# Xconnect over VRF Aware L2TPv3 in ASR1K

## Contents



## Introduction

This document describes how the Virtual Routing and Forwarding (VRF) can be used when you configure Layer 2 Tunneling Protocol (L2TP)v3 Xconnect over IP and Multiprotocol Label Switching (MPLS) network.

## Background Information

L2TP is the tunneling protocol used by Internet Service Providers (ISPs) in order to provide Virtual Private Network (VPN) in the dial access space over the internet.

It combines the best of Cisco's Layer 2 Forwarding (L2F) protocol and Microsoft's Point-to-Point Tunneling Protocol (PPTP). The main components of L2TP are L2TP Access Controller (LAC) and L2TP Network Server (LNS).

L2TP Access Controller: LAC is an access server connected to Public Switched Telephone Network (PSTN). The LAC is the initiator of incoming calls and the receiver of outgoing calls. It is connected to LNS over LAN or WAN.

L2TP Network Server: LNS is the network server for L2TP protcol where PPP sessions terminate and are authenticated. The LNS is the initiator of outgoing calls and the receiver of incoming calls.

L2TPv2 was designed to carry PPP traffic over IP networks. Network access equipment (DSL, cable modem or dial-up access interfaces) accepted PPP connections from subscribers and tunnelled the PPP sessions to the ISP over L2TP. The new version L2TPv3 is designed to carry any Layer 2 payload in addition to PPP which was the only payload that was supported by version 2.  Specifically, L2TPv3 defines the L2TP protocol for tunneling Layer 2 payloads over an IP core network with the use Layer 2 VPNs. Benefits of this feature include this:

- L2TPv3 simplifies deployment of VPNs
- L2TPv3 does not require MPLS
- L2TPv3 supports Layer 2 tunneling over IP for any payload

Here is the sample configuration of L2TPv3 pseudowire:

**1.enable**

**2.configureterminal**

**3.interface** type slot/port

**4.xconnect**peer-ip-address vcid**encapsulation l2tpv3pw-class**pw-class-name

Now take a look at how L2TPv3 Xconnect behaves when VRF is used. Here is the topology that is used for demonstration in which we Xconnect is configured between CPE and ASR1002 (IP) and ASR1004 (MPLS) with endpoints at ASR1000 in VRF (VRF Aware L2TPv3 is not supported on ASR1000 platform).

