestNetflowInlineRewrite

This test verifies the NetFlow lookup operation, the ACL permit and deny functionality, and the inline rewrite capabilities of the port ASIC. The test packet will undergo a NetFlow table lookup to obtain the rewrite information. The VLAN and the source and destination MAC addresses are rewritten when the packet reaches the targeted port.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on configuration of loopback port (for example, Spanning Tree Protocol).
Recommendation:	Schedule during downtime. Run this test during bootup only.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	All modules including supervisor engines.

TestNonDisruptiveLoopback

This test verifies the data path between the supervisor engine and the network ports of a module. In this test, a Layer 2 packet is flooded onto VLAN that contains a group of test ports. The test port group consists of one port per port ASIC channel. Each port in the test port group nondisruptively loops back the packet and directs it back to the supervisor engine's inband port. The ports in the test port group are tested in parallel.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Error disable a port after 10 consecutive failures. Error disable a channel if all of its ports failed the test in one test cycle. Reset the module after a failure of all channels.
Hardware support:	WS-X6148-FE-SFP, WS-X6148A-GE-TX, WS-X6148A-RJ-45.

TestNPLoopback

This test checks the data path of the ACE30 module for data path errors. This test runs at bootup, and the default configuration is a health-monitoring test that runs every 15 seconds. If TestNPLoopback fails, an SCP (Switch-module Configuration Protocol) message is sent to the ACE30 module indicating which network processors have failed. Upon receipt of the SCP message, ACE30 will take corrective action. If the TestNPLoopback test fails for ten consecutive times, the ACE30 module is reset.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY1.
Corrective action:	A syslog message is displayed to inform the ACE30 about the port(s) that failed the test on the failure code. Depending on the failure code, the ACE30 decides whether to take corrective action or not. The suggested action for ACE30 is to collect core dumps from all network processors and reset the ACE30 module.
Hardware support:	ACE30-MOD-K9.

TestTransceiverIntegrity

This security test is performed on the transceiver during transceiver online insertion and removal (OIR) or module bootup to make sure that the transceiver is supported.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Not applicable.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	Error disable the port.
Hardware support:	All modules with transceivers.

PFC Layer 2 Tests

- [TestBadBpduTrap](#), page 1-15
- [TestDontConditionalLearn](#), page 1-16
- [TestMatchCapture](#), page 1-16
- [TestNewIndexLearn](#), page 1-17

TestBadBpduTrap

This test is a combination of the TestTrap and the TestBadBpdu tests, which are described in the “[DFC Layer 2 Tests](#)” section on page 1-17.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	If you experience problems with the Layer 2 forwarding engine learning capability, run this test on-demand to verify the Layer 2 learning functionality. This test can also be used as a health-monitoring test.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines only.

TestDontConditionalLearn

This test is a combination of the TestDontLearn and the TestConditionalLearn tests, which are described in the [“DFC Layer 2 Tests” section on page 1-17](#).

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	If you experience problems with the Layer 2 forwarding engine learning capability, run this test on-demand to verify the Layer 2 learning functionality. This test can also be used as a health monitoring test.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestMatchCapture

This test is a combination of the TestProtocolMatchChannel and the TestCapture tests, which are described in the [“DFC Layer 2 Tests” section on page 1-17](#).

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Run this test on-demand to verify the Layer 2 learning functionality. This test can also be used as a health-monitoring test.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines only.

TestNewIndexLearn

This test is a combination of the TestNewLearn and the TestIndexLearn tests, which are described in the “DFC Layer 2 Tests” [section on page 1-17](#).

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	If you experience problems with the Layer 2 forwarding engine learning capability, run this test on-demand to verify the Layer 2 learning functionality. This test can also be used as a health-monitoring test.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines only.

DFC Layer 2 Tests

- [TestBadBpdu, page 1-18](#)
- [TestCapture, page 1-18](#)
- [TestConditionalLearn, page 1-19](#)
- [TestDontLearn, page 1-19](#)
- [TestIndexLearn, page 1-20](#)
- [TestNewLearn, page 1-20](#)
- [TestPortSecurity, page 1-21](#)
- [TestProtocolMatchChannel, page 1-21](#)
- [TestStaticEntry, page 1-22](#)
- [TestTrap, page 1-22](#)

TestBadBpdu

This test verifies the ability to trap or redirect packets to the switch processor. This test verifies that the Trap feature of the Layer 2 forwarding engine is working properly. When running the test on the supervisor engine, the diagnostic packet is sent from the supervisor engine's inband port and performs a packet lookup using the supervisor engine's Layer 2 forwarding engine. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The BPDU feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestCapture

This test verifies that the capture feature of Layer 2 forwarding engine is working properly. The capture functionality is used for multicast replication. When running the test on the supervisor engine, the diagnostic packet is sent from the supervisor engine's inband port and performs a packet lookup using the supervisor engine's Layer 2 forwarding engine. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The Capture feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	Schedule during downtime.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestConditionalLearn

This test verifies the ability to learn a Layer 2 source MAC address under specific conditions. When running the test on the supervisor engine, the diagnostic packet is sent from the supervisor engine's inband port and performs a packet lookup using the supervisor engine Layer 2 forwarding engine. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The Conditional Learn feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestDontLearn

This test verifies that new source MAC addresses are not populated in the MAC address table when they should not be learned. This test verifies that the "don't learn" feature of the Layer 2 forwarding engine is working properly. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine inband port through the switch fabric and looped back from one of the ports on the DFC-enabled module. The "don't learn" feature is verified during diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	Schedule during downtime.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestIndexLearn

This test ensures that existing MAC address table entries can be updated. This test verifies the Index Learn feature of the Layer 2 forwarding engine is working properly. When running the test on the supervisor engine, the diagnostic packet is sent from the supervisor engine's inband port and performs a packet lookup using the supervisor engine Layer 2 forwarding engine. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The Index Learn feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestNewLearn

This test verifies the Layer 2 source MAC address learning functionality of the Layer 2 forwarding engine. For supervisor engines, a diagnostic packet is sent from the supervisor engine inband port to verify that the Layer 2 forwarding engine is learning the new source MAC address from the diagnostic packet. For DFC-equipped modules, a diagnostic packet is sent from the supervisor engine inband port through the switch fabric and looped backed from one of the ports on the DFC-enabled module. The Layer 2 learning functionality is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestPortSecurity

This test verifies the ability to redirect packets to the CPU if a secure MAC address is transmitting the packets from a different port. For the supervisor engine, a diagnostic packet is sent from the supervisor engine's inband port and the port security feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine. For DFC-equipped modules, a diagnostic packet is sent from the supervisor engine inband port through the fabric and is looped back in one of the ports on the DFC-equipped module. The port security feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	None.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engine and DFC-equipped modules.

TestProtocolMatchChannel

This test verifies the ability to match specific Layer 2 protocols in the Layer 2 forwarding engine. When running the test on the supervisor engine, the diagnostic packet is sent from the supervisor engine's inband port and performs a packet lookup using the supervisor engine's Layer 2 forwarding engine. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The Match feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestStaticEntry

This test verifies the ability to populate static entries in the Layer 2 MAC address table. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The Static Entry feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestTrap

This test verifies the ability to trap or redirect packets to the switch processor. This test verifies that the Trap feature of the Layer 2 forwarding engine is working properly. When running the test on the supervisor engine, the diagnostic packet is sent from the supervisor engine's inband port and performs a packet lookup using the supervisor engine's Layer 2 forwarding engine. For DFC-equipped modules, the diagnostic packet is sent from the supervisor engine's inband port through the switch fabric and looped back from one of the DFC ports. The Trap feature is verified during the diagnostic packet lookup by the Layer 2 forwarding engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

PFC Layer 3 Tests

- [TestAclDeny, page 1-23](#)
- [TestAclPermit, page 1-23](#)
- [TestAclRedirect, page 1-24](#)

- [TestDQUP](#), page 1-24
- [TestInbandEdit](#), page 1-25
- [TestIPv4FibShortcut](#), page 1-25
- [TestIPv6FibShortcut](#), page 1-26
- [TestL3Capture2](#), page 1-26
- [TestMPLSFibShortcut](#), page 1-27
- [TestNATFibShortcut](#), page 1-27
- [TestNetflowShortcut](#), page 1-28
- [TestRBAcl](#), page 1-28

TestAclDeny

This test verifies that the ACL deny feature of the Layer 2 and Layer 3 forwarding engine is working properly. The test uses different ACL deny scenarios such as input, output, Layer 2 redirect, Layer 3 redirect, and Layer 3 bridges to determine whether or not the ACL deny feature is working properly.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Run this test on-demand.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Automatic ASIC reset for recovery.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestAclPermit

This test verifies that the ACL permit functionality is working properly. An ACL entry permitting a specific diagnostics packet is installed in the ACL TCAM. The corresponding diagnostic packet is sent from the supervisor engine and looked up by the Layer 3 forwarding engine to make sure that it hits the ACL TCAM entry and gets permitted and forwarded appropriately.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Run this test on-demand.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestAclRedirect

This test verifies the ACL redirect feature of the Layer 3 forwarding engine. This test verifies both ingress and egress Layer 3 redirect.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Run this test on-demand.
Default:	Off. This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestDQUP

This test verifies whether the DQUP and PUP packets can be generated when diagnostic packets hit QoS entry. This test receives the DQUP and PUP packets and ensures that the information in DQUP and PUP is correct.

Attributes	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Run this test on-demand if you suspect that DQUP and PUP is not working properly.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestInbandEdit

This test verifies the InbandEdit packets of the Layer 3 forwarding engine. One diagnostic InbandEdit packet is sent to create one diagnostic NetFlow entry and an adjacency entry, and one diagnostic packet is sent to ensure that the InbandEdit packet is forwarded according to the rewritten MAC and VLAN.

Attributes	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Test runs automatically on bootup. On-demand is also supported.
Default:	This test runs by default during bootup.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestIPv4FibShortcut

This test does the following:

- Verifies whether the IPv4 FIB forwarding of the Layer 3 forwarding engine is working properly. One diagnostic IPv4 FIB and an adjacency entry are installed, and a diagnostic packet is sent to make sure that the diagnostic packet is forwarded according to rewritten MAC and VLAN information.
- Verifies whether the FIB TCAM and adjacency devices are functional. One FIB entry is installed on each FIB TCAM device. A diagnostic packet is sent to make sure that the diagnostic packet is switched by the FIB TCAM entry installed on the TCAM device. This is not an exhaustive TCAM device test; only one entry is installed on each TCAM device.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Run this test on-demand to verify the Layer 3 forwarding functionality if you experience problems with the routing capability. This test can also be used as a health-monitoring test.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestIPv6FibShortcut

This test verifies that the IPV6 FIB forwarding of the Layer 3 forwarding engine is working properly. One diagnostic IPV6 FIB and an adjacency entry is installed, and a diagnostic IPv6 packet is sent to make sure the diagnostic packet is forwarded according to rewritten MAC and VLAN information.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Run this test on-demand to verify the Layer 3 forwarding functionality if you experience problems with the routing capability. This test can also be used as a health-monitoring test.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestL3Capture2

This test verifies that the Layer 3 capture (capture 2) feature of the Layer 3 forwarding engine is working properly. This capture feature is used for ACL logging and VACL logging. One diagnostic FIB and an adjacency entry with a capture 2 bit set is installed, and a diagnostic packet is sent to make sure that the diagnostic packet is forwarded according to the capture bit information.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test can also be used as a health-monitoring test. Use as a health-monitoring test if you are using ACL or VACL logging.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestMPLSFibShortcut

This test does the following:

- Verifies that the MPLS forwarding of the Layer 3 forwarding engine is working properly. One diagnostic MPLS FIB and an adjacency entry is installed, and a diagnostic MPLS packet is sent to make sure that the diagnostic packet is forwarded according to the MPLS label from the adjacency entry.
- Verifies the EoMPLS forwarding of the Layer 3 forwarding engine. One diagnostic EoMPLS Layer 2 FIB and an adjacency entry are installed and a diagnostic Layer 2 packet is sent to the forwarding engine to make sure it is forwarded accordingly with the MPLS labels and the encapsulated Layer 2 packet.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test can also be used as a health-monitoring test. Use as a health-monitoring test if you are routing MPLS traffic.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestNATFibShortcut

This test verifies the ability to rewrite a packet based on the NAT adjacency information (rewrite destination IP address). One diagnostic NAT FIB and an adjacency entry is installed, and the diagnostic packet is sent to make sure that the diagnostic packet is forwarded according to the rewritten IP address.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test can also be used as a health-monitoring test. Use as a health-monitoring test if the destination IP address is being rewritten (for example, if you are using NAT).
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestNetflowShortcut

This test verifies that the NetFlow forwarding functionality of the Layer 3 forwarding engine is working properly. One diagnostic NetFlow entry and an adjacency entry is installed, and a diagnostic packet is sent to make sure it is forwarded according to the rewritten MAC and VLAN information.

