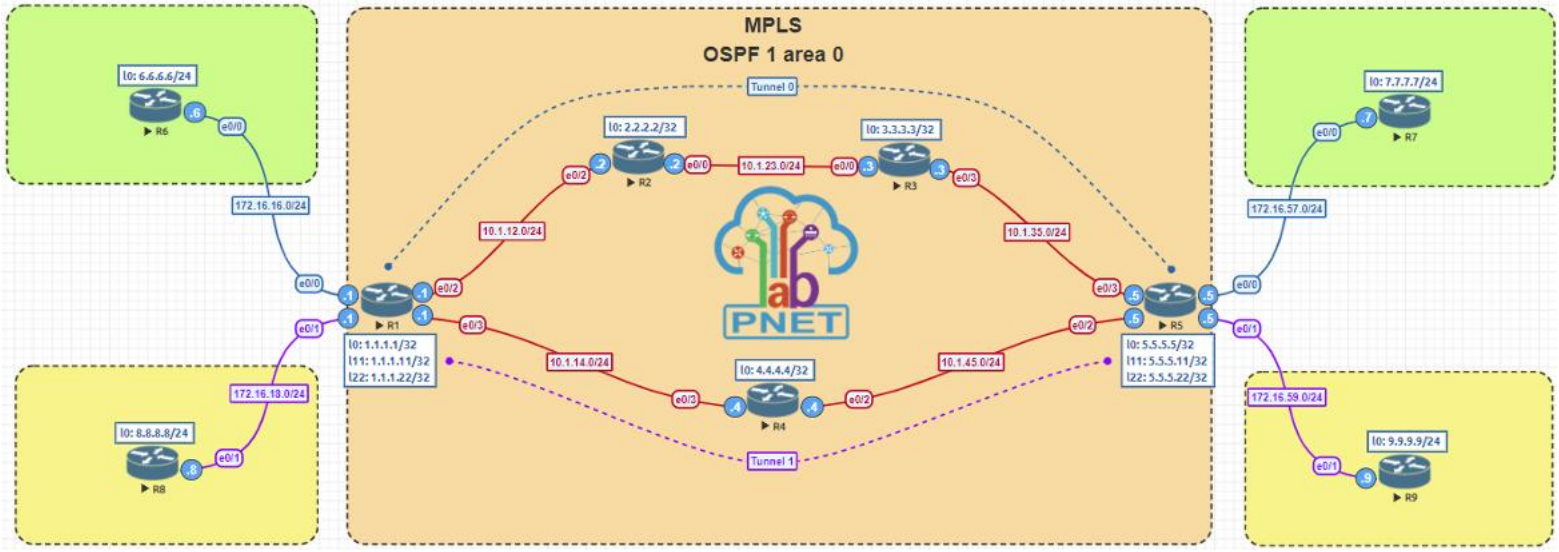
Per-VRF MPLS-TE Tunnels (MPLS Traffic Engineering)

Lab Topology:

Please use the following topology to complete this lab exercise:



Lab Objective:

The objective of this lab exercise is for you to learn and understand how can you configure Per-VRF MPLS TE Tunnels.

Recommendation:

Recommend do “MPLS basic LDP” lab to understand how LDP work in core domain before do this lab.
Reference: https://user.pnetlab.com/store/labs/detail?id=15960951497382

Task list:

Task 1: Assign the IP address to each interface on routers as shown in the topo.

Task 2: Configure OSPF to advertise all interfaces in the core domain and enable LDP on all internal interfaces.

Task 3: Create some VRFs and assign the interfaces that is facing the customer to the VRF. Then assign its IP address again.

Task 4: Configure IBGP on PE router (R1 and R5) to exchange VPNv4 routes.

Task 5: Create explicit path on PE router to determine the path for the traffic (Tunnel0: R1-R2-R3-R5 and Tunnel1: R1-R4-R5).

Task 6: Configure MPLS TE and setup tunnel to be used for TE.

Task 7: Setup routing between CEs in the same VPN domain (In this lab, use static routing for customer 1 and RIP for customer 2). Remember redistribute them into BGP and vice versa.

SOLUTION:

Task 1: Assign the IP address to each interface. Assign IP address to the interfaces of PE router which connect to CE isn't necessary because they will be remove after task 3.

