VyOS
The universal networking platform

Deployment Guide | Technical Doc

EVPN-VXLAN FOR L3VPN

April 2025

Introduction

This document is intended to be a contribution for VyOS Community, and a guide to show you how to configure on VyOS a L3VPN service using EVPN-VXLAN.

EVPN standards define two forwarding models that can be used to provide inter-subnet connectivity:

- **Asymmetric forwarding model**
- **Symmetric forwarding model**

Both models are supported by VyOS and will be explained and configured in this deployment guide.

An IP-VRF enables the communication between hosts residing in different subnets. Each subnet is represented by a bridging domain (BD) that is connected to the IP-VRF using a Layer-3 interface (SVI). In VyOS an SVI can either be a traditional bridge or a VLAN subinterface of a VLAN-aware bridge. The IP address allocated to that interface belongs to the subnet and is used as the default gateway for hosts in that subnet.

A PE maintains a virtual routing and forwarding (VRF) table for each provisioned VRF. The VRF stores the customer routes. An address resolution protocol (ARP) table is also maintained for each provisioned bridging domain to store the resolved IP-to-MAC address mapping for the IP-VRF.

Traffic between hosts belonging to the same subnet is forwarded based on the MAC forwarding database (FDB) maintained for each BD.

Traffic between hosts belonging to different subnets needs to be routed using the IP-VRF in addition to being forwarding in each of the ingress and egress bridge.

BGP-EVPN is the control plane for the transport of Ethernet frames, regardless of whether those frames are bridged or routed. An EVPN with integrated routing and bridging (EVPN-IRB) enables the use of EVPN to provide a inter-subnet connectivity between customers hosts.

In each PE provisioned, the bridging domain is bonded to the IP-VRF using a Layer-3 interface, referred to as an IRB interface or SVI. The SVI interface links the FDB of the BD to the VRF table of the IP-VRF and enables inter-subnet forwarding.

Multiple bridging domains may be connected to an IP-VRF and can therefore have multiple SVI interfaces. Traffic between hosts on the same subnet is forwarded by the PE by consulting the FDB table of the corresponding bridging domain. Traffic destined to a host in a different subnet needs to be routed based on the VRF table.

Layer-3 EVPN Asymmetric Forwarding Model

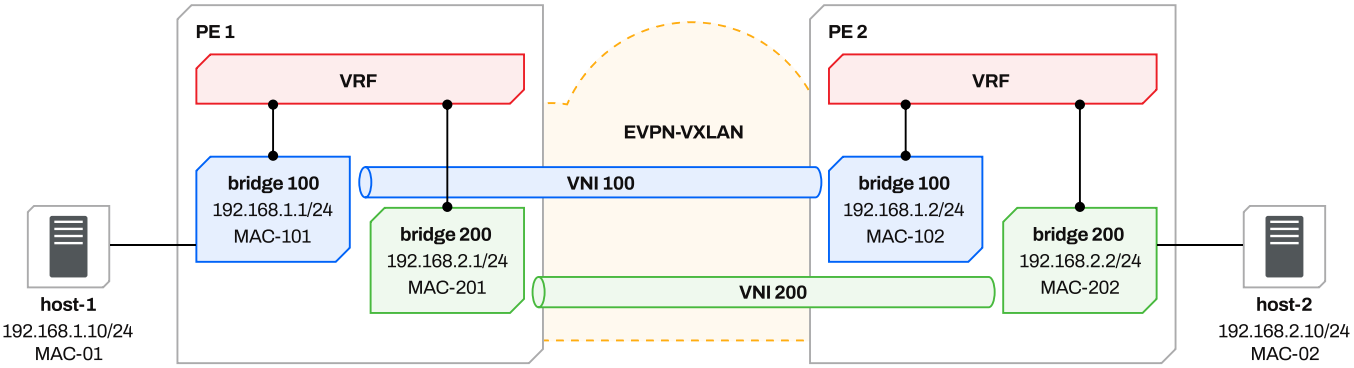

Figure-1: Asymmetric forwarding model

This model requires the creation of all bridging domains on all PE's participating in the EVPN-IRB, even if there are no connected hosts. PE1 consults its FDBs and its VRF table to perform L2 forwarding and L3 routing. The PE2 however only needs to consult its FDB to forward the traffic to its destination.

This model is simple to implement but is not scalable, since every participating PE needs to maintain FDB entries and ARP entries for all hosts in all bridging domains.

In the asymmetric model, all subnets are local in the VRF and there is no need to advertise IP-Prefix routes. MAC/IP routes are exchanged in the core for each provisioned bridging domain. Each PE advertises one MAC/IP route per connected host. The PE can also advertise one MAC/IP route per SVI interface.

Advertised Routes:

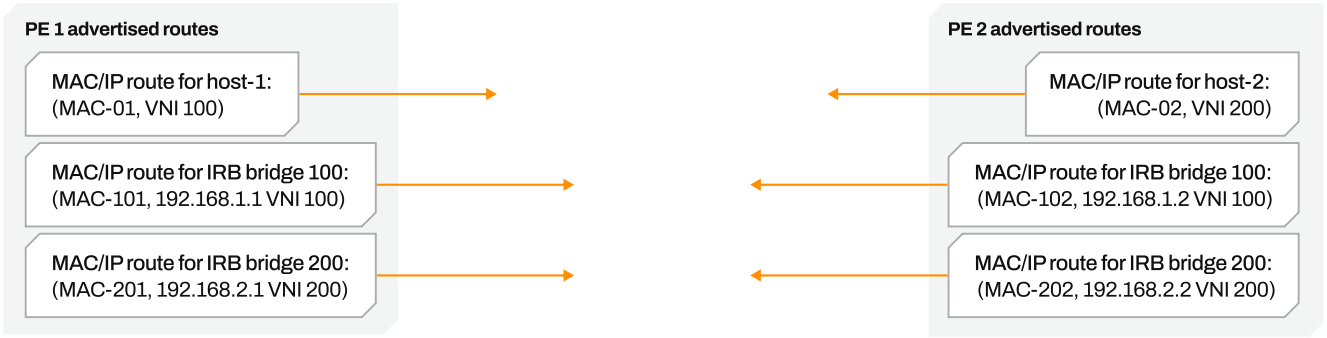

Figure-2: Asymmetric model – advertised routes

All traffic originated from host 1 to host 2 is forwarded as follows:

1. Host-1 is configured to use 192.168.1.1 as default gateway.
2. PE 1 consults FDB of bridge 100 and forwards the packet to the IRB bridge 100 interface.
3. PE 1 consults the VRF route table and determines that the next-hop is the local interface to bridge 200.
4. PE 1 consults FDB of bridge 200 and sends the packet to PE 2 over the VxLAN tunnel with VNI 200.
5. PE 2 decapsulates the packet, consults FDB of bridge 200, and forwards the packet to Host-2.

