



CHAPTER 21

Configuring Multiprotocol Label Switching

This chapter describes how to configure Multiprotocol Label Switching (MPLS) on a Catalyst 6500 series switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to these publications:

- The *Catalyst Supervisor Engine 32 PISA Cisco IOS Command Reference*, Release 12.2ZY, at this URL:
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2ZY/command/reference/cmdref.html>
 - The Release 12.2 publications at this URL:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_installation_and_configuration_guides_list.html
-

This chapter consists of these sections:

- [MPLS Label Switching](#), page 21-1
- [VPN Switching](#), page 21-9
- [Any Transport over MPLS](#), page 21-13

MPLS Label Switching

These sections describe MPLS label switching:

- [Understanding MPLS](#), page 21-2
- [Understanding MPLS Label Switching](#), page 21-2
- [Supported Hardware Features](#), page 21-4
- [Supported Cisco IOS Features](#), page 21-5
- [MPLS Guidelines and Restrictions](#), page 21-7
- [Configuring MPLS](#), page 21-7
- [MPLS Per-Label Load Balancing](#), page 21-7
- [MPLS Configuration Examples](#), page 21-8

Understanding MPLS

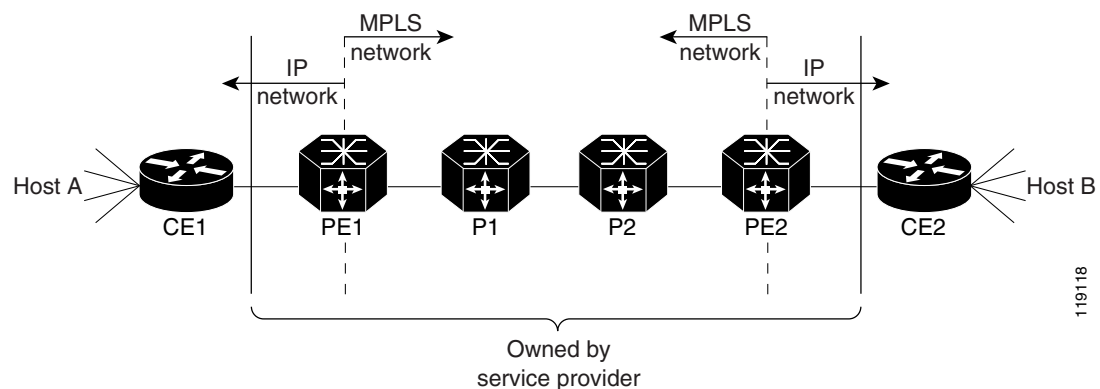
MPLS uses label switching to forward packets over various link-level technologies such as Packet-over-SONET (POS), Frame Relay, ATM, and Ethernet. Labels are assigned to packets based on groupings or forwarding equivalence classes (FECs). The label is added between the Layer 2 and the Layer 3 header.

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 21-1 shows an MPLS network of a service provider that connects two sites of a customer network.

Figure 21-1 MPLS Network



For additional information on MPLS, see this publication:

http://www.cisco.com/en/US/docs/ios/12_2/switch/configuration/guide/xcftagov_ps1835_TSD_Products_Configuration_Guide_Chapter.html

Understanding MPLS Label Switching

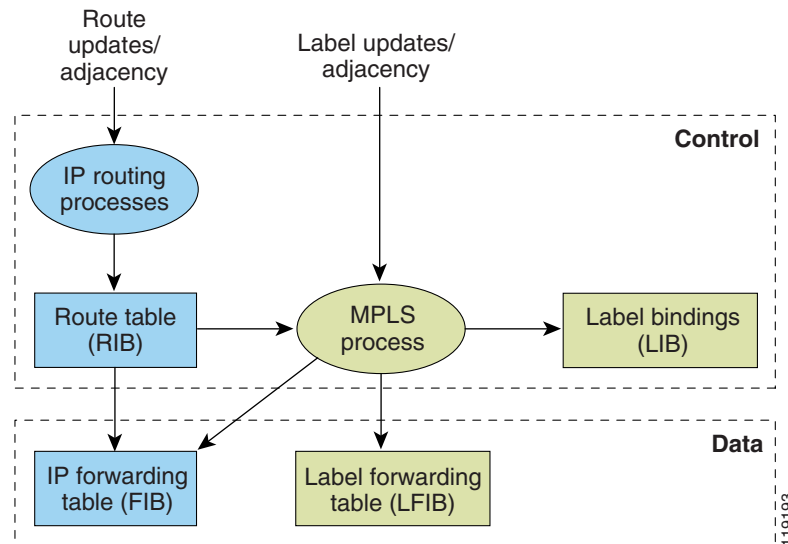
Supervisor Engine 32 PISA supports Multiprotocol Label Switching (MPLS) virtual private networks (VPNs), and Layer 2 Ethernet over MPLS (EoMPLS), with quality of service (QoS) and security.