<table>
<tr><td rowspan="3">Core</td><td>On R1
hostname R1
!
interface Loopback0
 ip address 1.1.1.1
255.255.255.255
interface Loopback11
 ip address 1.1.1.11
255.255.255.255
interface Loopback22
 ip address 1.1.1.22
255.255.255.255
interface Ethernet0/2
 ip address 10.1.12.1
255.255.255.0
interface Ethernet0/3
 ip address 10.1.14.1
255.255.255.0
!</td><td>On R5
hostname R5
!
interface Loopback0
 ip address 5.5.5.5
255.255.255.255
interface Loopback11
 ip address 5.5.5.11
255.255.255.255
interface Loopback22
 ip address 5.5.5.22
255.255.255.255
interface Ethernet0/2
 ip address 10.1.45.5
255.255.255.0
interface Ethernet0/3
 ip address 10.1.35.5
255.255.255.0
!</td></tr>
<tr><td>On R2
hostname R2
!
interface Loopback0
 ip address 2.2.2.2
255.255.255.255
interface Ethernet0/0
 ip address 10.1.23.2
255.255.255.0
interface Ethernet0/2
 ip address 10.1.12.2
255.255.255.0
!</td><td>On R3
hostname R3
!
interface Loopback0
 ip address 3.3.3.3
255.255.255.255
interface Ethernet0/0
 ip address 10.1.23.3
255.255.255.0
interface Ethernet0/3
 ip address 10.1.35.3
255.255.255.0
!</td></tr>
<tr><td colspan="2">On R4
hostname R4
!
interface Loopback0
 ip address 4.4.4.4 255.255.255.255
interface Ethernet0/2
 ip address 10.1.45.4 255.255.255.0
interface Ethernet0/3</td></tr>
</table>

	ip address 10.1.14.4 255.255.255.0 !	
Customer 1	**On R6** hostname R6 ! interface Loopback0 ip address 6.6.6.6 255.255.255.255 interface Ethernet0/0 ip address 172.16.16.6 255.255.255.0 !	**On R7** hostname R7 ! interface Loopback0 ip address 7.7.7.7 255.255.255.255 interface Ethernet0/0 ip address 172.16.57.7 255.255.255.0 !
Customer 2	**On R8** hostname R8 ! interface Loopback0 ip address 8.8.8.8 255.255.255.255 interface Ethernet0/1 ip address 172.16.18.8 255.255.255.0 !	**On R9** hostname R9 ! interface Loopback0 ip address 9.9.9.9 255.255.255.255 interface Ethernet0/1 ip address 172.16.59.9 255.255.255.0 !

Task 2: Configure OSPF in the core domain and enable LDP on all internal interfaces. Remember set ldp RID for R1 and R5 because we just advertise loopback 0 (isnt highest IP address) in OSPF.

On R1

```
mpls label range 100 199
!
interface Ethernet0/2
 mpls ip
interface Ethernet0/3
 mpls ip
!
router ospf 1
 network 1.1.1.1 0.0.0.0 area 0
 network 10.1.12.0 0.0.0.255 area 0
 network 10.1.14.0 0.0.0.255 area 0
!
mpls ldp router-id Loopback0 force
```

On R2

```
mpls label range 200 299
!
interface Ethernet0/0
```

```
 mpls ip
interface Ethernet0/2
 mpls ip
!
router ospf 1
 network 2.2.2.2 0.0.0.0 area 0
 network 10.1.12.0 0.0.0.255 area 0
 network 10.1.23.0 0.0.0.255 area 0
```

On R3

```
mpls label range 300 399
!
interface Ethernet0/0
 mpls ip
interface Ethernet0/3
 mpls ip
!
router ospf 1
 network 3.3.3.3 0.0.0.0 area 0
 network 10.1.23.0 0.0.0.255 area 0
 network 10.1.35.0 0.0.0.255 area 0
```

On R4

```
mpls label range 400 499
!
interface Ethernet0/2
 mpls ip
interface Ethernet0/3
 mpls ip
!
router ospf 1
 network 4.4.4.4 0.0.0.0 area 0
 network 10.1.14.0 0.0.0.255 area 0
 network 10.1.45.0 0.0.0.255 area 0
```

