

MPLS L3VPN Inter-AS Option B - VPNv4 EBGP Exchange

« MPLS L3VPN Inter-AS Option A - Back-to-Back VRF Exchange (</workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-a-back-to-back-vrf-exchange-Mjg2Nw%3D%3D>) | MPLS L3VPN Inter-AS Option C - Multihop VPNv4 EBGP Exchange (</workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-c-multihop-vpnv4-ebgp-exchange-Mjg2OQ%3D%3D>) »

Last updated: April 22, 2016

Note:

Initial Configuration & Diagrams: [Load the initial configuration files for the section named Inter AS L3VPN, 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 Inter AS L3VPN Diagram in order to complete this task.](#)

Task

- Disable the VLAN 40 subinterface of both R1 and XR1.
- Configure IGP routing and LDP in the first AS, which consists of R1, R2, and R3 as follows:
 - Enable OSPF Area 0 on the links between R1 & R3 and R2 & R3.
 - Enable OSPF Area 0 on the Loopback0 interfaces of R1, R2, & R3 as passive interfaces.
 - Enable LDP on the links between R1 & R3 and R2 & R3.
- Configure IGP routing and LDP in the second AS, which consists of R4, XR1, and XR2 as follows:
 - Use the following IS-IS NET addressing:
 - R4 – 49.0001.0000.0000.0004.00
 - XR1 – 49.0001.0000.0000.0019.00
 - XR2 – 49.0001.0000.0000.0020.00
 - Enable IS-IS Level 2 on the links between XR1 & R4 and R4 & XR2.
 - Advertise the Loopback0 interfaces of R4, XR1, and XR2 into IS-IS Level 2 as passive interfaces.
 - Enable LDP on the links between XR1 & R4 and R4 & XR2.
- Configure the following VRFs on PE routers R2 and XR2 follows:
 - VRF VPN_A:
 - Route Distinguisher: 101:201
 - Route Target Import: 101:201
 - Route Target Export: 101:201
 - Assign this VRF on the links in the 30.0.0.0 network on R2 and XR2.
 - VRF VPN_B:
 - Route Distinguisher: 102:202
 - Route Target Import: 102:202
 - Route Target Export: 102:202
 - Assign this VRF on the links in the 40.0.0.0 network on R2 and XR2.

- Configure RIPv2 routing for VRF VPN_A as follows:
 - Enable RIP between R9 & R2.
 - Enable RIP between R10 & XR2.
 - Advertise the Loopback0 networks of R9 & R10 into RIP.
- Configure EIGRP routing for VRF VPN_B as follows:
 - Use EIGRP Autonomous System 1.
 - Enable EIGRP between R7 & R2.
 - Enable EIGRP between R8 & XR2.
 - Advertise the Loopback0 networks of R7 & R8 into RIP.
- Configure VPNv4 BGP peerings as follows:
 - R1 and R2 should peer in AS 100 using each other's Loopback0 interfaces.
 - XR1 and XR2 should peer in AS 200 using each other's Loopback0 interfaces.
 - R1 and XR1 should peer EBGP.
- Redistribute between VPNv4 BGP and the VRF aware IGP processes on R2 and XR2.
- Once complete the following reachability should be achieved:
 - Customer routers R9 and R10 should have full IP reachability to each other's networks.
 - Customer routers R7 and R8 should have full IP reachability to each other's networks.
 - Traceroutes between these networks should indicate that separate Label Switch Paths are used within AS 100 and AS 200.

Configuration [Click to collapse](#)

```

R1:
interface Loopback0
  ip ospf 1 area 0
!
interface GigabitEthernet1.13
  ip ospf 1 area 0
  mpls ip
!
interface GigabitEthernet1.30
  mpls bgp forwarding
!
router bgp 100
  no bgp default ipv4-unicast
  no bgp default route-target filter
  neighbor 2.2.2.2 remote-as 100
  neighbor 2.2.2.2 update-source Loopback0
  neighbor 30.1.19.19 remote-as 200
!
address-family vpnv4
  neighbor 2.2.2.2 activate
  neighbor 2.2.2.2 send-community extended
  neighbor 2.2.2.2 next-hop-self
  neighbor 30.1.19.19 activate
  neighbor 30.1.19.19 send-community extended
exit-address-family
!
mpls ldp router-id Loopback0

R2:
vrf definition VPN_A
  rd 101:201
!
address-family ipv4
  route-target export 101:201
  route-target import 101:201
exit-address-family
!
vrf definition VPN_B
  rd 102:202
!
address-family ipv4
  route-target export 102:202
  route-target import 102:202
exit-address-family
!
interface Loopback0
  ip ospf 1 area 0
!
interface GigabitEthernet1.23
  ip ospf 1 area 0
  mpls ip
!
interface GigabitEthernet1.27

```

```

vrf forwarding VPN_B

ip address 40.2.7.2 255.255.255.0

!

interface GigabitEthernet1.29

vrf forwarding VPN_A

ip address 30.2.9.2 255.255.255.0

!

router eigrp 65535

!

address-family ipv4 vrf VPN_B

redistribute bgp 100

network 40.0.0.0

autonomous-system 1

exit-address-family

!

router ospf 1

passive-interface Loopback0

!

router rip

!

address-family ipv4 vrf VPN_A

redistribute bgp 100 metric 1

network 30.0.0.0

no auto-summary

version 2

exit-address-family

!

router bgp 100

no bgp default ipv4-unicast

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

neighbor 1.1.1.1 send-community extended

exit-address-family

!

address-family ipv4 vrf VPN_A

redistribute rip

exit-address-family

!

address-family ipv4 vrf VPN_B

redistribute eigrp 1

exit-address-family

!

mpls ldp router-id Loopback0

R3:

!

interface Loopback0

ip ospf 1 area 0

!

interface GigabitEthernet1.13

ip ospf 1 area 0

```

```

mpls ip
!
interface GigabitEthernet1.23
 ip ospf 1 area 0
 mpls ip
!
router ospf 1
 passive-interface Loopback0
!
mpls ldp router-id Loopback0

```

```

R4:
interface GigabitEthernet1.419
 ip router isis
 mpls ip
!
interface GigabitEthernet1.420
 ip router isis
 mpls ip
!
router isis
 net 49.0001.0000.0000.0004.00
 is-type level-2-only
 passive-interface Loopback0
!
mpls ldp router-id Loopback0

