



Ethernet over MPLS (EoMPLS)

- [Prerequisites for EoMPLS, page 9-1](#)
- [Restrictions for EoMPLS, page 9-2](#)
- [Information About EoMPLS, page 9-3](#)
- [Default Settings for EoMPLS, page 9-3](#)
- [How to Configure EoMPLS, page 9-4](#)



Note

- For complete syntax and usage information for the commands used in this chapter, see these publications:
http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
 - Cisco IOS Release 15.1SY supports only Ethernet interfaces. Cisco IOS Release 15.1SY does not support any WAN features or commands.
-



Tip

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

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

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

Prerequisites for EoMPLS

Before you configure EoMPLS, ensure that the network is configured as follows:

- Configure IP routing in the core so that the PE routers can reach each other through IP.
- Configure MPLS in the core so that a label switched path (LSP) exists between the PE routers.

Restrictions for EoMPLS

- EoMPLS in Cisco IOS Release 15.1SY does not support load balancing at the tunnel ingress; only one Interior Gateway Protocol (IGP) path is selected even if multiple IGP paths are available, but load balancing is available at the MPLS core.
- Ensure that the maximum transmission unit (MTU) of all intermediate links between endpoints is sufficient to carry the largest Layer 2 packet received.
- EoMPLS supports VLAN packets that conform to the IEEE 802.1Q standard. The 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames.
- When the QoS is enabled on a Layer 2 port, either 802.1q P bits or IP precedence bits can be preserved with the trusted configuration. However, by default the unpreserved bits are overwritten by the value of preserved bits. For instance, if you preserve the P bits, the IP precedence bits are overwritten with the value of the P bits. To preserve the IP precedence bits, use the **no platform qos rewrite ip dscp** command. The **no platform qos rewrite ip dscp** command is not compatible with the MPLS and MPLS VPN features.
- EoMPLS is not supported with private VLANs.
- The following restrictions apply to using trunks with EoMPLS:
 - To support Ethernet spanning tree bridge protocol data units (BPDUs) across an EoMPLS cloud, you must disable spanning tree for the Ethernet-over-MPLS VLAN. This ensures that the EoMPLS VLANs are carried only on the trunk to the customer switch. Otherwise, the BPDUs are not directed to the EoMPLS cloud.
 - The native VLAN of a trunk must not be configured as an EoMPLS VLAN.
- In Cisco IOS Release 15.1SY, all protocols (for example, CDP, VTP, BPDUs) are tunneled across the MPLS cloud without conditions.
- Unique VLANs are required across interfaces. You cannot use the same VLAN ID on different interfaces.
- EoMPLS tunnel destination route in the routing table and the CEF table must be a /32 address (host address where the mask is 255.255.255.255) to ensure that there is a label-switched path (LSP) from PE to PE.
- For a particular EoMPLS connection, both the ingress EoMPLS interface on the ingress PE and the egress EoMPLS interface on the egress PE have to be subinterfaces with dot1Q encapsulation or neither is a subinterface.
- 802.1Q in 802.1Q over EoMPLS is supported if the outgoing interface connecting to MPLS network is a port on an Layer 2 card.
- Shaping EoMPLS traffic is not supported if the egress interface connecting to an MPLS network is a Layer 2 LAN port (a mode known as PFC-based EoMPLS).
- EoMPLS based on a PFC does not perform any Layer 2 lookup to determine if the destination MAC address resides on the local or remote segment and does not perform any Layer 2 address learning (as traditional LAN bridging does).
- The AToM control word is not supported.
- Ethernet packets with hardware-level cyclic redundancy check (CRC) errors, framing errors, and runt packets are discarded on input.
- You must configure VLAN-based EoMPLS on subinterfaces.

- Port-based EoMPLS and VLAN-based EoMPLS are mutually exclusive. If you enable a main interface for port-to-port transport, you also cannot enter commands on a subinterface.
- EoMPLS is not supported on Layer 3 VLAN interfaces
- Flow Aware Transport pseudowire is not supported for Point-to-Point xconnect.
- Point-to-point EoMPLS works with a physical interface and subinterfaces.