On R5

```
mpls label range 500 599
!
interface Ethernet0/2
 mpls ip
interface Ethernet0/3
 mpls ip
!
router ospf 1
 network 5.5.5.5 0.0.0.0 area 0
```

```
 network 10.1.35.0 0.0.0.255 area 0
 network 10.1.45.0 0.0.0.255 area 0
!
mpls ldp router-id Loopback0 force
```

Verification

On R1

```
R1#show ip route
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 3 subnets
C        1.1.1.1 is directly connected, Loopback0
C        1.1.1.11 is directly connected, Loopback11
C        1.1.1.22 is directly connected, Loopback22
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/11] via 10.1.12.2, 00:00:53, Ethernet0/2
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/21] via 10.1.12.2, 00:00:43, Ethernet0/2
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/11] via 10.1.14.4, 00:00:43, Ethernet0/3
      5.0.0.0/32 is subnetted, 1 subnets
O        5.5.5.5 [110/21] via 10.1.14.4, 00:00:33, Ethernet0/3
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C        10.1.12.0/24 is directly connected, Ethernet0/2
L        10.1.12.1/32 is directly connected, Ethernet0/2
C        10.1.14.0/24 is directly connected, Ethernet0/3
L        10.1.14.1/32 is directly connected, Ethernet0/3
O        10.1.23.0/24 [110/20] via 10.1.12.2, 00:00:43, Ethernet0/2
O        10.1.35.0/24 [110/30] via 10.1.14.4, 00:00:33, Ethernet0/3
                      [110/30] via 10.1.12.2, 00:00:33, Ethernet0/2
O        10.1.45.0/24 [110/20] via 10.1.14.4, 00:00:33, Ethernet0/3
R1#show mpls ldp discovery
 Local LDP Identifier:
    1.1.1.1:0
    Discovery Sources:
    Interfaces:
        Ethernet0/2 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0
        Ethernet0/3 (ldp): xmit/recv
            LDP Id: 4.4.4.4:0
R1#show mpls ldp neighbor
    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 1.1.1.1:0
        TCP connection: 2.2.2.2.15336 - 1.1.1.1.646
        State: Oper; Msgs sent/rcvd: 16/14; Downstream
        Up time: 00:01:02
```

```
        LDP discovery sources:
          Ethernet0/2, Src IP addr: 10.1.12.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.12.2       2.2.2.2
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 1.1.1.1:0
        TCP connection: 4.4.4.4.15875 - 1.1.1.1.646
        State: Oper; Msgs sent/rcvd: 16/14; Downstream
        Up time: 00:00:58
        LDP discovery sources:
          Ethernet0/3, Src IP addr: 10.1.14.4
        Addresses bound to peer LDP Ident:
          10.1.45.4       10.1.14.4       4.4.4.4
R1#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing    Next
Hop
Label      Label      or Tunnel Id     Switched      interface
100        Pop Label  2.2.2.2/32       0             Et0/2
10.1.12.2
101        Pop Label  10.1.23.0/24     0             Et0/2
10.1.12.2
102        202        3.3.3.3/32       0             Et0/2
10.1.12.2
103        203        10.1.35.0/24     0             Et0/2
10.1.12.2
           403        10.1.35.0/24     0             Et0/3
10.1.14.4
104        Pop Label  4.4.4.4/32       0             Et0/3
10.1.14.4
105        Pop Label  10.1.45.0/24     0             Et0/3
10.1.14.4
106        406        5.5.5.5/32       0             Et0/3
10.1.14.4
```

Note: Do the same thing with R2, R3, R4, R5.

Task 3: Create some VRFs and assign the interfaces that is facing the customer to the VRF. Then assign its IP address again.