Attributes	Description
Disruptive or Nondisruptive:	Disruptive for looped back ports. The disruption is 500 ms.
Recommendation:	Run this test on-demand if you suspect that NetFlow is not working properly.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestRBACL

This test verifies the role based ACL (RBACL) feature of the Layer 3 forwarding engine. This test uses SGT and DGT instead of src_ip/dest_ip to get ACL lookup results.

Attributes	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Test runs automatically on bootup. On-demand and health-monitoring tests are also supported.
Default:	This test runs by default during bootup.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

DFC Layer 3 Tests

- [TestAclDeny, page 1-29](#)
- [TestAclPermit, page 1-29](#)
- [TestAclRedirect, page 1-30](#)
- [TestInbandEdit, page 1-30](#)
- [TestIPv4FibShortcut, page 1-31](#)
- [TestIPv6FibShortcut, page 1-31](#)
- [TestL3Capture2, page 1-32](#)
- [TestMPLSFibShortcut, page 1-32](#)
- [TestNATFibShortcut, page 1-33](#)

- [TestNetflowShortcut](#), page 1-33
- [TestRBACL](#), page 1-34

TestAcIDeny

This test verifies that the ACL deny feature of the Layer 2 and Layer 3 forwarding engine is working properly. The test uses different ACL deny scenarios such as input and output Layer 2 redirect, Layer 3 redirect, and Layer 3 bridges.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	Schedule during downtime if you are using ACLs.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestAcIPermit

This test verifies that the ACL permit functionality is working properly. An ACL entry permitting a specific diagnostics packet is installed in the ACL TCAM. The corresponding diagnostic packet is sent from the supervisor engine and is looked up by the Layer 3 forwarding engine to make sure it hits the ACL TCAM entry and gets permitted and forwarded correctly.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestAclRedirect

This test verifies the ACL redirect feature of the Layer 3 forwarding engine. This test verifies both ingress and egress Layer 3 redirect.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Run this test on-demand.
Default:	Off. This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestInbandEdit

This test verifies the InbandEdit packets of the Layer 3 forwarding engine. One diagnostic InbandEdit packet is sent to create one diagnostic NetFlow entry and an adjacency entry, and one diagnostic packet is sent to ensure that the InbandEdit packet is forwarded accordingly with the rewritten MAC and VLAN.

Attributes	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Test runs automatically on bootup. On-demand is also supported.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

TestIPv4FibShortcut

These tests do the following:

- Verifies whether the IPv4 FIB forwarding of the Layer 3 forwarding engine is working properly. One diagnostic IPv4 FIB and an adjacency entry is installed, and a diagnostic packet is sent to make sure that the diagnostic packet is forwarded according to rewritten MAC and VLAN information.
- Verifies whether the FIB TCAM and adjacency devices are functional. One FIB entry is installed on each FIB TCAM device. A diagnostic packet is sent to make sure that the diagnostic packet is switched by the FIB TCAM entry installed on the TCAM device. This is not an exhaustive TCAM device test; only one entry is installed on each TCAM device.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestIPv6FibShortcut

This test verifies that the IPv6 FIB forwarding functionality of the Layer 3 forwarding engine is working properly. One diagnostic IPv6 FIB and an adjacency entry is installed, and a diagnostic IPv6 packet is sent to make sure that the diagnostic packet is forwarded according to rewritten MAC and VLAN information.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestL3Capture2

This test verifies that the Layer 3 capture (capture 2) feature of the Layer 3 forwarding engine is working properly. This capture feature is used for ACL logging and VACL logging. One diagnostic FIB and an adjacency entry with a capture 2-bit set is installed, and a diagnostic packet is sent to make sure that the diagnostic packet is forwarded according to capture bit information.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestMPLSFibShortcut

This test does the following:

- Verifies that the MPLS forwarding of the Layer 3 forwarding engine is working properly. One diagnostic MPLS FIB and an adjacency entry is installed, and a diagnostic MPLS packet is sent to make sure that the diagnostic packet is forwarded according to the MPLS label from the adjacency entry.
- Verifies the EoMPLS forwarding of the Layer 3 forwarding engine. One diagnostic EoMPLS Layer 2 FIB and an adjacency entry are installed and a diagnostic Layer 2 packet is sent to the forwarding engine to make sure it is forwarded accordingly with the MPLS labels and the encapsulated Layer 2 packet.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestNATFibShortcut

This test verifies the ability to rewrite a packet based on NAT adjacency information, such as the rewrite destination IP address. One diagnostic NAT FIB and an adjacency entry is installed, and a diagnostic packet is sent to the forwarding engine to make sure the diagnostic packet is forwarded according to the rewritten IP address.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second. Duration of the disruption depends on the configuration of the looped-back port (for example, Spanning Tree Protocol).
Recommendation:	This test runs by default during bootup or after a reset or OIR.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestNetflowShortcut

This test verifies that the NetFlow forwarding functionality of the Layer 3 forwarding engine is working properly. One diagnostic NetFlow entry and an adjacency entry is installed, and a diagnostic packet is sent to make sure it is forwarded according to the rewritten MAC and VLAN information.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for looped-back ports. Disruption is typically less than one second.
Recommendation:	Run this test on-demand if you suspect that NetFlow is not working properly.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestRBACl

This test verifies the role based ACL (RBACL) feature of the Layer 3 forwarding engine. This test uses SGT and DGT instead of src_ip/dest_ip to get ACL lookup results.

Attributes	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Test runs automatically on bootup. On-demand and health-monitoring tests are also supported.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	DFC-equipped modules.

Replication Engine Tests

- [TestEgressSpan, page 1-34](#)
- [TestIngressSpan, page 1-35](#)
- [TestL3VlanMet, page 1-35](#)

TestEgressSpan

This test verifies that the egress SPAN replication functionality of the rewrite engine for both SPAN queues is working properly.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for both SPAN sessions. Disruption is typically less than one second.
Recommendation:	Run this test on-demand.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines, DFC-equipped modules.

TestIngressSpan

This test ensures that the port ASIC is able to tag packets for ingress SPAN. This test also verifies that the ingress SPAN operation of the rewrite engine for both SPAN queues is working properly.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive for both SPAN sessions. Also disruptive for the loopback port on modules. Duration of the disruption depends on the configuration of the loopback port (for example, Spanning Tree Protocol).
Recommendation:	Run this test on-demand.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestL3VlanMet

This test verifies that the multicast functionality of the replication engine is working properly. The replication engine is configured to perform multicast replication of a diagnostic packet onto two different VLANs. After the diagnostic packet is sent out from the supervisor engine's inband port, the test verifies that two packets are received back in the inband port on the two VLANs configured in the replication engine.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive for supervisor engines. Disruptive for DFC-equipped modules. Disruption is typically less than one second on looped-back ports.
Recommendation:	Run this test on-demand to test the multicast replication abilities of the replication engine.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

Fabric Tests

- [TestFabricCh0Health](#), page 1-36
- [TestFabricCh1Health](#), page 1-36
- [TestFabricExternalSnake](#), page 1-37
- [TestFabricFlowControlStatus](#), page 1-37
- [TestFabricInternalSnake](#), page 1-38

- [TestFabricVlanLoopback](#), page 1-38
- [TestFexFabricLinkStatus](#), page 1-39
- [TestSynchedFabChannel](#), page 1-39

TestFabricCh0Health

This test constantly monitors the health of the ingress and egress data paths for fabric channel 0 on 10-gigabit modules. The test runs every five seconds. Ten consecutive failures are treated as fatal and the module resets; three consecutive reset cycles may result in a fabric switchover.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not turn this test off. Use as a health-monitoring test.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	The module resets after 10 consecutive failures. Three consecutive resets powers down the module.
Hardware support:	WS-X6704-10GE.

TestFabricCh1Health

This test constantly monitors the health of the ingress and egress data paths for fabric channel 1 on 10-gigabit modules. The test runs every five seconds. Ten consecutive failures are treated as fatal and the module resets; three consecutive reset cycles might result in a fabric switchover.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not turn this test off. Use as a health-monitoring test.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	The module resets after 10 consecutive failures. Three consecutive failures resets powers down the module.
Hardware support:	WS-X6704-10GE module.

TestFabricExternalSnake

This test is executed for the chassis-active supervisor engine only during regular OIR bootup diagnostic testing stage. This test generates the test packet through the inband port of the supervisor engine and the test data path involves the port ASIC, the rewrite engine ASIC inside the supervisor engine, and the fabric ASIC.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	After the supervisor engine is online, this test is supported to run by on-demand diagnostics for chassis-active supervisor engine but not for chassis-standby supervisor engine.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	If the test fails, the supervisor engine is power reset because it is a major diagnostic error.
Hardware support:	Active Supervisor engines.

TestFabricFlowControlStatus

This test reads the switch fabric ASIC registers to detect flow-control status for each fabric channel. Flow-control events are logged into the diagnostic events queue. By default, this test is disabled as a health-monitoring test, and when enabled, this test runs every 15 seconds. This test reports per-slot or per-channel rate reduction, current fabric channel utilization, peak fabric-channel utilization, and SP CPU utilization in both ingress and egress directions.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Use as a health-monitoring test. Use this test when you suspect a problem with the fabric channel.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Flow control events are logged into the diagnostic event log.
Hardware support:	Supervisor engines.

TestFabricInternalSnake

This test is supported for a supervisor engine and module with a fabric switching ASIC. This test is executed by firmware INIT sequence code during bootup time when the firmware initializes the entire fabric ASIC.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	After the supervisor engine is online, this test is supported to run by on-demand diagnostics for chassis-standby supervisor engine but not for chassis-active supervisor engine. For modules with a fabric switching ASIC, this test is supported only for bootup diagnostic.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	If the test fails, firmware INIT sequence fails and the supervisor engine or module under test is power reset.
Hardware support:	Supervisor engines and fabric-enabled module.

TestFabricVlanLoopback

This test verifies the data path between the inband port of the module under test and the local fabric port responsible for switching traffic from and to the inband port through the per queue VLAN loopback feature provided by the hardware. When the test packet from the inband port arrives at the input queue of the local fabric port with matching VLAN to the pre-programmed VLAN loopback register, the test packet transverses the fabric, loops back to the output queue of the same fabric port, and forwards the test packet back to the inband port.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Test runs automatically on bootup. On-demand is also supported. Use this test to verify the data path between the local fabric channel and the inband port or use for debugging.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	Supervisor engines and DFC-equipped modules.