The PISA on the supervisor engine performs Layer 3 control-plane functions, including address resolution and routing protocols. The PISA processes information from the Routing and Label Distribution Protocols and builds the IP forwarding (FIB) table and the label forwarding (LFIB) table. The PISA distributes the information in both tables to the PFC3B.

The PFC3B receives the information and creates its own copies of the FIB and LFIB tables. Together, these tables comprise the FIB TCAM. The PFC3B looks 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 21-2 shows the various functional blocks that support MPLS label switching. 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 21-2 MPLS Forwarding, Control and Data Planes



IP to MPLS

At the ingress to the MPLS network, the PFC3B 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 PFC3B generates a result containing the imposition label(s) needed to switch the MPLS packet.



Note

If MPLS load sharing is configured, the adjacency may point to a load-balanced path. See [“Basic MPLS Load Balancing”](#) section on page 21-8.

MPLS to MPLS

At the core of an MPLS network, the PFC3B 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 PFC3B 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 PFC3B 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 PFC3B 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 PFC3B 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 PFC3B looks up the top label in the FIB TCAM and recirculates the packet. Then the PFC3B 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 PFC3B 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 PFC3B for additional processing.

Supported Hardware Features

The following hardware features are supported:

- 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 PFC3B provides the capability to recirculate the packets. See the “Recirculation” section on page 21-4.
- Configuration of MPLS switching is supported on VLAN interfaces with the **mpls ip** command.

Supported Cisco IOS Features

The following Cisco IOS software features are supported:



Note Multi-VPN Routing and Forwarding (VRF) for CE Routers (VRF Lite) is supported with the following features: IPv4 forwarding between VRFs interfaces, IPv4 ACLs, and IPv4 HSRP. IPv4 multicast is not supported.

- Multi-VRF for CE Routers (VRF Lite)—VRF-lite is a feature that enables a service provider to support two or more VPNs (using only VRF-based IPv4), where IP addresses can be overlapped among the VPNs. See this publication:
http://www.cisco.com/en/US/products/hw/routers/ps259/prod_bulletin09186a00800921d7.html.
- MPLS on Cisco routers—This feature provides basic MPLS support for imposing and removing labels on IP packets at label edge routers (LERs) and switching labels at label switch routers (LSRs). See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/mpls/config_library/12-2sx/mp-12-2sx-library.html.
- MPLS TE—MPLS traffic engineering software enables an MPLS backbone to replicate and expand upon the traffic engineering capabilities of Layer 2 ATM and Frame Relay networks. MPLS traffic engineering thereby makes traditional Layer 2 features available to Layer 3 traffic flows. For more information, see these publications:
http://www.cisco.com/en/US/docs/ios/12_2/switch/configuration/guide/xcftagc_ps1835_TSD_Products_Configuration_Guide_Chapter.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