On R1	**On R5**
ip vrf CusA rd 100:1 route-target export 100:10 route-target import 100:10 ! ip vrf CusB rd 100:2	ip vrf CusA rd 100:1 route-target export 100:10 route-target import 100:10 ! ip vrf CusB rd 100:2

route-target export 100:20 route-target import 100:20 ! interface Ethernet0/0 ip vrf forwarding CusA ip address 172.16.16.1 255.255.255.0 interface Ethernet0/1 ip vrf forwarding CusB ip address 172.16.18.1 255.255.255.0	route-target export 100:20 route-target import 100:20 ! interface Ethernet0/0 ip vrf forwarding CusA ip address 172.16.57.5 255.255.255.0 interface Ethernet0/1 ip vrf forwarding CusB ip address 172.16.59.5 255.255.255.0

Verification

On R1

```
R1#show ip int brief
Interface              IP-Address      OK? Method Status                Protocol
Ethernet0/0            172.16.16.1     YES manual up                    up
Ethernet0/1            172.16.18.1     YES manual up                    up
Ethernet0/2            10.1.12.1       YES NVRAM  up                    up
Ethernet0/3            10.1.14.1       YES NVRAM  up                    up
Loopback0              1.1.1.1         YES NVRAM  up                    up
Loopback11             1.1.1.11        YES NVRAM  up                    up
Loopback22             1.1.1.22        YES NVRAM  up                    up
R1#show vrf CusA
  Name                 Default RD            Protocols   Interfaces
  CusA                 100:1                 ipv4        Et0/0
R1#show vrf CusB
  Name                 Default RD            Protocols   Interfaces
  CusB                 100:2                 ipv4        Et0/1
```

Task 4: Configure IBGP on PE router (R1 and R5) to exchange VPNv4 routes.

On R1	**On R5**
router bgp 100 neighbor 5.5.5.5 remote-as 100 neighbor 5.5.5.5 update-source Loopback0 ! address-family vpnv4 neighbor 5.5.5.5 activate	router bgp 100 neighbor 1.1.1.1 remote-as 100 neighbor 1.1.1.1 update-source Loopback0 ! address-family vpnv4 neighbor 1.1.1.1 activate

Task 5: Setup routing between CEs in the same VPN domain (In this lab, use static routing for customer 1 and RIP for customer 2). Remember redistribute them into BGP and vice versa.

On R1 (RIP: network 172.16.18.0/24)	**On R5** (RIP: network 172.16.59/24)
router rip	router rip

version 2 no auto-summary address-family ipv4 vrf CusB redistribute bgp 100 network 172.16.0.0 ! router bgp 100 address-family ipv4 vrf CusA redistribute connected redistribute static address-family ipv4 vrf CusB redistribute rip ! ip route vrf CusA 6.6.6.0 255.255.255.0 172.16.16.6	version 2 no auto-summary address-family ipv4 vrf CusB redistribute bgp 100 network 172.16.0.0 ! router bgp 100 address-family ipv4 vrf CusA redistribute connected redistribute static address-family ipv4 vrf CusB redistribute rip ! ip route vrf CusA 7.7.7.0 255.255.255.0 172.16.57.7
On R6 ip route 0.0.0.0 0.0.0.0 172.16.16.1	**On R7** ip route 0.0.0.0 0.0.0.0 172.16.57.5
On R8 (RIP: network 172.16.18.0/24) router rip version 2 network 8.0.0.0 network 172.16.0.0 no auto-summary ip route 0.0.0.0 0.0.0.0 172.16.18.1	**On R9** (RIP: network 172.16.59/24) router rip version 2 network 9.0.0.0 network 172.16.0.0 no auto-summary ip route 0.0.0.0 0.0.0.0 172.16.59.5