Information About EoMPLS

- [AToM Overview, page 9-3](#)
- [EoMPLS Overview, page 9-3](#)

AToM Overview

Any Transport over MPLS (AToM) transports Layer 2 packets over an MPLS backbone. AToM uses a directed Label Distribution Protocol (LDP) session between edge routers for setting up and maintaining connections. Forwarding occurs through the use of two level labels that provide switching between the edge routers. The external label (tunnel label) routes the packet over the MPLS backbone to the egress PE at the ingress PE. The VC label is a demuxing label that determines the connection at the tunnel endpoint (the particular egress interface on the egress PE as well as the VLAN identifier for an Ethernet frame).

EoMPLS Overview

EoMPLS is one of the AToM transport types. EoMPLS works by encapsulating Ethernet PDUs in MPLS packets and forwarding them across the MPLS network. Each PDU is transported as a single packet. Cisco IOS Release 15.1SY supports two EoMPLS modes:

- VLAN mode—Transports Ethernet traffic from a source 802.1Q VLAN to a destination 802.1Q VLAN through a single VC over an MPLS network. VLAN mode uses VC type 5 as default (no dot1q tag) and VC type 4 (transport dot1 tag) if the remote PE does not support VC type 5 for subinterface (VLAN) based EoMPLS.
- Port mode—Allows all traffic on a port to share a single VC across an MPLS network. Port mode uses VC type 5.

**Note**

For both VLAN mode and port mode, EoMPLS in Cisco IOS Release 15.1SY does not allow local switching of packets between interfaces unless you use loopback interfaces.

LAN ports can receive Layer 2 traffic, impose labels, and switch the frames into the MPLS core.

Default Settings for EoMPLS

None.

How to Configure EoMPLS

- [Configuring VLAN-Based EoMPLS, page 9-4](#)
- [Configuring Port-Based EoMPLS, page 9-6](#)

Configuring VLAN-Based EoMPLS

To configure VLAN-based EoMPLS, perform this task on the provider edge (PE) routers:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.
Step 2	Router(config)# interface gigabitethernet <i>slot/interface.subinterface</i>	Specifies the Gigabit Ethernet subinterface. Make sure that the subinterface on the adjoining CE router is on the same VLAN as this PE router.
Step 3	Router(config-if)# encapsulation dot1q <i>vlan_id</i>	Enables the subinterface to accept 802.1Q VLAN packets. <ul style="list-style-type: none"> • The subinterfaces between the CE and PE routers that are running Ethernet over MPLS must be in the same subnet. • All other subinterfaces and backbone routers do not need to be on the same subnet.
Step 4	Router(config-if)# xconnect <i>peer_router_id vcid</i> encapsulation mpls	Binds the attachment circuit to a pseudowire VC. The syntax for this command is the same as for all other Layer 2 transports.

This is a VLAN-based EoMPLS configuration sample:

```
interface GigabitEthernet7/4.2
encapsulation dot1q 3
xconnect 13.13.13.13 3 encapsulation mpls
no shut
```



Note

The IP address is configured on subinterfaces of the CE devices.

To verify and display the configuration of Layer 2 VLAN transport over MPLS tunnels, perform the following:

- To display a single line for each VLAN, naming the VLAN, status, and ports, enter the **show vlan brief** command.

```
Router# show vlan brief
```

```
VLAN Name                Status    Ports
-----
1    default                 active
2    VLAN0002                active
3    VLAN0003                active
1002 fddi-default            act/unsup
1003 token-ring-default      act/unsup
1004 fddinet-default         act/unsup
1005 trnet-default           act/unsup
```

- To verify that the PE router endpoints have discovered each other, enter the **show mpls ldp discovery** command. When an PE router receives an LDP hello message from another PE router, it considers that router and the specified label space to be “discovered.”

```
Router# show mpls ldp discovery
Local LDP Identifier:
 13.13.13.13:0
Discovery Sources:
Interfaces:
  GE-WAN3/3 (ldp): xmit/rcv
    LDP Id: 12.12.12.12:0
Targeted Hellos:
 13.13.13.13 -> 11.11.11.11 (ldp): active/passive, xmit/rcv
    LDP Id: 11.11.11.11:0
```