- MPLS TE DiffServ Aware (DS-TE)—This feature provides extensions made to MPLS TE to make it DiffServ aware, allowing constraint-based routing of guaranteed traffic. See this publication:
http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsdserv3.html
- MPLS TE Forwarding Adjacency—This feature allows a network administrator to handle a traffic engineering, label-switched path (LSP) tunnel as a link in an Interior Gateway Protocol (IGP) network based on the Shortest Path First (SPF) algorithm. For information on forwarding adjacency with Intermediate System-to-Intermediate System (IS-IS) routing, see this publication:
http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fstefa_3.html
- MPLS TE Interarea Tunnels—This feature allows the router to establish MPLS TE tunnels that span multiple Interior Gateway Protocol (IGP) areas and levels, removing the restriction that had required the tunnel head-end and tail-end routers to be in the same area. See this publication:
http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsiarea3.html
- MPLS virtual private networks (VPNs)—This feature allows you to deploy scalable IPv4 Layer 3 VPN backbone services over a Cisco IOS network. See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/mp_13_vpns/configuration/12-2sx/mp-13-vpns-12-2sx-book.html
- MPLS VPN Carrier Supporting Carrier (CSC)—This feature enables one MPLS VPN-based service provider to allow other service providers to use a segment of its backbone network. See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/mp_ias_and_csc/configuration/12-2sx/mp-carrier-ldp-igp.html
- MPLS VPN Carrier Supporting Carrier IPv4 BGP Label Distribution—This feature allows you to configure your CSC network to enable Border Gateway Protocol (BGP) to transport routes and MPLS labels between the backbone carrier provider edge (PE) routers and the customer carrier customer edge (CE) routers. See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/mp_13_vpns/configuration/12-2sx/mp-13-vpns-12-2sx-book.html
- MPLS VPN Interautonomous System (InterAS) Support —This feature allows an MPLS VPN to span service providers and autonomous systems. See this publication:
http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/fsias24.html
- MPLS VPN Inter-AS IPv4 BGP label distribution—This feature enables you to set up a VPN service provider network so that the autonomous system boundary routers (ASBRs) exchange IPv4 routes with MPLS labels of the PE routers. See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/mp_ias_and_csc/configuration/12-2sx/mp-carrier-bgp.html
- MPLS VPN Hot Standby Router Protocol (HSRP) Support—This feature ensures that the HSRP virtual IP address is added to the correct IP routing table and not to the global routing table. See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/ipapp_fhrp/configuration/12-2sx/fhp-hsrp.html
- OSPF Sham-Link Support for MPLS VPN—This feature allows you to use a sham-link to connect VPN client sites that run the Open Shortest Path First (OSPF) protocol and share OSPF links in a MPLS VPN configuration. See this publication:
http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-sham-link.html
- Any Transport over MPLS (AToM)—Transports Layer 2 packets over an MPLS backbone. See the “Any Transport over MPLS” section on page 21-13.

MPLS Guidelines and Restrictions

When configuring MPLS, follow these guidelines and restrictions:

- The PFC3B supports up to 8 load-shared paths. Cisco IOS releases for other platforms support only 4 load-shared paths.
- The PFC3B supports MTU checking and fragmentation.



Note Fragmentation is supported with software (for IP to MPLS path). See the **mtu** command in the *Catalyst Supervisor Engine 32 PISA Cisco IOS Command Reference*, Release 12.2ZY.



Note For information on other limitations and restrictions, see the “[MPLS VPN Guidelines and Restrictions](#)” section on page 21-11 and the “[EoMPLS Guidelines and Restrictions](#)” section on page 21-14.

MPLS Supported Commands

The PFC3B 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/docs/ios/12_2/switch/command/reference/fswtch_r.html

Configuring MPLS

For information about configuring MPLS, see the Cisco IOS software documentation at the following URL:

http://www.cisco.com/en/US/docs/ios/12_2/switch/configuration/guide/xcftagov_ps1835_TSD_Products_Configuration_Guide_Chapter.html

MPLS Per-Label Load Balancing

The following sections provide information on basic MPLS, MLPS Layer 2 VPN, and MPLS Layer 3 VPN load balancing.

Basic MPLS Load Balancing

The maximum number of load balancing paths is 8. The PFC3B forwards MPLS labeled packets without explicit configuration. If the packet has three labels or less and the underlying packet is IPv4, then the PFC3B uses the source and destination IPv4 address. If the underlying packet is not IPv4 or more than three labels are present, the PFC3B parses down as deep as the fifth or lowest label and uses it for hashing.

MPLS Layer 2 VPN Load Balancing

Load balancing is based on the VC label in the MPLS core if the first nibble of the MAC address in the customer Ethernet frame is not 4.



Note

Load balancing is not supported at the ingress PE for Layer 2 VPNs.

MPLS Layer 3 VPN Load Balancing

MPLS Layer 3 VPN load balancing is similar to basic MPLS load balancing. For more information, see the [“Basic MPLS Load Balancing”](#) section on page 21-8.

