



Configuring MPLS and EoMPLS

This chapter describes how to configure multiprotocol label switching (MPLS) and Ethernet over MPLS (EoMPLS) on the Catalyst 3750 Metro switch. MPLS is a packet-switching technology that integrates link layer (Layer 2) switching with network layer (Layer 3) routing. With MPLS, data is transferred over any combination of Layer 2 technologies, using any Layer 3 protocol, with increased scalability. MPLS supports different routes between a source and destination over a router-based Internet backbone.

EoMPLS is tunneling mechanism that transports Layer 2 Ethernet frames over an MPLS network. You can connect two Layer 2 networks that are in different locations, without requiring bridges, routers, or switches at the locations. You enable the MPLS backbone to accept Layer 2 traffic by configuring the label-edge routers (LERs) at both ends of the MPLS backbone.

MPLS functionality is supported only on the enhanced-services (ES) ports; EoMPLS is supported on standard and ES ports.



Note

For more information about MPLS, see the “Multiprotocol Label Switching” section of the *Cisco IOS Switching Services Configuration Guide for Release 12.2*. For complete syntax and usage information for the MPLS commands used in this chapter, see the *Cisco IOS Switching Services Command Reference, Release 12.2*.

For more information about EoMPLS commands, see the command reference for this release.

This chapter contains these sections:

- [Understanding MPLS Services, page 37-2](#)
- [Understanding MPLS VPNs, page 37-3](#)
- [Configuring MPLS VPNs, page 37-6](#)
- [Understanding EoMPLS, page 37-12](#)
- [Enabling EoMPLS, page 37-15](#)
- [Configuring MPLS and EoMPLS QoS, page 37-18](#)
- [Monitoring and Maintaining MPLS and EoMPLS, page 37-22](#)

In Cisco IOS Release 12.2(25)SED or later, the switch supports hierarchical virtual private LAN service (H-VPLS) architecture to simulate LAN services over the MPLS network. The switch supports H-VPLS using IEEE 802.1Q tunneling or Ethernet over multiprotocol label switching (EoMPLS). For more information, see these software documents:

- For information about EoMPLS, see the “[Understanding EoMPLS](#)” section on page 37-12.
- For information about configuring EoMPLS, see the “[Enabling EoMPLS](#)” section on page 37-15 and the “[Configuring MPLS and EoMPLS QoS](#)” section on page 37-18.

- For information about 802.1Q tunneling, see the “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling” chapter.
- For information about configuring H-VPLS on Cisco 7600 routers, see the “Configuring Multiprotocol Label Switching on the Optical Services Modules” section in the *OSM Configuration Note, 12.2SX* at http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a00801e5c06.html

Understanding MPLS Services

In conventional Layer 3 forwarding, as a packet travels across the network, each router extracts the packet-forwarding information from the Layer 3 header and uses this information as an index for a routing table lookup to determine the packet's next hop. In most cases, the only relevant field in the header is the destination address field, but in some cases other header fields are also relevant. For this reason, each router through which the packet passes must analyze the packet header.

With MPLS, the Layer 3 header is analyzed only once and then is mapped into a fixed-length, unstructured value called a *label*. Many different headers can map to the same label, as long as those headers always result in the same choice of next hop. In effect, a label represents a *forwarding-equivalence class*—that is, a set of packets that can be very different but that are indistinguishable to the forwarding function.

The initial choice of label can be based exclusively on the contents of the Layer 3 header, or it can be based on policy, allowing forwarding decisions at subsequent hops to be based on policy. After a label is chosen, a short label header is put at the front of the Layer 3 packet and carried across the network as part of the packet. At subsequent hops through each MPLS router in the network, labels are exchanged, and the router uses MPLS forwarding-table lookups for the label to make forwarding decisions. It is not necessary to re-analyze the packet header. Because the label is a fixed length and unstructured, the MPLS forwarding-table lookup process is straightforward and fast.

Each label-switching router (LSR) in the network makes an independent, local decision as to which label value is used to represent which forwarding equivalence class. This association is known as a *label binding*. Each LSR informs its neighbors of the label bindings that it has made. When a labeled packet is sent from LSR A to neighboring LSR B, the label value carried by the packet is the label value that B assigned to represent the packet's forwarding equivalence class. Thus, the label value changes as the IP packet travels through the network.



Note

Because the Catalyst 3750 Metro switch is used as service-provider edge customer-located equipment (PE-CLE), rather than a service-provider core router, it does not normally operate as an LSR. The switch only performs label switching when it is connected to two different provider core routers over the ES ports to provide a redundant path. In this case, the switch uses QoS policies to classify MPLS packets on egress for label switching.

A label represents a forwarding-equivalence class, but it does not represent a particular path through the network. In general, the path through the network continues to be chosen by the existing Layer 3 routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Protocol (EIGRP), Intermediate-System-to-Intermediate-System (IS-IS), and Border Gateway Protocol (BGP). At each hop when a label is looked up, the choice of the next hop is determined by the dynamic routing algorithm.

Understanding MPLS VPNs

Using MPLS virtual private networks (VPNs) provides the capability to deploy and administer scalable Layer 3 VPN backbone services to business customers. A VPN is a secure IP-based network that shares resources on one or more physical networks. A VPN contains geographically dispersed sites that can communicate securely over a shared backbone.

VPN routes are distributed over the MPLS network by using multiprotocol BGP (MP-BGP), which also distributes the labels associated with each VPN route. MPLS VPN depends on VPN routing and forwarding (VRF) support to isolate the routing domains from each other. When routes are learned over an MPLS VPN, the switch learns the new route as a normal VRF route, except that the destination MAC address for the next hop is not the real address, but a specially formed address that contains an identifier that is allocated for the route. When an MPLS-VPN packet is received on a port, the switch looks up the labels in the routing table to determine what to do with the packet.

Each VPN is associated with one or more VPN VRF instances. A VRF includes routing and forwarding tables and rules that define the VPN membership of customer devices attached to the customer edge (CE) device. A customer site can be a member of multiple VPNs; however, a site can associate with only one VRF. A VRF has these elements:

- An IP routing table
- A Cisco Express Forwarding (CEF) table
- A set of interfaces that use the CEF forwarding table
- A set of rules and routing protocol parameters to control the information in the routing tables

A customer-site VRF contains all the routes available to the site from the VPNs to which it belongs. VPN routing information is stored in the IP routing table and the CEF table for each VRF. A separate set of tables is maintained for each VRF, which prevents information from being forwarded outside a VPN and prevents packets that are outside a VPN from being forwarded to a router within the VPN. Based on the routing information stored in the VRF IP routing table and the VRF CEF table, packets are forwarded to their destinations.

A PE router binds a label to each customer prefix that is learned from a CE device and includes the label in the network reachability information for the prefix that it advertises to other PE routers. When a PE router forwards a packet that is received from a CE device across the provider network, it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet, it examines the label and uses it to direct the packet to the correct CE device. A customer data-packet carries two levels of labels when traversing the backbone:

- The top label directs the packet to the correct PE router.
- The second label defines how that PE router should forward the packet to the CE device.