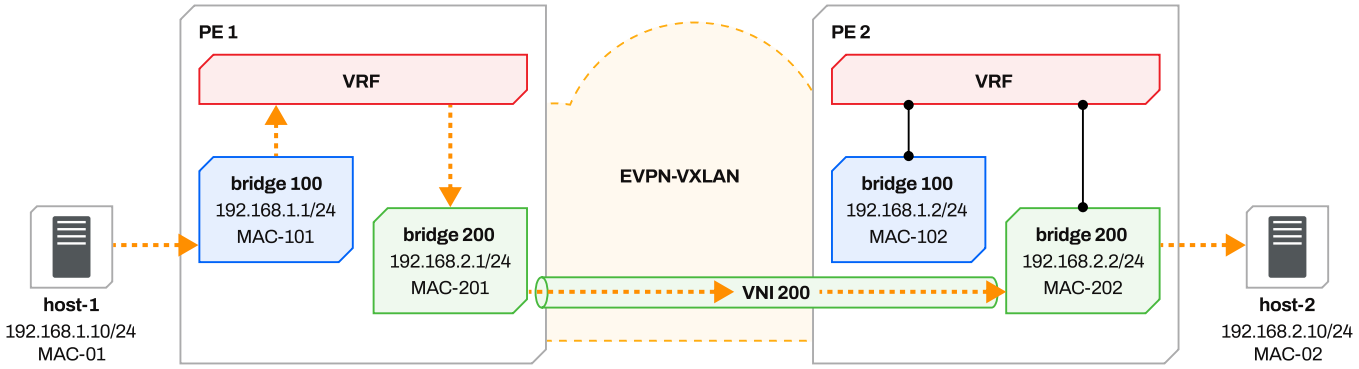

Figure-3: Asymmetric model – forwarding

Layer-3 EVPN Symmetric Interface-less VRF-to-VRF forwarding model

In the symmetric forwarding model, bridge domains are instantiated on a PE router only when it has directly attached hosts. In this model, PE routers advertise IP reachability to their local subnets using EVPN IP Prefix routes. This approach resembles traditional VRF-based services, where VRFs are interconnected directly over tunnels.

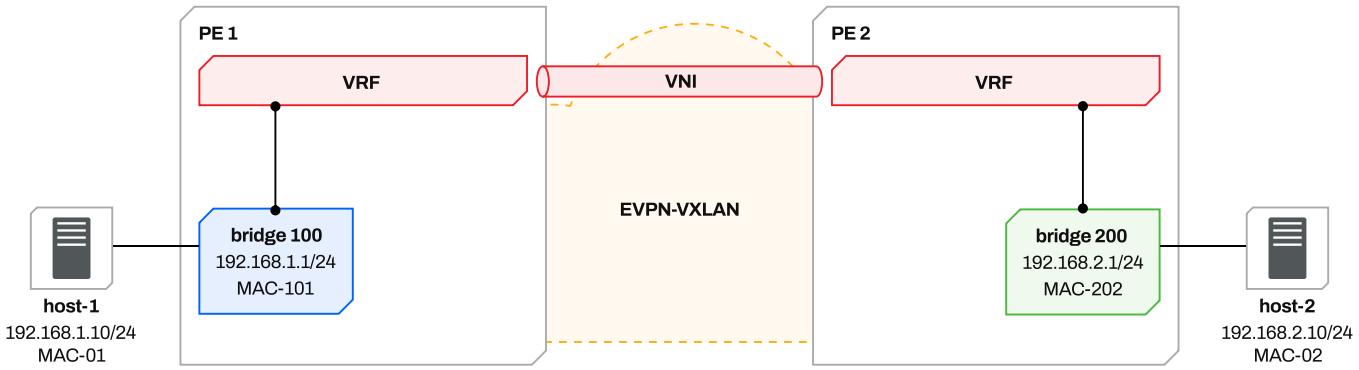

Figure-4: Symmetric forwarding model

The PE advertises an EVPN IP-Prefix route for any local route learned in the VRF. VRFs are connected directly over tunnels and local routes learned on the VRF are advertised using IP-Prefix routes. MAC/IP routes are not required.

Assuming that the different tenant subnets are not distributed across the core, the advertisement of IP-Prefix routes is sufficient for inter-subnet forwarding and there is no need to advertise any MAC/IP routes between PEs.

Advertised Routes:

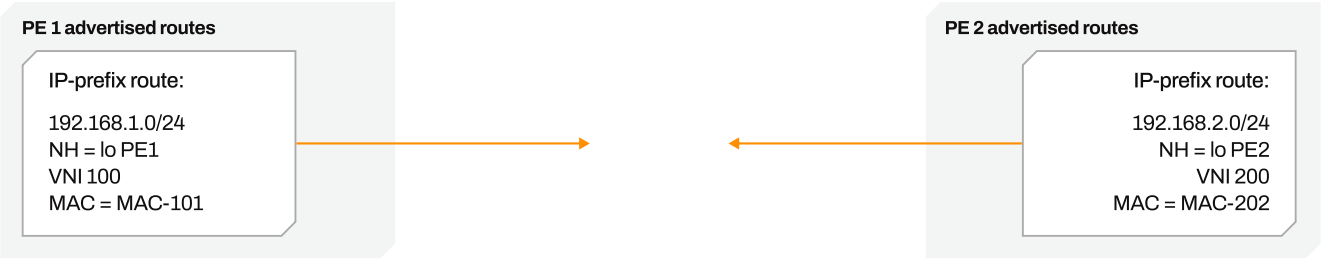

Figure-5: Symmetric model – advertised routes

For VXLAN encapsulation a VNI number and a router MAC address are included for data encapsulation. MAC address is provided in the Router MAC extended community.