MPLS Configuration Examples

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#sh ip ro 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#

Router# show mpls forwarding-table 88.0.0.0
Local  Outgoing  Prefix          Bytes tag  Outgoing   Next Hop
tag    tag or VC  or Tunnel Id    switched  interface
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 mpls 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 mpls 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 mpls cef adj ent 344105 de

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

```

VPN Switching

These sections describe VPN switching:

- [VPN Switching Operation, page 21-10](#)
- [MPLS VPN Guidelines and Restrictions, page 21-11](#)
- [MPLS VPN Supported Commands, page 21-11](#)
- [MPLS VPN Sample Configuration, page 21-12](#)

VPN Switching Operation

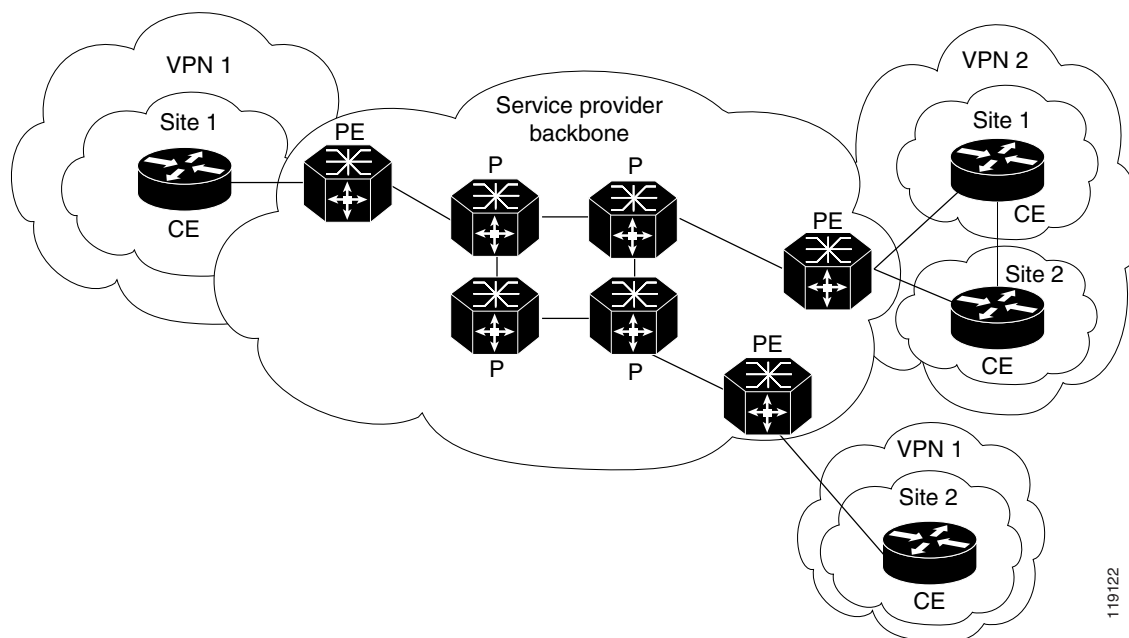
The IP VPN feature for MPLS allows a Cisco IOS network to deploy scalable IP Layer 3 VPN backbone services to multiple sites deployed on a shared infrastructure while also providing the same access or security policies as a private network. VPN based on MPLS technology provides the benefits of routing isolation and security, as well as simplified routing and better scalability.

Refer to the Cisco IOS software documentation for a conceptual MPLS VPN overview and configuration details at this URL:

http://www.cisco.com/en/US/docs/ios/12_2/switch/configuration/guide/xcftagov_ps1835_TSD_Products_Configuration_Guide_Chapter.html

A typical MPLS VPN network topology is shown in Figure 21-3.

Figure 21-3 VPNs with Service Provider Backbone



At the ingress PE, the PFC3B makes a forwarding decision based on the packet headers. The PFC3B contains a table that maps VLANs to VPNs. In the Catalyst 6500 series switch architecture, all physical ingress interfaces in the system are associated with a specific VPN. The PFC3B 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 line card 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 PFC3B 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 PFC3B must perform an IP lookup to determine the final destination. The IP lookup may require recirculation.

MPLS VPN Guidelines and Restrictions

When configuring MPLS VPN, follow these guidelines and restrictions:

- The PFC3B supports a total of 1024 VRFs per chassis with enhanced OSMs; using a nonenhanced OSM causes the system to default to 511 VRFs.
- VPNs are recirculated when the number of VPNs is over 511.

MPLS VPN Supported Commands

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/docs/ios/12_2/switch/command/reference/fswtch_r.html

Configuring MPLS VPN

For information on configuring MPLS VPN, refer to the “MPLS Virtual Private Networks” section at this URL:

http://www.cisco.com/en/US/docs/ios/12_2/switch/configuration/guide/xcftagc_ps1835_TSD_Products_Configuration_Guide_Chapter.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.

MPLS VPN Sample Configuration

This sample configuration shows LAN, OSM, and FlexWAN CE-facing interfaces. MPLS switching configuration is identical to configuration on other platforms.