```

```

R7:
router eigrp 1
 network 0.0.0.0
 no auto-summary

```

```

R8:
router eigrp 1
 network 0.0.0.0
 no auto-summary

```

```

R9:
router rip
 version 2
 network 9.0.0.0
 network 30.0.0.0
 no auto-summary

```

```

R10:
router rip
 version 2
 network 10.0.0.0
 network 30.0.0.0
 no auto-summary

```

```

XR1:
route-policy PASS
 pass

```

```

end-policy
!
router static
  address-family ipv4 unicast
    30.1.19.1/32 GigabitEthernet0/0/0/0.30
  !
  !
router isis 1
  is-type level-2-only
  net 49.0001.0000.0000.0019.00
  interface Loopback0
    passive
    address-family ipv4 unicast
  !
  !
  interface GigabitEthernet0/0/0/0.419
    address-family ipv4 unicast
  !
  !
  !
router bgp 200
  address-family vpnv4 unicast
    retain route-target all
  !
  neighbor 20.20.20.20
    remote-as 200
    update-source Loopback0
    address-family vpnv4 unicast
      next-hop-self
  !
  !
  neighbor 30.1.19.1
    remote-as 100
    address-family vpnv4 unicast
      route-policy PASS in
      route-policy PASS out
  !
  !
  !
mpls ldp
  router-id 19.19.19.19
  interface GigabitEthernet0/0/0/0.419
  !
  !
XR2:
vrf VPN_A
  address-family ipv4 unicast
    import route-target
      101:201
  !
  export route-target
    101:201
  !

```

```

!
!
vrf VPN_B
  address-family ipv4 unicast
    import route-target
      102:202
    !
  export route-target
    102:202
  !
!
!
interface GigabitEthernet0/0/0/0.820
  vrf VPN_B
  no ipv4 address
  ipv4 address 40.8.20.20 255.255.255.0
!
interface GigabitEthernet0/0/0/0.1020
  vrf VPN_A
  no ipv4 address
  ipv4 address 30.10.20.20 255.255.255.0
!
route-policy BGP_TO_RIP
  set rip-metric 1
end-policy
!
router isis 1
  is-type level-2-only
  net 49.0001.0000.0000.0020.00
  interface Loopback0
    passive
    address-family ipv4 unicast
  !
!
interface GigabitEthernet0/0/0/0.420
  address-family ipv4 unicast
!
!
!
router bgp 200
  address-family vpnv4 unicast
  !
  neighbor 19.19.19.19
  remote-as 200
  update-source Loopback0
  address-family vpnv4 unicast
  !
!
vrf VPN_A
  rd 101:201
  address-family ipv4 unicast
    redistribute rip
  !
!

```

```

vrf VPN_B
  rd 102:202
  address-family ipv4 unicast
    redistribute eigrp 1
  !
!
!
mpls ldp
  router-id 20.20.20.20
  interface GigabitEthernet0/0/0/0.420
  !
!
router eigrp 65535
  vrf VPN_B
  address-family ipv4
  autonomous-system 1
  redistribute bgp 200
  interface GigabitEthernet0/0/0/0.820
  !
!
!
!
router rip
  vrf VPN_A
  interface GigabitEthernet0/0/0/0.1020
  !
  redistribute bgp 200 route-policy BGP_TO_RIP
  !
!

```

Verification

MPLS L3VPN Inter-AS Option B, sometimes also called MPLS VPN Inter-AS with ASBRs Exchanging VPN-IPv4 Addresses, is similar to the previous Option A example with the exception that the Inter-AS link runs a single VPNv4 EBGp peering instead of multiple VRF aware IGP or BGP instances. Unlike the Option A example in which the providers did not run VPNv4 BGP with each other, Option B's VPNv4 EBGp peering between the providers means that the VPNv4 Route Distinguisher and Route Target fields have global significance between the MPLS Service Providers. This is evident based on the fact that the customer attached PE routers R2 and XR2, who do not have VPNv4 BGP peerings with each other, must agree on a standard RD and RT policy in order for routing exchange to be complete.

The final verification of this design, like any other L3VPN configuration, would be to test end-to-end connectivity between the customer sites. Once again just like in Option A the customer networks are oblivious to the fact that their traffic is transiting multiple providers. As long as the SP is able to provide end-to-end IP connectivity this is what the customer is most concerned with.

```
R9#show ip route rip
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
```

```
10.0.0.0/32 is subnetted, 1 subnets
R    10.10.10.10 [120/1] via 30.2.9.2, 00:00:15, GigabitEthernet1.29
30.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R    30.10.20.0/24 [120/1] via 30.2.9.2, 00:00:15, GigabitEthernet1.29
```

```
R9#ping 10.10.10.10 source loopback 0
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.10.10.10, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/24/32 ms
```

```
R9#traceroute 10.10.10.10 source loopback0
```

```
Type escape sequence to abort.
```

```
Tracing the route to 10.10.10.10
```

```
VRF info: (vrf in name/id, vrf out name/id)
```

```
 1 30.2.9.2 1 msec 1 msec 1 msec
 2 10.2.3.3 [MPLS: Labels 17/24 Exp 0] 26 msec 89 msec 31 msec
 3 10.1.3.1 [MPLS: Label 24 Exp 0] 32 msec 28 msec 32 msec
 4 30.1.19.19 [MPLS: Label 16015 Exp 0] 30 msec 30 msec 32 msec
 5 20.4.19.4 [MPLS: Labels 17/16003 Exp 0] 30 msec 30 msec 31 msec
 6 20.4.20.20 [MPLS: Label 16003 Exp 0] 23 msec 24 msec 24 msec
 7 30.10.20.10 20 msec * 23 msec
```

From inside the Service Provider network though, it begins to become clear that the design is now more complicated than the Service Providers simply treating each other as customer sites. Instead, the Service Providers now treat each other as VPNv4 EBGP peers, which means the design is more similar to normal IPv4 Unicast EBGP for regular unlabeled IPv4 transport over the Internet.