Traffic originated from host-1 to host-2 is forwarded as follows:

- **PE 1 consults its FDB and forwards traffic towards bridge 100 interface**
- **PE 1 consults its VRF, encapsulates the packet, and forwards it towards PE 2**
- **PE 2 identifies the VRF using the VNI, consults its VRF, and forwards the packet towards the local bridge**
- **PE 2 performs an FDB lookup and forwards the packet to the local host-2**

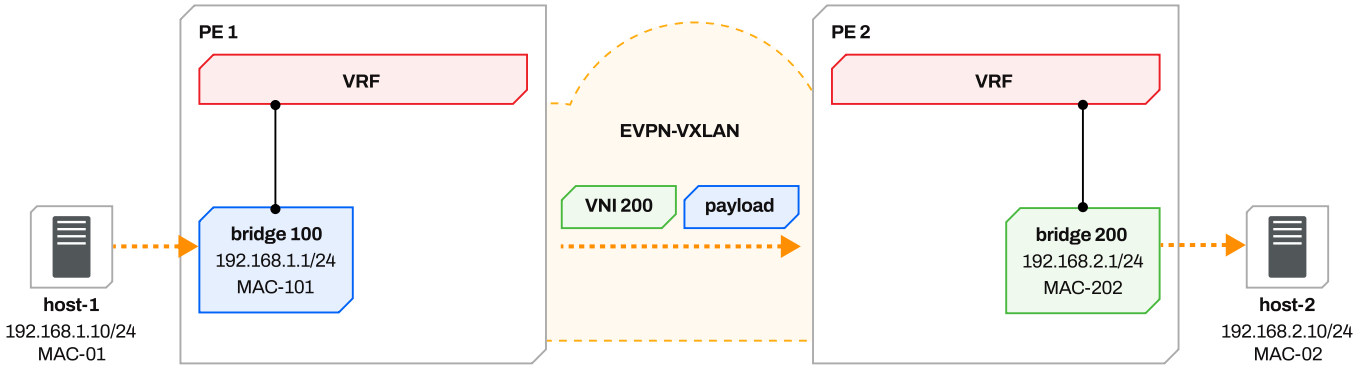

Figure-6: Symmetric model – forwarding

Scenario | Topology

In this lab setup, we're using a simple topology with two PE routers, connecting two remote customer sites to the IP backbone. The IP backbone is running OSPFv2 as IGP for loopbacks redistribution. In a real-world deployment, the IP backbone may be of any design and scale, as long as there is IP reachability between the loopback interfaces of the PEs. The underlying IGP can be any routing protocol, such as OSPF or IS-IS. P1 acts as route-reflector for bgp l2vpn address-famfily. The PE routers establish an iBGP session with the Route Reflector (RR) for the exchange of EVPN routes, using their loopback interfaces as the transport addresses. In a real-world scenario, of course, the RR would not be in the data path and could be either a physical or virtual router located somewhere in the IP backbone.

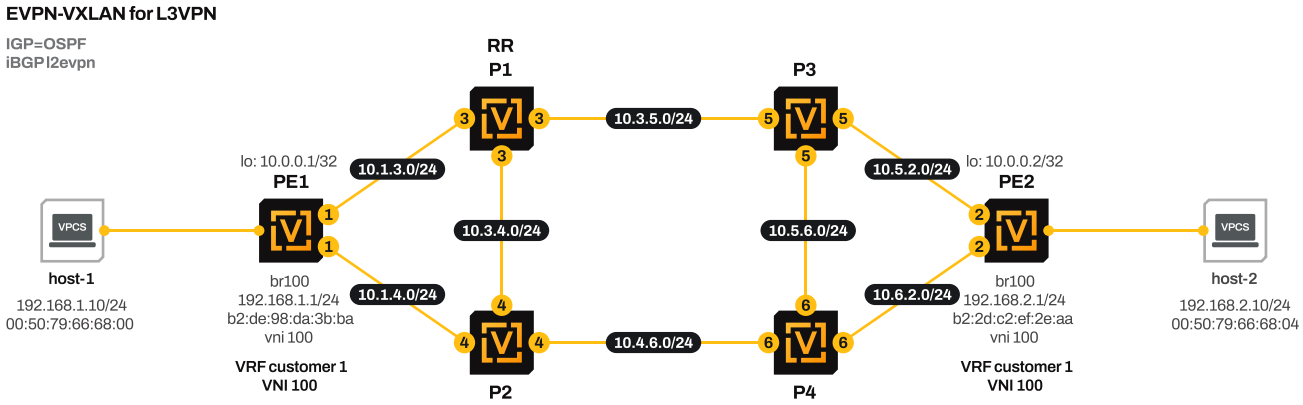

Figure-7: lab topology

Configurations and Deployment

Layer-3 EVPN Asymmetric Model

As seen in Figure-8, in our example we are implementig one customer VRF which has two subnets in two remote sites.

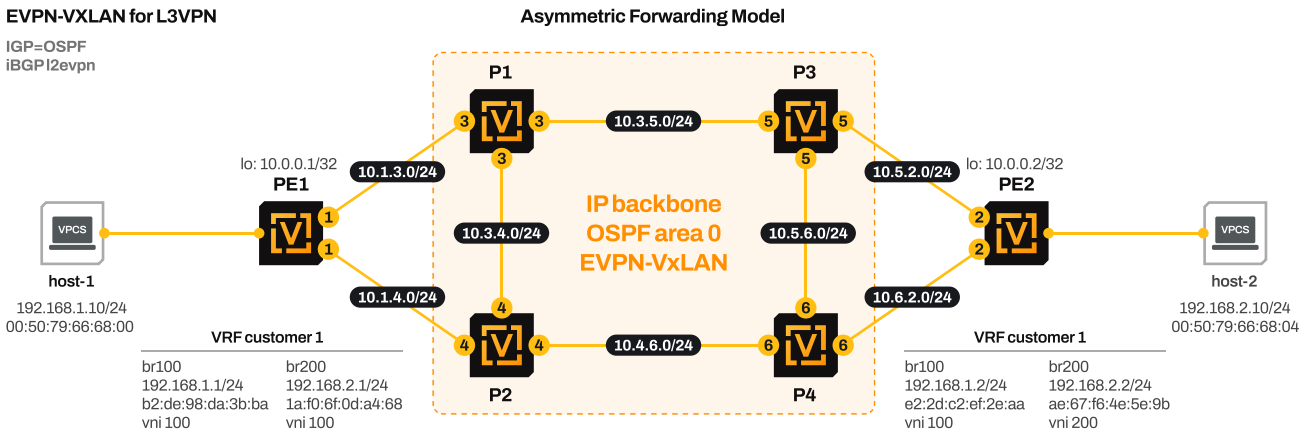

Figure-8: L3VPN asymmetric model deployment

Let's see a simple configuration for the IP Backbone.