```

!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
mls mpls tunnel-recir
!
interface Loopback0
  ip address 10.4.4.4 255.255.255.255
!
interface GigabitEthernet4/2
  description Catalyst link to P2
  no ip address
  mls qos trust dscp
!
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
  mls qos trust dscp
!
interface GigabitEthernet7/3.73
  encapsulation dot1Q 73
  ip vrf forwarding blues
  ip address 10.19.7.1 255.255.255.0
!
interface POS8/1
  description OSM link to CE3
  ip vrf forwarding blues
  ip address 10.19.8.1 255.255.255.252
  encapsulation ppp
  mls qos trust dscp
  pos scramble-atm
  pos flag c2 22
!
interface POS9/0/0
  description FlexWAN link to CE1
  ip vrf forwarding blues
  ip address 10.19.9.1 255.255.255.252
  encapsulation ppp
  pos scramble-atm
  pos flag c2 22
!
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 vpv4
  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
!

```

Any Transport over MPLS

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).

AToM supports the following like-to-like transport types:

- Ethernet over MPLS (EoMPLS) (VLAN mode and port mode)
- Frame Relay over MPLS with DLCI-to-DLCI connections
- ATM AAL5 over MPLS
- ATM Cell Relay over MPLS



Note Additional AToM types are planned in future releases.

The PFC3B supports both hardware-based EoMPLS as well as OSM-, FlexWAN, or FlexWAN2-based EoMPLS. For more information, see this URL:

http://www.cisco.com/en/US/docs/routers/7600/install_config/12.2SX_OSM_config/mpls.html

For information on requirements for Supervisor Engine 2-based EoMPLS, see this URL:

http://www.cisco.com/en/US/docs/routers/7600/install_config/12.2SX_OSM_config/mpls.html#Supervisor_Engine_2-Based_EoMPLS

For information on other AToM implementations (ATM AAL5 over MPLS, ATM Cell Relay over MPLS, Frame Relay over MPLS), see this URL:

http://www.cisco.com/en/US/docs/routers/7600/install_config/12.2SX_OSM_config/mpls.html#Any_Transport_over_MPLS

These sections describe AToM:

- [AToM Load Balancing, page 21-14](#)
- [Understanding EoMPLS, page 21-14](#)
- [EoMPLS Guidelines and Restrictions, page 21-14](#)
- [Configuring EoMPLS, page 21-16](#)

AToM Load Balancing

EoMPLS 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.

Understanding EoMPLS

EoMPLS is one of the AToM transport types. AToM transports Layer 2 packets over a MPLS backbone using a directed 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 works by encapsulating Ethernet PDUs in MPLS packets and forwarding them across the MPLS network. Each PDU is transported as a single packet.



Note

Use OSM-based EoMPLS when you want local Layer 2 switching and EoMPLS on the same VLAN. You need to configure EoMPLS on the SVI; the core-facing card must be an OSM. When local Layer 2 switching is not required, use PFC3B-based EoMPLS configured on the subinterface or physical interface.

EoMPLS Guidelines and Restrictions

When configuring EoMPLS, follow these guidelines and restrictions:

- 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.
- If QoS is disabled globally, both the 802.1p and IP precedence bits are preserved. 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. The PFC3B provides a new command that allows you to trust the P bits while preserving the IP precedence bits. To preserve the IP precedence bits, use the **no mls qos rewrite ip dscp** command.

**Note**

The **no mls qos rewrite ip dscp** command is not compatible with the MPLS and MPLS VPN features. See [Chapter 38, “Configuring PFC QoS.”](#)

**Note**

Do not use the **no mls qos rewrite ip dscp** command if you have PFC3B-based EoMPLS and PFX-based EoMPLS services in the same system.

- 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 the supervisor engine 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 directed to the supervisor engine and not to the EoMPLS cloud.
 - The native VLAN of a trunk must not be configured as an EoMPLS VLAN.
- All protocols (for example, CDP, VTP, BPDUs) are tunneled across the MPLS cloud without conditions.
- ISL encapsulation is not supported for the interface that receives EoMPLS packets.
- 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 PFC3B-based EoMPLS).
- EoMPLS 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). This functionality (local switching) is available only when using OSM and FlexWAN modules as uplinks.
- In previous releases of AToM, the command used to configure AToM circuits was **mpls l2 transport route**. This command has been replaced with the **xconnect** command. You can use the **xconnect** command to configure EoMPLS circuits.
- The AToM control word is not supported.
- EoMPLS is not supported on Layer 3 VLAN interfaces.
- Point-to-point EoMPLS works with a physical interface and subinterfaces.

Configuring EoMPLS

These sections describe how to configure EoMPLS:

- [Prerequisites, page 21-16](#)
- [Configuring VLAN-Based EoMPLS, page 21-16](#)
- [Configuring Port-Based EoMPLS, page 21-19](#)

Prerequisites

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.

EoMPLS works by encapsulating Ethernet PDUs in MPLS packets and forwarding them across the MPLS network. Each PDU is transported as a single packet. Two methods are available to configure EoMPLS:

- 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 does not allow local switching of packets between interfaces unless you use loopback ports.
- A system can have both an OSM or FlexWAN configuration and EoMPLS configuration enabled at the same time. Cisco supports this configuration but does not recommend it. Unless the uplinks to the MPLS core are through OSM or FlexWAN-enabled interfaces, OSM or FlexWAN-based EoMPLS connections will not be active; this causes packets for OSM or FlexWAN-based EoMPLS arriving on non-WAN interfaces to be dropped. For information on WAN (FlexWAN and OSM) EoMPLS, see this URL:

http://www.cisco.com/en/US/docs/routers/7600/install_config/12.2SX_OSM_config/mpls.html

LAN ports can receive Layer 2 traffic, impose labels, and switch the frames into the MPLS core without using an OSM or FlexWAN module.

You can configure an OSM or a FlexWAN module to face the core of MPLS network and use either the OSM configuration, the FlexWAN configuration, or the PFC3B configuration. For more information, see this URL:

http://www.cisco.com/en/US/docs/routers/7600/install_config/12.2SX_OSM_config/mpls.html

Configuring VLAN-Based EoMPLS

When configuring VLAN-based EoMPLS, follow these guidelines and restrictions:

- 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.

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.

Verifying the Configuration

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 make sure 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 make sure 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 ensure 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  \
                                           0          Gi2/1      23.2.0.1
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
Router#
```