TestFexFabricLinkStatus

This test monitors the IA parent-client link and reports failure if the link port loses its connection to the associated IA client.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive health monitoring test.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.1(2)SY
Corrective action:	Displays the “%CONST_DIAG-6-RSL_LINK_DOWN: RSL link on module <i>number</i> port <i>number</i> is down syslog” message.
Hardware support:	WS-X6904-40G-2T

TestSynchedFabChannel

This test periodically checks the fabric synchronization status for both the module and the fabric. This test is available only for fabric-enabled modules. This test is not a packet-switching test so it does not involve the data path. This test sends an SCP control message to the module and fabric to query the synchronization status.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not turn this test off. Use as a health-monitoring test.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	The module resets after five consecutive failures. Three consecutive reset cycles results in the module powering down. A fabric switchover may be triggered, depending on the type of failure.
Hardware support:	All fabric-enabled modules.

Exhaustive Memory Tests

- [TestAclQosTcam](#), page 1-40
- [TestAsicMemory](#), page 1-40
- [TestEarlMemOnBootup](#), page 1-41



Note

Because the supervisor engine must be rebooted after running memory tests, run memory tests on the other modules before running them on the supervisor engine. For more information about running on-demand online diagnostic tests see the “[Configuring On-Demand Online Diagnostics](#)” section on page 16-4.

TestAclQosTcam

This test tests all the bits and checks the location of both ACL and QoS TCAMs on the PFC.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive. Disruption is approximately one hour.
Recommendation:	Use this test only if you suspect a problem with the hardware or before putting the hardware into a live network. Do not run any traffic in the background on the module that you are testing. The supervisor engine must be rebooted after running this test.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	All modules including supervisor engines.

TestAsicMemory

This test uses an algorithm to test the memory on a module.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive. Disruption is approximately one hour.
Recommendation:	Use this test only if you suspect a problem with the hardware or before putting the hardware into a live network. Do not run any traffic in the background on the module that you are testing. The supervisor engine must be rebooted after running this test.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	All modules including supervisor engines.

TestEarlMemOnBootup

This test runs on bootup and tests all the bits and locations of EARL memories supported by the Generic Memory Testing Logic (GMTL). EARL memories are tested by drivers during bootup initialization. This test retrieves and displays the bootup test results from the drivers.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Not applicable.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Supervisor engines and DFC-equipped modules.

Service Module Tests

- [TestPcLoopback, page 1-41](#)
- [TestPortASICLoopback, page 1-42](#)

TestPcLoopback

This test verifies the longest datapath between the supervisor and the NAM service module. A packet is sent from the supervisor to the module and is looped back by the PC to the supervisor engine.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test runs automatically during bootup.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	WS-SVC-NAM-1, WS-SVC-NAM-2.

TestPortASICLoopback

This test verifies the health of the ASIC ports on the NAM service module. A packet is sent from the supervisor engine and looped back at the ASIC.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test runs automatically during bootup.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	WS-SVC-NAM-1, WS-SVC-NAM-2.

Stress Tests

- [TestEobcStressPing](#), page 1-42
- [TestMicroburst](#), page 1-43
- [TestNVRAMBatteryMonitor](#), page 1-43
- [TestTrafficStress](#), page 1-44

TestEobcStressPing

This test stresses a module's EOBC link with the supervisor engine. The test is started when the supervisor engine initiates a number of sweep-ping processes (the default is one). The sweep-ping process pings the module with 20,000 SCP-ping packets. The test passes if all 20,000 packets respond before each packet-ping timeout, which is two seconds. If unsuccessful, the test allows five retries to account for traffic bursts on the EOBC bus during the test.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive. Disruption is several minutes.
Recommendation:	Use this test to qualify hardware before installing it in your network.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Supervisor engines.

TestMicroburst

This test monitors packet microbursts in the port ASICs and logs them to SEA unless consecutive failures reach the threshold.

Attribute	Description
Disruptive or Nondisruptive:	NonDisruptive.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Supervisor Engine 2T and DFC-equipped switching modules.

TestNVRAMBatteryMonitor

This test monitors the NVRAM battery status of the supervisor engine and logs every battery status change to OBFL. A syslog is printed if the battery voltage remains below a certain threshold for three consecutive days. Use the **show logging onboard** command and look for 'NVRAM battery power' to check change history of NVRAM battery status.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This is an automatic health monitor test that runs every hour. It does not run at bootup and cannot be run on-demand.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	This test monitors the NVRAM battery status of the supervisor engine and logs battery status changes to on-board failure logging (OBFL). In the show logging onboard command output, the “NVRAM battery power” information is the NVRAM battery status history. After 72 consecutive test failures, a syslog advises module replacement.
Hardware support:	Supervisor Engine 2T.

TestTrafficStress

This test stress tests the switch and the installed modules by configuring all of the ports on the modules into pairs, which then pass packets between each other. After allowing the packets to pass through the switch for a predetermined period, the test verifies that the packets are not dropped.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive. Disruption is several minutes.
Recommendation:	Use this test to qualify hardware before installing it in your network.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Supervisor engines.

General Tests

- [ScheduleSwitchover](#), page 1-44
- [TestCFRW](#), page 1-45
- [TestFirmwareDiagStatus](#), page 1-45
- [TestOBFL](#), page 1-45
- [TestRwEngineOverSubscription](#), page 1-46
- [TestVDB](#), page 1-46

ScheduleSwitchover

This test allows you to trigger a switchover at any time using the online diagnostics scheduling capability.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Schedule this test during downtime to test the ability of the standby supervisor engine to take over after a switchover.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	None
Hardware support:	Supervisor engines.

TestCFRW

This test verifies the CompactFlash disk or disks on the supervisor engine. This test is performed during system boot-up or whenever a disk is inserted. A 128-byte temporary file is written to each disk present in the slot and read back. The content read back is checked and the temporary file is deleted. You can also execute this test from the CLI.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not disable. No traffic is affected.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Format or replace the failed CompactFlash.
Hardware support:	Removable CompactFlash devices.

TestFirmwareDiagStatus

This test displays the results of the power-on diagnostic tests run by the firmware during the module bootup.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test can only be run at bootup.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide.
Hardware support:	All modules.

TestOBFL

This test verifies the on-board failure logging capabilities. During this test a diagnostic message is logged on the module.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test is run automatically during bootup and cannot be run on-demand.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Supervisor engines, DFC-equipped switching modules, WS-SVC-WISM2.

TestRwEngineOverSubscription

This is a health-monitoring test that is not enabled by default. This test runs on the module every one second and checks if the rewrite engine gets oversubscribed by retrieving drop counters and generates a syslog message if the drops exceed the set threshold.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test is run only as a health-monitoring test.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Supervisor engines, DFC-equipped modules.

TestVDB

This test is available on PoE-equipped modules. This test queries the result of diagnostic tests that run on the PoE daughter card.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	This test is run automatically during bootup.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	Modules with a PoE daughter card.

Critical Recovery Tests

- [TestTxPathMonitoring](#), page 1-47



Note

These tests are also considered critical recovery tests:

- [TestFabricCh0Health](#), page 1-36
- [TestFabricCh1Health](#), page 1-36
- [TestSynchedFabChannel](#), page 1-39

TestTxPathMonitoring

This test sends index-directed packets periodically to each port on the supervisor engine and supported modules to verify ASIC synchronization and correct any related problems. The test runs every two seconds.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive.
Recommendation:	Do not change the default settings.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable (self-recovering).
Hardware support:	Supervisor engines and DFC-equipped modules.

ViSN Tests

- [TestRslHm, page 1-47](#)
- [TestVSActiveToStandbyLoopback, page 1-48](#)
- [TestVslBridgeLink, page 1-48](#)
- [TestVslLocalLoopback, page 1-49](#)
- [TestVslStatus, page 1-49](#)

TestRslHm

This test monitors the data and control links between the remote switch and core switches. A diagnostic packet is sent from the supervisor engine inband port on the remote switch to the supervisor engine inband port on the core switch and is pinged back along the reverse data path. This tests each RSL link between the remote switch and both active and standby core switches.

Attribute	Description
Disruptive or Nondisruptive:	Nondisruptive health monitoring test.
Recommendation:	Do not disable.
Default:	On.
Initial Release:	15.0(1)SY.
Corrective action:	None. See the system message guide for more information.
Hardware support:	VSL-capable modules.

TestVSActiveToStandbyLoopback

This test is the only GOLD test that tests the full data path across the virtual switch links. This test selects an uplink port in the standby virtual switch supervisor engine as the loopback point and sends the VLAN flood packet from the active virtual switch supervisor engine inband port to the system. Due to the configuration of the FPOE and LTL VLAN flood region for all VSL modules and VSL interfaces in the active and standby virtual switch, the packet goes across VSL and arrives at the uplink port of the standby virtual switch supervisor engines, and loopbacks from there. The packet comes back to the inband port of the active supervisor engine due to the preconfiguration of FPOE and LTL in the standby and active virtual switches. In case of a test failure, the error check is executed for SP CPU, fabric flow control, and other errors in both active and standby virtual switches.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	Disable all health monitoring tests before executing this test. This test is run only for on-demand diagnostic testing.
Default:	Off.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	VSL-capable modules.

TestVslBridgeLink

This test provides diagnostic coverage for VSL-capable modules and the supervisor engine during module bootup. The data path of this test picks only one port corresponding to the local and remote bridge inband port as the loopback points. A diagnostic packet is sent from the inband port of the supervisor engine to the loopback points on the VSL module, and the packet traverses the bridge link between two fabric data path complexes to verify the hardware bridge link functionality.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test is run automatically during bootup.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	VSL-capable modules.

TestVslLocalLoopback

This test verifies the hardware functionality of each port on the VSL module before the VSL link interface is up. The data path of this test is constrained with the VSL module. A diagnostic packet is sent from the local inband port of the VSL module to each port to run a loopback test. This test is run only during module bootup.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test is run automatically during bootup and cannot be run on-demand.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	VSL-capable modules.

TestVslStatus

This test reports the status change detected by the VSLP protocol. When any link problem is detected by the VSLP protocol, the status of the link is changed and the result is updated accordingly. This test also triggers the loopback test to check the hardware status requested by the VSLP protocol.

Attribute	Description
Disruptive or Nondisruptive:	Disruptive.
Recommendation:	This test is effective once the VSL modules are online.
Default:	This test runs by default during bootup or after a reset or OIR.
Initial Release:	15.0(1)SY.
Corrective action:	Not applicable.
Hardware support:	VSL-capable modules.