IP backbone

P1

```
set interfaces loopback lo address '10.0.0.3/32'

set interfaces ethernet eth0 address '10.1.3.3/24'
set interfaces ethernet eth0 description 'to_PE1'
set interfaces ethernet eth1 address '10.3.5.3/24'
set interfaces ethernet eth1 description 'to_P3'
set interfaces ethernet eth2 address '10.3.4.3/24'
set interfaces ethernet eth2 description 'to_P2'

set protocols ospf area 0 network '10.1.3.0/24'
set protocols ospf area 0 network '10.3.4.0/24'
set protocols ospf area 0 network '10.3.5.0/24'
set protocols ospf interface eth0 network 'point-to-point'
set protocols ospf interface eth1 network 'point-to-point'
set protocols ospf interface eth2 network 'point-to-point'
set protocols ospf parameters router-id '10.0.0.3'

# similar configuration for P2, P3 and P4

# bgp configuration just for P1 acting as route-reflector

set protocols bgp neighbor 10.0.0.1 peer-group 'PEs'
set protocols bgp neighbor 10.0.0.2 peer-group 'PEs'
se set protocols bgp parameters router-id '10.0.0.3'
set protocols bgp peer-group PEs address-family l2vpn-evpn route-reflector-client
set protocols bgp peer-group PEs remote-as 'internal'
set protocols bgp peer-group PEs update-source 'lo'
set protocols bgp system-as '65000'
```

PE1

```
set interfaces loopback lo address '10.0.0.1/32'

set interfaces ethernet eth0 address '10.1.3.1/24'
set interfaces ethernet eth0 description 'to_P1'
set interfaces ethernet eth1 address '10.1.4.1/24'
set interfaces ethernet eth1 description 'to_P2'
set interfaces ethernet eth2 description 'customer-1_site-1'

set protocols ospf area 0 network '10.1.3.0/24'
set protocols ospf area 0 network '10.1.4.0/24'
set protocols ospf interface eth0 network 'point-to-point'
set protocols ospf interface eth1 network 'point-to-point'
set protocols ospf parameters router-id '10.0.0.1'
set protocols ospf redistribute connected

set protocols bgp address-family l2vpn-evpn advertise ipv4 unicast
set protocols bgp address-family l2vpn-evpn advertise-all-vni
set protocols bgp neighbor 10.0.0.3 peer-group 'RR'
set protocols bgp parameters router-id '10.0.0.1'
set protocols bgp peer-group RR address-family l2vpn-evpn
set protocols bgp peer-group RR remote-as 'internal'
set protocols bgp peer-group RR update-source 'lo'
set protocols bgp system-as '65000'
```

PE2

```
set interfaces loopback lo address '10.0.0.2/32'

set interfaces ethernet eth0 address '10.2.5.2/24'
set interfaces ethernet eth0 description 'to_P3'
set interfaces ethernet eth1 address '10.2.6.2/24'
set interfaces ethernet eth1 description 'to_P4'
set interfaces ethernet eth2 description 'customer-1_site-2'

set protocols ospf area 0 network '10.2.5.0/24'
set protocols ospf area 0 network '10.2.6.0/24'
set protocols ospf interface eth0 network 'point-to-point'
set protocols ospf interface eth1 network 'point-to-point'
set protocols ospf parameters router-id '10.0.0.2'
set protocols ospf redistribute connected

set protocols bgp address-family l2vpn-evpn advertise ipv4 unicast
set protocols bgp address-family l2vpn-evpn advertise-all-vni
set protocols bgp neighbor 10.0.0.3 peer-group 'RR'
set protocols bgp parameters router-id '10.0.0.2'
set protocols bgp peer-group RR address-family l2vpn-evpn
set protocols bgp peer-group RR remote-as 'internal'
set protocols bgp peer-group RR update-source 'lo'
set protocols bgp system-as '65000'
```

Overlay – customer VRFs

Our customer “customer-1” is assigned an individual VRF. Every PE that provides access to the customer needs to have VRF and all bridging domians configured, bridge 100 and bridge 200 in our case.

PE1

```
set interfaces bridge br100 address '192.168.1.1/24'
set interfaces bridge br100 description 'customer-1_BD-100'
set interfaces bridge br100 member interface eth2
set interfaces bridge br100 member interface vxlan100
set interfaces bridge br100 vrf 'customer-1'

set interfaces bridge br200 address '192.168.2.1/24'
set interfaces bridge br200 description 'customer-1_BD-200'
set interfaces bridge br200 member interface vxlan200
set interfaces bridge br200 vrf 'customer-1'

set interfaces vxlan vxlan100 mtu '1500'
set interfaces vxlan vxlan100 parameters nolearning
set interfaces vxlan vxlan100 port '4789'
set interfaces vxlan vxlan100 source-address '10.0.0.1'
set interfaces vxlan vxlan100 vni '100'

set interfaces vxlan vxlan200 mtu '1500'
set interfaces vxlan vxlan200 parameters nolearning
set interfaces vxlan vxlan200 port '4789'
set interfaces vxlan vxlan200 source-address '10.0.0.1'
set interfaces vxlan vxlan200 vni '200'

set vrf name customer-1 protocols bgp address-family ipv4-unicast redistribute connected
set vrf name customer-1 protocols bgp address-family l2vpn-evpn advertise ipv4 unicast
set vrf name customer-1 protocols bgp system-as '65000'
set vrf name customer-1 table '100'
```

PE2

```
set interfaces bridge br100 address '192.168.1.2/24'
set interfaces bridge br100 description 'customer-1'
set interfaces bridge br100 member interface vxlan100
set interfaces bridge br100 vrf 'customer-1'

set interfaces bridge br200 address '192.168.2.2/24'
set interfaces bridge br200 description 'customer-1_BD-200'
set interfaces bridge br200 member interface eth2
set interfaces bridge br200 member interface vxlan200
set interfaces bridge br200 vrf 'customer-1'

set interfaces vxlan vxlan100 mtu '1500'
set interfaces vxlan vxlan100 parameters nolearning
set interfaces vxlan vxlan100 port '4789'
set interfaces vxlan vxlan100 source-address '10.0.0.2'
set interfaces vxlan vxlan100 vni '100'

set interfaces vxlan vxlan200 mtu '1500'
set interfaces vxlan vxlan200 parameters nolearning
set interfaces vxlan vxlan200 port '4789'
set interfaces vxlan vxlan200 source-address '10.0.0.2'
set interfaces vxlan vxlan200 vni '200'

set vrf name customer-1 protocols bgp address-family ipv4-unicast redistribute connected
set vrf name customer-1 protocols bgp address-family l2vpn-evpn advertise ipv4 unicast
set vrf name customer-1 protocols bgp system-as '65000'
set vrf name customer-1 table '100'
```

Layer-3 EVPN Symmetric Interface-less VRF-to-VRF model

As seen in Figure-9 in this example we are implementig one customer VRF who has two subnets connected to two remote sites.

