Layer 2 VPN Architectures

A complete guide to understanding, designing, and deploying Layer 2 VPN technologies and pseudowire emulation applications

Wei Luo, CCIE® No. 13,291
Carlos Pignataro, CCIE No. 4619
Dmitry Bokotey, CCIE No. 4460
Anthony Chan, CCIE No. 10,266

ciscopress.com

Layer 2 VPN Architectures

Wei Luo, CCIE No. 13,291
Carlos Pignataro, CCIE No. 4619
Dmitry Bokotey, CCIE No. 4460
Anthony Chan, CCIE No. 10,266

Cisco Press

800 East 96th Street
Indianapolis, IN 46240 USA

Example 7-32 **show controllers eompls forwarding-table** *Command*

```
SanFran

LC-CON0#show controllers eompls forwarding-table 0 100

Port # 0, VLAN-ID # 100, Table-index 100
EoMPLS configured: 1
tag_rew_ptr             = D001BB58
Leaf entry?     = 1
FCR index       = 20
           **tagrew_psa_addr    = 0006ED60
           **tagrew_vir_addr    = 7006ED60
           **tagrew_phy_addr    = F006ED60
        [0-7] loq 8800 mtu 4458  oq 4000 ai 3 oi 04019110 (encaps size 4)
        cw-size 4 vlanid-rew 200
        gather A30 (bufhdr size 32 EoMPLS (Control Word) Imposition profile 81)
        2 tag: 18 18
        counters 1182, 10 reported 1182, 10.
    Local OutputQ (Unicast):    Slot:2  Port:0  RED queue:0  COS queue:0
    Output Q (Unicast):         Port:0          RED queue:0  COS queue:0

_______________________________________________________________________________

NewYork

LC-CON0#show controllers eompls forwarding-table 0 200

Port # 0, VLAN-ID # 200, Table-index 200
EoMPLS configured: 1
tag_rew_ptr             = D0027B90
Leaf entry?     = 1
FCR index       = 20
           **tagrew_psa_addr    = 0009EE40
           **tagrew_vir_addr    = 7009EE40
           **tagrew_phy_addr    = F009EE40
        [0-7] loq 9400 mtu 4458  oq 4000 ai 8 oi 84000002 (encaps size 4)
        cw-size 4 vlanid-rew 100
        gather A30 (bufhdr size 32 EoMPLS (Control Word) Imposition profile 81)
        2 tag: 17 18
        counters 1182, 10 reported 1182, 10.
    Local OutputQ (Unicast):    Slot:5  Port:0  RED queue:0  COS queue:0
    Output Q (Unicast):         Port:0          RED queue:0  COS queue:0
```

NOTE Other platforms that do not require manual configuration do not provide VLAN ID rewrite information in their output.

Port VLAN ID Inconsistency Issue

In EoMPLS using CE switches, Spanning Tree Protocol (described in Chapter 4) still runs between the customer end devices, and it is transported across the MPLS core backbone. In a VLAN ID rewrite scenario, Port VLAN ID (PVID) inconsistency stems from the Per VLAN Spanning Tree + (PVST+) BPDU being received on a different VLAN than it was originated. Therefore, when the trunk port on Oakland receives a PVST+ BPDU from the Albany's STP of VLAN 200 with a tag of VLAN 200, you get an error message as soon as the circuit comes up:

```
%SPANTREE-2-BLOCK_PVID_LOCAL : Blocking [chars] on [chars]. Inconsistent local
vlan.
```

The listed STP instance is that of the native VLAN ID of the listed interface. The first [chars] is the interface, and the second [chars] is the STP instance. As a result, the trunk port is held in a blocking state (for both VLAN 100 and VLAN 200) until the inconsistency is resolved.

To unblock the interface, change the VLAN IDs on the CEs so that they match. When this is not possible and VLAN ID rewrite is required, you must turn off the STP. This alternative opens a door to bridging loops; therefore, you should use it with extreme caution.

Case Study 7-7: Map to Pseudowire

For EoMPLS configuration, you might choose to configure a pseudowire class template that consists of configuration settings used by all attachment circuits that are bound to the class. Pseudowire was introduced in Chapters 2, "Pseudowire Emulation Framework and Standards," and 6 and is discussed in further detail in the advanced configuration case studies of Chapter 9, "Advanced AToM Case Studies."