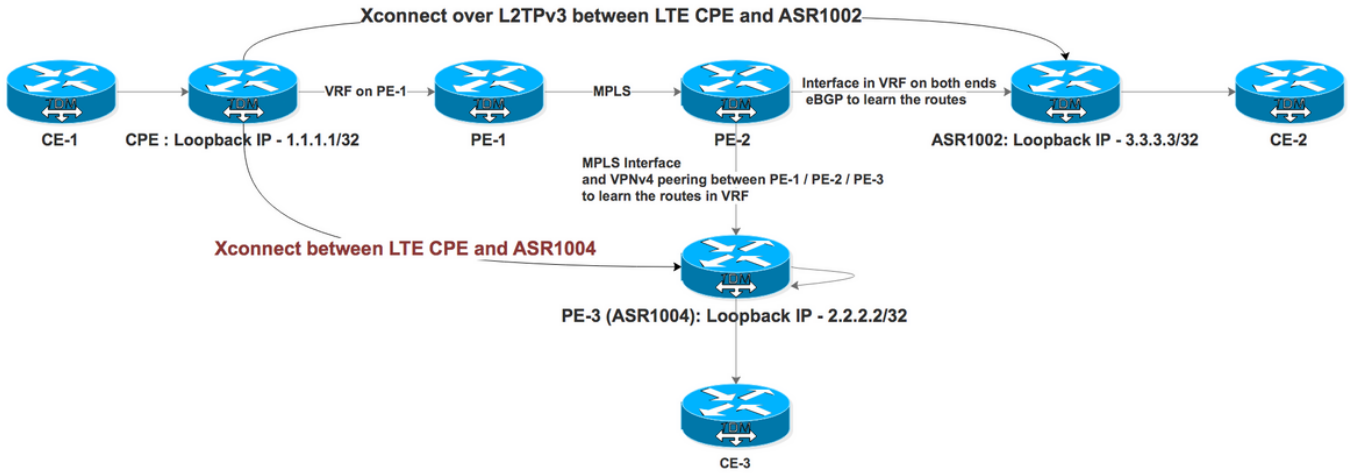


# Test Case I: L2TPv3 Xconnect over IP network with Endpoints in VRF

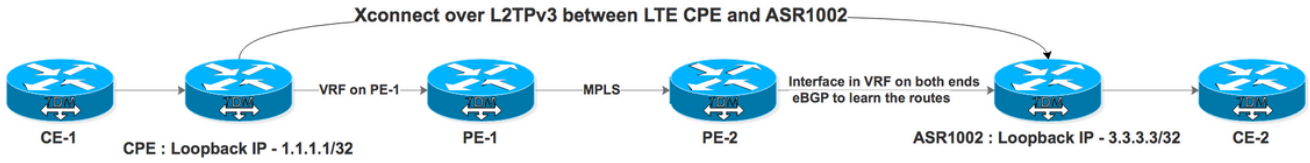


---

PE-1 and PE-2 make the MPLS network for ISP. CPE is connected to PE-1 over VRF and ASR1002 is connected to PE-2 over VRF. ASR1002 also has VRF on the interface connected to PE-2. The reachability of CPE loopback from ASR1002 is via VRF over IP interface.

Configuration on CPE for Xconnect towards ASR1002:

```
interface FastEthernet4.2381

encapsulation dot1Q 2381

xconnect 3.3.3.3 2381 encapsulation l2tpv3 pw-class PSEUDO_CLASS >>>>>>>>>>> Xconnect with
ASR1002
```

```
pseudowire-class PSEUDO_CLASS

encapsulation l2tpv3

interworking vlan

protocol l2tpv3 L2TP_CLASS

ip local interface Loopback0

ip tos reflect


l2tp-class L2TP_CLASS

authentication

password cisco


interface Gigabit0/1

ip address 192.168.8.190 255.255.255.0

end


Interface Loopback0

ip address 1.1.1.1 255.255.255.255

end


ip route 0.0.0.0 0.0.0.0 192.168.8.1 >>>>>>>>>>>>>>> Default route towards PE-1
```

Working Configuration on ASR1002:

```
interface GigabitEthernet0/0/0.906 ————————————————> Interface connected to PE-2 is in VRF

encapsulation dot1Q 906

ip vrf forwarding L2TP_VRF

ip address 10.1.1.1 255.255.255.252


interface GigabitEthernet0/0/1.2381

encapsulation dot1Q 2381

xconnect 1.1.1.1 2381 encapsulation l2tpv3 pw-class PSEUDO_CLASS


pseudowire-class PSEUDO_CLASS

encapsulation l2tpv3
```

```
interworking vlan

protocol l2tpv3 L2TP_CLASS

ip local interface Loopback11



l2tp-class L2TP_CLASS

authentication

password cisco


interface Loopback11

ip vrf forwarding L2TP_VRF ————————————————————> Source is in VRF

ip address 3.3.3.3 255.255.255.255


router bgp 1


address-family ipv4 vrf L2TP_VRF

redistribute connected

neighbor 10.1.1.2 remote-as 2 ————————————————> eBGP with PE-2 in VRF

neighbor 10.1.1.2 activate

neighbor 10.1.1.2 soft-reconfiguration inbound

exit-address-family
```

VRF L2TP_VRF:

```
B       1.1.1.1/32 [20/0] via 10.1.1.2, 1d ——————————> Xconnect end point learned via eBGP in
VRF
```

Let us now check the status of Xconnect on CPE:

```
CPE #sh xconnect all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

 UP=Up       DN=Down            AD=Admin Down      IA=Inactive

 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                        S1 Segment 2                        S2
```

```
------+--------------------------------+--+--------------------------------+--

UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp 3.3.3.3:2381             DOWN
           Interworking: vlan                   Session ID: 1906980494
                                                 Tunnel ID: 2886222725
                                                Protocol State: DOWN
                                                Remote Circuit State: DOWN
                                                pw-class: PSEUDO_CLASS_VLAN
```