Verification

On R1

```
R1#show ip route vrf CusA
      6.0.0.0/24 is subnetted, 1 subnets
S        6.6.6.0 [1/0] via 172.16.16.6
      7.0.0.0/24 is subnetted, 1 subnets
B        7.7.7.0 [200/0] via 5.5.5.5, 00:00:07
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.16.0/24 is directly connected, Ethernet0/0
L        172.16.16.1/32 is directly connected, Ethernet0/0
B        172.16.57.0/24 [200/0] via 5.5.5.5, 00:00:07
R1#show ip route vrf CusB
      8.0.0.0/24 is subnetted, 1 subnets
R        8.8.8.0 [120/1] via 172.16.18.8, 00:00:17, Ethernet0/1
      9.0.0.0/24 is subnetted, 1 subnets
B        9.9.9.0 [200/1] via 5.5.5.5, 00:00:43
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.18.0/24 is directly connected, Ethernet0/1
```

```
L        172.16.18.1/32 is directly connected, Ethernet0/1
B        172.16.59.0/24 [200/0] via 5.5.5.5, 00:00:43
R1#show ip bgp vpnv4 all
      Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf CusA)
*>  6.6.6.0/24       172.16.16.6              0         32768 ?
*>i 7.7.7.0/24       5.5.5.5                  0    100      0 ?
*>  172.16.16.0/24   0.0.0.0                  0         32768 ?
*>i 172.16.57.0/24   5.5.5.5                  0    100      0 ?
Route Distinguisher: 100:2 (default for vrf CusB)
*>  8.8.8.0/24       172.16.18.8              1         32768 ?
*>i 9.9.9.0/24       5.5.5.5                  1    100      0 ?
*>  172.16.18.0/24   0.0.0.0                  0         32768 ?
*>i 172.16.59.0/24   5.5.5.5                  0    100      0 ?
R1#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next
Hop
Label      Label      or Tunnel Id     Switched      interface
100        Pop Label  2.2.2.2/32       0             Et0/2
10.1.12.2
101        Pop Label  10.1.23.0/24     0             Et0/2
10.1.12.2
102        202        3.3.3.3/32       0             Et0/2
10.1.12.2
103        203        10.1.35.0/24     0             Et0/2
10.1.12.2
           403        10.1.35.0/24     0             Et0/3
10.1.14.4
104        Pop Label  4.4.4.4/32       0             Et0/3
10.1.14.4
105        Pop Label  10.1.45.0/24     0             Et0/3
10.1.14.4
106        406        5.5.5.5/32       0             Et0/3
10.1.14.4
107        No Label   172.16.16.0/24[V]   \
                                       0             aggregate/CusA
108        No Label   6.6.6.0/24[V]    0             Et0/0
172.16.16.6
109        No Label   172.16.18.0/24[V]   \
                                       0             aggregate/CusB
110        No Label   8.8.8.0/24[V]    0             Et0/1
172.16.18.8
```

On R6

```
R6#ping 7.7.7.7 source l0
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
Packet sent with a source address of 6.6.6.6
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
R6#ping 9.9.9.9 source l0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 9.9.9.9, timeout is 2 seconds:
Packet sent with a source address of 6.6.6.6
U.U.U
Success rate is 0 percent (0/5)
R6#trace 7.7.7.7 source l0
Type escape sequence to abort.
Tracing the route to 7.7.7.7
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.16.1 1 msec 1 msec 1 msec
  2 10.1.14.4 [MPLS: Labels 406/508 Exp 0] 5 msec 5 msec 5 msec
  3 172.16.57.5 [MPLS: Label 508 Exp 0] 3 msec 4 msec 4 msec
  4 172.16.57.7 5 msec 6 msec *
```