Tip

For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

[Participate in the Technical Documentation Ideas forum](#)

INDEX

Numerics

4K VLANs (support for 4,096 VLANs) [26-2](#)

802.1AE Tagging [70-2](#)

802.1Q

Layer 2 protocol tunneling

See Layer 2 protocol tunneling

mapping to ISL VLANs [26-7](#)

trunks [21-4](#)

restrictions [21-2](#)

tunneling

configuration guidelines [29-1](#)

configuring tunnel ports [29-6](#)

overview [29-4](#)

802.1Q Ethertype

specifying custom [21-15](#)

802.1X [83-1](#)

802.1x accounting [83-43](#)

802.3ad

See LACP

802.3af [20-2](#)

802.3at [20-2](#)

802.3x Flow Control [10-9](#)

A

AAA

fail policy [83-8, 84-5](#)

AAA (authentication, authorization, and accounting). See also port-based authentication. [83-6, 84-2](#)

aaa accounting dot1x command [83-44](#)

aaa accounting system command [83-44](#)

abbreviating commands [2-6](#)

access, restricting MIB [86-10](#)

access control entries and lists [69-1](#)

access-enable host timeout (not supported) [69-4](#)

access port, configuring [21-14](#)

access rights [86-9](#)

access setup, example [86-11](#)

accounting

with 802.1x [83-43](#)

with IEEE 802.1x [83-16](#)

ACEs and ACLs [69-1](#)

ACLs

downloadable [84-2](#)

downloadable (dACLs) [83-24](#)

Filter-ID [83-25](#)

per-user [83-24](#)

port

defined [73-2](#)

redirect URL [83-25](#)

static sharing [83-25](#)

acronyms, list of [1-1](#)

activating lawful intercept [86-8](#)

admin function (mediation device) [86-7, 86-8](#)

administration, definition [86-6](#)

advertisements, VTP [25-4](#)

aggregate label [37-2, 37-5](#)

aggregate policing [63-4](#)

aging time

accelerated

for MSTP [31-45](#)

maximum

for MSTP [31-45, 31-46](#)

alarms

major [15-7](#)

minor [15-7](#)

Allow DHCP Option 82 on Untrusted Port

configuring [78-10](#)

understanding [78-5](#)

any transport over MPLS (AToM) [39-3](#)

Ethernet over MPLS [39-3](#)

ARP ACL [69-12](#)

ARP spoofing [80-3](#)

AToM [39-3](#)

audience [1-i](#)

authentication control-direction command [83-53](#)
 authentication event command [83-45](#)
 authentication failed VLAN
 See restricted VLAN
 authentication open comand [83-15](#)
 authentication password, VTP [25-5](#)
 authentication periodic command [83-38, 83-50](#)
 authentication port-control command [83-45](#)
 authentication timer reauthenticate command [83-38](#)
 authorized ports with 802.1X [83-12](#)
 auto enablement [83-30](#)
 automatic QoS
 configuration guidelines and restrictions [66-2](#)
 macros [66-4](#)
 overview [66-2](#)
 AutoQoS [66-1](#)

B

BackboneFast
 See STP BackboneFast
 backup interfaces
 See Flex Links
 binding database, DHCP snooping
 See DHCP snooping binding database
 binding table, DHCP snooping
 See DHCP snooping binding database
 blocking state, STP [31-8](#)
 BPDU
 RSTP format [31-16](#)
 BPDU guard
 See STP BPDU guard
 BPDUs
 Bridge Assurance [32-5](#)
 Shared Spanning Tree Protocol (SSTP) [32-20](#)
 Bridge Assurance
 description [32-4 to 32-6](#)
 inconsistent state [32-5](#)
 supported protocols and link types [32-5](#)

bridge domain
 configuring [41-8](#)
 bridge groups [35-1](#)
 bridge ID
 See STP bridge ID
 bridge priority, STP [31-34](#)
 bridge protocol data units
 see BPDUs
 bridging [35-1](#)
 broadcast storms
 see traffic-storm control

C

CALEA, See Communications Assistance for Law Enforcement Act (CALEA)

CDP

host presence detection [83-14, 85-4](#)
 to configure Cisco phones [19-3](#)

CEF

configuring
 RP [33-5](#)
 supervisor engine [33-4](#)
 examples [33-3](#)
 Layer 3 switching [33-2](#)
 packet rewrite [33-2](#)

channel-group group

command [12-9, 12-13, 12-14, 12-15, 23-9, 23-14, 23-15, 23-16](#)
 command example [23-9, 23-15](#)

Cisco Discovery Protocol

See CDP

Cisco Emergency Responder [19-4](#)

Cisco EnergyWise [13-1](#)

Cisco Express Forwarding [37-3](#)

CISCO-IP-TAP-MIB

citapStreamVRF [86-2](#)
 overview [86-8](#)
 restricting access to [86-10, 86-11](#)

- CISCO-TAP2-MIB
 - accessing [86-9](#)
 - overview [86-8](#)
 - restricting access to [86-10, 86-11](#)
- CISP [83-30](#)
- CIST regional root
 - See MSTP
- CIST root
 - See MSTP
- class command [63-9](#)
- class map configuration [63-8, 64-11](#)
- clear authentication sessions command [83-40](#)
- clear counters command [10-12](#)
- clear dot1x command [83-40](#)
- clear interface command [10-13](#)
- CLI
 - accessing [2-1](#)
 - backing out one level [2-6](#)
 - console configuration mode [2-5](#)
 - getting list of commands [2-6](#)
 - global configuration mode [2-5](#)
 - history substitution [2-4](#)
 - interface configuration mode [2-5](#)
 - privileged EXEC mode [2-5](#)
 - ROM monitor [2-7](#)
 - software basics [2-4](#)
- Client Information Signalling Protocol
 - See CISP
- collection function [86-6](#)
- command line processing [2-3](#)
- commands, getting list of [2-6](#)
- Communications Assistance for Law Enforcement Act
 - CALEA for Voice [86-5](#)
 - lawful intercept [86-4](#)
- community ports [27-7](#)
- community VLANs [27-6, 27-7](#)
- configuration example
 - EoMPLS port mode [39-4, 39-7](#)
 - EoMPLS VLAN mode [39-4](#)
- configuration guidelines
 - EVCs [41-2](#)
- configuring [63-9, 64-11](#)
 - lawful intercept [86-10, 86-11, 86-12](#)
 - SNMP [86-10](#)
- console configuration mode [2-5](#)
- content IAP [86-6](#)
- control plane policing
 - See CoPP
- CoPP [77-1](#)
 - applying QoS service policy to control plane [77-5](#)
 - configuring
 - ACLs to match traffic [77-5](#)
 - enabling MLS QoS [77-5](#)
 - packet classification criteria [77-5](#)
 - service-policy map [77-5](#)
 - control plane configuration mode
 - entering [77-5](#)
 - displaying
 - dynamic information [77-9](#)
 - number of conforming bytes and packets [77-9](#)
 - rate information [77-9](#)
 - entering control plane configuration mode [77-5](#)
 - monitoring statistics [77-9](#)
 - overview [77-3](#)
 - packet classification guidelines [77-2](#)
 - traffic classification
 - defining [77-6](#)
 - guidelines [77-7](#)
 - overview [77-6](#)
 - sample ACLs [77-8](#)
 - sample classes [77-6](#)
- CoS
 - override priority [19-6, 20-5](#)
- counters
 - clearing interface [10-12, 10-13](#)
- critical authentication [83-8](#)
- critical authentication, IEEE 802.1x [83-47](#)
- CSCsr62404 [10-9](#)

[CSCtx75254](#) [5-2](#)
[cTap2MediationDebug](#) notification [86-12](#)
[cTap2MediationNewIndex](#) object [86-8](#)
[cTap2MediationTable](#) [86-8](#)
[cTap2MediationTimedOut](#) notification [86-12](#)
[cTap2MIBActive](#) notification [86-12](#)
[cTap2StreamDebug](#) notification [86-12](#)
[cTap2StreamTable](#) [86-8](#)

D

[dACL](#)
 See [ACLs, downloadable](#) [83-24](#)
[dCEF](#) [33-4](#)
[debug](#) commands
 [IP MMLS](#) [44-31](#)
[DEC](#) spanning-tree protocol [35-1](#)
[default](#) configuration
 [802.1X](#) [83-31, 84-7](#)
 [dynamic ARP inspection](#) [80-6](#)
 [EVCs](#) [41-9](#)
 [Flex Links](#) [22-4](#)
 [IP MMLS](#) [44-15](#)
 [MSTP](#) [31-26](#)
 [MVR](#) [47-5](#)
 [UDLD](#) [11-4](#)
 [voice VLAN](#) [19-4](#)
 [VTP](#) [25-9](#)
[default](#) VLAN [21-10](#)
[denial of service](#) protection [76-1](#)
[DHCP](#) binding database
 See [DHCP snooping binding database](#)
[DHCP](#) binding table
 See [DHCP snooping binding database](#)
[DHCP](#) option 82
 [circuit ID](#) suboption [78-7](#)
 overview [78-5](#)
 [packet format](#), suboption
 [circuit ID](#) [78-7](#)
 [remote ID](#) [78-7](#)
 [remote ID](#) suboption [78-7](#)
[DHCP](#) option 82 allow on untrusted port [78-10](#)
[DHCP](#) snooping
 [802.1X](#) data insertion [83-15](#)
 [binding](#) database
 See [DHCP snooping binding database](#)
 [configuration](#) guidelines [78-8](#)
 [configuring](#) [78-9](#)
 [default](#) configuration [78-8](#)
 [displaying](#) binding tables [78-18](#)
 [enabling](#) [78-9, 78-10, 78-11, 78-12, 78-13, 78-14](#)
 [enabling](#) the database agent [78-14](#)
 [message exchange](#) process [78-6](#)
 [monitoring](#) [79-5, 79-6](#)
 [option 82](#) data insertion [78-5](#)
 overview [78-3](#)
 [Snooping](#) database agent [78-7](#)
[DHCP](#) snooping binding database
 [described](#) [78-5](#)
 [entries](#) [78-5](#)
[DHCP](#) snooping binding table
 See [DHCP snooping binding database](#)
[DHCP](#) Snooping Database Agent
 [adding](#) to the database (example) [78-18](#)
 [enabling](#) (example) [78-15](#)
 overview [78-7](#)
 [reading](#) from a TFTP file (example) [78-17](#)
[DHCP](#) snooping increased bindings limit [78-14](#)
[DiffServ](#)
 [configuring](#) short pipe mode [67-30](#)
 [configuring](#) uniform mode [67-34](#)
 [short pipe](#) mode [67-27](#)
 [uniform](#) mode [67-28](#)
[DiffServ](#) tunneling modes [67-4](#)
[Disabling](#) PIM Snooping Designated Router Flooding [46-6](#)
[distributed](#) Cisco Express Forwarding
 See [dCEF](#)

- distributed egress SPAN 56-10, 56-15
 - DNS, See Domain Name System
 - DNS, see Domain Name System
 - documentation, related 1-i
 - Domain Name System 86-2
 - DoS protection 76-1
 - monitoring packet drop statistics
 - using monitor session commands 76-10
 - using VACL capture 76-12
 - QoS ACLs 76-2
 - security ACLs 76-2
 - uRPF check 76-6
 - dot1x initialize interface command 83-39
 - dot1x max-reauth-req command 83-43
 - dot1x max-req command 83-42
 - dot1x pae authenticator command 83-34
 - dot1x re-authenticate interface command 83-39
 - dot1x timeout quiet-period command 83-41
 - DSCP-based queue mapping 65-14
 - duplex command 10-5, 10-6
 - duplex mode
 - autonegotiation status 10-6
 - configuring interface 10-4
 - dynamic ARP inspection
 - ARP cache poisoning 80-3
 - ARP requests, described 80-3
 - ARP spoofing attack 80-3
 - configuration guidelines 80-2
 - configuring
 - log buffer 80-13, 80-15
 - logging system messages 80-14
 - rate limit for incoming ARP packets 80-5, 80-10
 - default configuration 80-6
 - denial-of-service attacks, preventing 80-10
 - described 80-3
 - DHCP snooping binding database 80-4
 - displaying
 - ARP ACLs 80-15
 - configuration and operating state 80-15
 - trust state and rate limit 80-15
 - error-disabled state for exceeding rate limit 80-5
 - function of 80-4
 - interface trust states 80-4
 - log buffer
 - configuring 80-13, 80-15
 - logging of dropped packets, described 80-6
 - logging system messages
 - configuring 80-14
 - man-in-the middle attack, described 80-4
 - network security issues and interface trust states 80-4
 - priority of ARP ACLs and DHCP snooping entries 80-6
 - rate limiting of ARP packets
 - configuring 80-10
 - described 80-5
 - error-disabled state 80-5
 - validation checks, performing 80-11
 - Dynamic Host Configuration Protocol snooping 78-1
-
- ## E
- EAC 70-2
 - EAPOL. See also port-based authentication. 83-6
 - eFSU, See Enhanced Fast Software Upgrade (eFSU)
 - egress SPAN 56-10
 - electronic traffic, monitoring 86-7
 - enable mode 2-5
 - enable sticky secure MAC address 85-8
 - enabling
 - IP MMLS
 - on router interfaces 44-16
 - lawful intercept 86-8
 - SNMP notifications 86-12
 - Endpoint Admission Control (EAC) 70-2
 - EnergyWise 13-1
 - enhanced Fast Software Upgrade (eFSU)
 - aborting (issu abortversion command) 5-13