In the below output we see that the final PE routers R2 and XR2 learn the VPNv4 BGP routes via their Inter-AS PEs R1 and XR1, however transitive attributes such as the AS-Path, the VPNv4 Route Distinguisher, and the VPNv4 Route Target values remain.

```
R2#show bgp vpnv4 unicast all
BGP table version is 21, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

```

Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 101:201 (default for vrf VPN_A)
*> 9.9.9.9/32     30.2.9.9          1          32768 ?
*>i 10.10.10.10/32 1.1.1.1           0    100    0 200 ?
*> 30.2.9.0/24    0.0.0.0           0          32768 ?
*>i 30.10.20.0/24  1.1.1.1           0    100    0 200 ?
Route Distinguisher: 102:202 (default for vrf VPN_B)
*> 7.7.7.7/32     40.2.7.7          130816     32768 ?
*>i 8.8.8.8/32     1.1.1.1           0    100    0 200 ?
*> 40.2.7.0/24    0.0.0.0           0          32768 ?
*>i 40.8.20.0/24  1.1.1.1           0    100    0 200 ?
```

```
R2#show bgp vpnv4 unicast all 8.8.8.8
BGP routing table entry for 102:202:8.8.8.8/32, version 19
Paths: (1 available, best #1, table VPN_B)
Not advertised to any peer
Refresh Epoch 1
200
1.1.1.1 (metric 3) (via default) from 1.1.1.1 (1.1.1.1)
Origin incomplete, metric 0, localpref 100, valid, internal, best
Extended Community: RT:102:202 0x8800:32768:0 0x8801:1:128256
0x8802:65281:2560 0x8803:1:1500
mpls labels in/out nlabel/26
rx pathid: 0, tx pathid: 0x0
```

```
RP/0/0/CPU0:XR2#show bgp vpnv4 unicast
```

```
Wed May 13 01:24:16.081 UTC
BGP router identifier 20.20.20.20, local AS number 200
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 19
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 101:201 (default for vrf VPN_A)
*>i9.9.9.9/32     19.19.19.19       100    0 100 ?
*> 10.10.10.10/32 30.10.20.10        1          32768 ?
*>i30.2.9.0/24    19.19.19.19       100    0 100 ?
*> 30.10.20.0/24  0.0.0.0           0          32768 ?
Route Distinguisher: 102:202 (default for vrf VPN_B)
```

```

*>i7.7.7.7/32      19.19.19.19      100      0 100 ?
*> 8.8.8.8/32      40.8.20.8        2570240   32768 ?
*>i40.2.7.0/24     19.19.19.19      100      0 100 ?
*> 40.8.20.0/24    0.0.0.0          0         32768 ?

```

Processed 8 prefixes, 8 paths

```
RP/0/0/CPU0:XR2#show bgp vpnv4 unicast vrf VPN_A 9.9.9.9
```

Wed May 13 01:25:22.566 UTC

BGP routing table entry for 9.9.9.9/32, Route Distinguisher: 101:201

Versions:

Process	bRIB/RIB	SendTblVer
Speaker	14	14

Last Modified: May 13 01:18:17.625 for 00:07:05

Paths: (1 available, best #1)

Not advertised to any peer

Path #1: Received by speaker 0

Not advertised to any peer

100

19.19.19.19 (metric 20) from 19.19.19.19 (19.19.19.19)

Received Label 16011

Origin incomplete, localpref 100, valid, internal, best, group-best, import-candidate, imported

Received Path ID 0, Local Path ID 1, version 14

Extended community: RT:101:201

Source VRF: VPN_A, Source Route Distinguisher: 101:201

Since the PE routers R2 and XR2 do not have an IGP route to the transit link between R1 and XR1, these edge routers are changing the next-hop value to their own local Loopback.

Changing the BGP next-hop value of a VPNv4 route has an additional special significance, because the next-hop value is what's used to determine which LDP derived Transport Label to use when encapsulating MPLS traffic towards the destination. In this specific design, the VPNv4 next-hop value is changing three times for every route that is originated by the customer facing PEs.

When R2 originates the VPNv4 route 7.7.7.7/32 towards R1, it sets the next-hop value to its own local Loopback. This means that R1 will use the transport label towards 2.2.2.2 in order to label switch traffic towards 7.7.7.7/32. Notice that an "In" label has been allocated locally by R1, and is what will be advertised towards XR1.

```
R1#show bgp vpnv4 unicast all 7.7.7.7/32
```

BGP routing table entry for 102:202:7.7.7.7/32, version 4

Paths: (1 available, best #1, no table)

Advertised to update-groups:

2

Refresh Epoch 1

Local

2.2.2.2 (metric 3) (via default) from 2.2.2.2 (2.2.2.2)

Origin incomplete, metric 130816, localpref 100, valid, internal, best

Extended Community: RT:102:202 Cost:pre-bestpath:128:130816

0x8800:32768:0 0x8801:1:128256 0x8802:65281:2560 0x8803:65281:1500

0x8806:0:117901063

mpls labels in/out 22/21

rx pathid: 0, tx pathid: 0x0

When a VPNv4 router without local VRFs inserts itself in the forwarding path by changing the next-hop to self, LFIB state is created for such paths. This forwarding state is what is then used for packet forwarding in hardware. Not only does R1 have LFIB entries for the LDP derived labels, but additional entries have been created for all VPN routes for which R1 set itself as a next-hop. Notice that local label 22 has an entry, the VPN "in" label which is advertised to XR1 for 102:202:7.7.7.7/32. The "out" label from the VPN route, 21, is the label that R2 originated for the route.

```
R1#show mpls forwarding-table
```

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Switched	Label	Outgoing interface	Next Hop
16	Pop Label	3.3.3.3/32	0		Gi1.13	10.1.3.3
17	Pop Label	10.2.3.0/24	0		Gi1.13	10.1.3.3
18	16	2.2.2.2/32	0		Gi1.13	10.1.3.3
19	Pop Label	30.1.19.19/32	0		Gi1.30	30.1.19.19
20	19	101:201:9.9.9.9/32	0			
			3728		Gi1.13	10.1.3.3
21	20	101:201:30.2.9.0/24	0		Gi1.13	10.1.3.3
22	21	102:202:7.7.7.7/32	0		Gi1.13	10.1.3.3
23	22	102:202:40.2.7.0/24	0		Gi1.13	10.1.3.3
24	16015	101:201:10.10.10.10/32	0			
			1912		Gi1.30	30.1.19.19
25	16016	101:201:30.10.20.0/24	0		Gi1.30	30.1.19.19
26	16017	102:202:8.8.8.8/32	0		Gi1.30	30.1.19.19
27	16018	102:202:40.8.20.0/24	0		Gi1.30	30.1.19.19

When R1 sends the route to its VPNv4 EBGp neighbor, XR1, the next-hop value is changed to the local peering address. Specifically this is the address 30.1.19.1, as seen below. Also notice that the VPNv4 MPLS Label is changing every time the next-hop value changes.