VPN Benefits

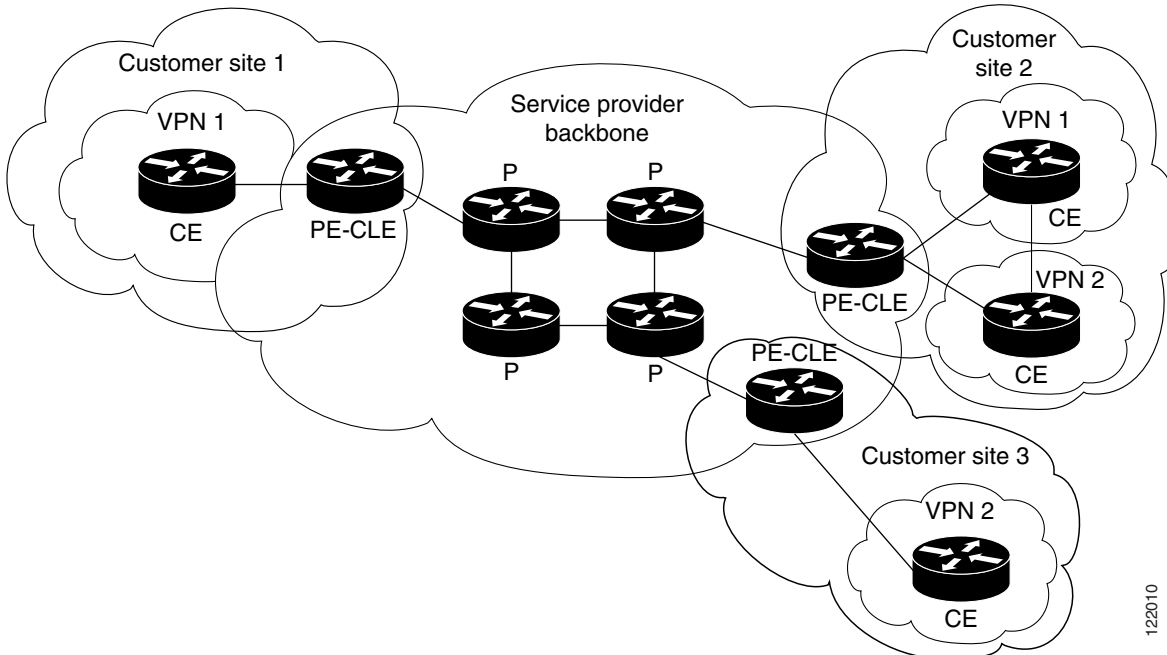
MPLS VPNs allow service providers to deploy scalable VPNs and build the foundation to deliver value-added services, including:

- Connectionless service—MPLS VPNs are connectionless, which means that no prior action is required to establish communication between hosts. A connectionless VPN does not require tunnels and encryption for network privacy.
- Centralized service—MPLS VPNs are seen as private intranets, which allows delivery of targeted IP services to a group of users represented by a VPN.

- Scalability—MPLS-based VPNs use the peer model and Layer 3 connectionless architecture to leverage a highly scalable solution. The peer model requires a customer site to act as a peer to one PE router as opposed to all other customer provider-edge or CE devices that are members of the VPN. The PE routers maintain VPN routes for those VPNs who are members. Routers in the core network do not maintain any VPN routes.
- Security—MPLS VPNs offer the same level of security as connection-oriented VPNs. Packets from one VPN do not inadvertently go to another VPN. Security provided at the edge of a provider network ensures that packets received from a customer are placed on the correct VPN; security provided at the backbone ensures that VPN traffic is kept separate.
- Easy to create—Because MPLS VPNs are connectionless, no specific point-to-point connection maps or topologies are required, and you can add sites to intranets and extranets to form closed user groups.
- Flexible addressing—Customers can continue to use their present address spaces without network address translation (NAT) because the MPLS VPN provides a public and private view of the address. A NAT is required only if two VPNs with overlapping address spaces want to communicate.
- Straightforward migration—You can build MPLS VPNs over multiple network architectures. Migration for the end customer is simplified because the CE router is not required to support MPLS, so no customer's intranet modifications are needed.
- MPLS VPN also provides increased BGP functionality.

Figure 37-1 shows an example of a VPN with a service-provider backbone network, PE-CLE routers, and CE devices.

Figure 37-1 VPNs with a Service-Provider Backbone



Each VPN contains customer devices attached to the CE devices. The customer devices use VPNs to exchange information between devices, and the provider routers (P) are not aware of the VPNs.


Figure 37-2 shows five customer sites communicating within three VPNs. The VPNs can communicate with these sites:

VPN1: Sites 2 and 4

VPN2: Sites 1, 3, and 4

VPN3: Sites 1, 3, and 5

Figure 37-2 Customer Sites with VPNs



Distribution of VPN Routing Information

The distribution of VPN routing information is controlled through the use of VPN route target communities, implemented by BGP extended communities. VPN routing information is distributed in this manner:

- When a VPN route learned from a CE device is added to the BGP process, a list of VPN route target extended community attributes is associated with it. The attribute values are obtained from an export list of route targets associated with the VRF from which the route was learned.
- An import list of route target extended communities is also associated with each VRF. The import list defines route target extended community attributes that a route must have in order for the route to be imported into the VRF. For example, if the import list for a particular VRF includes route target communities A, B, and C, then any VPN route that carries any of those route target extended communities—A, B, or C—is imported into the VRF.

A PE router can learn an IP prefix from a CE device by static configuration, through a BGP session with the CE device, or through the routing information protocol (RIP) exchange with the CE router. The IP prefix is a member of the IPv4 address family. After it learns the IP prefix, the PE router converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family and uniquely identifies the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses.

BGP distributes reachability information for VPN-IPv4 prefixes for each VPN. BGP communication takes place at two levels: within IP domains, known as autonomous systems (internal BGP or IBGP), and between autonomous systems (external BGP or EBGp). PE-to-PE sessions are IBGP sessions, and PE-CE sessions are EBGp sessions.

BGP propagates reachability information for VPN-IPv4 prefixes among PE routers by using the BGP multiprotocol extensions, which define support for address families other than IPv4. It does this in a way that ensures that the routes for a given VPN are learned only by other members of that VPN, which enables members of the VPN to communicate with each other.

Configuring MPLS VPNs

This section includes this information about configuring MPLS VPNs on a Catalyst 3750 Metro switch used as a PE router:

- [Default MPLS Configuration, page 37-6](#)
- [MPLS VPN Configuration Guidelines, page 37-6](#)

These sections describe the required tasks:

- [Enabling MPLS, page 37-7](#)
- [Defining VPNs, page 37-8](#)
- [Configuring BGP Routing Sessions, page 37-9](#)
- [Configuring PE-to-PE Routing Sessions, page 37-9](#)

You must also configure a PE to CE routing session. These sections provide example configurations:

- [Configuring BGP PE-to-CE Routing Sessions, page 37-10](#)
- [Configuring RIP PE-to-CE Routing Sessions, page 37-10](#)
- [Configuring Static Route PE-to-CE Routing Sessions, page 37-11](#)

For an example of packet flow in an MPLS VPN, see the “[Packet Flow in an MPLS VPN](#)” section on [page 37-11](#).