- accepting the new software version 5-12
- committing the new software to standby RP (issu commitversion command) 5-12
- displaying maximum outage time for module 5-10
- error handling 5-5
- forcing a switchover (issu runversion command) 5-10
- issu loadversion command 5-8
- loading new software onto standby RP 5-8
- memory reservation on module 5-5
- memory reservation on module, prohibiting 5-4
- OIR not supported 5-3
- operation 5-3
- outage times 5-4
- performing 5-5
- steps 5-6
- usage guidelines and limitations 5-2
- verifying redundancy mode 5-7
- environmental monitoring
 - LED indications 15-7
 - SNMP traps 15-7
 - supervisor engine and switching modules 15-7
 - Syslog messages 15-7
 - using CLI commands 15-1
- EOBC
 - for MAC address table synchronization 21-3
- EoMPLS 39-3
 - configuring 39-4
 - configuring VLAN mode 39-3
 - guidelines and restrictions 39-2
 - port mode 39-3
 - VLAN mode 39-3
- ERSPAN 56-1
- EtherChannel
 - channel-group group
 - command 12-9, 12-14, 12-15, 23-9, 23-14, 23-15, 23-16
 - command example 23-9, 23-15
 - configuration guidelines 4-31, 23-2
 - configuring
 - Layer 2 23-9
 - configuring (tasks) 4-31, 23-7
 - interface port-channel
 - command example 23-8
 - interface port-channel (command) 12-7, 12-8, 23-8
 - lacp system-priority
 - command example 23-11
 - Layer 2
 - configuring 23-9, 23-15
 - load balancing
 - configuring 23-11
 - understanding 23-7
 - Min-Links 23-13, 23-14
 - modes 23-4
 - PAgP
 - understanding 23-5
 - port-channel interfaces 23-7
 - port-channel load-balance
 - command 23-11
 - command example 23-12
 - STP 23-7
 - understanding 4-4, 23-3
- EtherChannel Guard
 - See STP EtherChannel Guard
- Ethernet
 - setting port duplex 10-10
- Ethernet flow point
 - See EFP
- Ethernet over MPLS (EoMPLS) configuration
 - EoMPLS port mode 39-6
 - EoMPLS VLAN mode 39-4
- Ethernet Virtual Connection
 - See EVC
- EVC
 - broadcast domain 41-4
 - configuration guidelines 41-2
 - default configuration 41-9
 - supported features 41-2
- EXP mutation 67-4

extended range VLANs [26-2](#)

See VLANs

extended system ID

MSTP [31-39](#)

Extensible Authentication Protocol over LAN. See EAPOL.

F

fall-back bridging [35-1](#)

fast link notification

on VSL failure [4-15](#)

fiber-optic, detecting unidirectional links [11-1](#)

FIB TCAM [37-3](#)

figure

lawful intercept overview [86-5](#)

Flex Links [22-1](#)

configuration guidelines [22-2](#)

configuring [22-4](#)

default configuration [22-4](#)

description [22-2](#)

monitoring [22-6](#)

flex links

interface preemption [22-3](#)

flow control [10-9](#)

forward-delay time

MSTP [31-45](#)

forward-delay time, STP [31-35](#)

frame distribution

See EtherChannel load balancing

G

get requests [86-7, 86-8, 86-11](#)

global configuration mode [2-5](#)

guest VLAN and 802.1x [83-19](#)

H

hardware Layer 3 switching

guidelines [33-2](#)

hello time

MSTP [31-44](#)

hello time, STP [31-35](#)

High Capacity Power Supply Support [14-4](#)

history

CLI [2-4](#)

host mode

see port-based authentication

host ports

kinds of [27-7](#)

host presence CDP message [19-4, 83-14](#)

host presence TLV message [85-4](#)

http

[//www-tac.cisco.com/Teams/ks/c3/xmlkvery.php?srId=612293409](http://www-tac.cisco.com/Teams/ks/c3/xmlkvery.php?srId=612293409) [23-3](#)

I

IAP

content IAP [86-6](#)

definition [86-6](#)

content IAP [86-6](#)

identification IAP [86-6](#)

types of

ICMP unreachable messages [69-2](#)

ID IAP [86-6](#)

IEEE 802.1Q Ethertype

specifying custom [21-15](#)

IEEE 802.1Q Tagging on a Per-Port Basis [29-7](#)

IEEE 802.1w

See RSTP

IEEE 802.1x

accounting [83-16, 83-43](#)

authentication failed VLAN [83-20](#)

critical ports [83-21](#)

- DHCP snooping [83-15](#)
- guest VLAN [83-19](#)
- MAC authentication bypass [83-26](#)
- network admission control Layer 2 validation [83-27](#)
- port security interoperability [83-23](#)
- RADIUS-supplied session timeout [83-38](#)
- voice VLAN [83-22](#)
- wake-on-LAN support [83-28](#)
- IEEE 802.3ad
 - See LACP
- IEEE 802.3af [20-2](#)
- IEEE 802.3at [20-2](#)
- IEEE 802.3x Flow Control [10-9](#)
- IEEE bridging protocol [35-1](#)
- IGMP [45-1](#)
 - configuration guidelines [52-9](#)
 - enabling [45-9](#)
 - join messages [45-3](#)
 - leave processing
 - enabling [45-13](#)
 - queries [45-4](#)
 - query interval
 - configuring [45-12](#)
 - snooping
 - fast leave [45-6](#)
 - joining multicast group [45-3, 48-4](#)
 - leaving multicast group [45-5, 48-4](#)
 - understanding [45-3, 48-3](#)
 - snooping querier
 - enabling [45-9](#)
 - understanding [45-3, 48-3](#)
- IGMPv3 [44-26](#)
- IGMP v3lite [44-26](#)
- ignore port trust [63-11](#)
- inaccessible authentication bypass [83-21](#)
- ingress SPAN [56-10](#)
- intercept access point
 - See IAP
- intercept-related information (IRI) [86-6, 86-7](#)
- intercepts, multiple [86-6](#)
- interface
 - configuration mode [2-5](#)
 - Layer 2 modes [21-4](#)
 - number [10-2](#)
- interface port-channel
 - command example [23-8](#)
- interface port-channel (command) [12-7, 12-8, 23-8](#)
- interfaces
 - configuring, duplex mode [10-3](#)
 - configuring, speed [10-3](#)
 - configuring, overview [10-2](#)
 - counters, clearing [10-12, 10-13](#)
 - displaying information about [10-12](#)
 - maintaining [10-12](#)
 - monitoring [10-12](#)
 - range of [10-2](#)
 - restarting [10-13](#)
 - shutting down
 - task [10-13](#)
- interfaces command [10-2](#)
- interfaces range command [55-3](#)
- interfaces range macro command [10-2](#)
- internal VLANs [26-3](#)
- Internet Group Management Protocol [45-1, 48-1](#)
- IP accounting, IP MMLS and [44-2](#)
- IP CEF
 - topology (figure) [33-4](#)
- ip flow-export source command [58-3, 58-4, 58-5](#)
- ip local policy route-map command [34-5](#)
- IP MMLS
 - cache, overview [44-4](#)
 - configuration guideline [44-1](#)
 - debug commands [44-31](#)
 - default configuration [44-15](#)
 - enabling
 - on router interfaces [44-16](#)
 - Layer 3 MLS cache [44-4](#)
 - overview [44-2](#)

packet rewrite [44-5](#)

router

enabling globally [44-16](#)

enabling on interfaces [44-16](#)

PIM, enabling [44-16](#)

IP multicast

IGMP snooping and [45-8](#)

MLDv2 snooping and [52-9](#)

overview [45-2, 48-2, 49-2](#)

IP multicast MLS

See IP MMLS

ip multicast-routing command

enabling IP multicast [44-16](#)

IP phone

configuring [19-5](#)

ip pim command

enabling IP PIM [44-16](#)

ip policy route-map command [34-5](#)

IP Source Guard [79-1](#)

configuring [79-3](#)

configuring on private VLANs [79-5](#)

displaying [79-5, 79-6](#)

overview [79-2](#)

IP unnumbered [35-1](#)

IPv4 Multicast VPN [50-1](#)

IPv6 Multicast Layer 3 Switching [51-1](#)

IPv6 QoS [61-4](#)

ISL trunks [21-4](#)

isolated port [27-7](#)

isolated VLANs [27-6, 27-7](#)

J

join messages, IGMP [45-3](#)

jumbo frames [10-6](#)

K

keyboard shortcuts [2-3](#)

L

label edge router [37-2](#)

label switched path [39-1](#)

label switch router [37-2, 37-4](#)

LACP

system ID [23-6](#)

Law Enforcement Agency (LEA) [86-4](#)

lawful intercept

admin function [86-7, 86-8](#)

collection function [86-6](#)

configuring [86-10, 86-11, 86-12](#)

enabling [86-8](#)

IRI [86-6](#)

mediation device [86-5](#)

overview [86-4, 86-5](#)

prerequisites [86-1](#)

processing [86-7](#)

security considerations [86-9](#)

SNMP notifications [86-12](#)

lawful intercept processing [86-7](#)

Layer 2

configuring interfaces [21-5](#)

access port [21-14](#)

trunk [21-8](#)

defaults [21-5](#)

interface modes [21-4](#)

show interfaces [10-8, 10-9, 21-6, 21-13](#)

switching

understanding [21-2](#)

trunks

understanding [21-4](#)

VLAN

interface assignment [26-6](#)

Layer 2 Interfaces

- configuring 21-1
- Layer 2 protocol tunneling
 - configuring Layer 2 tunnels 30-3
 - overview 30-2
- Layer 2 Traceroute 59-1
- Layer 2 traceroute
 - and ARP 59-2
 - and CDP 59-1
 - described 59-2
 - IP addresses and subnets 59-2
 - MAC addresses and VLANs 59-2
 - multicast traffic 59-2
 - multiple devices on a port 59-2
 - unicast traffic 59-2
 - usage guidelines 59-1
- Layer 3
 - IP MMLS and MLS cache 44-4
- Layer 3 switched packet rewrite
 - CEF 33-2
- Layer 3 switching
 - CEF 33-2
- Layer 4 port operations (ACLs) 69-2
- leave processing, IGMP
 - enabling 45-13
- leave processing, MLDv2
 - enabling 52-12
- LSRs 67-2, 67-6, 67-7
- Link Failure
 - detecting unidirectional 31-25
- link negotiation 10-5
- link redundancy
 - See Flex Links
- LLDP-MED
 - configuring
 - TLVs 20-8
- load deferral
 - MEC traffic recovery 4-6
- Local Egress Replication 44-19
- logical operation unit