- To verify that the label distribution session has been established, enter the **show mpls ldp neighbor** command. The third line of the output shows that the state of the LDP session is operational and shows that messages are being sent and received.

```
Router# show mpls ldp neighbor
Peer LDP Ident: 12.12.12.12:0; Local LDP Ident 13.13.13.13:0
TCP connection: 12.12.12.12.646 - 13.13.13.13.11010
State: Oper; Msgs sent/rcvd: 1649/1640; Downstream
Up time: 23:42:45
LDP discovery sources:
  GE-WAN3/3, Src IP addr: 34.0.0.2
Addresses bound to peer LDP Ident:
 23.2.1.14      37.0.0.2      12.12.12.12      34.0.0.2
 99.0.0.1
Peer LDP Ident: 11.11.11.11:0; Local LDP Ident 13.13.13.13:0
TCP connection: 11.11.11.11.646 - 13.13.13.13.11013
State: Oper; Msgs sent/rcvd: 1650/1653; Downstream
Up time: 23:42:29
LDP discovery sources:
  Targeted Hello 13.13.13.13 -> 11.11.11.11, active, passive
Addresses bound to peer LDP Ident:
 11.11.11.11      37.0.0.1      23.2.1.13
```

- To verify that the label forwarding table is built correctly, enter the **show mpls forwarding-table** command to verify that a label has been learned for the remote PE and that the label is going from the correct interface to the correct next-hop.

```
Router# show mpls forwarding-table
Local  Outgoing  Prefix          Bytes tag  Outgoing     Next Hop
tag    tag or VC  or Tunnel Id    switched  interface
16     Untagged  223.255.254.254/32  \
20     Untagged  12ckt(2)        133093    V12         point2point
21     Untagged  12ckt(3)        185497    V13         point2point
24     Pop tag   37.0.0.0/8      0         GE3/3       34.0.0.2
25     17       11.11.11.11/32  0         GE3/3       34.0.0.2
26     Pop tag   12.12.12.12/32  0         GE3/3       34.0.0.2
```

The output shows the following data:

- Local tag—Label assigned by this router.
- Outgoing tag or VC—Label assigned by next hop.
- Prefix or Tunnel Id—Address or tunnel to which packets with this label are going.
- Bytes tag switched— Number of bytes switched out with this incoming label.
- Outgoing interface—Interface through which packets with this label are sent.

- Next Hop—IP address of neighbor that assigned the outgoing label.
- To display the state of the currently routed VCs, enter the **show mpls l2transport vc** command.

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

Local intf	Local circuit	Dest address	VC ID	Status
V12	Eth VLAN 2	11.11.11.11	2	UP
V13	Eth VLAN 3	11.11.11.11	3	UP

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

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

```
Local interface: 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

To support 802.1Q-in-802.1Q traffic and Ethernet traffic over EoMPLS in Cisco IOS Release 15.1SY, configure port-based EoMPLS by performing this task:

	Command	Purpose
Step 1	Router# configure terminal	Enters global configuration mode.

	Command	Purpose
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</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.

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

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

```
Local intf      Local circuit      Dest address      VC ID      Status
-----
Gi8/48         Ethernet           75.0.78.1         1          UP
Gi7/11.2000    Eth VLAN 2000     75.0.78.1         2000       UP
```

```
Router# show run interface gigabitethernet 8/48
Building configuration...
```

```
Current configuration : 86 bytes
!
interface GigabitEthernet8/48
 no ip address
 xconnect 75.0.78.1 1 encapsulation mpls
end
```

```
Router# show run interface gigabitethernet 7/11
Building configuration...
```

```
Current configuration : 118 bytes
!
interface GigabitEthernet7/11
 description Traffic-Generator
 no ip address
 logging event link-status
 speed nonegotiate
end
```

```
Router# show run int gigabitethernet 7/11.2000
Building configuration...
```

```
Current configuration : 112 bytes
!
interface GigabitEthernet7/11.2000
 encapsulation dot1Q 2000
 xconnect 75.0.78.1 2000 encapsulation mpls
end
```

```
Router# show mpls l2transport vc 1 detail
```

```
Local interface: Gi7/47 up, line protocol up, Ethernet up
 Destination address: 75.0.80.1, VC ID: 1, VC status: up
 Tunnel label: 5704, next hop 75.0.83.1
 Output interface: Te8/3, imposed label stack {5704 10038}
 Create time: 00:30:33, last status change time: 00:00:43
 Signaling protocol: LDP, peer 75.0.80.1:0 up
 MPLS VC labels: local 10579, remote 10038
 Group ID: local 155, remote 116
 MTU: local 1500, remote 1500
 Remote interface description:
 Sequencing: receive disabled, send disabled
```