Default MPLS Configuration

By default, label switching of IPv4 packets along normally routed paths is globally enabled. MPLS forwarding of IPv4 packets is disabled by default on interfaces.

If no distribution protocol is explicitly configured by the **mpls label protocol** global configuration command, tag distribution protocol (TDP) is the default label distribution protocol for the switch. Cisco recommends that you configure label distribution protocol (LDP) for MPLS.

If no protocol is explicitly configured for an interface, the default label distribution protocol for the switch is used. By default, the labels of all destinations are advertised to all LDP neighbors.

No VRFs are defined. The default routing table for an interface is the global routing table.

MPLS VPN Configuration Guidelines

MPLS requires that CEF is enabled on the switch. CEF is enabled by default. For more information about CEF, see the “[Configuring Cisco Express Forwarding](#)” section on [page 34-82](#).

The switch must connect to the MPLS network through an ES port. MPLS configuration is supported only on ES ports.

Do not configure VLAN mapping on an interface configured for MPLS.

The switch supports a total of 26 VRFs and VPNs.

VRFs are not compatible with the PBR template. If you configure the PBR template by entering the **sdm prefer routing-pbr** command, any preconfigured VRFs are removed from the configuration. PBR and VRFs cannot function on the same switch.

Enabling MPLS

To use MPLS in a network, such as the one shown in [Figure 37-1](#), MPLS must be globally enabled and explicitly configured on the PE-CLE routers.

Beginning in privileged EXEC mode, follow these steps to incrementally deploy MPLS through a network, assuming that packets to all destination prefixes should be label-switched:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip cef	Enable CEF on the device if it has been disabled.
Step 3	mpls ip	Enable MPLS forwarding of IPv4 packets on normally routed paths if it has been disabled.
Step 4	mpls label protocol ldp	Set the label protocol on the switch to LDP. The default protocol is TDP.
Step 5	mpls ldp advertise-labels [for prefix-access-list [to peer-access-list]]	Enable MPLS label advertising on the switch. If no keywords are included, there are no restrictions on which labels are advertised. <ul style="list-style-type: none"> • (Optional) for prefix-access-list—Specify which destinations should have their labels advertised. • (Optional) to peer-access-list—Specify which LDP neighbors should receive label advertisements. A label switch router (LSR) is identified by its router ID, which consists of the first 4 bytes of its 6-byte LDP identifier.
Step 6	interface interface-id	Enter interface configuration mode, and specify the Layer 3 ES interface connected to the MPLS network.
Step 7	mpls ip	Enable MPLS forwarding of IPv4 packets along normally routed paths for the interface.
Step 8	end	Return to privileged EXEC mode.
Step 9	show mpls forwarding table show mpls interfaces	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Repeat these steps for every PE-CLE router in the network and the appropriate interfaces until all routers and connected interfaces are enabled for MPLS.

Use the **no mpls ip** global configuration command to disable MPLS on the switch. Use the **no mpls label protocol ldp** global configuration command to return to the default TDP.

Defining VPNs

Beginning in privileged EXEC mode, follow these steps to define VPN routing instances on the PE-CLE router:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing (required only if routing is disabled).
Step 3	ip vrf <i>vrf-name</i>	Enter VRF configuration mode, and define the VPN routing instance by assigning a VRF name.
Step 4	rd <i>route-distinguisher</i>	Create a VRF table by specifying a route distinguisher. Enter either an AS number and an arbitrary number (xxx:y) or an IP address and arbitrary number (A.B.C.D:y).
Step 5	route-target { export import both } <i>route-target-ext-community</i>	Create a list of import, export, or import and export route target communities for the specified VRF. Enter either an AS system number and an arbitrary number (xxx:y) or an IP address and an arbitrary number (A.B.C.D:y). The <i>route-target-ext-community</i> should be the same as the <i>route-distinguisher</i> entered in Step 4.
Step 6	import map <i>route-map</i>	(Optional) Associate the specified import route map with the VRF.
Step 7	export map <i>route-map</i>	(Optional) Associate the specified export route map with the VRF.
Step 8	exit	Return to global configuration mode.
Step 9	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 ES or VLAN interface to be associated with the VRF.
Step 10	ip vrf forwarding <i>vrf-name</i>	Associate the Layer 3 interface with the VRF.
Step 11	end	Return to privileged EXEC mode.
Step 12	show ip vrf	Display the defined VRFs and interfaces.
Step 13	show ip route vrf show ip cef vrf <i>vrf-name</i>	Display the IP routing table for a VRF. Display the CEF forwarding table associated with a VRF.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip vrf** *vrf-name* global configuration command to delete a VRF and remove all interfaces from it. Use the **no ip vrf forwarding** interface configuration command to remove an interface from a VRF.

Configuring BGP Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE-CLE router to configure BGP routing sessions in a provider network:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing (required only if routing is disabled).
Step 3	router bgp <i>autonomous-system-number</i>	Enable a BGP routing process, assign it the AS number passed to the other BGP routers, and enter router configuration mode. The AS number can be from 1 to 65535, with 64512 to 65535 designated as private autonomous numbers.
Step 4	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>	Specify a neighbor IP address or BGP peer group that identifies it to the local autonomous system. The AS number can be from 1 to 65535.
Step 5	neighbor <i>ip-address</i> activate	Activate the advertisement of the IPv4 address family.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip bgp neighbor	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system* global configuration command to delete the BGP routing session.

Configuring PE-to-PE Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE-CLE router to configure a PE-to-PE routing session in a provider network that uses IBGP:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Enter router configuration mode.
Step 3	address-family vpnv4 [unicast]	Enter address family configuration mode to configure routing sessions that use standard VPNv4 address prefixes. (Optional) unicast —Specify VPNv4 unicast address prefixes.
Step 4	neighbor <i>ip-address</i> remote-as <i>as-number</i>	Define an IBGP session between the PE routers.
Step 5	neighbor <i>ip-address</i> activate	Activate the advertisement of the IPv4 address family.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip bgp [ipv4] [neighbors] [vpnv4]	Verify BGP configuration. Display information about all BGP IPv4 prefixes.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system* global configuration command to delete the BGP routing session.