On R8

```
R8#ping 9.9.9.9 source l0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 9.9.9.9, timeout is 2 seconds:
Packet sent with a source address of 8.8.8.8
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Task 6: Create explicit path on PE router to determine the path for the traffic (Tunnel0: R1-R2-R3-R5 and Tunnel1: R1-R4-R5).

On R1	On R5
ip explicit-path name Tunnel0 enable next-address 2.2.2.2 next-address 3.3.3.3 next-address 5.5.5.5 ! ip explicit-path name Tunnel1 enable next-address 4.4.4.4 next-address 5.5.5.5	ip explicit-path name Tunnel0 enable next-address 3.3.3.3 next-address 2.2.2.2 next-address 1.1.1.1 ! ip explicit-path name Tunnel1 enable next-address 4.4.4.4 next-address 1.1.1.1

Task 7: Configure MPLS TE and setup tunnel to be used for TE.

1. Enable MPLS-TE on global mode and on each interfaces belong to core domain and configure MPLS-TE on OSPF.

On R1	On R5

mpls traffic-eng tunnels ! interface Ethernet0/2 mpls traffic-eng tunnels interface Ethernet0/3 mpls traffic-eng tunnels ! router ospf 1 mpls traffic-eng router-id Loopback0 mpls traffic-eng area 0	mpls traffic-eng tunnels ! interface Ethernet0/2 mpls traffic-eng tunnels interface Ethernet0/3 mpls traffic-eng tunnels ! router ospf 1 mpls traffic-eng router-id Loopback0 mpls traffic-eng area 0
On R2 mpls traffic-eng tunnels ! interface Ethernet0/0 mpls traffic-eng tunnels interface Ethernet0/2 mpls traffic-eng tunnels ! router ospf 1 mpls traffic-eng router-id Loopback0 mpls traffic-eng area 0	**On R3** mpls traffic-eng tunnels ! interface Ethernet0/0 mpls traffic-eng tunnels interface Ethernet0/3 mpls traffic-eng tunnels ! router ospf 1 mpls traffic-eng router-id Loopback0 mpls traffic-eng area 0
On R4 mpls traffic-eng tunnels ! interface Ethernet0/2 mpls traffic-eng tunnels interface Ethernet0/3 mpls traffic-eng tunnels ! router ospf 1 mpls traffic-eng router-id Loopback0 mpls traffic-eng area 0	

2. Setup tunnel to be used for TE. Remember enter bgp next-hop loopback on vrf mode.

On R1 mpls traffic-eng reoptimize timers frequency 10 ! ip vrf CusA bgp next-hop Loopback11 ! ip vrf CusB bgp next-hop Loopback22	**On R5** mpls traffic-eng reoptimize timers frequency 10 ! ip vrf CusA bgp next-hop Loopback11 ! ip vrf CusB bgp next-hop Loopback22

! interface Tunnel0 ip unnumbered Loopback0 tunnel mode mpls traffic-eng tunnel destination 5.5.5.5 tunnel mpls traffic-eng autoroute announce tunnel mpls traffic-eng path-option 1 explicit name Tunnel0 tunnel mpls traffic-eng path-option 2 dynamic ! interface Tunnel1 ip unnumbered Loopback0 tunnel mode mpls traffic-eng tunnel destination 5.5.5.5 tunnel mpls traffic-eng autoroute announce tunnel mpls traffic-eng path-option 1 explicit name Tunnel1 ! ip route 5.5.5.11 255.255.255.255 Tunnel0 ip route 5.5.5.22 255.255.255.255 Tunnel1	! interface Tunnel0 ip unnumbered Loopback0 tunnel mode mpls traffic-eng tunnel destination 1.1.1.1 tunnel mpls traffic-eng autoroute announce tunnel mpls traffic-eng path-option 1 explicit name Tunnel0 tunnel mpls traffic-eng path-option 2 dynamic ! interface Tunnel1 ip unnumbered Loopback0 tunnel mode mpls traffic-eng tunnel destination 1.1.1.1 tunnel mpls traffic-eng autoroute announce tunnel mpls traffic-eng path-option 1 explicit name Tunnel1 ! ip route 1.1.1.11 255.255.255.255 Tunnel0 ip route 1.1.1.22 255.255.255.255 Tunnel1

Verification

On R6

```
R6#trace 7.7.7.7 source l0
Type escape sequence to abort.
Tracing the route to 7.7.7.7
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.16.1 2 msec 0 msec 1 msec
  2 10.1.12.2 [MPLS: Labels 207/508 Exp 0] 1 msec 4 msec 5 msec
  3 10.1.23.3 [MPLS: Labels 308/508 Exp 0] 3 msec 3 msec 2 msec
  4 172.16.57.5 [MPLS: Label 508 Exp 0] 1 msec 4 msec 3 msec
  5 172.16.57.7 4 msec 2 msec *
```

On R1

```
R1#show ip route vrf CusA
      6.0.0.0/24 is subnetted, 1 subnets
S        6.6.6.0 [1/0] via 172.16.16.6
      7.0.0.0/24 is subnetted, 1 subnets
B        7.7.7.0 [200/0] via 5.5.5.11, 00:29:05
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.16.0/24 is directly connected, Ethernet0/0
```

```
L        172.16.16.1/32 is directly connected, Ethernet0/0
B        172.16.57.0/24 [200/0] via 5.5.5.11, 00:29:05
R1#show ip route vrf CusB
      8.0.0.0/24 is subnetted, 1 subnets
R        8.8.8.0 [120/1] via 172.16.18.8, 00:00:21, Ethernet0/1
      9.0.0.0/24 is subnetted, 1 subnets
B        9.9.9.0 [200/1] via 5.5.5.22, 00:02:14
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.18.0/24 is directly connected, Ethernet0/1
L        172.16.18.1/32 is directly connected, Ethernet0/1
B        172.16.59.0/24 [200/0] via 5.5.5.22, 00:02:14
```