```
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast rd 102:202 7.7.7.7/32
```

```
Wed May 13 02:01:35.687 UTC
```

```
BGP routing table entry for 7.7.7.7/32, Route Distinguisher: 102:202
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
```

```
Speaker          4         4
```

```
Local Label: 16013
```

```
Last Modified: May 13 01:16:26.451 for 00:45:09
```

```
Paths: (1 available, best #1)
```

```
Advertised to peers (in unique update groups):
```

```
20.20.20.20
```

```
Path #1: Received by speaker 0
```

```
Advertised to peers (in unique update groups):
```

```
20.20.20.20
```

```
100
```

```
30.1.19.1 from 30.1.19.1 (1.1.1.1)
```

```
Received Label 22
```

```
Origin incomplete, localpref 100, valid, external, best, group-best, import-candidate, not-in-vrf
```

```
Received Path ID 0, Local Path ID 1, version 4
```

```
Extended community: EIGRP route-info:0x8000:0 EIGRP AD:1:128256 EIGRP RHB:255:1:2560 EIGRP LM:0xff:1:1500 0x8806:0x00:0x00:0x07:0x07:
```

```
0x07:0x07 RT:102:202
```

XR1 also creates LFIB state for all the VPN routes for which it is setting itself as the next-hop. Notice that the "out" or "received" label for 102:202:7.7.7.7/32 is 22.

```
RP/0/0/CPU0:XR1#show mpls forwarding
```

```
Wed May 13 02:19:08.005 UTC
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16000	17	20.20.20.20/32	Gi0/0/0/0.419	20.4.19.4	9851
16001	Pop	20.4.20.0/24	Gi0/0/0/0.419	20.4.19.4	0
16003	Pop	4.4.4.4/32	Gi0/0/0/0.419	20.4.19.4	7272
16007	Pop	30.1.19.1/32	Gi0/0/0/0.30	30.1.19.1	10574
16011	20	101:201:9.9.9.9/32	Gi0/0/0/0.30	30.1.19.1	3244
16012	21	101:201:30.2.9.0/24	\		
			Gi0/0/0/0.30	30.1.19.1	0
16013	22	102:202:7.7.7.7/32	Gi0/0/0/0.30	30.1.19.1	0
16014	23	102:202:40.2.7.0/24	\		
			Gi0/0/0/0.30	30.1.19.1	0
16015	16003	101:201:10.10.10.10/32	\		
				20.20.20.20	2584
16016	16008	101:201:30.10.20.0/24	\		
				20.20.20.20	0
16017	16005	102:202:8.8.8.8/32		20.20.20.20	0
16018	16009	102:202:40.8.20.0/24	\		
				20.20.20.20	0

A very important point about this design is that IOS XR cannot label switch traffic towards a next-hop that is not learned via a /32 host route. This means that in order for XR1 to label switch traffic towards the connected link to R1, it must have a /32 host route to the address 30.1.19.1. This is the reason that XR1 has a static route locally configured in its routing table, as seen below.

```

RP/0/0/CPU0:XR1#show run router static
Wed May 13 02:04:01.147 UTC

router static

  address-family ipv4 unicast

    30.1.19.1/32 GigabitEthernet0/0/0.30

  !

!

RP/0/0/CPU0:XR1#show route static
Wed May 13 02:04:15.376 UTC

S   30.1.19.1/32 is directly connected, 00:49:57, GigabitEthernet0/0/0.30

```

Note that IOS does this automatically when the 'mpls bgp forwarding' command is enabled, as it is on R1's interface towards XR1. This command is automatically generated when an EBGp VPNv4 session is established.

```

R1#show run int Gi1.30
Building configuration...

Current configuration : 118 bytes
!
interface GigabitEthernet1.30
  encapsulation dot1Q 30
  ip address 30.1.19.1 255.255.255.0
  mpls bgp forwarding
end

R1#show ip route connected | include GigabitEthernet1.30
C       30.1.19.0/24 is directly connected, GigabitEthernet1.30
L       30.1.19.1/32 is directly connected, GigabitEthernet1.30
C       30.1.19.19/32 is directly connected, GigabitEthernet1.30

R1#show ip route 30.1.19.19
Routing entry for 30.1.19.19/32

  Known via "connected", distance 0, metric 0 ( connected, via interface )

  Routing Descriptor Blocks:

    * directly connected, via GigabitEthernet1.30

      Route metric is 0, traffic share count is 1

```

XR1 continues to advertise the route on to XR2, again updating the next-hop value of the VPNv4 route, which is due to the **next-hop-self** command under the VPNv4 BGP process. Note again that since the next-hop value changed, so did the MPLS Label.

```

RP/0/0/CPU0:XR2#show bgp vpnv4 unicast rd 102:202 7.7.7.7/32

Wed May 13 02:07:55.391 UTC

BGP routing table entry for 7.7.7.7/32, Route Distinguisher: 102:202

Versions:

  Process          bRIB/RIB  SendTblVer
  Speaker          16        16

Last Modified: May 13 01:18:17.625 for 00:49:37

Paths: (1 available, best #1)

  Not advertised to any peer

  Path #1: Received by speaker 0

  Not advertised to any peer

  100

  19.19.19.19 (metric 20) from 19.19.19.19 (19.19.19.19)

    Received Label 16013

    Origin incomplete, localpref 100, valid, internal, best, group-best, import-candidate, imported

    Received Path ID 0, Local Path ID 1, version 16

    Extended community: EIGRP route-info:0x8000:0 EIGRP AD:1:128256 EIGRP RHB:255:1:2560 EIGRP LM:0xff:1:1500 0x8806:0x00:0x00:0x07:0x07:
0x07:0x07 RT:102:202

    Source VRF: VPN_B, Source Route Distinguisher: 102:202