It says Segment 2 is down, which means the path from CPE to ASR1002 is having an issue. However, we are able to ping the Endpoint. The debugs on CPE shows that tunnel to Endpoint is failed or there is no route to endpoint.

```
CPE #sh xconnect all de
Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
 UP=Up       DN=Down            AD=Admin Down      IA=Inactive
 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware

XC ST  Segment 1                        S1 Segment 2                         S2
------+--------------------------------+--+--------------------------------+--

UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp 3.3.3.3:2381             DOWN
           Interworking: vlan                   Session ID: 1906980494
                                                 Tunnel ID: 2886222725
                                                Protocol State: DOWN
                                                Remote Circuit State: DOWN
                                                pw-class: PSEUDO_CLASS_VLAN
```

The main issue here is that Endpoint is reachable via VRF on ASR1002. The Xconnect endpoint needs to be in Global Routing Table for it to come up. Let us now configure a route for CPE Loopback 1.1.1.1/32 in global pointing to interface GigabitEthernet0/0/0.906 which is itself in VRF.

```
 ip route 1.1.1.1 255.255.255.255 GigabitEthernet0/0/0.906 10.1.1.2

             S        1.1.1.1/32 [1/0] via 10.1.1.2, GigabitEthernet0/0/0.906
```

Once the dummy static route is configured, Xconnect comes up. You can also point it to Null0.

This is a workaround to let the router believe that Endpoint is reachable via Global not VRF and is just used for Control Plane. The actual data plane traffic will be via VRF only.

Here are the ping results with and without VRF:

```
 ip route 1.1.1.1 255.255.255.255 GigabitEthernet0/0/0.906 10.1.1.2
```

```
             S        1.1.1.1/32 [1/0] via 10.1.1.2, GigabitEthernet0/0/0.906
```

Status of Xconnect on CPE:

```
CPE #sh xconnect all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

 UP=Up       DN=Down            AD=Admin Down      IA=Inactive

 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                          S2

------+---------------------------------+--+---------------------------------+--


UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp3.3.3.3:2381             UP

           Interworking: vlan                   Session ID: 1906980494

                                                 Tunnel ID: 2886222725

                                                Protocol State: UP

                                                Remote Circuit State: UP

                                                pw-class: PSEUDO_CLASS_VLAN
```