Configuring BGP PE-to-CE Routing Sessions

Beginning in privileged EXEC mode, follow these steps on the PE-CLE router to configure a BGP PE-to-CE routing session:

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>router bgp <i>autonomous-system-number</i></code>	Configure the BGP routing process with the AS number passed to other BGP routers, and enter router configuration mode.
Step 3	<code>address-family ipv4 [unicast] vrf <i>vrf-name</i></code>	Define EPGP parameters for PE-to-CE routing sessions, and enter VRF address-family configuration mode. Note The default is <i>off</i> for auto-summary and synchronization in the VRF address-family configuration mode.
Step 4	<code>neighbor <i>address</i> remote-as <i>as-number</i></code>	Define an EBGP session between PE and CE routers.
Step 5	<code>neighbor <i>address</i> activate</code>	Activate the advertisement of the IPv4 address family.
Step 6	<code>end</code>	Return to privileged EXEC mode.
Step 7	<code>show ip bgp [ipv4] [neighbors]</code>	Verify BGP configuration. Display information about all BGP IPv4 prefixes.
Step 8	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

Use the `no router bgp as-number` global configuration command to delete the BGP routing session.

Configuring RIP PE-to-CE Routing Sessions



Note

You can also use the OSPF routing protocol for PE-to-CE routing sessions.

Beginning in privileged EXEC mode, follow these steps on the PE-CLE router to configure RIP PE-to-CE routing:

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>router rip</code>	Enable RIP routing, and enter router configuration mode.
Step 3	<code>address-family ipv4 [unicast] vrf <i>-name</i></code>	Define RIP parameters for PE to CE routing sessions, and enter VRF address-family configuration mode. Note The default is <i>off</i> for auto-summary and synchronization in the VRF address-family configuration mode.
Step 4	<code>network <i>prefix</i></code>	Enable RIP on the PE-to-CE link.
Step 5	<code>end</code>	Return to privileged EXEC mode.
Step 6	<code>show ip rip database [<i>network-prefix</i>]</code>	Verify the configuration.
Step 7	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

Use the `no router rip` global configuration command to disable RIP routing.

Configuring Static Route PE-to-CE Routing Sessions

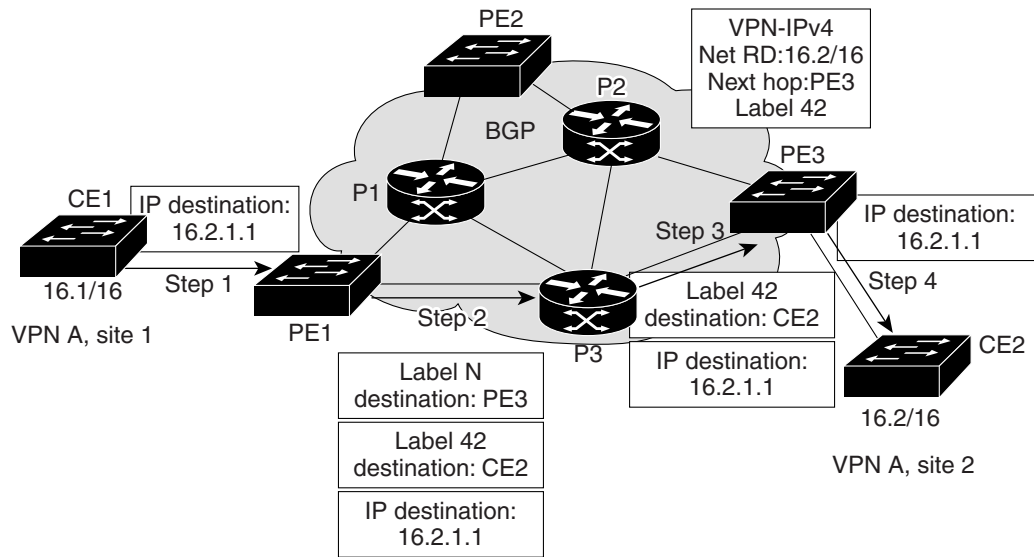
Beginning in privileged EXEC mode, follow these steps on the PE-CLE router to configure static routing:

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>ip route vrf vrf-name prefix mask</code>	Define the VRF static routing table to use for PE-to-CE sessions.
Step 3	<code>router bgp autonomous-system-number</code>	Enter the BGP routing process AS number, and enter router configuration mode.
Step 4	<code>address-family ipv4 [unicast] vrf vrf-name</code>	Define static route parameters for every PE-to-CE routing session, and enter VRF address-family configuration mode. Note The default is <i>off</i> for auto-summary and synchronization in the VRF address-family configuration mode.
Step 5	<code>redistribute static</code>	Redistribute VRF static routes into the VRF BGP table.
Step 6	<code>redistribute connected</code>	Redistribute directly connected networks into the VRF BGP table.
Step 7	<code>end</code>	Return to privileged EXEC mode.
Step 8	<code>show ip bgp [ipv4]</code>	Verify the configuration.
Step 9	<code>copy running-config startup-config</code>	(Optional) Save your entries in the configuration file.

Packet Flow in an MPLS VPN

Figure 37-3 is an example of packet flow between two customer sites in an MPLS VPN network.

Figure 37-3 Sample MPLS VPN Packet Flow



A customer (Fast Ethernet) port on switch PE1 is configured for routed operation in a VPN. The port uses static routing or a routing protocol (RIP, OSPF, EIGRP, or BGP) to forward packets. MP-BGP is configured over the PE1 switch ES port with a route distinguisher (RD) that is associated with the customer's VPN. MP-BGP is configured to redistribute the routes and their associated VPN labels over the ES port that is using this RD.

The packet flow follows these steps:

-
- Step 1** Provider-edge switch PE1 (which could be a Catalyst 3750 Metro switch) receives a packet from the customer switch at site 1. The switch determines from the lookup table that the VRF is a VLAN running MPLS and uses the MPLS lookup table to determine what to do with the packet. The MPLS lookup table contains the peer LSR as the destination MAC address and the local interface as the source MAC address.
 - Step 2** PE1 finds a BGP route with the appropriate next hop and labels, adds the appropriate labels to the packet, and forwards the packet out of the ES port to the next hop router (P3).
 - Step 3** The P3 router receives the packet and forwards it over the MPLS-VPN network, based on the packet's top label—the interior gateway protocol (IGP) label—and then removes the top label.
 - Step 4** PE3 receives the packet, removes the MPLS encapsulation, and forwards the packet by using the VRF interface associated with the VPN label contained in the packet that has the customer-edge switch CE2 as the destination.
-

Understanding EoMPLS

Any Transport over MPLS (AToM) is a solution for transporting Layer 2 packets over an MPLS network, allowing service providers to use the MPLS network to provide connectivity between customer sites with existing Layer 2 networks. Instead of separate networks with network management environments, service providers can use the MPLS network to transport all types of traffic for different customers. The Catalyst 3750 Metro switch supports EoMPLS, a subset of AToM that uses a tunneling mechanism to carry Layer 2 Ethernet traffic.