- See LOU
- loop guard
 - See STP loop guard
- LOU
 - description 69-3
 - determining maximum number of 69-3
- LSRs 67-2, 67-6

M

- mab command 83-45, 83-50
- MAC address-based blocking 72-1
- MAC address table notification 21-7
- MAC authentication bypass. See also port-based authentication. 83-26
- MAC move (port security) 85-3
- macros 3-1
 - See Smartports macros
- MACSec 70-2
- magic packet 83-28
- mapping 802.1Q VLANs to ISL VLANs 26-7
- markdown
 - see QoS markdown
- match ip address command 34-4
- match length command 34-4
- maximum aging time
 - MSTP 31-45
- maximum aging time, STP 31-36
- maximum hop count, MSTP 31-46
- MEC
 - configuration 4-49
 - described 4-15
 - failure 4-16
 - port load share deferral 4-17
- mediation device
 - admin function 86-7, 86-8
 - definition 86-5
 - description 86-5
- MIBs

- CISCO-IP-TAP-MIB 86-2, 86-8, 86-10
- CISCO-TAP2-MIB 86-8, 86-9, 86-10
- SNMP-COMMUNITY-MIB 86-9
- SNMP-USM-MIB 86-4, 86-9
- SNMP-VACM-MIB 86-4, 86-9
- microflow policing 63-4
- Mini Protocol Analyzer 60-1
- Min-Links 23-13
- MLD
 - report 52-5
- MLD snooping
 - query interval
 - configuring 52-10
- MLDv1 52-2
- MLDv2 52-1
 - enabling 52-11
 - leave processing
 - enabling 52-12
 - queries 52-6
 - snooping
 - fast leave 52-8
 - joining multicast group 52-5
 - leaving multicast group 52-7
 - understanding 52-3
 - snooping querier
 - enabling 52-10
 - understanding 52-3
- MLDv2 Snooping 52-1
- monitoring
 - Flex Links 22-6
 - MVR 47-8
 - private VLANs 27-16
- monitoring electronic traffic 86-7
- MPLS 37-1, 37-2
 - aggregate label 37-2
 - any transport over MPLS 39-3
 - basic configuration 37-9
 - core 37-4
 - DiffServ Tunneling Modes 67-26
 - egress 37-4
 - experimental field 67-3
 - hardware features 37-5
 - ingress 37-4
 - IP to MPLS path 37-4
 - labels 37-2
 - MPLS to IP path 37-4
 - MPLS to MPLS path 37-4
 - nonaggregate lable 37-2
 - QoS default configuration 67-13
 - restrictions 37-1
 - VPN 67-11
 - VPN guidelines and restrictions 38-2
- MPLS QoS
 - Classification 67-2
 - Class of Service 67-2
 - commands 67-15
 - configuring a class map 67-17
 - configuring a policy map 67-20
 - configuring egress EXP mutation 67-24
 - configuring EXP Value Maps 67-25
 - Differentiated Services Code Point 67-2
 - displaying a policy map 67-24
 - E-LSP 67-2
 - EXP bits 67-2
 - features 67-2
 - IP Precedence 67-2
 - QoS Tags 67-2
 - queueing-only mode 67-17
- MPLS QoS configuration
 - class map to classify MPLS packets 67-17
- MPLS supported commands 37-2
- MPLS VPN
 - limitations and restrictions 38-2
- MQC 62-1
- MST
 - interoperation with Rapid PVST+ 32-20
 - root bridge 32-20
- MSTP

- boundary ports
 - configuration guidelines 31-2
 - described 31-22
- CIST, described 31-19
- CIST regional root 31-20
- CIST root 31-21
- configuration guidelines 31-2
- configuring
 - forward-delay time 31-45
 - hello time 31-44
 - link type for rapid convergence 31-46
 - maximum aging time 31-45
 - maximum hop count 31-46
 - MST region 31-38
 - neighbor type 31-46
 - path cost 31-42
 - port priority 31-41
 - root switch 31-39
 - secondary root switch 31-40
 - switch priority 31-43
- CST
 - defined 31-19
 - operations between regions 31-20
- default configuration 31-26
- displaying status 31-47
- enabling the mode 31-38
- extended system ID
 - effects on root switch 31-39
 - effects on secondary root switch 31-40
 - unexpected behavior 31-39
- IEEE 802.1s
 - implementation 31-23
 - port role naming change 31-23
 - terminology 31-21
- interoperability with IEEE 802.1D
 - described 31-24
 - restarting migration process 31-47
- IST
 - defined 31-19
 - master 31-20
 - operations within a region 31-20
- mapping VLANs to MST instance 31-38
- MST region
 - CIST 31-19
 - configuring 31-38
 - described 31-19
 - hop-count mechanism 31-22
 - IST 31-19
 - supported spanning-tree instances 31-19
- overview 31-18
- root switch
 - configuring 31-39
 - effects of extended system ID 31-39
 - unexpected behavior 31-39
 - status, displaying 31-47
- MTU size (default) 26-3
- multiauthentication (multiauth). See also port-based authentication. 83-15
- multicast
 - IGMP snooping and 45-8
 - MLDv2 snooping and 52-9
 - non-RPF 44-7
 - overview 45-2, 48-2, 49-2
 - PIM snooping 46-4
- multicast flood blocking 82-1
- multicast groups
 - joining 45-3, 48-4
 - leaving 45-5, 52-7
- multicast groups, IPv6
 - joining 52-5
- Multicast Listener Discovery version 2 52-1
- Multicast Replication Mode Detection enhancement 44-18
- multicast storms
 - see traffic-storm control
- multicast television application 47-3
- multicast VLAN 47-2
- Multicast VLAN Registration 47-1

multichassis EtherChannel

see MEC 4-15

Multidomain Authentication (MDA). See also port-based authentication. 83-14

Multilayer MAC ACL QoS Filtering 69-9

multiple path RPF check 76-7

Multiple Spanning Tree

See MST

MUX-UNI Support 37-7

MUX-UNI support 37-7

MVAP (Multi-VLAN Access Port). See also port-based authentication. 83-22

MVR

and IGMPv3 47-2

configuring interfaces 47-6

default configuration 47-5

example application 47-3

in the switch stack 47-5

monitoring 47-8

multicast television application 47-3

restrictions 47-1

setting global parameters 47-6

N

NAC

agentless audit support 83-27

critical authentication 83-21, 83-47

IEEE 802.1x authentication using a RADIUS server 83-50

IEEE 802.1x validation using RADIUS server 83-50

inaccessible authentication bypass 83-47

Layer 2 IEEE 802.1x validation 83-50

Layer 2 IEEE802.1x validation 83-27

native VLAN 21-11

NDAC 70-2

NEAT

configuring 83-54

overview 83-30

NetFlow

table, displaying entries 33-5

Network Device Admission Control (NDAC) 70-2

Network Edge Access Topology

See NEAT

network ports

Bridge Assurance 32-5

description 32-2

nonaggregate label 37-2, 37-5

non-RPF multicast 44-7

normal-range VLANs

See VLANs

notifications, See SNMP notifications

NSF with SSO does not support IPv6 multicast traffic. 7-1, 8-1

O

OIR 10-11

online diagnostics

CompactFlash disk verification 1-45

configuring 16-2

datapath verification 1-15

diagnostic sanity check 16-29

egress datapath test 1-5

error counter test 1-5

interrupt counter test 1-5

memory tests 16-29

overview 16-2

running tests 16-7

test descriptions 1-1

understanding 16-2

online diagnostic tests 1-1

online insertion and removal

See OIR

out-f-band MAC address table synchronization

configuring 21-6

in a VSS 4-2

out of profile

see QoS out of profile

P

packet capture 60-2

packet rewrite

CEF 33-2

IP MMLS and 44-5

packets

multicast 73-6

PAgP

understanding 23-5

path cost

MSTP 31-42

PBACLs 69-5

PBF 74-4

PBR (policy-based routing)

configuration (example) 34-7

enabling 34-4

peer inconsistent state

in PVST simulation 32-20

per-port VTP enable and disable 25-16

PFC

recirculation 37-5

PIM, IP MMLS and 44-16

PIM snooping

designated router flooding 46-6

enabling globally 46-5

enabling in a VLAN 46-5

overview 46-4

platform aging command

configuring IP MLS 53-3, 53-4

platform ip multicast command

enabling IP MMLS 44-17 to 44-27

PoE 20-2

Cisco prestandard 20-3

IEEE 802.3af 20-2

IEEE 802.3at 20-2

PoE management 20-3

power policing 20-4

power use measurement 20-4

police command 63-13, 63-14

policy-based ACLs (PBACLs) 69-5

policy-based forwarding (PBF) 75-2

policy-based routing (PBR)

configuring 34-1

policy map 63-9, 64-11

attaching to an interface 63-17, 64-17, 76-6

policy-map command 63-9

port ACLs

defined 73-2

port ACLs (PACLs) 73-1

Port Aggregation Protocol

see PAgP

port-based authentication

AAA authorization 83-33

accounting 83-16

configuring 83-43

authentication server

defined 83-7, 84-3

RADIUS server 83-7

client, defined 83-7, 84-3

configuration guidelines 83-2, 84-1

configuring

guest VLAN 83-45

inaccessible authentication bypass 83-47

initializing authentication of a client 83-39

manual reauthentication of a client 83-39

RADIUS server 83-35, 84-10

RADIUS server parameters on the switch 83-34, 84-9

restricted VLAN 83-46

switch-to-authentication-server retransmission time 83-42

switch-to-client EAP-request frame retransmission time 83-41

switch-to-client frame-retransmission number 83-42, 83-43

switch-to-client retransmission time 83-41

- user distribution 83-44
 - VLAN group assignment 83-44
- default configuration 83-31, 84-7
- described 83-6
- device roles 83-7, 84-3
- DHCP snooping 83-15
- DHCP snooping and insertion 78-6
- displaying statistics 83-57, 84-15
- EAPOL-start frame 83-10
- EAP-request/identity frame 83-10
- EAP-response/identity frame 83-10
- enabling
 - 802.1X authentication 83-33, 83-34, 84-9
 - periodic reauthentication 83-38
- encapsulation 83-7
- guest VLAN
 - configuration guidelines 83-19, 83-20
 - described 83-19
- host mode 83-13
- inaccessible authentication bypass
 - configuring 83-47
 - described 83-21
 - guidelines 83-4
- initiation and message exchange 83-10
- MAC authentication bypass 83-26
- magic packet 83-28
- method lists 83-33
- modes 83-13
- multiauth mode, described 83-15
- multidomain authentication mode, described 83-14
- multiple-hosts mode, described 83-13
- ports
 - authorization state and dot1x port-control command 83-12
 - authorized and unauthorized 83-12
 - critical 83-21
 - voice VLAN 83-22
- port security
 - and voice VLAN 83-23
 - described 83-23
 - interactions 83-23
 - multiple-hosts mode 83-13
 - pre-authentication open access 83-15, 83-36
 - resetting to default values 83-57
 - supplicant, defined 83-7
 - switch
 - as proxy 83-7, 84-3
 - RADIUS client 83-7
 - switch supplicant
 - configuring 83-54
 - overview 83-30
 - user distribution
 - configuring 83-44
 - described 83-18
 - guidelines 83-4
- VLAN assignment
 - AAA authorization 83-33
 - characteristics 83-17
 - configuration tasks 83-18
 - described 83-17
- VLAN group
 - guidelines 83-4
- voice VLAN
 - described 83-22
 - PVID 83-22
 - VVID 83-22
 - wake-on-LAN, described 83-28
- port-based QoS features
 - see QoS
- port-channel
 - see EtherChannel
- port-channel load-balance
 - command 23-11
 - command example 23-11, 23-12
- port-channel load-defer command 4-49
- port-channel port load-defer command 4-49
- port cost, STP 31-32
- port debounce timer