```

The end result of this is that when a traceroute is sent between the customer sites it can be seen that the VPN Label changes three times during transit. The LFIB state created on R1 and XR1 allows for the "stitching" of the LSPs to occur.

```

R8#traceroute 7.7.7.7 source loopback 0

Type escape sequence to abort.

Tracing the route to 7.7.7.7

VRF info: (vrf in name/id, vrf out name/id)

 1 40.8.20.20 2 msec 1 msec 1 msec
 2 20.4.20.4 [MPLS: Labels 16/ 16013 Exp 0] 24 msec 20 msec 20 msec
 3 20.4.19.19 [MPLS: Label 16013 Exp 0] 20 msec 20 msec 27 msec
 4 30.1.19.1 [MPLS: Label 22 Exp 0] 43 msec 21 msec 28 msec
 5 10.1.3.3 [MPLS: Labels 16/ 21 Exp 0] 32 msec 32 msec 28 msec
 6 40.2.7.2 [MPLS: Label 21 Exp 0] 24 msec 20 msec 20 msec
 7 40.2.7.7 20 msec * 24 msec

```

From R8 to R7 the VPN label 16013 terminates at XR1, a second LSP via VPN label 22 terminates at R1, and a third LSP via VPN label 21 terminates at R2. This can likewise be seen on the reverse path, with label value 26 from R2 to R1, label value 16017 from R1 to XR1, and label value 16005 from XR1 to XR2.

```

R7#traceroute 8.8.8.8 source loopback 0

Type escape sequence to abort.

Tracing the route to 8.8.8.8

VRF info: (vrf in name/id, vrf out name/id)

 1 40.2.7.2 4 msec 1 msec 1 msec
 2 10.2.3.3 [MPLS: Labels 17/ 26 Exp 0] 27 msec 25 msec 24 msec
 3 10.1.3.1 [MPLS: Label 26 Exp 0] 23 msec 24 msec 28 msec
 4 30.1.19.19 [MPLS: Label 16017 Exp 0] 19 msec 24 msec 20 msec
 5 20.4.19.4 [MPLS: Labels 17/ 16005 Exp 0] 28 msec 32 msec 31 msec
 6 20.4.20.20 [MPLS: Label 16005 Exp 0] 28 msec 20 msec 24 msec
 7 40.8.20.8 20 msec * 24 msec

```

Keep in mind that this also means that the traffic on the Inter-AS link between R1 and XR1 is now MPLS Labeled, where in the previous Inter-AS Option A example the traffic was native IPv4. This can have a potential impact when trying to implement features such as QoS, which require different configuration to classify packets as MPLS vs. native IPv4 packets.

Another caveat of this configuration is the need for the command **no bgp default route-target filter** in regular IOS on R1 and **retain route-target all** in IOS XR on XR1. These are needed because the VPNv4 peers do not have the VRFs VPN_A and VPN_B configured locally, which means that they will discard the VPNv4 routes when received from their peers because they do not have a matching Route Target Import Policy. The potential problem of this can be seen below with the default behavior of the route target filter.

```
R1#debug bgp vpnv4 unicast update
BGP updates debugging is on for address family: VPNv4 Unicast
R1#config t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#router bgp 100
R1(config-router)#bgp default route-target filter
R1(config-router)#end
R1#clear bgp vpnv4 unicast * in
BGP: nbr_topo global 2.2.2.2 VPNv4 Unicast:base (0x7FF7B88F7BE0:1) rcvd Refresh Start-of-RIB
BGP: nbr_topo global 2.2.2.2 VPNv4 Unicast:base (0x7FF7B88F7BE0:1) refresh_epoch is 2
BGP(4): 2.2.2.2 rcvd UPDATE w/ attr: nexthop 2.2.2.2, origin ?, localpref 100, metric 130816, extended community RT:102:202 Cost:pre-bestpat
h:128:130816 0x8800:32768:0 0x8801:1:128256 0x8802:65281:2560 0x8803:65281:1500 0x8806:0:117901063
BGP(4): 2.2.2.2 rcvd 102:202:7.7.7.7/32, label 21 -- DENIED due to: extended community not supported;
BGP(4): no valid path for 102:202:7.7.7.7/32
BGP(4): 2.2.2.2 rcvd UPDATE w/ attr: nexthop 2.2.2.2, origin ?, localpref 100, metric 0, extended community RT:102:202 Cost:pre-bestpath:12
8:2816 0x8800:32768:0 0x8801:1:256 0x8802:65280:2560 0x8803:65281:1500 0x8806:0:671221506
BGP(4): 2.2.2.2 rcvd 102:202:40.2.7.0/24, label 22 -- DENIED due to: extended community not supported;
BGP(4): no valid path for 102:202:40.2.7.0/24
BGP(4): 2.2.2.2 rcvd UPDATE w/ attr: nexthop 2.2.2.2, origin ?, localpref 100, metric 0, extended community RT:101:201
BGP(4): 2.2.2.2 rcvd 101:201:30.2.9.0/24, label 20 -- DENIED due to: extended community not supported;
BGP(4): no valid path for 101:201:30.2.9.0/24
BGP(4): 2.2.2.2 rcvd UPDATE w/ attr: nexthop 2.2.2.2, origin ?, localpref 100, metric 1, extended community RT:101:201
BGP(4): 2.2.2.2 rcvd 101:201:9.9.9.9/32, label 19 -- DENIED due to: extended community not supported;
BGP(4): no valid path for 101:201:9.9.9.9/32
BGP: nbr_topo global 2.2.2.2 VPNv4 Unicast:base (0x7FF7B88F7BE0:1) rcvd Refresh End-of-RIB
```

Since R1 now drops inbound advertisements from its VPNv4 peers R2 and XR1, the end-to-end control plane is now broken.

```

R7#show ip route eigrp

R7#

R9#show ip route rip

R9#

R2#show bgp vpnv4 unicast all
BGP table version is 27, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 101:201 (default for vrf VPN_A)					
*> 9.9.9.9/32	30.2.9.9	1		32768	?
*> 30.2.9.0/24	0.0.0.0	0		32768	?
Route Distinguisher: 102:202 (default for vrf VPN_B)					
*> 7.7.7.7/32	40.2.7.7	130816		32768	?
*> 40.2.7.0/24	0.0.0.0	0		32768	?

```

RP/0/0/CPU0:XR2#show bgp vpnv4 unicast
Wed May 13 02:30:18.139 UTC
BGP router identifier 20.20.20.20, local AS number 200
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 23
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard

```

```

Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 101:201 (default for vrf VPN_A)					
*> 10.10.10.10/32	30.10.20.10	1		32768	?
*> 30.10.20.0/24	0.0.0.0	0		32768	?
Route Distinguisher: 102:202 (default for vrf VPN_B)					
*> 8.8.8.8/32	40.8.20.8	2570240		32768	?
*> 40.8.20.0/24	0.0.0.0	0		32768	?