Example 7-33 shows configuration of the VC 100 in Ethernet port mode.

Example 7-33 *Pseudowire Class Configuration*

```
hostname SanFran
!
pseudowire-class ethernet-port
  encapsulation mpls

int GigabitEthernet1/4
  xconnect 192.168.1.103 100 pw-class ethernet-port
!
```

Common Troubleshooting Techniques

This section introduces you to the most common troubleshooting techniques for PEs in EoMPLS. You first learn the commands and outputs for the Cisco router PEs and then learn to troubleshoot Cisco 7600 series switches.

Troubleshooting EoMPLS on Routers

The first common step in troubleshooting problems is attempting to discover failure by verifying the status of a VC by issuing the **show mpls l2transport vc** command.

Three conditions must be met so that the VC is UP:

- The disposition interfaces are programmed if the VC has been configured and the CE interface is UP.
- If the IGP label exists, it can be implicit null in a back-to-back configuration.
- The imposition interface is programmed if the disposition interface is programmed and you have a remote VC label and an IGP label (LSP to the peer).

If the status field is marked DOWN (that is, the VC is not ready to carry traffic between the two VC endpoints), as shown in the output of Example 7-34, execute the **show mpls l2transport vc detail** command seen in Example 7-35 for more in-depth information.

Example 7-34 **show mpls l2transport vc** *Command*

```
NewYork#show mpls l2transport vc

Local intf     Local circuit            Dest address    VC ID      Status
-------------  ----------------------- --------------- ---------- ----------
Et0/0          Ethernet                 192.168.1.102   100        DOWN
```

Example 7-35 **show mpls l2transport vc detail** *Command*

```
NewYork#show mpls l2transport vc detail
Local interface: Et0/0 up, line protocol up, Ethernet up
  Destination address: 192.168.1.102, VC ID: 100, VC status: down
    Preferred path: not configured
    Default path: active
    Tunnel label: 16, next hop point2point
    Output interface: Se5/0, imposed label stack {16 16}
  Create time: 00:18:10, last status change time: 00:03:51
  Signaling protocol: LDP, peer 192.168.1.102:0 up
    MPLS VC labels: local 16, remote 16
    Group ID: local 0, remote 0
    MTU: local 1500, remote unknown
    Remote interface description:
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 0, send 0
    byte totals:   receive 0, send 0
    packet drops:  receive 0, send 78
```

Table 7-4 describes some of the significant fields of the **show mpls l2transport vc detail** command output.

Table 7-4 **show mpls l2transport vc detail** *Command Output Fields*

Field	Description
Destination address	The IP address of the remote router specified for this VC as part of the **mpls l2transport route** or **xconnect** command.
VC ID	The virtual circuit identifier assigned to the interface on the router.
VC status	The status of the VC. The status can be one of the following: **UP**—The VC is in a state in which it can carry traffic between the two VC endpoints. A VC is UP when both imposition and disposition interfaces are programmed. **DOWN**—The VC is not ready to carry traffic between the two VC endpoints. **ADMIN DOWN**—A user has disabled the VC.
Tunnel label	An IGP label that routes the packet over the MPLS backbone to the destination router with the egress interface.
Output interface	The interface on the remote router that has been enabled to transmit and receive Layer 2 packets.
Imposed label stack	A summary of the MPLS label stack that directs the VC to the PE router.
Signaling protocol	The type of protocol that sends the MPLS labels. The output also shows the status of the peer router.
MPLS VC labels	The local VC label is a disposition label, which determines the egress interface of an arriving packet from the MPLS backbone. The remote VC label is a disposition VC label of the remote peer router.
MTU	The maximum transmission unit specified for the local and remote interfaces.
Sequencing	This field describes whether sequencing of out-of-order packets is enabled or disabled.

In the **show mpls l2transport vc detail** command output, pay attention to the remote unknown next to the local MTU. One of the possible causes is a remote interface down or an MTU mismatch. Verify to make sure that MTU on each side is the same. If an EoMPLS tunnel is still down after this and you cannot pass traffic, perform another check by issuing the **show mpls forwarding-table** command, as demonstrated in Example 7-36.