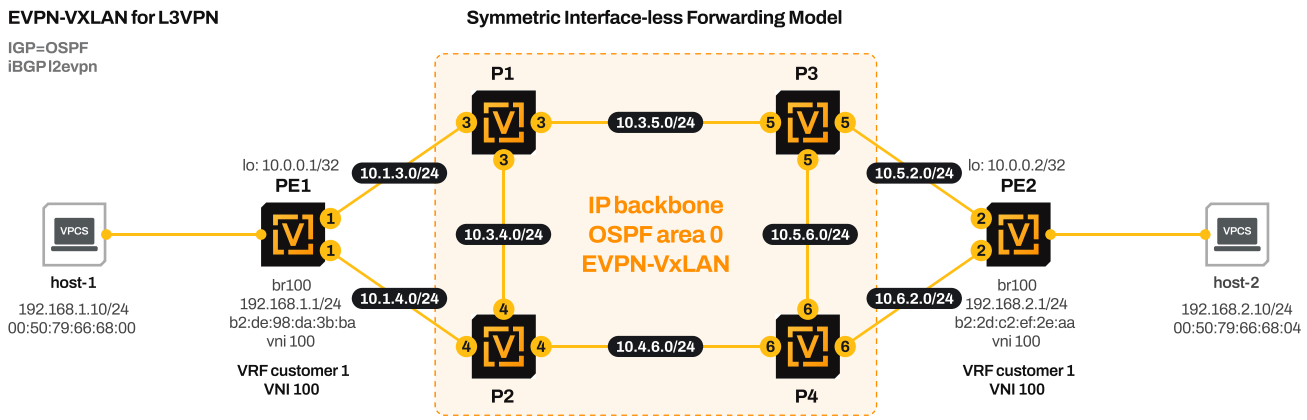

Figure-9: L3VPN symmetric model deployment

The IP backbone setup remains unchanged from the asymmetric model example. There's no need to modify the backbone configuration when switching to the symmetric model.

Overlay – customer VRFs

Our customer "customer-1" is assigned an individual VRF that would support having other customers with overlapping address ranges. Every PE that provides access to the customer needs to have VRF and VNI configured. Although it is possible to use different VNI values for VRFs in PE's, for the sake of simplicity we'll use VNI 100 for both, this way we can use route-target auto-derive and there's no need for explicitly configuration.

PE1

```
set interfaces bridge br100 address '192.168.1.1/24'
set interfaces bridge br100 description 'customer-1'
set interfaces bridge br100 member interface eth2
set interfaces bridge br100 member interface vxlan100
set interfaces bridge br100 vrf 'customer-1'

set interfaces vxlan vxlan100 mtu '1500'
set interfaces vxlan vxlan100 parameters nolearning
set interfaces vxlan vxlan100 port '4789'
set interfaces vxlan vxlan100 source-address '10.0.0.1'
set interfaces vxlan vxlan100 vni '100'

set vrf name customer-1 protocols bgp address-family ipv4-unicast redistribute connected
set vrf name customer-1 protocols bgp address-family l2vpn-evpn advertise ipv4 unicast
set vrf name customer-1 protocols bgp system-as '65000'
set vrf name customer-1 table '100'
set vrf name customer-1 vni '100'
```

PE2

```
set interfaces bridge br100 address '192.168.2.1/24'
set interfaces bridge br100 description 'customer-1'
set interfaces bridge br100 member interface eth2
set interfaces bridge br100 member interface vxlan100
set interfaces bridge br100 vrf 'customer-1'

set interfaces vxlan vxlan100 mtu '1500'
set interfaces vxlan vxlan100 parameters nolearning
set interfaces vxlan vxlan100 port '4789'
set interfaces vxlan vxlan100 source-address '10.0.0.2'
set interfaces vxlan vxlan100 vni '100'

set vrf name customer-1 protocols bgp address-family ipv4-unicast redistribute connected
set vrf name customer-1 protocols bgp address-family l2vpn-evpn advertise ipv4 unicast
set vrf name customer-1 protocols bgp system-as '65000'
set vrf name customer-1 table '100'
set vrf name customer-1 vni '100'
```

Validation and Troubleshooting

Layer-3 EVPN Asymmetric Model

1. Verify all OSPF sessions are in full state

```
vyos@PE1# run show ip ospf neighbor

Neighbor ID     Pri State       Up Time      Dead Time Address   Interface           RXmtL RqstL DBsmL
10.0.0.3          1 Full/-      1h53m10s     36.743s 10.1.3.3    eth0:10.1.3.1       0     0     0
10.0.0.4          1 Full/-      1h53m05s     38.230s 10.1.4.4    eth1:10.1.4.1       0     0     0
```

2. Verify that remote PE loopback is known via OSPF

```
vyos@PE1# run show ip route 10.0.0.2
Routing entry for 10.0.0.2/32
  Known via "ospf", distance 110, metric 20, best
  Last update 01:55:03 ago
  * 10.1.3.3, via eth0, weight 1
  * 10.1.4.4, via eth1, weight 1

[edit]
vyos@PE1#
```

3. Verify PE reachability

```
vyos@PE1# run ping 10.0.0.2 source-address 10.0.0.1 count 4
PING 10.0.0.2 (10.0.0.2) from 10.0.0.1 : 56(84) bytes of data.
64 bytes from 10.0.0.2: icmp_seq=1 ttl=62 time=2.45 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=62 time=2.26 ms
64 bytes from 10.0.0.2: icmp_seq=3 ttl=62 time=2.81 ms
64 bytes from 10.0.0.2: icmp_seq=4 ttl=62 time=2.82 ms

--- 10.0.0.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3005ms
rtt min/avg/max/mdev = 2.258/2.584/2.824/0.240 ms
[edit]
vyos@PE1#
```