Processed 4 prefixes, 4 paths

From the above output we see that the final PE routers of R2 and XR2 only know about VPNv4 routes that they are locally originating. All Inter-AS exchanged VPNv4 routes are no longer learned. The same would be true if XR1 reverted back to its default filtering behavior.

```

RP/0/0/CPU0:XR1#debug bgp update vpnv4 unicast in
RP/0/0/CPU0:XR1(config)#no router bgp 100
RP/0/0/CPU0:XR1(config)#router bgp 200
RP/0/0/CPU0:XR1(config-bgp)#address-family vpnv4 unicast
RP/0/0/CPU0:XR1(config-bgp-af)#no retain route-target all
RP/0/0/CPU0:XR1(config-bgp-af)#commit
Wed May 13 02:32:31.570 UTC

RP/0/0/CPU0:XR1#clear bgp vpnv4 unicast * soft in
RP/0/0/CPU0:XR1#RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: UPDATE from 20.20.20.20 contains nh 20.20.20.20/32, gw_afi 0, flags 0x0, nlr_i_afi 4
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: NH-Validate-Create: addr=20.20.20.20/32, len=12, nlr_i_afi=4, nbr=20.20.20.20, gw_afi=0, gwlen=4, gwaddrlen=32:: nhout=0x10905220, validity=1, attrwdrflags=0x00000000
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: --bgp4_rcv_attributes--: END: nbr=20.20.20.20:: msg=0x1002fdcc/91, uprlen=72, attrb1=0x1002fde3/68, ipv4reachlen=0, msginpath=0x3ef0bf0, asloopcheck=1, attrwdrfl=0x00000000:: samecluster=0, local_as_prepend=0, attr_wdr_flags 0x00000000, myascount=0:: rcvdata=0x1002fe27/0, errptr=0x1002fe1c/11
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received UPDATE from 20.20.20.20 with attributes:
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): nexthop 20.20.20.20/32, origin ?, localpref 100, metric 0, extended community RT:101:201
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received prefix 2ASN:101:201:30.10.20.0/24 (path ID: none) with MPLS label 16008 from neighbor 20.20.20.20
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Prefix 2ASN:101:201:30.10.20.0/24 (path ID: none) received from 20.20.20.20 DENIED RT extended community is not imported locally
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: UPDATE from 20.20.20.20 contains nh 20.20.20.20/32, gw_afi 0, flags 0x0, nlr_i_afi 4
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: NH-Validate-Create: addr=20.20.20.20/32, len=12, nlr_i_afi=4, nbr=20.20.20.20, gw_afi=0, gwlen=4, gwaddrlen=32:: nhout=0x10905220, validity=1, attrwdrflags=0x00000000
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: --bgp4_rcv_attributes--: END: nbr=20.20.20.20:: msg=0x1002fb64/92, uprlen=73, attrb1=0x1002fb7b/69, ipv4reachlen=0, msginpath=0x3ef0bf0, asloopcheck=1, attrwdrfl=0x00000000:: samecluster=0, local_as_prepend=0, attr_wdr_flags 0x00000000, myascount=0:: rcvdata=0x1002fbc0/0, errptr=0x1002fbb5/11
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received UPDATE from 20.20.20.20 with attributes:
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): nexthop 20.20.20.20/32, origin ?, localpref 100, metric 1, extended community RT:101:201
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received prefix 2ASN:101:201:10.10.10.10/32 (path ID: none) with MPLS label 16003 from neighbor 20.20.20.20
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Prefix 2ASN:101:201:10.10.10.10/32 (path ID: none) received from 20.20.20.20 DENIED RT extended community is not imported locally
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: UPDATE from 20.20.20.20 contains nh 20.20.20.20/32, gw_afi 0, flags 0x0, nlr_i_afi 4
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: NH-Validate-Create: addr=20.20.20.20/32, len=12, nlr_i_afi=4, nbr=20.20.20.20, gw_afi=0, gwlen=4, gwaddrlen=32:: nhout=0x10905220, validity=1, attrwdrflags=0x00000000
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: --bgp4_rcv_attributes--: END: nbr=20.20.20.20:: msg=0x10037cb0/131, uprlen=112, attrb1=0x10037cc7/108, ipv4reachlen=0, msginpath=0x3ef0bf0, asloopcheck=1, attrwdrfl=0x00000000:: samecluster=0, local_as_prepend=0, attr_wdr_flags 0x00000000, myascount=0:: rcvdata=0x10037d33/0, errptr=0x10037d00/51
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received UPDATE from 20.20.20.20 with attributes:
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): nexthop 20.20.20.20/32, origin ?, localpref 100, metric 0, extended community COST:128:128:10240 EIGRP route-info:0x8000:0 EIGRP AD:1:256 EIGRP RHB:255:0:2560 EIGRP LM:0x0:1:1500 RT:102:202
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received prefix 2ASN:102:202:40.8.20.0/24 (path ID: none) with MPLS label 16009 from neighbor 20.20.20.20
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Prefix 2ASN:102:202:40.8.20.0/24 (path ID: none) received from 20.20.20.20 DENIED RT extended community is not imported locally
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: UPDATE from 20.20.20.20 contains nh 20.20.20.20/32, gw_afi 0, flags 0x0, nlr_i_afi 4

```

i 4

```
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: NH-Validate-Create: addr=20.20.20.20/32, len=12, nlrifa=4, nbr=20.20.20.20, gw  
afi=0, gwlen=4, gwaddrlen=32:: nhout=0x10905220, validity=1, attrwdrflags=0x00000000  
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr]: --bgp4_rcv_attributes--: END: nbr=20.20.20.20:: msg=0x10037e48/132, uprlen=113,  
attrbl=0x10037e5f/109, ipv4reachlen=0, msginpath=0x3ef0bf0, asloopcheck=1, attrwdrfl=0x00000000:: samecluster=0, loc  
al_as_prepend=0, attr_wdr_flags 0x00000000, myascount=0:: rcvdata=0x10037ecc/0, errptr=0x10037e99/51  
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received UPDATE from 20.20.20.20 with attributes:  
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): nexthop 20.20.20.20/32, origin ?, localpref 100, metric 2570240, extende  
d community COST:128:128:2570240 EIGRP route-info:0x8000:0 EIGRP AD:1:128256 EIGRP RHB:255:1:2560 EIGRP LM:0x0:1:150  
0 RT:102:202  
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Received prefix 2ASN:102:202:8.8.8.8/32 (path ID: none) with MPLS label  
16005 from neighbor 20.20.20.20  
RP/0/0/CPU0:May 13 02:35:09.209 : bgp[1047]: [default-rtr] (vpn4u): Prefix 2ASN:102:202:8.8.8.8/32 (path ID: none) received from 20.20.20.20  
DENIED RT extended community is not imported locally
```