EoMPLS encapsulates Ethernet frames in MPLS packets and forwards them across the MPLS network. Each frame is transported as a single packet, and the PE routers connected to the backbone add and remove labels as appropriate for packet encapsulation:

- The ingress PE router receives an Ethernet frame and encapsulates the packet by removing the preamble, the start of frame delimiter (SFD), and the frame check sequence (FCS). The rest of the packet header is not changed.
- The ingress PE router adds a point-to-point virtual connection (VC) label and a label switched path (LSP) tunnel label for normal MPLS routing through the MPLS backbone.
- The network core routers use the LSP tunnel label to move the packet through the MPLS backbone and do not distinguish Ethernet traffic from any other types of packets in the MPLS backbone.
- At the other end of the MPLS backbone, the egress PE router receives the packet and de-encapsulates the packet by removing the LSP tunnel label if one is present. The PE router also removes the VC label from the packet.
- The PE router updates the header, if necessary, and sends the packet out the appropriate interface to the destination switch.

The MPLS backbone uses the tunnel labels to transport the packet between the PE routers. The egress PE router uses the VC label to select the outgoing interface for the Ethernet packet. EoMPLS tunnels are unidirectional; for bidirectional EoMPLS, you need to configure one tunnel in each direction.

The point-to-point VC requires you to configure VC endpoints at the two PE routers. Only the PE routers at the ingress and egress points of the MPLS backbone know about the VCs dedicated to transporting Layer 2 traffic. Other routers do not have table entries for these VCs.

This section includes additional information about these topics:

- [Interaction with Other Features](#), page 37-13
- [EoMPLS Limitations](#), page 37-14

Interaction with Other Features

This section describes how EoMPLS interacts other features. It includes these sections:

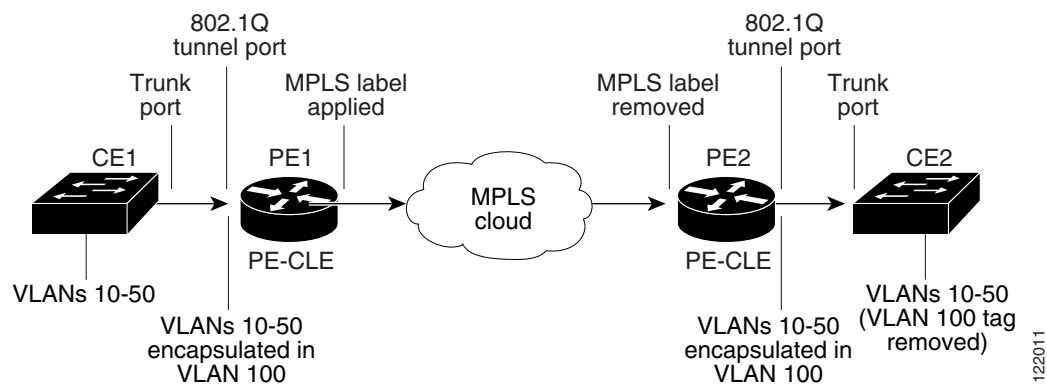
- [EoMPLS and 802.1Q Tunneling](#), page 37-13
- [EoMPLS and Layer 2 Tunneling](#), page 37-14
- [EoMPLS and QoS](#), page 37-14

EoMPLS and 802.1Q Tunneling

IEEE 802.1Q tunneling enables service providers to use a single VLAN to support customers who have multiple VLANs, while preserving customer VLAN IDs and segregating traffic in different VLANs. For more information about 802.1Q tunneling, see [Chapter 16, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

[Figure 37-4](#) is an example configuration where 802.1Q-tunneled traffic is forwarded using EoMPLS over an MPLS network. To support 802.1Q tunneling in a topology where a Layer 2 device connects to an MPLS network through a switch functioning as a PE-CLE, the ingress LAN port on the PE-CLE that receives the 802.1Q tunnel-encapsulated traffic (PE1) is configured as a tunnel port that accepts VLAN 100 traffic. On PE1, the interface is configured for port-based EoMPLS forwarding, with PE2 as the destination IP address. When packets from VLANs 10 to 50 arrive from CE1, they are encapsulated in VLAN 100 and sent to the PE1 egress port that is connected to the MPLS network. At the egress port, an MPLS tag is added to the frame header before it is mapped to a VC and forwarded to the next MPLS PE-CLE (PE2).

Figure 37-4 EoMPLS Example



By entering the **mpls l2transport route** or the **xconnect** interface configuration command on either a VLAN for VLAN-based EoMPLS or an Ethernet port for port-based EoMPLS, you can configure an EoMPLS tunnel to forward traffic based on either the customer VLAN or the Ethernet port.

- To forward 802.1Q tunnel-encapsulated traffic through the MPLS core to a specific recipient on the other side of the MPLS network, configure port-based EoMPLS.
- To forward 802.1Q tunnel-encapsulated traffic from an access device to a PE router, configure VLAN-based EoMPLS.

EoMPLS and Layer 2 Tunneling

Layer 2 protocol tunneling over an EoMPLS link allows CDP, STP, and VTP protocol data units (PDUs) to be tunneled through an MPLS network. To support Layer 2 protocol tunneling when the Layer 2 device connects to an MPLS network through a switch functioning as a PE, you configure the ingress port on the PE that receives the Layer 2 protocol traffic as a tunnel port. The Layer 2 protocol traffic is encapsulated before it is forwarded over the MPLS network. For more information about Layer 2 protocol tunneling, see [Chapter 16, “Configuring IEEE 802.1Q and Layer 2 Protocol Tunneling.”](#)

EoMPLS and QoS

EoMPLS supports QoS by using three experimental bits in a label to determine the priority of packets. To support QoS between label edge routers (LERs), you set the experimental bits in both the VC and tunnel labels. EoMPLS QoS classification occurs on ingress, and you can only match on Layer 3 parameters (such as IP or DSCP), not Layer 2 parameters (CoS). See the [“Configuring MPLS and EoMPLS QoS” section on page 37-18](#) for more information about EoMPLS and QoS.

EoMPLS Limitations

These restrictions apply to EoMPLS:

- EoMPLS requires that at least one of the two ES ports be configured for MPLS. Therefore, if you want to run EoMPLS on an ES port, you can only configure it on the ES port that is not configured for MPLS.
- MTU—EoMPLS does not support packet fragmentation and reassembly. Therefore, the maximum transmission unit (MTU) of all intermediate links between endpoints must be sufficient to carry the largest Layer 2 VLAN received. The ingress and egress PE routers must have the same MTU value.
- Address Format—All loopback addresses on PE routers must be configured with 32-bit masks to ensure proper operation of MPLS forwarding. OSPF requires the use of loopback addresses.
- Packet Format—EoMPLS supports VLAN packets that conform to the IEEE 802.1Q standard. ISL encapsulation is not supported between PE and CE routers.
- The maximum number of VLANs using EoMPLS on a switch is 1005.
- Layer 2 connection restrictions:
 - You cannot have a direct Layer 2 connection between PE routers with EoMPLS.
 - You cannot have more than one Layer 2 connection between routers if those routers are configured to transport Ethernet VLANs over the MPLS backbone. Adding a second Layer 2 connection causes the spanning-tree state to constantly toggle if you disable spanning tree on the peer router.