4. Verify that BGP session are established

```
vyos@PE1# run show bgp summary

L2VPN EVPN Summary (VRF default):
BGP router identifier 10.0.0.1, local AS number 65000 vrf-id 0
BGP table version 0
RIB entries 7, using 672 bytes of memory
Peers 1, using 20 KiB of memory
Peer groups 1, using 64 bytes of memory

Neighbor   V   AS   MsgRcvd   MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd  PfxSnt Desc
10.0.0.3   4   65000    129       125        3    0    0 01:59:30            4       4 FRRouting/9.1.3

Total number of neighbors 1
[edit]
vyos@PE1#
```

5. Verify evpn routes advertised to RR

```
vyos@PE1# run show bgp l2vpn evpn neighbors 10.0.0.3 advertised-routes
BGP table version is 0, local router ID is 10.0.0.1
Default local pref 100, local AS 65000
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 10.0.0.1:2
 *> [3]:[0]:[32]:[10.0.0.1]
                                  100  32768 i
Route Distinguisher: 10.0.0.1:3
 *> [2]:[0]:[48]:[00:50:79:66:68:00]
                                  100  32768 i
 *> [2]:[0]:[48]:[00:50:79:66:68:00]:[32]:[192.168.1.10]
                                  100  32768 i
 *> [3]:[0]:[32]:[10.0.0.1]
                                  100  32768 i

Total number of prefixes 4
[edit]
vyos@PE1#
```

6. EVPN routes. Just type-2 MAC/IP and type-3 IMET routes are displayed, as expected for this model.

```
vyos@PE1# run show bgp l2vpn evpn route
BGP table version is 1, local router ID is 10.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN type-1 prefix: [1]:[EthTag]:[ESI]:[IPlen]:[VTEP-IP]:[Frag-id]
EVPN type-2 prefix: [2]:[EthTag]:[MAClen]:[MAC]:[IPlen]:[IP]
EVPN type-3 prefix: [3]:[EthTag]:[IPlen]:[OrigIP]
EVPN type-4 prefix: [4]:[ESI]:[IPlen]:[OrigIP]
EVPN type-5 prefix: [5]:[EthTag]:[IPlen]:[IP]

   Network          Next Hop            Metric LocPrf Weight Path
                    Extended Community
Route Distinguisher: 10.0.0.1:2
 *> [3]:[0]:[32]:[10.0.0.1]
                    10.0.0.1                           32768 i
                    ET:8 RT:65000:200
Route Distinguisher: 10.0.0.1:3
 *> [2]:[0]:[48]:[00:50:79:66:68:00]
                    10.0.0.1                           32768 i
                    ET:8 RT:65000:100
 *> [2]:[0]:[48]:[00:50:79:66:68:00]:[32]:[192.168.1.10]   <- Type-2 MAC/IP route for local connected host-2
                    10.0.0.1                           32768 i
                    ET:8 RT:65000:100
 *> [3]:[0]:[32]:[10.0.0.1]
                    10.0.0.1                           32768 i
                    ET:8 RT:65000:100
Route Distinguisher: 10.0.0.2:2
 *>i[2]:[0]:[48]:[00:50:79:66:68:01]
                    10.0.0.2                 0    100      0 i
                    RT:65000:200 ET:8
 *>i[2]:[0]:[48]:[00:50:79:66:68:01]:[32]:[192.168.2.10]   <- Type-2 MAC/IP route for local connected host-2
                    10.0.0.2                 0    100      0 i
                    RT:65000:200 ET:8
```

```
 *>i[3]:[0]:[32]:[10.0.0.2]
                    10.0.0.2                0    100      0 i
                    RT:65000:200 ET:8
Route Distinguisher: 10.0.0.2:3
 *>i[3]:[0]:[32]:[10.0.0.2]
                    10.0.0.2                0    100      0 i
                    RT:65000:100 ET:8

Displayed 8 prefixes (8 paths)
[edit]
vyos@PE1#
```

7. EVPN routes in detail.

```
vyos@PE1# run show bgp l2vpn evpn route detail
...

Route Distinguisher: 10.0.0.2:2
BGP routing table entry for 10.0.0.2:2:[2]:[0]:[48]:[00:50:79:66:68:01]
Paths: (1 available, best #1)
  Not advertised to any peer
  Route [2]:[0]:[48]:[00:50:79:66:68:01] VNI 200
  Local
    10.0.0.2 (metric 20) from 10.0.0.3 (10.0.0.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best (First path received)
      Extended Community: RT:65000:200 ET:8
      Originator: 10.0.0.2, Cluster list: 10.0.0.3
      Last update: Mon Apr 14 12:39:33 2025
BGP routing table entry for 10.0.0.2:2:[2]:[0]:[48]:[00:50:79:66:68:01]:[32]:[192.168.2.10]
Paths: (1 available, best #1)
  Not advertised to any peer
  Route [2]:[0]:[48]:[00:50:79:66:68:01]:[32]:[192.168.2.10] VNI 200
  Local
    10.0.0.2 (metric 20) from 10.0.0.3 (10.0.0.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best (First path received)
      Extended Community: RT:65000:200 ET:8
      Originator: 10.0.0.2, Cluster list: 10.0.0.3
      Last update: Mon Apr 14 12:39:33 2025
...

vyos@PE1#
```

8. Show routes customer-1 VRF. Both prefixes are known as directly connected

```
vyos@PE1# run show ip route vrf customer-1
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
       T - Table, v - VNC, V - VNC-Direct, A - Babel, F - PBR,
       f - OpenFabric,
       > - selected route, * - FIB route, q - queued, r - rejected, b - backup
       t - trapped, o - offload failure

VRF customer-1:
C>* 192.168.1.0/24 is directly connected, br100, 02:16:12
C>* 192.168.2.0/24 is directly connected, br200, 02:16:04
[edit]
vyos@PE1#
```

9. br200 maintain FDB entry and ARP entry for remote host-2.

```
vyos@PE1# run show bridge br200 fdb
Interface    Mac address        State      Flags
-----------  -----------------  ---------  -----------------
...

vxlan200     00:50:79:66:68:01             self,extern_learn


vyos@PE1# run show arp interface br200
Address       Interface    Link layer address    State
------------  -----------  --------------------  -------
192.168.2.10  br200        00:50:79:66:68:01     NOARP
[edit]
vyos@PE1#
```