# Test Case II: L2TPv3 Xconnect over MPLS network with Endpoints in VRF

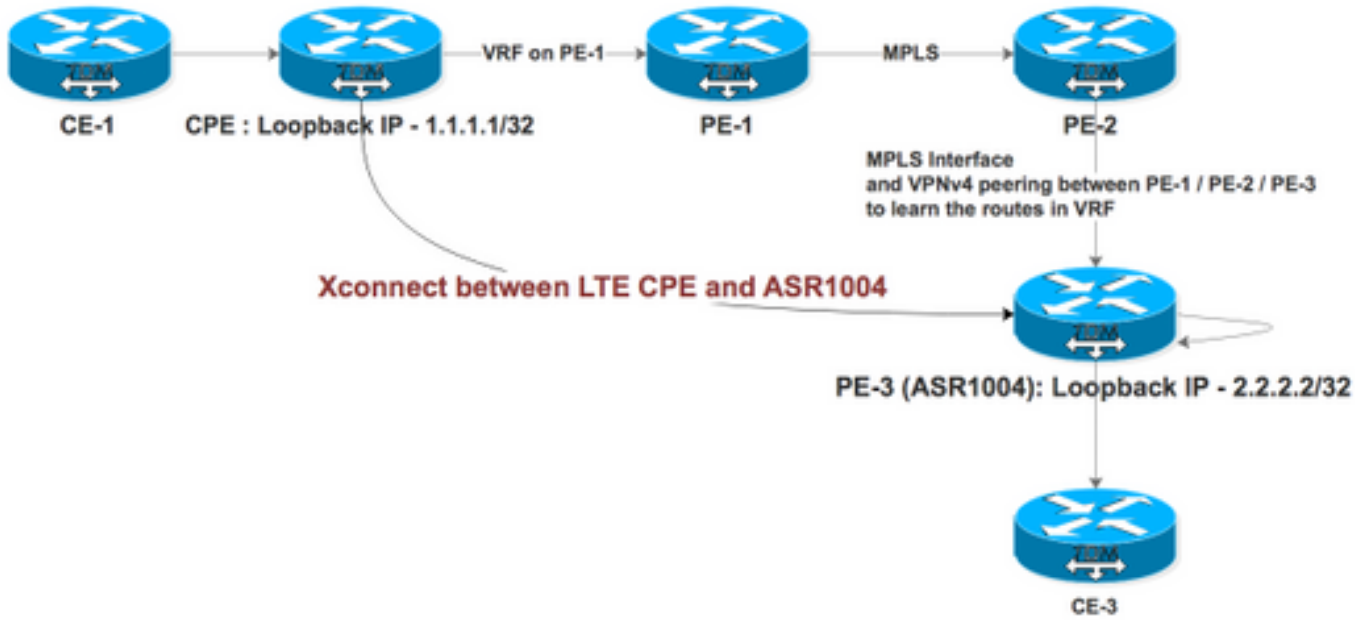


---

PE-1, PE-2 and PE-3 make the MPLS network for ISP with PE-2 acting as Route Reflector (RR). CPE is connected to PE-1 over VRF and ASR1004 is connected to PE-2 with MPLS enabled on the interface. ASR1004 also has VRF in which it is supposed to receive the VPNv4 routes from PE-1 via RR. The reachability of CPE looback from ASR1004 is via VRF over MPLS interface.

Configuration on CPE for Xconnect towards ASR1004:

```
CPE #sh xconnect all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

 UP=Up       DN=Down            AD=Admin Down      IA=Inactive

 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                          S2

------+--------------------------------+--+--------------------------------+--


UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp3.3.3.3:2381             UP

           Interworking: vlan                   Session ID: 1906980494

                                                 Tunnel ID: 2886222725

                                                Protocol State: UP

                                                Remote Circuit State: UP

                                                pw-class: PSEUDO_CLASS_VLAN
```

Configuration on ASR1004:

```
CPE #sh xconnect all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

 UP=Up       DN=Down            AD=Admin Down      IA=Inactive

 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                         S2

------+---------------------------------+--+---------------------------------+--


UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp3.3.3.3:2381             UP

           Interworking: vlan                   Session ID: 1906980494

                                                 Tunnel ID: 2886222725

                                                Protocol State: UP

                                                Remote Circuit State: UP

                                                pw-class: PSEUDO_CLASS_VLAN
```

Route entry for Xconnect End Point:

```
ASR1004#sh ip rou vrf L2TP_VRF 1.1.1.1 .   ————————————————————> Xconnect End Point also learned
via VRF

Routing Table: L2TP_VRF
Routing entry for 1.1.1.1/32
Known via "bgp 2", distance 200, metric 0, type internal
Last update from 11.11.11.11 6d17h ago
Routing Descriptor Blocks:
* 11.11.11.11 (default), from 22.22.22.22, 6d17h ago
Route metric is 0, traffic share count is 1
AS Hops 0
MPLS label: 18
MPLS Flags: MPLS Required


We observed that Segment 2 was continuously flapping on both ends.


ASR1004#sh xc all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

  UP=Up       DN=Down            AD=Admin Down      IA=Inactive

  SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                         S2
```

```
------+---------------------------------+--+---------------------------------+--

DN pri   ac Gi0/0/1:2380(Eth VLAN)       UP l2tp 1.1.1.1:2380           DN  >>>>>>>>>>>>>>

            Interworking: vlan                   Session ID: 2543426569

                                                 Tunnel ID: 3352120314

                                                 Protocol State: DOWN

                                                 Remote Circuit State: DOWN

                                                 pw-class: PSEUDO_CLASS_VLAN

ASR1004#sh xc all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

  UP=Up       DN=Down            AD=Admin Down      IA=Inactive

  SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                         S2

------+---------------------------------+--+---------------------------------+--

UP pri   ac Gi0/0/1:2380(Eth VLAN)       UP l2tp 1.1.1,1:2380           UP >>>>>>>>>>>>>>>>

            Interworking: vlan                   Session ID: 2543426569

                                                 Tunnel ID: 3352120314

                                                 Protocol State: UP

                                                 Remote Circuit State: UP

                                                 pw-class: PSEUDO_CLASS_VLAN
```

