

MPLS L2 VPN - VLAN Based Point-to-Point EoMPLS Service

« [MPLS L2 VPN - Port Based Point-to-Point Service using EVC \(/workbook/view/service-provider-v4/task/mpls-l2-vpn-port-based-point-to-point-service-using-enc-Mjk0Mw%3D%3D\)](#) | [MPLS L2 VPN - VPLS - Multipoint Service \(/workbook/view/service-provider-v4/task/mpls-l2-vpn-vpls-multipoint-service-Mjk0NQ%3D%3D\)](#) »

Last updated: April 22, 2016

»
CONTENTS

Note:

Initial Configuration & Diagrams: [Load the initial configuration files for the section named L2VPN VLAN Based, which can be found in CCIE SPv4 Topology Diagrams & Initial Configurations \(<http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs>\).](#) [Refer to the L2VPN VLAN Based Diagram in order to complete this task.](#)

A Note L2VPN Platform Support: [L2VPN data plane forwarding is not currently supported in XRv, however the control plan functionality is. Throughout the L2VPN sections of the workbook, we will configure the different variations of L2VPN on the XRv platform to explore IOS-XR specific syntax, and to validate the protocol mechanics via the control-plane state.](#)

Task

- Configure R2, R4, R5 and XR1 for EoMPLS MPLS as follows:
 - The attachment circuits are the Ethernet links connecting to R1, R7, R8 and XR2 respectively.
 - R2 and XR1 should form a pseudowire between their Loopback0 interfaces.
 - Extend VLAN 100 between R1 and XR2.
 - R4 and R5 should form a pseudowire between their Loopback0 interfaces.
 - Extend VLAN 200 between R7 and R8.
 - No other VLANs should be allowed over each pseudowire.
- Once complete, R7 and R8 should form an OSPFv2 adjacency and have IP reachability to each other's Loopback0 networks.
- Note that R1 and XR1 will not form an OSPFv2 adjacency due to the XRv L2VPN limitation.

Configuration [Click to collapse](#)

```
!  
! Using EVC Config  
!  
R2:  
interface GigabitEthernet2  
  no ip address  
  negotiation auto  
  no keepalive  
  service instance 10 ethernet  
    encapsulation dot1q 100  
    xconnect 19.19.19.19 219 encapsulation mpls  
  !  
R4:  
interface GigabitEthernet2  
  no ip address  
  negotiation auto  
  no keepalive  
  service instance 4 ethernet  
    encapsulation dot1q 200  
    xconnect 5.5.5.5 45 encapsulation mpls  
  !  
R5:  
interface GigabitEthernet2  
  no ip address  
  speed 1000  
  no negotiation auto  
  no keepalive  
  service instance 5 ethernet  
    encapsulation dot1q 200  
    xconnect 4.4.4.4 45 encapsulation mpls  
  !  
XR1:  
interface GigabitEthernet0/0/0/1.100 l2transport  
  dot1q vlan 100  
  !  
l2vpn  
pw-class VC_219  
  encapsulation mpls  
  no transport-mode  
  transport-mode vlan  
  !  
  !  
xconnect group GROUP1  
p2p XR1_R2  
  interface GigabitEthernet0/0/0/1.100  
  neighbor ipv4 2.2.2.2 pw-id 219  
  pw-class VC_219  
  !  
  !
```

```
!  
!  
!  
! Non EVC Config  
!  
R2:  
interface GigabitEthernet2.100  
  encapsulation dot1q 100  
  xconnect 19.19.19 219 encapsulation mpls  
!  
R4:  
interface GigabitEthernet2.200  
  encapsulation dot1q 200  
  xconnect 5.5.5.5 45 encapsulation mpls  
!  
R5:  
interface GigabitEthernet2  
  encapsulation dot1q 200  
  xconnect 4.4.4.4 45 encapsulation mpls  
!  
XR1:  
interface GigabitEthernet0/0/0/1.100 l2transport  
  dot1q vlan 100  
!  
l2vpn  
pw-class VC_219  
  encapsulation mpls  
  no transport-mode  
  transport-mode vlan  
!  
!  
xconnect group GROUP1  
  p2p XR1_R2  
  interface GigabitEthernet0/0/0/1.100  
  neighbor ipv4 2.2.2.2 pw-id 219  
  pw-class VC_219  
!  
!  
!  
!
```

Verification

Just like the configuration of Point-to-Point pseudowires, the configuration can be done using the standard "legacy" syntax or using the EVC framework. In this configuration, we are leveraging what is referred to as "VC-Type 4", which is a VLAN based pseudowire. In the previous examples, any frames sent into the port would have been transmitted over the pseudowire, regardless of what VLAN tag was

<https://t.me/learningnets>

on the frame. The PE simply receives the frames on its access circuit and forwards them onto the pseudowire. That service is referred to as "VC Type 5" or Port Based pseudowire. In a VLAN based pseudowire, a single VLAN is extended.