- EoMPLS and trunking restrictions:
 - To support Ethernet spanning-tree bridge protocol data units (BPDUs) across an EoMPLS backbone, you must disable spanning tree for the EoMPLS VLAN. This ensures that the EoMPLS VLANs are carried only on the trunk to the customer switch.
 - The native VLAN of a trunk cannot be an EoMPLS VLAN.
- You can enable EoMPLS on 802.1Q interfaces by using the **mpls l2transport route** or the **xconnect** interface configuration command.
- Do not configure VLAN mapping on an interface configured for EoMPLS.
- Do not configure EoMPLS on private-VLAN interfaces.

Enabling EoMPLS

This section includes this information about configuring EoMPLS on a switch used as a PE router:

- [Default EoMPLS Configuration, page 37-15](#)
- [EoMPLS Configuration Guidelines, page 37-15](#)
- [Configuring EoMPLS, page 37-16](#)
- [Packet Flow in an EoMPLS Network, page 37-17](#)

Default EoMPLS Configuration

By default, EoMPLS is not configured.

The **mpls ldp router-id** command is disabled. No virtual connections are configured.

EoMPLS Configuration Guidelines

When you configure EoMPLS, you must follow these guidelines:

- EoMPLS requires that at least one of the two ES ports be configured for MPLS. Therefore, if you want to run EoMPLS on an ES port, you can only configure it on the ES port that is not configured for MPLS.
- Before enabling EoMPLS, you must enable dynamic MPLS labeling by using the **mpls ip** interface configuration command on all paths between the imposition and disposition LERs. MPLS is globally enabled by default.
- For VLAN-based EoMPLS, you must configure VLANs on the switch.
- EoMPLS operation between two PE routers requires an LDP session between the routers. The IP address used by each router as its LDP router ID must be IP-reachable by the other. Use the optional **mpls ldp router-id** global configuration command to control the selection of the LDP router ID by specifying the interface whose IP address should be used.
 - If the specified interface is up and has an IP address, you can use the command without the optional **force** keyword. When the router ID is selected, that IP address is selected as the router ID.

- If the specified interface is not up or does not have an IP address, use the **force** keyword with the command to ensure that the IP address of the specified interface is used when that interface is brought up.
- Both PE routers require a loopback address that you can use to create a VC between the routers. When you use OSPF as the interior gateway protocol, you must configure all loopback addresses on PE routers with 32-bit masks to ensure proper operation of MPLS forwarding between the PE routers.
- Do not configure EoMPLS on an interface configured for VLAN mapping.

Configuring EoMPLS

You configure VLAN-based EoMPLS on a VLAN interface. When VLAN-based EoMPLS is enabled, the switch associates the tunnel and VC labels based on the VLAN ID. You use the same commands to enable port-based EoMPLS on an ES interface.

Beginning in privileged EXEC mode, follow these steps on the PE-CLE routers to configure EoMPLS to transport Layer 2 packets between two endpoints:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mpls label protocol ldp	Enable LDP for all interfaces. By default, TDP is enabled. This command causes all interfaces to use LDP.
Step 3	interface loopback0	Enter interface configuration mode for a loopback interface.
Step 4	ip address <i>ip-address subnet mask</i>	Assign an IP address to the loopback interface.
Step 5	exit	Return to global configuration mode
Step 6	mpls ldp router-id loopback0 force	(Optional) Force the IP address of loopback interface 0 to be used as the router ID.
Step 7	interface <i>interface-id</i>	Enter a Layer 3 VLAN (for VLAN-based EoMPLS) or the interface ID of an ES port (for port-based EoMPLS), and enter interface configuration mode.
Step 8	mpls l2transport route <i>destination vc-id</i> or xconnect <i>destination vc-id encapsulation mpls</i>	Configure the interface to transport the Layer 2 VLAN packets over MPLS. <ul style="list-style-type: none"> • <i>destination</i>—The IP address of the PE router at the other end of the VC. • <i>vc-id</i>—A unique value defined for the VC. The VC-ID connects the endpoints of the VC and must be the same on both ends of the VC. The range is from 1 to 4294967295.
Step 9	end	Return to privileged EXEC mode.
Step 10	show mpls l2transport vc	Verify the configuration.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no mpls l2transport route *destination vc-id*** or **no xconnect *destination vc-id encapsulation mpls*** interface command to delete the EoMPLS tunnel.

This example shows how to configure an EoMPLS tunnel between switch PE1's VLAN 3 interface and PE2's VLAN 4 interface.

PE1 has an IP address 10.0.0.1/32, and PE2 has IP address 20.0.0.1/32. Both PE routers are configured with an MPLS connection to the MPLS core. The VC ID is 123.

Enter these commands on the PE1 switch:

```
Switch(config)# interface loopback0
Switch(config-if)# ip address 10.10.10.10 255.255.255.255
Switch(config-if)# exit
Switch(config)# interface vlan 3
Switch(config-if)# mpls 12transport route 20.0.0.1 123
```

Enter these commands on the PE2 switch:

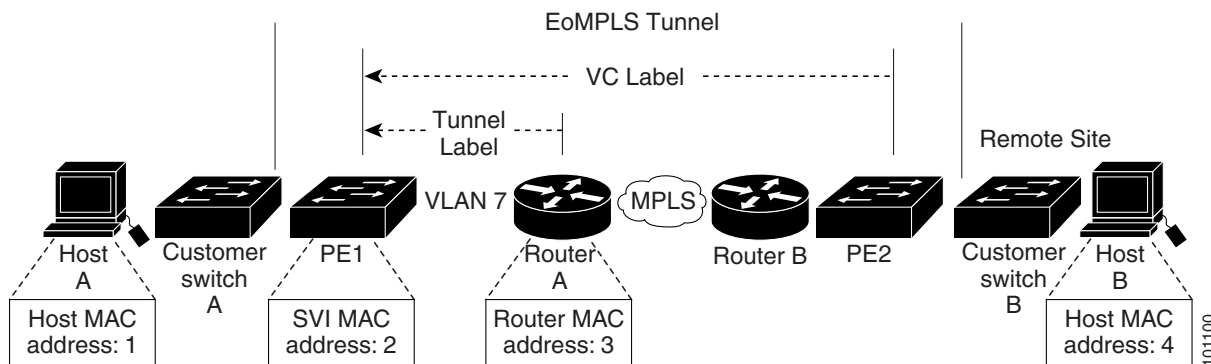
```
Switch(config)# interface loopback0
Switch(config-if)# ip address 20.20.20.20 255.255.255.255
Switch(config-if)# exit
Switch(config)# interface vlan 4
Switch(config-if)# mpls 12transport route 10.0.0.1 123
```

Packet Flow in an EoMPLS Network

Figure 37-5 is an example of packet flow in an EoMPLS network. A customer port on PE1 is configured for a per-port EoMPLS tunnel to a remote customer port on PE2. This allows the two physically separated customer switches (A and B) connected to these ports to appear as if they are directly connected on the same physical LAN.