An alternate solution to this problem would be to locally configure the VRFs on the edge routers R1 and XR1, or to configure them as VPNv4 Route Reflectors, as Route Reflectors do not filter prefixes based on the Route Target extended community. These possible solutions can be seen below.

```
R1#show bgp vpnv4 unicast all
```

```
R1#config t
```

Enter configuration commands, one per line. End with CNTL/Z.

```
R1(config)#vrf definition VPN_A
```

```
R1(config-vrf)# rd 101:201
```

```
R1(config-vrf)# !
```

```
R1(config-vrf)# address-family ipv4
```

```
R1(config-vrf-af)# route-target export 101:201
```

```
R1(config-vrf-af)# route-target import 101:201
```

```
R1(config-vrf-af)# exit-address-family
```

```
R1(config-vrf)#
```

```
R1(config-vrf)#vrf definition VPN_B
```

```
R1(config-vrf)# rd 102:202
```

```
R1(config-vrf)# !
```

```
R1(config-vrf)# address-family ipv4
```

```
R1(config-vrf-af)# route-target export 102:202
```

```
R1(config-vrf-af)# route-target import 102:202
```

```
R1(config-vrf-af)# exit-address-family
```

```
R1(config-vrf-af)#end
```

```
R1#
```

```
RP/0/0/CPU0:XR1#conf t
```

```
Wed May 13 02:37:57.637 UTC
```

```
RP/0/0/CPU0:XR1(config)#router bgp 200
```

```
RP/0/0/CPU0:XR1(config-bgp)#neighbor 20.20.20.20
```

```
RP/0/0/CPU0:XR1(config-bgp-nbr)#address-family vpnv4 unicast
```

```
RP/0/0/CPU0:XR1(config-bgp-nbr-af)#route-reflector-client
```

```
RP/0/0/CPU0:XR1(config-bgp-nbr-af)#commit
```

```
Wed May 13 02:38:20.896 UTC
```

```
RP/0/0/CPU0:May 13 02:38:21.306 : config[65729]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by user 'admin'. Use 'show configuration commit changes 100000215' to view the changes.
```

```
RP/0/0/CPU0:XR1(config-bgp-nbr-af)#RP/0/0/CPU0:May 13 02:38:22.956 : bgp[1047]: %ROUTING-BGP-5-ADJCHANGE : neighbor 20.20.20.20 Down - RR client configuration changed (VRF: default) (AS: 200)
```

```
RP/0/0/CPU0:May 13 02:38:25.625 : bgp[1047]: %ROUTING-BGP-5-ADJCHANGE : neighbor 20.20.20.20 Up (VRF: default) (AS: 200)
```

```
R1#show bgp vpnv4 unicast all
```

```
BGP table version is 33, local router ID is 1.1.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
          r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
```

```
          x best-external, a additional-path, c RIB-compressed,
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
RPKI validation codes: V valid, I invalid, N Not found
```

```
          Network                  Next Hop                  Metric LocPrf Weight Path
```

```
Route Distinguisher: 101:201 (default for vrf VPN_A)
```

```
*>i 9.9.9.9/32          2.2.2.2                  1    100      0 ?
```

```
*> 10.10.10.10/32      30.1.19.19              0      0      0 200 ?
```

```
*>i 30.2.9.0/24         2.2.2.2                  0    100      0 ?
```

```
*> 30.10.20.0/24      30.1.19.19              0      0      0 200 ?
```

```
Route Distinguisher: 102:202 (default for vrf VPN_B)
```

```
*>i 7.7.7.7/32          2.2.2.2                 130816  100      0 ?
```

```
*> 8.8.8.8/32          30.1.19.19              0      0      0 200 ?
```

```
*>i 40.2.7.0/24      2.2.2.2          0      100      0 ?
*> 40.8.20.0/24     30.1.19.19      0 200 ?
```

Notice that now the LFIB state on R1 now shows an "V" - this is referencing the fact that the routes are in a VRF. The downside to this technique is that additional resources are consumed because R1 now has to install all of the routes in each VRF's RIB / FIB - even though they will never be used for forwarding towards locally attached customers. With the previous solution, only LFIB state was created, and no routes were installed (there were no VRFs to install the routes into).

R1#show mpls forwarding-table

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Switched	Label	Outgoing interface	Next Hop
16	Pop Label	3.3.3.3/32	0		Gi1.13	10.1.3.3
17	Pop Label	10.2.3.0/24	0		Gi1.13	10.1.3.3
18	16	2.2.2.2/32	0		Gi1.13	10.1.3.3
19	Pop Label	30.1.19.19/32	0		Gi1.30	30.1.19.19
20	20	30.2.9.0/24[V]	0		Gi1.13	10.1.3.3
21	21	7.7.7.7/32[V]	3098		Gi1.13	10.1.3.3
22	19	9.9.9.9/32[V]	1824		Gi1.13	10.1.3.3
23	22	40.2.7.0/24[V]	0		Gi1.13	10.1.3.3
24	16016	30.10.20.0/24[V]	0		Gi1.30	30.1.19.19
25	16017	8.8.8.8/32[V]	1302		Gi1.30	30.1.19.19
26	16015	10.10.10.10/32[V]	\			
			3030		Gi1.30	30.1.19.19
27	16018	40.8.20.0/24[V]	0		Gi1.30	30.1.19.19

R1#show mpls forwarding-table vrf VPN_B 7.7.7.7 detail

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Switched	Label	Outgoing interface	Next Hop
21	21	7.7.7.7/32[V]	3098		Gi1.13	10.1.3.3