Validation for this pseudowire is similar to the previous ones. The control-plane signaling is the same, thus the key items to check for are the status and labels allocated over the Targeted LDP session.

```
R2#show mpls l2transport vc detail
```

```
Local interface: Gi2 up, line protocol up, Eth VLAN 100 up
  Destination address: 19.19.19.19, VC ID: 219, VC status: up
    Output interface: Gi1.23, imposed label stack {40 16005}
    Preferred path: not configured
    Default path: active
    Next hop: 20.2.3.3
  Create time: 00:01:06, last status change time: 00:01:01
    Last label FSM state change time: 00:01:01
    Last peer autosense occurred at: 00:01:01
  Signaling protocol: LDP, peer 19.19.19.19:0 up
    Targeted Hello: 2.2.2.2(LDP Id) -> 19.19.19.19, LDP is UP
    Graceful restart: not configured and not enabled
    Non stop routing: not configured and not enabled
    Status TLV support (local/remote) : enabled/supported
      LDP route watch : enabled
      Label/status state machine : established, LruRru
    Last local dataplane status rcvd: No fault
    Last BFD dataplane status rcvd: Not sent
    Last BFD peer monitor status rcvd: No fault
    Last local AC circuit status rcvd: No fault
    Last local AC circuit status sent: No fault
    Last local PW i/f circ status rcvd: No fault
    Last local LDP TLV status sent: No fault
    Last remote LDP TLV status rcvd: No fault
    Last remote LDP ADJ status rcvd: No fault
  MPLS VC labels: local 42, remote 16005
  Group ID: local 0, remote 512
  MTU: local 1500, remote 1500
  Remote interface description: GigabitEthernet0_0_0_1.100
  Sequencing: receive disabled, send disabled
  Control Word: Off (configured: autosense)
  SSO Descriptor: 19.19.19.19/219, local label: 42
  Dataplane:
    SSM segment/switch IDs: 12317/4123 (used), PWID: 7
  VC statistics:
    transit packet totals: receive 0, send 6
    transit byte totals: receive 0, send 720
    transit packet drops: receive 0, seq error 0, send 0
```

```
R2#show mpls l2transport binding
```

```
Destination Address: 19.19.19.19, VC ID: 219
  Local Label: 42
    Cbit: 0, VC Type: Eth VLAN, GroupID: 0
    MTU: 1500, Interface Desc: n/a
    VCCV: CC Type: CW [1], RA [2], TTL [3]
      CV Type: LSPV [2]
  Remote Label: 16005
    Cbit: 0, VC Type: Eth VLAN, GroupID: 512
    MTU: 1500, Interface Desc: GigabitEthernet0_0_0_1.100
    VCCV: CC Type: RA [2], TTL [3]
      CV Type: LSPV [2]
```

Note that due to a limitation in the XRv implementation, the L2VPN control-plane mechanics will succeed, but data-plane forwarding will not. R1 and XR2 will not have reachability due to this reason.

```
RP/0/0/CPU0:XR1#show l2vpn xconnect detail
```

```
Thu Jun 25 01:06:01.392 UTC
```

```
Group GROUP1, XC XR1_R2, state is up; Interworking none
```

```
AC: GigabitEthernet0/0/0/1.100, state is up
```

```
Type VLAN; Num Ranges: 1
```

```
VLAN ranges: [100, 100]
```

```
MTU 1500; XC ID 0x1; interworking none
```

```
Statistics:
```

```
  packets: received 0, sent 0
```

```
  bytes: received 0, sent 0
```

```
  drops: illegal VLAN 0, illegal length 0
```

```
PW: neighbor 2.2.2.2, PW ID 219, state is up ( established )
```

```
PW class VC_219, XC ID 0xff000001
```

```
Encapsulation MPLS, protocol LDP
```

```
Source address 19.19.19.19
```

```
PW type Ethernet VLAN, control word disabled, interworking none
```

```
PW backup disable delay 0 sec
```

```
Sequencing not set
```

```
PW Status TLV in use
```