```

VC statistics:
  packet totals: receive 26, send 0
  byte totals:   receive 13546, send 0
  packet drops:  receive 0, send 0

```

To display the VC type:

```
Router# remote command switch show mpls l2transport vc 1 de
```

```

Local interface: GigabitEthernet7/47, Ethernet
Destination address: 75.0.80.1, VC ID: 1
VC status: receive UP, send DOWN
VC type: receive 5, send 5
  Tunnel label: not ready, destination not in LFIB
  Output interface: unknown, imposed label stack {}
  MPLS VC label: local 10579, remote 10038
Linecard VC statistics:
  packet totals: receive: 0 send: 0
  byte totals:   receive: 0 send: 0
  packet drops: receive: 0 send: 0
Control flags:
  receive 1, send: 31
!

```

To verify and display the configuration of Layer 2 VLAN transport over MPLS tunnels, perform the following:

- To display a single line for each VLAN, naming the VLAN, status, and ports, enter the **show vlan brief** command.

```
Router# show vlan brief
```

VLAN Name	Status	Ports
1 default	active	
2 VLAN0002	active	Gil/4
1002 fddi-default	act/unsup	
1003 token-ring-default	act/unsup	
1004 fddinet-default	act/unsup	
1005 trnet-default	act/unsup	

- To verify that the PE router endpoints have discovered each other, enter the **show mpls ldp discovery** command. When an PE router receives an LDP Hello message from another PE router, it considers that router and the specified label space to be “discovered.”

```

Router# show mpls ldp discovery
Local LDP Identifier:
  13.13.13.13:0
Discovery Sources:
Interfaces:
  GE-WAN3/3 (ldp): xmit/recv
    LDP Id: 12.12.12.12:0
Targeted Hellos:
  13.13.13.13 -> 11.11.11.11 (ldp): active/passive, xmit/recv
    LDP Id: 11.11.11.11:0

```

- To verify that the label distribution session has been established, enter the **show mpls ldp neighbor** command. The third line of the output shows that the state of the LDP session is operational and shows that messages are being sent and received.

```

Router# show mpls ldp neighbor
Peer LDP Ident: 12.12.12.12:0; Local LDP Ident 13.13.13.13:0
TCP connection: 12.12.12.12.646 - 13.13.13.13.11010
State: Oper; Msgs sent/rcvd: 1715/1706; Downstream
Up time: 1d00h

```

```

LDP discovery sources:
  GE-WAN3/3, Src IP addr: 34.0.0.2
Addresses bound to peer LDP Ident:
  23.2.1.14      37.0.0.2      12.12.12.12    34.0.0.2
  99.0.0.1
Peer LDP Ident: 11.11.11.11:0; Local LDP Ident 13.13.13.13:0
TCP connection: 11.11.11.11.646 - 13.13.13.13.11013
State: Oper; Msgs sent/rcvd: 1724/1730; Downstream
Up time: 1d00h
LDP discovery sources:
  Targeted Hello 13.13.13.13 -> 11.11.11.11, active, passive
Addresses bound to peer LDP Ident:
  11.11.11.11    37.0.0.1      23.2.1.13

```

- To verify that the label forwarding table is built correctly, enter the **show mpls forwarding-table** command.

```

Router# show mpls forwarding-table
Local  Outgoing  Prefix          Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id    switched  interface
16     Untagged   223.255.254.254/32  \
                                           0          Gi2/1      23.2.0.1
20     Untagged   12ckt(2)        55146580  V12        point2point
24     Pop tag    37.0.0.0/8      0         GE3/3      34.0.0.2
25     17        11.11.11.11/32  0         GE3/3      34.0.0.2
26     Pop tag    12.12.12.12/32  0         GE3/3      34.0.0.2

```

- The output displays the following data:
 - Local tag—Label assigned by this router.
 - Outgoing tag or VC—Label assigned by next hop.
 - Prefix or Tunnel Id—Address or tunnel to which packets with this label are going.
 - Bytes tag switched— Number of bytes switched out with this incoming label.
 - Outgoing interface—Interface through which packets with this label are sent.
 - Next Hop—IP address of neighbor that assigned the outgoing label.
- To display the state of the currently routed VCs, enter the **show mpls l2transport vc** command:

```

Router# show mpls l2transport vc
Local intf    Local circuit    Dest address    VC ID    Status
-----
V12          Eth VLAN 2      11.11.11.11    2        UP

```

**Tip**

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

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

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