MAC/Encaps=18/26, MRU=1496, Label Stack{16 21}

0050569E6E6A0050569E5913810000D8847 000100000015000

VPN route: VPN_B

No output feature configured

R1#show ip route vrf VPN_B 7.7.7.7

Routing Table: VPN_B

Routing entry for 7.7.7.7/32

Known via "bgp 100", distance 200, metric 130816, type internal

Last update from 2.2.2.2 00:11:37 ago

Routing Descriptor Blocks:

* 2.2.2.2 (default), from 2.2.2.2, 00:11:37 ago

Route metric is 130816, traffic share count is 1

AS Hops 0

MPLS label: 21

MPLS Flags: MPLS Required

RP/0/0/CPU0:XR1#show bgp vpnv4 unicast

Wed May 13 02:39:33.381 UTC

BGP router identifier 19.19.19.19, local AS number 200

BGP generic scan interval 60 secs

BGP table state: Active

Table ID: 0x0 RD version: 0

BGP main routing table version 25

BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best

i - internal, r RIB-failure, S stale, N Nexthop-discard

Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path

```

Route Distinguisher: 101:201
*> 9.9.9.9/32      30.1.19.1          0 100 ?
*>i10.10.10.10/32  20.20.20.20        1 100  0 ?
*> 30.2.9.0/24     30.1.19.1          0 100 ?
*>i30.10.20.0/24   20.20.20.20        0 100  0 ?

Route Distinguisher: 102:202
*> 7.7.7.7/32      30.1.19.1          0 100 ?
*>i8.8.8.8/32      20.20.20.20        2570240 100 0 ?
*> 40.2.7.0/24     30.1.19.1          0 100 ?
*>i40.8.20.0/24    20.20.20.20        0 100  0 ?

```

Processed 8 prefixes, 8 paths

R2#show bgp vpnv4 unicast all

```

BGP table version is 35, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 101:201 (default for vrf VPN_A)					
*> 9.9.9.9/32	30.2.9.9	1		32768	?
*>i 10.10.10.10/32	1.1.1.1	0	100	0	200 ?
*> 30.2.9.0/24	0.0.0.0	0		32768	?
*>i 30.10.20.0/24	1.1.1.1	0	100	0	200 ?
Route Distinguisher: 102:202 (default for vrf VPN_B)					
*> 7.7.7.7/32	40.2.7.7	130816		32768	?
*>i 8.8.8.8/32	1.1.1.1	0	100	0	200 ?
*> 40.2.7.0/24	0.0.0.0	0		32768	?
*>i 40.8.20.0/24	1.1.1.1	0	100	0	200 ?

RP/0/0/CPU0:XR2#show bgp vpnv4 unicast

Wed May 13 02:40:20.188 UTC

BGP router identifier 20.20.20.20 local AS number 200

BGP generic scan interval 60 secs

BGP table state: Active

Table ID: 0x0 RD version: 0

BGP main routing table version 43

BGP scan interval 60 secs

```

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 101:201 (default for vrf VPN_A)					
*>i9.9.9.9/32	19.19.19.19		100	0	100 ?
*> 10.10.10.10/32	30.10.20.10	1		32768	?
*>i30.2.9.0/24	19.19.19.19		100	0	100 ?
*> 30.10.20.0/24	0.0.0.0	0		32768	?
Route Distinguisher: 102:202 (default for vrf VPN_B)					
*>i7.7.7.7/32	19.19.19.19		100	0	100 ?
*> 8.8.8.8/32	40.8.20.8	2570240		32768	?

```
*>i40.2.7.0/24      19.19.19.19      100      0 100 ?
*> 40.8.20.0/24     0.0.0.0          0        32768 ?
```

Processed 8 prefixes, 8 paths

R9#show ip route rip

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

10.0.0.0/32 is subnetted, 1 subnets

R 10.10.10.10 [120/1] via 30.2.9.2, 00:00:10, GigabitEthernet1.29

30.0.0.0/8 is variably subnetted, 3 subnets, 2 masks

R 30.10.20.0/24 [120/1] via 30.2.9.2, 00:00:10, GigabitEthernet1.29

R10#traceroute 9.9.9.9 source loopback 0

Type escape sequence to abort.

Tracing the route to 9.9.9.9

VRF info: (vrf in name/id, vrf out name/id)

```
 1 30.10.20.20 4 msec 2 msec 0 msec
 2 20.4.20.4 [MPLS: Labels 16/16011 Exp 0] 27 msec 26 msec 24 msec
 3 20.4.19.19 [MPLS: Label 16011 Exp 0] 24 msec 21 msec 23 msec
 4 30.1.19.1 [MPLS: Label 22 Exp 0] 24 msec 32 msec 29 msec
 5 10.1.3.3 [MPLS: Labels 16/19 Exp 0] 27 msec 28 msec 29 msec
 6 30.2.9.2 [MPLS: Label 19 Exp 0] 22 msec 24 msec 24 msec
 7 30.2.9.9 24 msec * 26 msec
```

R7#traceroute 8.8.8.8 source loopback 0

Type escape sequence to abort.

Tracing the route to 8.8.8.8

VRF info: (vrf in name/id, vrf out name/id)

```
 1 40.2.7.2 3 msec 1 msec 0 msec
 2 10.2.3.3 [MPLS: Labels 17/25 Exp 0] 26 msec 23 msec 24 msec
 3 10.1.3.1 [MPLS: Label 25 Exp 0] 28 msec 24 msec 24 msec
 4 30.1.19.19 [MPLS: Label 16017 Exp 0] 24 msec 24 msec 24 msec
 5 20.4.19.4 [MPLS: Labels 17/16005 Exp 0] 30 msec 29 msec 31 msec
 6 20.4.20.20 [MPLS: Label 16005 Exp 0] 28 msec 23 msec 24 msec
 7 40.8.20.8 24 msec * 26 msec
```

This Inter-AS Layer 3 VPNs design is documented under Section 10.b in RFC 4364
(<https://tools.ietf.org/html/rfc4364#section-10>)

« MPLS L3VPN Inter-AS Option A - Back-to-Back VRF Exchange (/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-a-back-to-back-vrf-exchange-Mjg2Nw%3D%3D) | MPLS L3VPN Inter-AS Option C - Multihop VPNv4 EBGP Exchange (/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-c-multihop-vpnv4-ebgp-exchange-Mjg2OQ%3D%3D) »

<https://t.me/learningnets>