MPLS	Local	Remote
-----	-----	-----
Label	16005	42
Group ID	0x200	unknown
Interface	GigabitEthernet0/0/0/1.100	unknown
MTU	1500	1500
Control word disabled		disabled
PW type	Ethernet VLAN	Ethernet VLAN
VCCV CV type 0x2		0x2
	(LSP ping verification)	(LSP ping verification)
VCCV CC type 0x6		0x6
	(router alert label)	(router alert label)
	(TTL expiry)	(TTL expiry)
-----	-----	-----

```
Incoming Status (PW Status TLV):
```

```
  Status code: 0x0 (Up) in Notification message
```

```
Outgoing Status (PW Status TLV):
```

```
  Status code: 0x0 (Up) in Notification message
```

```
MIB cpwVcIndex: 4278190081
```

```
Create time: 25/06/2015 00:49:04 (00:16:57 ago)
```

```
Last time status changed: 25/06/2015 01:01:25 (00:04:36 ago)
```

```
Statistics:
```

```
  packets: received 0, sent 0
```

```
  bytes: received 0, sent 0
```

R7 and R8 have connectivity using their VLAN200 sub-interfaces.

```
R7#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
8.8.8.8	1	FULL/DR	00:00:36	10.0.200.8	GigabitEthernet2.200

```
R7#show ip route ospf
```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default, U - per-user static route

o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP

a - application route

+ - replicated route, % - next hop override

Gateway of last resort is not set

8.0.0.0/32 is subnetted, 1 subnets

```
0      8.8.8.8 [110/2] via 10.0.200.8, 00:06:52, GigabitEthernet2.200
```

```
R7#ping 8.8.8.8 source 7.7.7.7
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:

Packet sent with a source address of 7.7.7.7

```
!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/7/21 ms

Lets shut down the link between R4 and R5 and perform a packet capture on R6. This will allow us to see a two label stack along with the VLAN tag being transported across the pseudowire between R4 and R5. Note that currently R4 and R5 are using their Gig1.45 link to reach each other. So MPLS encapsulated packets between R4 and R5 will not carry a Transport label, due to the PHP process.

```
R4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R4(config)#interface Gig1.45
R4(config-subif)#shut
R4(config-subif)#end

R4#
%OSPF-5-ADJCHG: Process 1, Nbr 5.5.5.5 on GigabitEthernet1.45 from FULL to DOWN, Neighbor Down: Interface down or detached
%SYS-5-CONFIG_I: Configured from console by console

R6#monitor capture CAP match any interface Gig1.46 both
R6#monitor capture CAP start
%BUFCAP-6-ENABLE: Capture Point CAP enabled.

R7#ping 8.8.8.8 source 7.7.7.7
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
Packet sent with a source address of 7.7.7.7
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/20 ms

R6#monitor capture CAP stop
R6#
%BUFCAP-6-DISABLE: Capture Point CAP disabled.

R6#monitor capture CAP export ftp://cisco:cisco@169.254.254.1/l2vpn.eompls.vlan.001.pcap
Writing l2vpn.eompls.vlan.001.pcap
Exported Successfully
```

Notice that VLAN tag 200 is on the frame being transported over the core.

We EVC, we can actually remove the VLAN tag as the frames traverse the pseudowire, and have the VLAN tag pushed back onto the frame as it exits the Access Circuit.

```
R4:
interface GigabitEthernet2
 service instance 4 ethernet
  rewrite ingress tag pop 1 symmetric

R5:
interface GigabitEthernet2
 service instance 5 ethernet
  rewrite ingress tag pop 1 symmetric
!
```

Notice that after applying this config the frames over the pseudowire don't carry a user VLAN.

```
R6#no monitor capture CAP
R6#monitor capture CAP match any interface Gig1.46 both
R6#monitor capture CAP start
%BUFCAP-6-ENABLE: Capture Point CAP enabled.

R7#ping 8.8.8.8 source 7.7.7.7
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
Packet sent with a source address of 7.7.7.7
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/20 ms

R6#monitor capture CAP stop
R6#
%BUFCAP-6-DISABLE: Capture Point CAP disabled.

R6#monitor capture CAP export ftp://cisco:cisco@169.254.254.1/l2vpn.eompls.vlan.rewrite.001.pcap
Writing l2vpn.eompls.vlan.rewrite.001.pcap
Exported Successfully
```

So although there is no user VLAN on the frames through the core, the EFP's Advanced Frame Manipulation stage of the pipeline pushes the correct VLAN back onto the frame as its sent out of the Access Circuit. As mentioned during the previous section, the `symmetric` argument allows for this reverse action to occur.

« MPLS L2 VPN - Port Based Point-to-Point Service using EVC (/workbook/view/service-provider-v4/task/mpls-l2-vpn-port-based-point-to-point-service-using-enc-Mjk0Mw%3D%3D) | MPLS L2 VPN - VPLS - Multipoint Service (/workbook/view/service-provider-v4/task/mpls-l2-vpn-vpls-multipoint-service-Mjk0NQ%3D%3D) »