Logs from CPE:

```
CPE#sh xconnect all de

Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State

 UP=Up       DN=Down            AD=Admin Down      IA=Inactive

 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                         S2

------+---------------------------------+--+---------------------------------+--

DN pri   ac Fa4.2380:2380(Eth VLAN)      UP l2tp 2.2.2.2:2380            DN ---------------à
Flapping with ASR1004

           Interworking: vlan                  Session ID: 3434660693

                                                Tunnel ID: 1760690853
```

```
                                          Protocol State: DOWN
                                          Remote Circuit State: DOWN
                                          pw-class: PSEUDO_CLASS
UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp 3.3.3.3:2381             UP ------------------
---à Stable with ASR1002
           Interworking: vlan                  Session ID: 1906980494
                                          Tunnel ID: 2886222725
                                          Protocol State: UP
                                          Remote Circuit State: UP
                                          pw-class: PSEUDO_CLASS


CPE#sh l2tp session


L2TP Session Information Total tunnels 2 sessions 2


LocID      RemID      TunID      Username, Intf/      State  Last Chg Uniq ID
                                 Vcid, Circuit
2714490989 3697021268 1760690853 2380, Fa4.2380:2380  est   00:00:03 0  ————————> Flapping with
ASR1004
1906980494 2361475239 2886222725 2381, Fa4.2381:2381  est   15:37:06 0       ————————> Stable
with ASR1002
```

You cannot configure a static route in this case as exit interface is MPLS enabled interface. As a workaround, there are two interfaces looped back to each other and configured one in VRF with other in global. Then configured a static route in global pointing towards VRF interface, with this Xconnect became stable.

```
CPE#sh xconnect all de
Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
 UP=Up       DN=Down            AD=Admin Down      IA=Inactive
 SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware


XC ST  Segment 1                         S1 Segment 2                          S2
------+---------------------------------+--+---------------------------------+--
DN pri   ac Fa4.2380:2380(Eth VLAN)      UP l2tp 2.2.2.2:2380           DN --------------à
Flapping with ASR1004
           Interworking: vlan                  Session ID: 3434660693
```

```
                                          Tunnel ID: 1760690853

                                          Protocol State: DOWN

                                          Remote Circuit State: DOWN

                                          pw-class: PSEUDO_CLASS

UP pri   ac Fa4.2381:2381(Eth VLAN)      UP l2tp 3.3.3.3:2381             UP ------------------
---à Stable with ASR1002

           Interworking: vlan                   Session ID: 1906980494

                                          Tunnel ID: 2886222725

                                          Protocol State: UP

                                          Remote Circuit State: UP

                                          pw-class: PSEUDO_CLASS


CPE#sh l2tp session


L2TP Session Information Total tunnels 2 sessions 2


LocID      RemID      TunID      Username, Intf/      State  Last Chg Uniq ID

                                 Vcid, Circuit

2714490989 3697021268 1760690853 2380, Fa4.2380:2380  est   00:00:03 0  ————————> Flapping with
ASR1004

1906980494 2361475239 2886222725 2381, Fa4.2381:2381  est   15:37:06 0       ————————> Stable
with ASR1002
```

L2TP Session Information Total tunnels 2 sessions 2:

```
LocID      RemID      TunID      Username, Intf/      State  Last Chg Uniq ID

                                 Vcid, Circuit

2714490989 3697021268 1760690853 2380, Fa4.2380:2380  est   00:20:03 0

1906980494 2361475239 2886222725 2381, Fa4.2381:2381  est   15:37:06 0
```

The traffic flow is seen as in case of ASR1004:

- When traffic comes from CPE on ASR1004, it comes in MPLS interface Gi0/0/1 and gets switched directly to Gi0/0/0 Access port.
- When traffic comes from Access Port Gi0/0/0, it takes the looped path of Gi0/0/0 -> Gi0/0/2 -> Gi0/0/3 -> Gi0/0/1.

The main issue with this workaround is for QFP utilization on ASR1000 platform as packet processing is done twice:

```
LocID      RemID      TunID      Username, Intf/      State  Last Chg Uniq ID
                                 Vcid, Circuit
2714490989 3697021268 1760690853 2380, Fa4.2380:2380  est   00:20:03 0
1906980494 2361475239 2886222725 2381, Fa4.2381:2381  est   15:37:06 0
```

This behavior is documented in Doc Bug: CSCvi42964