10. Finally, letś try a ping test and capture traffic from host-1 to host-2:

```
host-1> ping 192.168.2.10

84 bytes from 192.168.2.10 icmp_seq=1 ttl=62 time=4.436 ms
84 bytes from 192.168.2.10 icmp_seq=2 ttl=62 time=3.991 ms
84 bytes from 192.168.2.10 icmp_seq=3 ttl=62 time=3.972 ms
84 bytes from 192.168.2.10 icmp_seq=4 ttl=62 time=3.910 ms
84 bytes from 192.168.2.10 icmp_seq=5 ttl=62 time=4.206 ms

host-1>
```

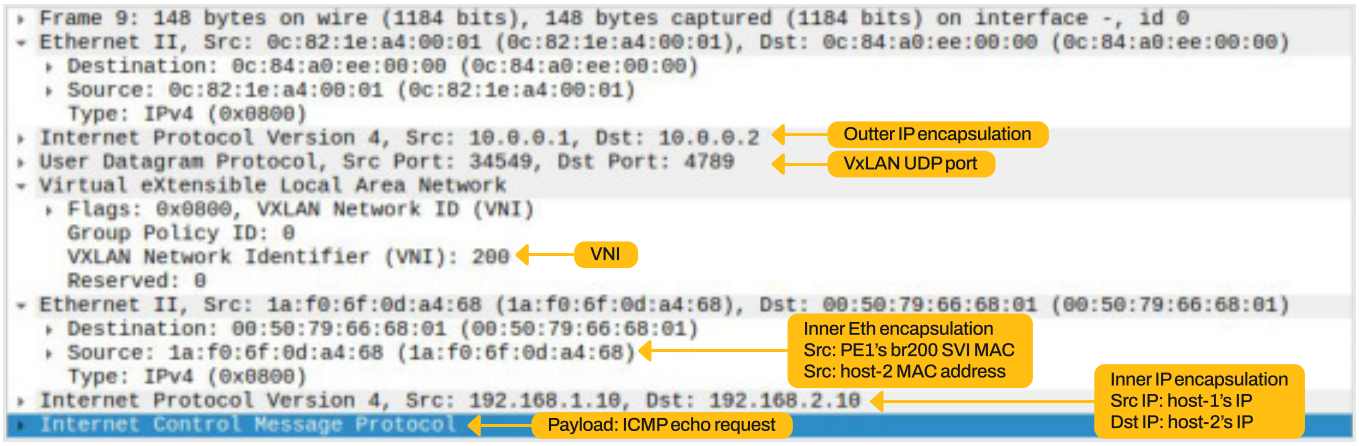

Figure-10: L3VPN asymmetric model deployment – packet capture

Layer-3 EVPN Symmetric Interface-less VRF-to-VRF model

1. Verify all OSPF sessions are in full state

```
vyos@PE1# run show ip ospf neighbor

Neighbor ID     Pri State     Up Time    Dead Time Address     Interface            RXmtL RqstL DBsmL
10.0.0.3          1 Full/-    1h05m08s     39.656s 10.1.3.3    eth0:10.1.3.1            0     0     0
10.0.0.4          1 Full/-    1h05m04s     30.437s 10.1.4.4    eth1:10.1.4.1            0     0     0

[edit]
vyos@PE1#
```

2. Verify that remote PE loopback is known via OSPF

```
vyos@PE1# run show ip route 10.0.0.2
Routing entry for 10.0.0.2/32
  Known via "ospf", distance 110, metric 20, best
  Last update 01:05:48 ago
  * 10.1.3.3, via eth0, weight 1
  * 10.1.4.4, via eth1, weight 1

[edit]
vyos@PE1#
```

3. Verify PE reachability

```
vyos@PE1# run ping 10.0.0.2 source-address 10.0.0.1 count 4
PING 10.0.0.2 (10.0.0.2) from 10.0.0.1 : 56(84) bytes of data.
64 bytes from 10.0.0.2: icmp_seq=1 ttl=62 time=2.74 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=62 time=2.59 ms
64 bytes from 10.0.0.2: icmp_seq=3 ttl=62 time=2.66 ms
64 bytes from 10.0.0.2: icmp_seq=4 ttl=62 time=2.45 ms

--- 10.0.0.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3006ms
rtt min/avg/max/mdev = 2.454/2.610/2.743/0.105 ms
[edit]
vyos@PE1#
```

4. Verify that BGP session are established

```
vyos@PE1# run show bgp summary
L2VPN EVPN Summary (VRF default):
BGP router identifier 10.0.0.1, local AS number 65000 vrf-id 0
BGP table version 0
RIB entries 3, using 288 bytes of memory
Peers 1, using 20 KiB of memory
Peer groups 1, using 64 bytes of memory

Neighbor     V      AS   MsgRcvd MsgSent TblVer  InQ OutQ  Up/Down State/PfxRcd PfxSnt Desc
10.0.0.3     4   65000        32      31      1    0    0 00:27:40            1      1 FRRouting/9.1.3

Total number of neighbors 1
[edit]
vyos@PE1#
```

5. Verify EVPN routes advertised to RR

```
vyos@PE1# run show bgp l2vpn evpn neighbors 10.0.0.3 advertised-routes

BGP table version is 0, local router ID is 10.0.0.1
Default local pref 100, local AS 65000
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 192.168.1.1:2
 *> [5]:[0]:[24]:[192.168.1.0]
                           0    100  32768 ?

Total number of prefixes 1
[edit]
vyos@PE1#
```

6. Let's see evpn routes. Just type-5 routes are displayed, as expected for this model.

```
vyos@PE1# run show bgp l2vpn evpn route
BGP table version is 1, local router ID is 10.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN type-1 prefix: [1]:[EthTag]:[ESI]:[IPlen]:[VTEP-IP]:[Frag-id]
EVPN type-2 prefix: [2]:[EthTag]:[MAClen]:[MAC]:[IPlen]:[IP]
EVPN type-3 prefix: [3]:[EthTag]:[IPlen]:[OrigIP]
EVPN type-4 prefix: [4]:[ESI]:[IPlen]:[OrigIP]
EVPN type-5 prefix: [5]:[EthTag]:[IPlen]:[IP]

   Network          Next Hop            Metric LocPrf Weight Path
                    Extended Community
Route Distinguisher: 192.168.1.1:2
 *> [5]:[0]:[24]:[192.168.1.0]
                    10.0.0.1                 0         32768 ?
                    ET:8 RT:65000:100 Rmac:b2:de:98:da:3b:ba
Route Distinguisher: 192.168.2.1:2
 *>i[5]:[0]:[24]:[192.168.2.0]
                    10.0.0.2                 0    100      0 ?
                    RT:65000:100 ET:8 Rmac:e2:2d:c2:ef:2e:aa

Displayed 2 prefixes (2 paths)
[edit]
vyos@PE1#
```