The EoMPLS tunnel is configured with the IP address of Switch B and a VC ID that is associated with the remote customer port. PE1 establishes a tunnel LSP with PE2 by using a label advertised with LDP by Router A, which is connected to the ES port of PE1. PE1 then establishes a targeted LDP session to PE2 to advertise the VC label associated with the VC ID. When PE2 is configured with the EoMPLS tunnel, it also establishes a targeted LDP session to advertise the VC label it associated to the VC ID. This establishes an EoMPLS tunnel between the two ES ports on switch PE1 and switch PE2.

Figure 37-5 Sample EoMPLS Packet Flow



Assume that Host A is connected to the customer switch on VLAN 3 that has a trunk port connected to PE1 configured for 802.1Q tagging. Host A sends a packet to Host B, using the specific values of MAC addresses, labels, and VLANs shown in the figure. The customer switch tags the host packet and forwards it over the trunk port to PE1.

The tagged packet is received on the CE port that is configured for per-port EoMPLS tunneling. The PE1 switch examines the packet headers and looks at the tables stored in the switch to determine what to do with the packet. Because the port is configured for per-port EoMPLS tunneling, the switch does not remove any VLAN tags that are in the packet, but assigns the packet to an internal VLAN. Only the customer port and the ES port are configured with that internal VLAN, which makes the PE1 ES port the only possible destination for the packet.

The ES port encapsulates the packet header with the tunnel label and the VC label and forwards the packet to the next hop, in this case Router A, for it to be sent across the MPLS network.

The router receives the packet and forwards it over the MPLS network to the remote PE2 switch. PE2 removes the MPLS encapsulation and sends the packet out of the port associated with the VC label. Customer Switch B removes the final VLAN tag and forwards the packet to the remote host B.

VLAN-based EoMPLS packet flow is basically the same as port-based EoMPLS, except that the customer VLAN is used instead of an internal VLAN. The PE1 switch looks up the customer VLAN ID to determine that the packet is forwarded to the ES port, where the packets is again examined and encapsulated with the tunnel label and VC label based on the EoMPLS for that VLAN.

Configuring MPLS and EoMPLS QoS

Quality of service (QoS) in MPLS and EoMPLS enables network administrators to provide differentiated types of service across an MPLS network. Each packet can receive the particular kind of service specified by the packet QoS. To preserve QoS IP precedence bits, you must globally disable QoS.

After you enable QoS, you can preserve Differentiated Services Code Point (DSCP) or IP precedence bits by using a trusted configuration at the interface level. For more information, see the [“Configuring Ingress Classification by Using Port Trust States” section on page 32-49](#). However, the unreserved bits are automatically overwritten by the value of the preserved bits. For example, if you preserve the DSCP bits, the IP precedence and CoS bits are overwritten with the value of the DSCP bits. You can also set MPLS and EoMPLS QoS priority by using 3 experimental bits in the MPLS label to determine the priority of packets.



Note

The switch supports only DSCP and IP precedence classification for MPLS and EoMPLS.

This section contains this information:

- [Understanding MPLS QoS, page 37-18](#)
- [Enabling MPLS and EoMPLS QoS, page 37-20](#)

Understanding MPLS QoS

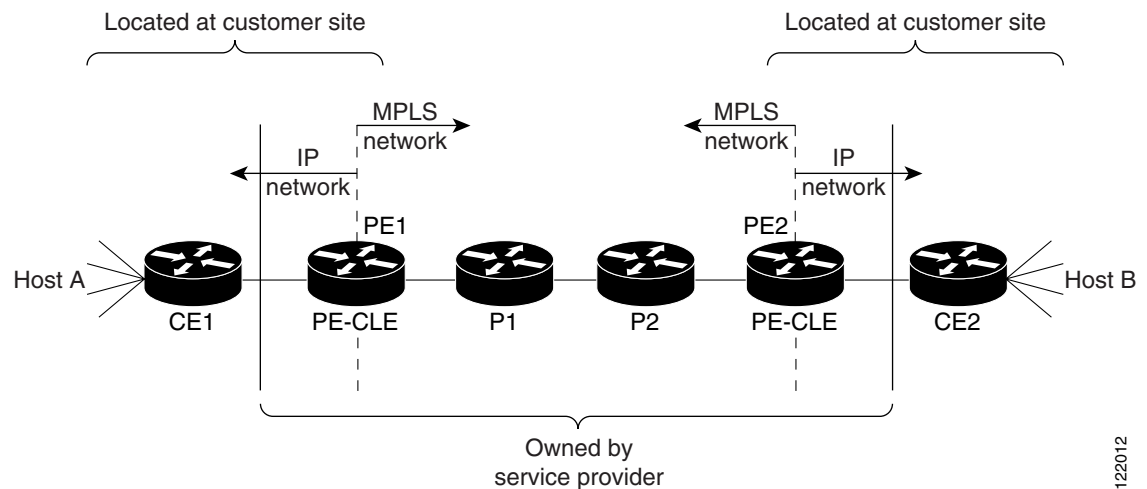
Service in an MPLS network can be specified in different ways, for example, by using the IP precedence bit settings in IP packets. When you send IP packets from one site to another, the IP precedence field (the first 3 bits of the DSCP field in the header of an IP packet) specifies the QoS. Based on the IP precedence marking, the packet is given the desired treatment such as latency or bandwidth. If the network is an MPLS network, the IP precedence bits are copied into the MPLS EXP field at the edge of the network.

A service provider might also want to set a QoS value for an MPLS packet to a different value. Instead of overwriting the value in the IP precedence field that belongs to a customer, the service provider can set the MPLS experimental field. The IP header remains available for the customer's use, and the QoS of an IP packet is not changed as the packet travels through the MPLS network.

By choosing different values for the MPLS experimental field, you can mark packets based on their characteristics, such as rate or type, so that packets have the priority that they require during periods of congestion.

Figure 37-6 shows an MPLS network that connects two sites of an IP network that belongs to a customer.

Figure 37-6 MPLS Network Connecting Two Customer Sites



PE1 and PE2 are customer-located routers at the boundaries between the MPLS network and the IP network and are the ingress and egress PE devices. CE1 and CE2 are customer edge devices. P1 and P2 are service provider routers within the core of the service-provider network.

Packets arrive as IP packets at PE1, the ingress PE-CLE router, and PE1 sends the packets to the MPLS network as MPLS packets. Within the service-provider network, there is no IP precedence field for the queuing mechanism to look at because the packets are MPLS packets. The packets remain as MPLS packets until they arrive at PE2, the egress PE-CLE router. PE2 removes the label from each packet and forwards the packets as IP packets.

MPLS QoS enables service providers to classify packets according to their type, input interface, and other factors by setting (marking) each packet within the MPLS experimental field without changing the IP precedence or DSCP field. You can use the IP Precedence or DSCP bits to specify the QoS for an IP packet and use the MPLS experimental bits to specify the QoS for an MPLS packet. In an MPLS network, configure the MPLS experimental field value at PE1 (the ingress router) to set the QoS value in the MPLS packet.