- disabling [10-10](#)
 - displaying [10-10](#)
 - enabling [10-10](#)
- PortFast
 - edge ports [32-2](#)
 - network ports [32-2](#)
 - See STP PortFast
- PortFast Edge BPDU filtering
 - See STP PortFast Edge BPDU filtering
- PortFast port types
 - description [32-2, 32-2 to ??](#)
 - edge [32-2](#)
 - network [32-2](#)
- port mode [39-3](#)
- port negotiation [10-5](#)
- port priority
 - MSTP [31-41](#)
- port priority, STP [31-31](#)
- ports
 - setting the debounce timer [10-10](#)
- port security
 - aging [85-9, 85-10](#)
 - configuring [85-4](#)
 - described [85-3](#)
 - displaying [85-10](#)
 - enable sticky secure MAC address [85-8](#)
 - sticky MAC address [85-3](#)
 - violations [85-3](#)
- Port Security is supported on trunks [85-2, 85-5, 85-7, 85-9](#)
- port security MAC move [85-3](#)
- port security on PVLAN ports [85-2](#)
- Port Security with Sticky Secure MAC Addresses [85-3](#)
- power management
 - enabling/disabling redundancy [14-2](#)
 - overview [14-1](#)
 - powering modules up or down [14-3](#)
 - power policing [20-8](#)
- power negotiation
 - through LLDP [20-8](#)
- Power over Ethernet [20-2](#)
- power over ethernet [20-2](#)
- pre-authentication open access. See port-based authentication.
- preemption, default configuration [22-4](#)
- preemption delay, default configuration [22-4](#)
- prerequisites for lawful intercept [86-1](#)
- primary links [22-2](#)
- primary VLANs [27-6](#)
- priority
 - overriding CoS [19-6, 20-5](#)
- private hosts [28-1](#)
- private hosts feature
 - configuration guidelines [28-1](#)
 - configuring (detailed steps) [28-9](#)
 - configuring (summary) [28-8](#)
 - multicast operation [28-4](#)
 - overview [28-4](#)
 - port ACLs (PACLs) [28-7](#)
 - port types [28-5, 28-6](#)
 - protocol-independent MAC ACLs [28-4](#)
 - restricting traffic flow with PACLs [28-5](#)
 - spoofing protection [28-3](#)
- private VLANs [27-1](#)
 - across multiple switches [27-9](#)
 - and SVIs [27-10](#)
 - benefits of [27-5](#)
 - community VLANs [27-6, 27-7](#)
 - configuration guidelines [27-2, 27-4, 27-10](#)
 - configuring [27-10](#)
 - host ports [27-14](#)
 - promiscuous ports [27-15](#)
 - routing secondary VLAN ingress traffic [27-13](#)
 - secondary VLANs with primary VLANs [27-12](#)
 - VLANs as private [27-11](#)
 - end station access to [27-8](#)
 - IP addressing [27-8](#)
 - isolated VLANs [27-6, 27-7](#)
 - monitoring [27-16](#)

- ports
 - community [27-7](#)
 - configuration guidelines [27-4](#)
 - isolated [27-7](#)
 - promiscuous [27-7](#)
- primary VLANs [27-6](#)
- secondary VLANs [27-6](#)
- subdomains [27-5](#)
- traffic in [27-10](#)
- privileged EXEC mode [2-5](#)
- promiscuous ports [27-7](#)
- protocol tunneling
 - See Layer 2 protocol tunneling [30-2](#)
- PVRST
 - See Rapid-PVST [31-3](#)
- PVST
 - description [31-3](#)
- PVST simulation
 - description [32-20](#)
 - peer inconsistent state [32-20](#)
 - root bridge [32-20](#)

Q

- QoS
 - auto-QoS
 - enabling for VoIP [66-4](#)
 - IPv6 [61-4](#)
 - See also automatic QoS [66-1](#)
- QoS CoS
 - port value, configuring [65-2](#)
- QoS default configuration [68-2](#)
- QoS DSCP
 - maps, configuring [65-7](#)
- QoS mapping
 - CoS values to DSCP values [65-4](#), [65-7](#)
 - DSCP markdown values [65-8](#), [67-14](#)
 - DSCP mutation [65-3](#), [67-25](#)
 - DSCP values to CoS values [65-9](#)

- IP precedence values to DSCP values [65-7](#)
- QoS markdown [63-4](#)
- QoS out of profile [63-4](#)
- QoS policing rule
 - aggregate [63-4](#)
 - microflow [63-4](#)
- QoS port
 - trust state [65-10](#)
- QoS port-based or VLAN-based [65-12](#)
- QoS receive queue [65-18](#)
- QoS statistics data export [68-2](#)
 - configuring [68-2](#)
 - configuring destination host [68-7](#)
 - configuring time interval [68-6](#), [68-8](#)
- QoS transmit queues [64-6](#), [65-15](#), [65-16](#)
- QoS VLAN-based or port-based [65-12](#)
- queries, IGMP [45-4](#)
- queries, MLDv2 [52-6](#)

R

- RADIUS [78-6](#)
- RADIUS. See also port-based authentication. [83-7](#)
- range
 - command [55-3](#)
 - macro [10-2](#)
- rapid convergence [31-14](#)
- Rapid-PVST
 - enabling [31-36](#)
- Rapid PVST+
 - interoperation with MST [32-20](#)
- Rapid-PVST+
 - overview [31-3](#)
- Rapid Spanning Tree
 - See RSTP
- Rapid Spanning Tree Protocol
 - See RSTP
- receive queues
 - see QoS receive queues

- recirculation [37-5](#)
- redirect URLs
 - described [83-25](#)
- reduced MAC address [31-3](#)
- redundancy (RPR+) [9-1](#)
 - configuring [9-4](#)
 - configuring supervisor engine [9-2](#)
 - displaying supervisor engine configuration [9-4](#)
- related documentation [1-i](#)
- Remote Authentication Dial-In User Service. See RADIUS.
- report, MLD [52-5](#)
- reserved-range VLANs
 - See VLANs
- restricted VLAN
 - configuring [83-46](#)
 - described [83-20](#)
 - using with IEEE 802.1x [83-20](#)
- restricting MIB access [86-10, 86-11](#)
- rewrite, packet
 - CEF [33-2](#)
 - IP MMLS [44-5](#)
- RIF cache monitoring [10-12](#)
- ROM monitor
 - CLI [2-7](#)
- root bridge
 - MST [32-20](#)
 - PVST simulation [32-20](#)
- root bridge, STP [31-29](#)
- root guard
 - See STP root guard
- root switch
 - MSTP [31-39](#)
- route-map (IP) command [34-4](#)
- route maps
 - defining [34-4](#)
- router guard [49-1](#)
- RPF
 - failure [44-7](#)
 - non-RPF multicast [44-7](#)
- RPR and RPR+ support IPv6 multicast traffic [9-1](#)
- RSTP
 - active topology [31-13](#)
 - BPDU
 - format [31-16](#)
 - processing [31-17](#)
 - designated port, defined [31-13](#)
 - designated switch, defined [31-13](#)
 - interoperability with IEEE 802.1D
 - described [31-24](#)
 - restarting migration process [31-47](#)
 - topology changes [31-17](#)
 - overview [31-13](#)
 - port roles
 - described [31-13](#)
 - synchronized [31-15](#)
 - proposal-agreement handshake process [31-14](#)
 - rapid convergence
 - described [31-14](#)
 - edge ports and Port Fast [31-14](#)
 - point-to-point links [31-14, 31-46](#)
 - root ports [31-14](#)
 - root port, defined [31-13](#)
 - See also MSTP

S

- secondary VLANs [27-6](#)
- Secure MAC Address Aging Type [85-10](#)
- security
 - configuring [71-1](#)
- security, port [85-3](#)
- security considerations [86-9](#)
- Security Exchange Protocol (SXP) [70-2](#)
- Security Group Access Control List (SGACL) [70-2](#)
- Security Group Tag (SGT) [70-2](#)
- serial interfaces
 - clearing [10-13](#)

- synchronous
 - maintaining 10-13
- service instance
 - configuration mode 41-5
 - creating 41-4
 - defined 41-4
- service-policy input command 63-17, 64-17, 65-4, 65-6, 67-25, 76-6
- service-provider network, MSTP and RSTP 31-18
- set default interface command 34-4
- set interface command 34-4
- set ip default next-hop command 34-4
- set ip df command
 - PBR 34-4
- set ip next-hop command 34-4
- set ip precedence command
 - PBR 34-4
- set ip vrf command
 - PBR 34-4
- set power redundancy enable/disable command 14-2
- set requests 86-7, 86-8, 86-11
- setting up lawful intercept 86-7
- SGACL 70-2
- SGT 70-2
- short pipe mode
 - configuring 67-30
- show authentication command 83-58
- show catalyst6000 chassis-mac-address command 31-4
- show dot1x interface command 83-39
- show eobc command 10-12
- show history command 2-4
- show ibc command 10-12
- show interfaces command 10-8, 10-9, 10-12, 21-6, 21-13
 - clearing interface counters 10-12
 - displaying, speed and duplex mode 10-6
- show ip local policy command 34-5
- show mab command 83-61
- show module command 9-4
- show platform aging command 53-4
- show platform entry command 33-5
- show platform ip multicast group command
 - displaying IP MMLS group 44-27
- show platform ip multicast interface command
 - displaying IP MMLS interface 44-27
- show platform ip multicast source command
 - displaying IP MMLS source 44-27
- show platform ip multicast statistics command
 - displaying IP MMLS statistics 44-27
- show platform ip multicast summary
 - displaying IP MMLS configuration 44-27
- show protocols command 10-12
- show rif command 10-12
- show running-config command 10-12
 - displaying ACLs 73-7, 73-8
- show version command 10-12
- shutdown command 10-13
- shutdown interfaces
 - result 10-13
- slot number, description 10-2
- smart port macros 3-1
 - configuration guidelines 3-2
- Smartports macros
 - applying global parameter values 3-14
 - applying macros 3-14
 - creating 3-13
 - default configuration 3-4
 - defined 3-4
 - displaying 3-15
 - tracing 3-2
- SNMP
 - configuring 86-10
 - default view 86-9
 - get and set requests 86-7, 86-8, 86-11
 - notifications 86-9, 86-12
- SNMP-COMMUNITY-MIB 86-9
- SNMP-USM-MIB 86-4, 86-9
- SNMP-VACM-MIB 86-4, 86-9
- snooping