Type-5 IP-prefix route for local VRF prefix

Encapsulation Tunnel: 8-VxLAN
Auto-derived route-target: 65000:100
Router MAC: PE1 br100's SVI MAC reuired for VxLAV data encapsulation

Type-5 IP-prefix route for remore VRF prefix

Next-hop 10.0.0.2

Encapsulation Tunnel: 8-VxLAN
Auto-derived route-target: 65000:100
Router MAC: PE2 br100's SVI MAC reuired for VxLAV data encapsulation

7. Now, let's see EVPN routes in detail.

```
vyos@PE1# run show bgp l2vpn evpn route detail
Route Distinguisher: 192.168.1.1:2
BGP routing table entry for 192.168.1.1:2:[5]:[0]:[24]:[192.168.1.0]
Paths: (1 available, best #1)
  Advertised to non peer-group peers:
  10.0.0.3
  Route [5]:[0]:[24]:[192.168.1.0] VNI 100
  Local
    10.0.0.1 from 0.0.0.0 (10.0.0.1)
      Origin incomplete, metric 0, weight 32768, valid, sourced, local, best (First path received)
      Extended Community: ET:8 RT:65000:100 Rmac:b2:de:98:da:3b:ba
```

```
      Last update: Fri Apr 11 21:51:56 2025
Route Distinguisher: 192.168.2.1:2
BGP routing table entry for 192.168.2.1:2:[5]:[0]:[24]:[192.168.2.0]
Paths: (1 available, best #1)
  Not advertised to any peer
  Route [5]:[0]:[24]:[192.168.2.0] VNI 100
  Local
    10.0.0.2 (metric 20) from 10.0.0.3 (10.0.0.2)
      Origin incomplete, metric 0, localpref 100, valid, internal, best (First path received)
      Extended Community: RT:65000:100 ET:8 Rmac:e2:2d:c2:ef:2e:aa
      Originator: 10.0.0.2, Cluster list: 10.0.0.3
      Last update: Fri Apr 11 21:52:16 2025

Displayed 2 prefixes (2 paths)
[edit]
vyos@PE1#
```

8. Show routes customer-1 VRF

```
vyos@PE1# run show ip route vrf customer-1
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
       T - Table, v - VNC, V - VNC-Direct, A - Babel, F - PBR,
       f - OpenFabric,
       > - selected route, * - FIB route, q - queued, r - rejected, b - backup
       t - trapped, o - offload failure

VRF customer-1:
C>* 192.168.1.0/24 is directly connected, br100, 00:02:09
B>* 192.168.2.0/24 [200/0] via 10.0.0.2, br100 onlink, weight 1, 00:01:43
[edit]
vyos@PE1#
```

9. Finally, letś try a ping test and capture traffic from host-1 to host-2:

```
host-1> ping 192.168.2.10
84 bytes from 192.168.2.10 icmp_seq=1 ttl=62 time=4.436 ms
84 bytes from 192.168.2.10 icmp_seq=2 ttl=62 time=3.991 ms
84 bytes from 192.168.2.10 icmp_seq=3 ttl=62 time=3.972 ms
84 bytes from 192.168.2.10 icmp_seq=4 ttl=62 time=3.910 ms
84 bytes from 192.168.2.10 icmp_seq=5 ttl=62 time=4.206 ms
host-1>
```

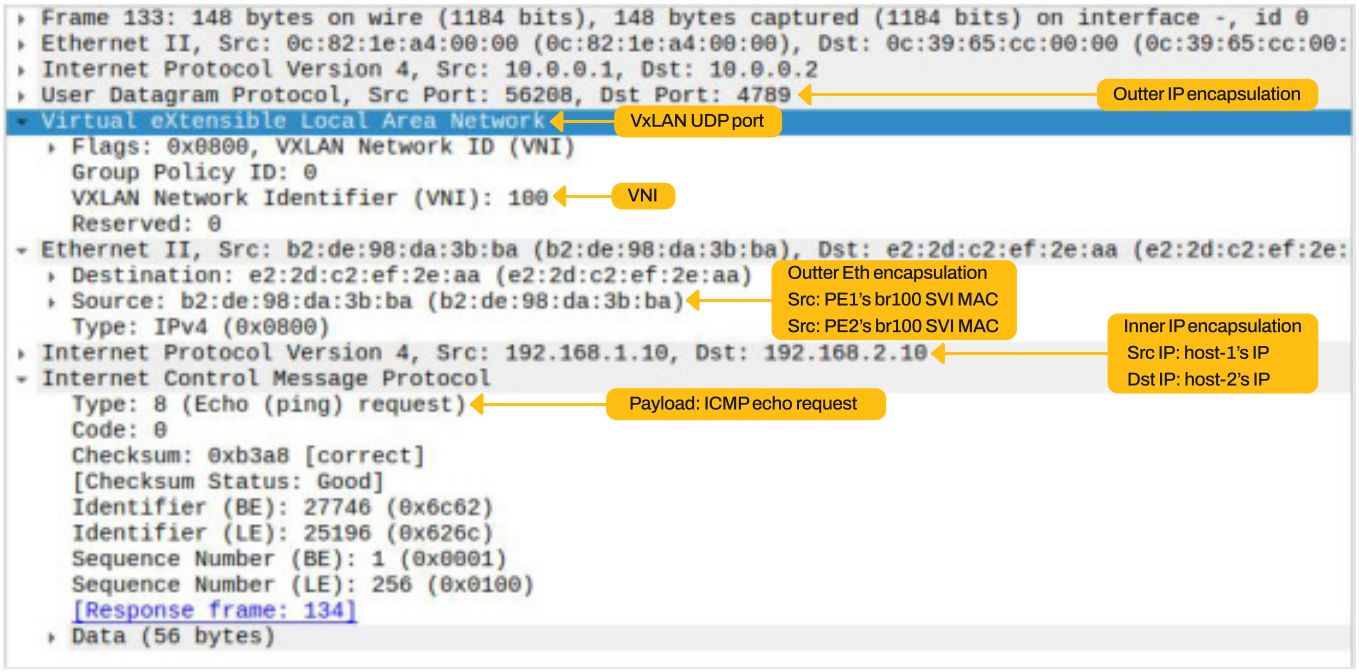

Figure-11: L3VPN symmetric model deployment – packet capture