It is important to assign the correct priority to a packet. The priority of a packet affects how the packet is treated during periods of congestion. For example, service providers have service-level agreements with customers that specify how much traffic the service provider has agreed to deliver. To comply with the agreement, the customer must not send more than the agreed-upon rate. Packets are considered to be in-rate or out-of-rate. If there is congestion in the network, out-of-rate packets might be dropped more aggressively.

122012

Enabling MPLS and EoMPLS QoS

This section describes how to configure MPLS QoS on the ingress PE router. It includes these topics:

- [Default MPLS and EoMPLS QoS Configuration, page 37-20](#)
- [Setting the Priority of Packets with Experimental Bits, page 37-20](#)

For more information about QoS, see [Chapter 32, “Configuring QoS.”](#)

Default MPLS and EoMPLS QoS Configuration

QoS is disabled. Packets are not modified, and the CoS, DSCP, and IP precedence values in the packet are not changed. Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

The default behavior for the VLAN-based EoMPLS packets is to relay the 802.1p bits into the EXP bits of the VC and tunnel labels. The default behavior for the port-based EoMPLS packets is to use a value of 0 in the EXP bits of the VC and tunnel labels. You can change the default behavior for VLAN- or port-based EoMPLS by applying a hierarchical QoS policy on an ES port.

When QoS is enabled with the `mls qos` global configuration command and all other QoS settings are at their defaults, traffic is classified as best effort (the DSCP value is set to 0) without any policing. No policy maps are configured.



Note

For MPLS and EoMPLS QoS, you can match only Layer 3 parameters (IP or DSCP values), not Layer 2 (CoS values).

Setting the Priority of Packets with Experimental Bits

MPLS and EoMPLS provide QoS on the ingress router by using 3 experimental bits in a label to determine the priority of packets. To support QoS between LERs, set the experimental bits in both the VC and tunnel labels. If you do not assign values to the experimental bits, the priority bits in the 802.1Q header tag control information field are written into the experimental bit fields.

The process includes these steps on the ingress router:

- Configure a class map to classify IP packets according to their DSCP or IP precedence classification.



Note

The switch supports only DSCP and IP precedence classification for MPLS and EoMPLS.

- Configure a policy map to mark MPLS packets (write their classification into the MPLS experimental field).
- Configure the input interface to attach the service policy.

Beginning in privileged EXEC mode, follow these steps to set the experimental bits for EoMPLS or MPLS QoS:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS globally. QoS runs from the default settings described in Chapter 32, “Configuring QoS.”
Step 3	class-map <i>class-map-name</i>	Specify the name of the traffic class, and enter class-map configuration mode.
Step 4	match { ip dscp <i>dscp-list</i> ip precedence <i>ip-precedence-list</i> }	Specify the matching criteria for 802.1Q packets. <ul style="list-style-type: none"> ip dscp <i>dscp-list</i>—list of up to eight IP DSCP values to match against incoming packets. The range is 0 to 63. ip precedence <i>ip-precedence-list</i>—list of up to eight IP-precedence values to match against incoming packets. Separate each value with a space. The range is 0 to 7. <p>Note The switch does not support the cos and vlan keywords for MPLS and EoMPLS.</p>
Step 5	exit	Return to global configuration mode.
Step 6	policy-map <i>policy-map-name</i>	Specify the name of the traffic policy to configure, and enter policy-map configuration mode.
Step 7	class <i>class-name</i>	Specify the name of the predefined traffic class configured with the class-map command, and enter policy-map class configuration mode.
Step 8	set mpls experimental <i>exp-number</i>	Specify the value to which the MPLS bits are set if the packets match the specified policy map. The range is 0 to 7.
Step 9	exit	Return to policy-map configuration mode.
Step 10	exit	Return to global configuration mode.
Step 11	interface <i>interface-id</i>	Enter the interface ID, and enter interface configuration mode. The interface should be the ES egress port of the ingress router.
Step 12	service-policy output <i>policy-map-name</i>	Attach the specified policy map to the output interface.
Step 13	end	Return to privileged EXEC mode.
Step 14	show policy-map [policy-map-name [class <i>class-map-name</i>]] show policy-map interface <i>interface-id</i>	Verify the configuration.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an existing policy map, use the **no policy-map** *policy-map-name* global configuration command. To delete an existing class, use the **no class** *class-name* policy-map configuration command.

This example shows how to use class and policy maps to configure different experimental bit settings for DSCP and IP precedence for MPLS QoS.

```
Switch(config)# class-map match-all gold-class
Switch(config-cmap)# match ip dscp 1
Switch(config-cmap)# exit
Switch(config)# class-map match-all silver-class
```

```

Switch(config-cmap) # match ip precedence 2
Switch(config-cmap) # exit

Switch(config) # policy-map out-policy
Switch(config-pmap) # class gold-class
Switch(config-pmap-c) # set mpls experimental 5
Switch(config-pmap-c) # exit
Switch(config-pmap) # class silver-class
Switch(config-pmap-c) # set mpls experimental 4
Switch(config-pmap-c) # exit

Switch(config) # interface gigabitethernet1/1/1
Switch(config-if) # service-policy output out-policy
Switch(config-if) # end

```

Monitoring and Maintaining MPLS and EoMPLS

To clear MPLS counters or display MPLS and EoMPLS information, use the privileged EXEC commands in [Table 37-1](#).

Table 37-1 Commands for Displaying MPLS and EoMPLS Information

Command	Purpose
<code>clear mpls counters</code>	Clear MPLS forwarding counters.
<code>show mpls forwarding-table</code>	Display the contents of the MPLS label forwarding information base (LFIB).
<code>show mpls interfaces</code>	Display information about interfaces that have been configured for label switching.
<code>show mpls ip binding</code>	Display specified information about label bindings learned by LDP.
<code>show mpls l2transport vc [detail] [summary]</code>	Display detailed or summary information about the EoMPLS virtual connections that have been enabled to route Layer 2 packets on a provider-edge device.
<code>show mpls l2transport vc [vc-id] [vc-id-min - vc-id-max]</code>	Display information about the specified VC or range of VCs. The range is from 1 to 4294967295.
<code>show mpls label range</code>	Display the range of local labels available for use on packet interfaces.
<code>show mpls ldp bindings</code>	Display the contents of the label information base (LIB).
<code>show mpls ldp discovery</code>	Display the status of the LDP discovery process.
<code>show mpls ldp neighbor</code>	Display the status of LDP sessions.
<code>show mpls ldp parameters</code>	Display current LDP parameters.
<code>show mpls prefix-map</code>	Show the prefix map used to assign a QoS map to network prefixes that match a standard IP access list.
<code>show mpls ldp backoff</code>	Display information about the configured session setup backoff parameters and any potential LDP peers with which session setup attempts are being throttled.