- See IGMP snooping
- software
 - upgrading router 5-5
- source specific multicast with IGMPv3, IGMP v3lite, and URD 44-26
- SPAN
 - configuration guidelines 56-2
 - configuring 56-12
 - sources 56-16, 56-19, 56-21, 56-22, 56-24, 56-25, 56-26, 56-28
 - VLAN filtering 56-30
 - destination port support on EtherChannels 56-12, 56-19, 56-22, 56-24, 56-25, 56-29
 - distributed egress 56-10, 56-15
 - modules that disable for ERSPAN 56-7
 - input packets with don't learn option
 - ERSPAN 56-28, 56-29
 - local SPAN 56-17, 56-18, 56-19
 - RSPAN 56-22, 56-23, 56-25
 - understanding 56-12
 - local SPAN egress session increase 56-3, 56-16
 - overview 56-7
- SPAN Destination Port Permit Lists 56-15
- spanning-tree backbonefast
 - command 32-15, 32-16
 - command example 32-15, 32-16
- spanning-tree cost
 - command 31-33
 - command example 31-33
- spanning-tree portfast
 - command 32-2, 32-3, 32-4
 - command example 32-3, 32-4
- spanning-tree portfast bpduguard
 - command 32-8
- spanning-tree port-priority
 - command 31-31
- spanning-tree protocol for bridging 35-1
- spanning-tree uplinkfast
 - command 32-13
 - command example 32-13
- spanning-tree vlan
 - command 31-27, 31-29, 31-30, 31-31, 32-8, 32-17
 - command example 31-28, 31-29, 31-30, 31-31
- spanning-tree vlan cost
 - command 31-33
- spanning-tree vlan forward-time
 - command 31-35
 - command example 31-35
- spanning-tree vlan hello-time
 - command 31-35
 - command example 31-35
- spanning-tree vlan max-age
 - command 31-36
 - command example 31-36
- spanning-tree vlan port-priority
 - command 31-31
 - command example 31-32
- spanning-tree vlan priority
 - command 31-34
 - command example 31-34
- speed
 - configuring interface 10-4
- speed command 10-4
- speed mode
 - autonegotiation status 10-6
- standards, lawful intercept 86-4
- standby links 22-2
- static sharing
 - description 83-25
- statistics
 - 802.1X 83-57, 84-15
- sticky ARP 76-9
- sticky MAC address 85-3
- Sticky secure MAC addresses 85-8, 85-9
- storm control
 - see traffic-storm control
- STP
 - configuring 31-26
 - bridge priority 31-34

- enabling [31-27, 31-28](#)
 - forward-delay time [31-35](#)
 - hello time [31-35](#)
 - maximum aging time [31-36](#)
 - port cost [31-32](#)
 - port priority [31-31](#)
 - root bridge [31-29](#)
 - secondary root switch [31-30](#)
- defaults [31-25](#)
- EtherChannel [23-7](#)
- normal ports [32-3](#)
- understanding [31-2](#)
 - 802.1Q Trunks [31-12](#)
 - Blocking State [31-8](#)
 - BPDU s [31-4](#)
 - disabled state [31-12](#)
 - forwarding state [31-11](#)
 - learning state [31-10](#)
 - listening state [31-9](#)
 - overview [31-3](#)
 - port states [31-6](#)
 - protocol timers [31-5](#)
 - root bridge election [31-5](#)
 - topology [31-5](#)
- STP BackboneFast
 - configuring [32-15](#)
 - figure
 - adding a switch [32-18](#)
 - spanning-tree backbonefast
 - command [32-15, 32-16](#)
 - command example [32-15, 32-16](#)
 - understanding [32-13](#)
- STP BPDU Guard
 - configuring [32-7](#)
 - spanning-tree portfast bpdu-guard
 - command [32-8](#)
 - understanding [32-7](#)
- STP bridge ID [31-3](#)
- STP EtherChannel guard [32-16](#)
- STP extensions
 - description [?? to 32-20](#)
- STP loop guard
 - configuring [32-19](#)
 - overview [32-17](#)
- STP PortFast
 - BPDU filter
 - configuring [32-10](#)
 - BPDU filtering [32-9](#)
 - configuring [32-2](#)
 - spanning-tree portfast
 - command [32-2, 32-3, 32-4](#)
 - command example [32-3, 32-4](#)
 - understanding [32-2](#)
- STP port types
 - normal [32-3](#)
- STP root guard [32-17](#)
- STP UplinkFast
 - configuring [32-12](#)
 - spanning-tree uplinkfast
 - command [32-13](#)
 - command example [32-13](#)
 - understanding [32-11](#)
- subdomains, private VLAN [27-5](#)
- supervisor engine
 - environmental monitoring [15-1](#)
 - redundancy [9-1](#)
 - synchronizing configurations [9-4](#)
- supervisor engine redundancy
 - configuring [9-2](#)
- supervisor engines
 - displaying redundancy configuration [9-4](#)
- supplicant [83-7](#)
- surveillance [86-7](#)
- Switched Port Analyzer [56-1](#)
- switch fabric functionality [18-1](#)
 - monitoring [18-2](#)
- switchport
 - configuring [21-14](#)

- example 21-13
- show interfaces 10-8, 10-9, 21-6, 21-13
- switchport access vlan 21-6, 21-7, 21-10, 21-14
 - example 21-15
- switchport mode access 21-4, 21-6, 21-7, 21-14
 - example 21-15
- switchport mode dynamic 21-9
- switchport mode dynamic auto 21-4
- switchport mode dynamic desirable 21-4
 - default 21-5
 - example 21-13
- switchport mode trunk 21-4, 21-9
- switchport nonegotiate 21-4
- switchport trunk allowed vlan 21-11
- switchport trunk encapsulation 21-7, 21-9
- switchport trunk encapsulation dot1q
 - example 21-13
- switchport trunk encapsulation negotiate
 - default 21-5
- switchport trunk native vlan 21-11
- switchport trunk pruning vlan 21-12
- switch priority
 - MSTP 31-43
- switch TopN reports
 - foreground execution 58-2
 - running 58-3
 - viewing 58-3
- SXP 70-2
- system event archive (SEA) 54-1

T

- TDR
 - checking cable connectivity 10-14
 - enabling and disabling test 10-14
 - guidelines 10-14
- Telnet
 - accessing CLI 2-2
- Time Domain Reflectometer 10-14
- TLV
 - host presence detection 19-4, 83-14, 85-4
- traceroute, Layer 2
 - and ARP 59-2
 - and CDP 59-1
 - described 59-2
 - IP addresses and subnets 59-2
 - MAC addresses and VLANs 59-2
 - multicast traffic 59-2
 - multiple devices on a port 59-2
 - unicast traffic 59-2
 - usage guidelines 59-1
- traffic-storm control
 - command
 - broadcast 81-4
 - described 81-2
 - monitoring 81-5
 - thresholds 81-2
- traffic suppression
 - see traffic-storm control
- transmit queues
 - see QoS transmit queues
- traps, see SNMP notifications
- trunks 21-4
 - 802.1Q Restrictions 21-2
 - allowed VLANs 21-11
 - configuring 21-8
 - default interface configuration 21-6
 - default VLAN 21-10
 - different VTP domains 21-4
 - native VLAN 21-11
 - to non-DTP device 21-4
 - VLAN 1 minimization 21-12
- trusted boundary 19-6
- trusted boundary (extended trust for CDP devices) 19-4
- tunneling 67-4, 67-26
- tunneling, 802.1Q
 - See 802.1Q 29-4
- type length value

See TLV

U

UDE

- configuration 36-5
- overview 36-4

UDE and UDLR 36-1

UDLD

- default configuration 11-4
- enabling
 - globally 11-5
 - on ports 11-5, 11-6
- overview 11-2

UDLR 36-1

- back channel 36-3
- configuration 36-6
- tunnel

(example) 36-7

ARP and NHRP 36-4

UDLR (unidirectional link routing) 36-1

UDP port for SNMP notifications 86-12

UMFB 82-2

unauthorized ports with 802.1X 83-12

unicast storms

- see traffic-storm control

Unidirectional Ethernet 36-1

unidirectional ethernet

- example of setting 36-5

UniDirectional Link Detection Protocol

- see UDLD

uniform mode

- configuring 67-34

unknown multicast flood blocking

- See UMFB

unknown unicast and multicast flood blocking 82-1

unknown unicast flood blocking

- See UUFRL

unknown unicast flood rate-limiting

See UUFRL

UplinkFast

See STP UplinkFast

URD 44-26

User-Based Rate Limiting 63-6, 63-15

user EXEC mode 2-5

UUFRL 82-2

UUFRL 82-2

V

VACLs 74-2

configuring

examples 74-5

Layer 3 VLAN interfaces 74-5

Layer 4 port operations 69-2

logging

configuration example 74-8

configuring 74-7

restrictions 74-7

MAC address based 74-2

multicast packets 73-6

SVIs 74-5

WAN interfaces 74-2

vlan

command 26-5, 26-6, 56-20

command example 26-6

VLAN Access Control Lists

See VACLs

VLAN-based QoS filtering 69-10

VLAN-bridge spanning-tree protocol 35-1

vlan database

command 26-5, 26-6, 56-20

vlan group command 83-44

VLAN locking 26-4

vlan mapping dot1q

command 26-8, 26-9

VLAN maps

applying 73-8

- VLAN mode [39-3](#)
- VLAN port provisioning verification [26-4](#)
- VLANs
 - allowed on trunk [21-11](#)
 - configuration guidelines [26-2](#)
 - configuring [26-1](#)
 - configuring (tasks) [26-4](#)
 - defaults [26-3](#)
 - extended range [26-3](#)
 - interface assignment [26-6](#)
 - multicast [47-2](#)
 - name (default) [26-3](#)
 - normal range [26-3](#)
 - reserved range [26-3](#)
 - support for 4,096 VLANs [26-2](#)
 - token ring [26-3](#)
 - trunks
 - understanding [21-4](#)
 - understanding [26-2](#)
 - VLAN 1 minimization [21-12](#)
 - VTP domain [26-4](#)
- VLAN translation
 - command example [26-9](#)
- voice VLAN
 - Cisco 7960 phone, port connections [19-2](#)
 - configuration guidelines [19-1](#)
 - configuring IP phone for data traffic
 - override CoS of incoming frame [19-6, 20-5](#)
 - configuring ports for voice traffic in
 - 802.1Q frames [19-5](#)
 - connecting to an IP phone [19-5](#)
 - default configuration [19-4](#)
 - overview [19-2](#)
- voice VLAN. See also port-based authentication. [83-22](#)
- VPN
 - configuration example [38-4](#)
 - guidelines and restrictions [38-2](#)
- VPN supported commands [38-2](#)
- VPN switching [38-1](#)
- VSS
 - dual-active detection
 - Enhanced PAgP, advantages [4-27](#)
 - Enhanced PAgP, description [4-28](#)
 - enhanced PAgP, description [4-50](#)
 - fast-hello, advantages [4-27](#)
 - fast-hello, description [4-28](#)
 - VSLP fast-hello, configuration [4-51](#)
- VSS Quad-Sup SSO (V4SO) [4-9](#)
- VTP
 - advertisements [25-4, 25-5](#)
 - client, configuring [25-15](#)
 - configuration guidelines [25-1](#)
 - default configuration [25-9](#)
 - disabling [25-15](#)
 - domains [25-3](#)
 - VLANs [26-4](#)
 - modes
 - client [25-4](#)
 - server [25-4](#)
 - transparent [25-4](#)
 - monitoring [25-17](#)
 - overview [25-2](#)
 - per-port enable and disable [25-16](#)
 - pruning
 - configuration [21-12](#)
 - configuring [25-12](#)
 - overview [25-7](#)
 - server, configuring [25-15](#)
 - statistics [25-17](#)
 - transparent mode, configuring [25-15](#)
 - version 2
 - enabling [25-13](#)
 - overview [25-5](#)
 - version 3
 - enabling [25-13](#)
 - overview [25-6](#)
 - server type, configuring [25-11](#)

W

wake-on-LAN. See also port-based authentication. [83-28](#)

web-based authentication

- AAA fail policy [84-5](#)

- description [84-2](#)

wiretaps [86-4](#)