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 view 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
V13	Eth VLAN 3	11.11.11.11	3	UP

To see detailed information about each VC, add the keyword **detail**.

```
Router# show mpls l2transport vc detail
```

```
Local interface: V12 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: V13 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

When configuring port-based EoMPLS, follow these guidelines and restrictions:

- 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.
- 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.

To support 802.1Q-in-802.1Q traffic and Ethernet traffic over EoMPLS, 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 <i>slot/interface</i>	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 <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.

The following is an example of a port-based configuration:

```
!
EoMPLS:
```

```
router# show mpls l2transport vc
```

```
Local intf      Local circuit    Dest address     VC ID           Status
-----
Fa8/48         Ethernet        75.0.78.1       1               UP
Gi7/11.2000    Eth VLAN 2000   75.0.78.1       2000            UP
```

Port-Based EoMPLS Config:

```
router# show run interface f8/48
Building configuration...
```

```
Current configuration : 86 bytes
!
interface FastEthernet8/48
 no ip address
 xconnect 75.0.78.1 1 encapsulation mpls
end
```

```
Sub-Interface Based Mode:
router# show run interface g7/11
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 g7/11.2000
Building configuration...

Current configuration : 112 bytes
!
interface GigabitEthernet7/11.2000
 encapsulation dot1q 2000
 xconnect 75.0.78.1 2000 encapsulation mpls
end
```

```

kb7606# 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 obtain the VC type:

```

kb7606# 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
!

```

Verifying the Configuration

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 make sure 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 make sure 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 make sure 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  \
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 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 view 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

```