Example 7-36 **show mpls forwarding-table** *Command*

```
NewYork#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
17     Untagged    10.1.1.0/24       0          Se5/0      point2point
18     Untagged    192.168.1.101/32  0          Se5/0      point2point
19     Untagged    192.168.1.102/32  0          Se5/0      point2point
20     Untagged    l2ckt(100)        4592       Et0/0.100  point2point
```

The Untagged result in the Outgoing tag or VC field indicates that an MPLS label might not be exchanged between the PE and P (Denver) routers.

A couple possible causes exist. Either the **mpls ip** has not been enabled per interface, or CEF is disabled (or not enabled) on the P or PE router. To verify, issue the **show mpls ldp discovery** command, as in Example 7-37.

Example 7-37 **show mpls ldp discovery** *Command*

```
NewYork#show mpls ldp discovery
 Local LDP Identifier:
    192.168.1.103:0
    Discovery Sources:
    Targeted Hellos:
        192.168.1.103 -> 192.168.1.102 (ldp): active/passive, xmit/recv
            LDP Id: 192.168.1.102:0
```

The output shows whether you have a direct LDP session open between directly connected MPLS-enabled interfaces. By observing this behavior, you can conclude that MPLS indeed was not enabled per interface facing the core. To solve this problem, enable MPLS on an interface or check whether CEF is enabled.

Two common issues result when the circuit does not come up:

- The remote port is down or not configured.
- The MTU is mismatched.

Examples 7-38 and 7-39 display the output of the **show mpls l2transport vc** *vcid* **detail** command with the two conditions, respectively.

Example 7-38 *Remote Port Down or Not Configured*

```
NewYork#show mpls l2transport vc 10 detail
Local interface: FastEternet0/0.10 up, line protocol up, Eth VLAN 10 up
  Destination address: 192.168.1.102, VC ID: 10, VC status: down
    Tunnel label: not ready
    Output interface: unknown, imposed label stack {}
  Create time: 22:31:53, last status change time: 04:02:56
  Signaling protocol: LDP, peer 192.168.1.102:0 up
    MPLS VC labels: local 19, remote unassigned
    Group ID: local 0, remote unknown
    MTU: local 1500, remote unknown
    Remote interface description:
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 1650, send 1743
    byte totals:   receive 552557, send 550044
    packet drops:  receive 0, send 7
```

Example 7-39 *MTU Mismatch*

```
NewYork#show mpls l2transport vc 10 detail
Local interface: FastEternet0/0.10 up, line protocol up, Eth VLAN 10 up
  Destination address: 192.168.1.102, VC ID: 10, VC status: down
    Tunnel label: not ready
    Output interface: unknown, imposed label stack {}
  Create time: 22:36:10, last status change time: 00:00:20
  Signaling protocol: LDP, peer 192.168.1.102:0 up
    MPLS VC labels: local 19, remote 21
    Group ID: local 0, remote 0
    MTU: local 1500, remote 1000
    Remote interface description: *** To SanFran ***
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 1880, send 1901
    byte totals:   receive 168476, send 155436
    packet drops:  receive 0, send 13
```

The highlighted portions of Examples 7-38 and 7-39 call attention to the faulty conditions. Compare the output to output of the same command in an operational environment from Example 7-40.

Example 7-40 *Working Example*

```
NewYork#show mpls l2transport vc 10 detail
Local interface: FastEternet0/0.10 up, line protocol up, Eth VLAN 10 up
  Destination address: 192.168.1.102, VC ID: 10, VC status: up
    Preferred path: not configured
    Default path: active
    Tunnel label: 17, next hop 10.1.1.202
    Output interface: Et1/0, imposed label stack {17 21}
  Create time: 23:06:37, last status change time: 00:30:47
  Signaling protocol: LDP, peer 192.168.1.102:0 up
    MPLS VC labels: local 19, remote 21
    Group ID: local 0, remote 0
    MTU: local 1500, remote 1500
    Remote interface description: *** To SanFran ***
  Sequencing: receive disabled, send disabled
  VC statistics:
    packet totals: receive 1683, send 1777
    byte totals:   receive 565455, send 563328
    packet drops:  receive 0, send 7
```

Example 7-41 presents the verification and configuration sequence of enabling MPLS on the Serial 5/0 interface.