MPLS L3VPN Inter-AS Option A - Back-to-Back VRF Exchange

« Multi-VRF CE (VRF Lite) (/workbook/view/service-provider-v4/task/multi-vrf-ce-vrf-lite-Mjg2Ng%3D%3D) | MPLS L3VPN Inter-AS Option B - VPNv4 EBGP Exchange (/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-b-vpnv4-ebgp-exchange-Mjg2OA%3D%3D) »

Last updated: April 23, 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

- 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
    - XR1 – 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 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.
- Configure the following VRFs on PE routers R1 and R2 in AS 100 as follows:
  - VRF VPN_A:
    - Route Distinguisher: 100:1
    - Route Target Import: 100:1
    - Route Target Export: 100:1
    - Assign this VRF on the links in the 30.0.0.0 network on R1 and R2.
  - VRF VPN_B:
    - Route Distinguisher: 100:2
    - Route Target Import: 100:2

        - Route Target Export: 100:2
        - Assign this VRF on the links in the 40.0.0.0 network on R1 and R2.
- Configure the following VRFs on PE routers XR1 and XR2 in AS 200 as follows:
    - VRF VPN_A:
        - Route Distinguisher: 200:1
        - Route Target Import: 200:1
        - Route Target Export: 200:1
        - Assign this VRF on the links in the 30.0.0.0 network on XR1 and XR2.
    - VRF VPN_B:
        - Route Distinguisher: 200:2
        - Route Target Import: 200:2
        - Route Target Export: 200:2
        - Assign this VRF on the links in the 40.0.0.0 network on XR1 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 EIGRP.
- Redistribute between VPNv4 BGP and the VRF aware IGP processes on the PE routers R1, R2, XR1, 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:
vrf definition VPN_A
 rd 100:1
 !
 address-family ipv4
 route-target export 100:1
 route-target import 100:1
 exit-address-family
!
vrf definition VPN_B
  rd 100:2
  !
 address-family ipv4
 route-target export 100:2
 route-target import 100:2
 exit-address-family
!
interface Loopback0
  ip ospf 1 area 0
!
interface GigabitEthernet1.13
  ip ospf 1 area 0
  mpls ip
!
interface GigabitEthernet1.30
  vrf forwarding VPN_A
  ip address 30.1.19.1 255.255.255.0
!
interface GigabitEthernet1.40
  vrf forwarding VPN_B
  ip address 40.1.19.1 255.255.255.0
!
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 2.2.2.2 remote-as 100
  neighbor 2.2.2.2 update-source Loopback0
  !
  address-family vpnv4
   neighbor 2.2.2.2 activate
   neighbor 2.2.2.2 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


R2:
vrf definition VPN_A
 rd 100:1
 !
 address-family ipv4
 route-target export 100:1
 route-target import 100:1
 exit-address-family
!
vrf definition VPN_B
 rd 100:2
 !
 address-family ipv4
 route-target export 100:2
 route-target import 100:2
 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:
vrf VPN_A
  address-family ipv4 unicast
   import route-target
    200:1
   !
   export route-target
    200:1
  !
 !
!
vrf VPN_B
  address-family ipv4 unicast
```

```
   import route-target
    200:2
   !
   export route-target
    200:2
  !
 !
!
interface GigabitEthernet0/0/0/0.30
  vrf VPN_A
  no ipv4 address
  ipv4 address 30.1.19.19 255.255.255.0
!
interface GigabitEthernet0/0/0/0.40
  vrf VPN_B
  no ipv4 address
  ipv4 address 40.1.19.19 255.255.255.0
!
route-policy BGP_TO_RIP
  set rip-metric 1
end-policy
!
router isis 1
 is-type level-2-only
 net 49.0001.0000.0000.0019.00
 interface Loopback0
  passive
  address-family ipv4 unicast
 !
!
interface GigabitEthernet0/0/0/0.419
  address-family ipv4 unicast
  !
 !
!
router bgp 200
 address-family vpnv4 unicast
 !
 neighbor 20.20.20.20
  remote-as 200
  update-source Loopback0
  address-family vpnv4 unicast
  !
!
vrf VPN_A
 rd 200:1
 address-family ipv4 unicast
  redistribute rip
 !
!
vrf VPN_B
 rd 200:2
 address-family ipv4 unicast
  redistribute eigrp 1
```

```
  !
 !
!
mpls ldp
 router-id 19.19.19.19
 interface GigabitEthernet0/0/0/0.419
 !
!
router eigrp 65535
 vrf VPN_B
  address-family ipv4
   autonomous-system 1
   redistribute bgp 200
   interface GigabitEthernet0/0/0/0.40
   !
  !
 !
!
router rip
 vrf VPN_A
interface GigabitEthernet0/0/0/0.30
!
redistribute bgp 200 route-policy BGP_TO_RIP
!
!
end


XR2:
vrf VPN_A
  address-family ipv4 unicast
   import route-target
    200:1
   !
   export route-target
    200:1
  !
 !
!
vrf VPN_B
 address-family ipv4 unicast
  import route-target
   200:2
  !
  export route-target
   200:2
  !
 !
!
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 200:1
 address-family ipv4 unicast
  redistribute rip
 !
!
vrf VPN_B
 rd 200:2
 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

This example demonstrates one of the ways that a Service Provider can extend a customer's Layer 3 MPLS VPN through another Service Provider, giving the customer transparent transport between their sites. This specific design is known as MPLS L3VPN Inter-AS Option A, or sometimes called Back-to-Back VRF Exchange. This is the simplest of all of the Inter-AS designs, because the different Service Providers essentially treat each other as just another customer site.

In this example the Inter-AS link between R1 and XR1 has one subinterfaces per- VRF that the Service Providers wants to exchange. Each of them treat these links as if they are CE facing links, with normal VRF membership, VRF aware routing configuration, and redistribution of learned routes into VPNv4 BGP and back. The disadvantage of using this design is that scalability becomes an issue as the number of VRFs that the Service Providers want to exchange increase. This scaling problem is solved with the later options B and C, and in the hybrid option AB.

Verification of this design works just like any other MPLS L3VPN design covered up to this point. The PE routers R1 and R2 learn routes from the CE routers XR1, R7, and R9 and redistribute them into VPNv4 BGP. R1 and R2 agree on a Route Target Import and Export policy that will allow them to accept each other's routes for VRF VPN_A and VRF VPN_B.

```
R2#show ip route vrf VPN_A

Routing Table: VPN_A
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

      9.0.0.0/32 is subnetted, 1 subnets
R        9.9.9.9 [120/1] via 30.2.9.9, 00:00:13, GigabitEthernet1.29
      10.0.0.0/32 is subnetted, 1 subnets
B        10.10.10.10 [200/1] via 1.1.1.1, 00:00:23
      30.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B        30.1.19.0/24 [200/0] via 1.1.1.1, 00:00:23
C        30.2.9.0/24 is directly connected, GigabitEthernet1.29
L        30.2.9.2/32 is directly connected, GigabitEthernet1.29
B        30.10.20.0/24 [200/1] via 1.1.1.1, 00:00:23

R1#show ip route vrf VPN_A

Routing Table: VPN_A
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

      9.0.0.0/32 is subnetted, 1 subnets
B        9.9.9.9 [200/1] via 2.2.2.2, 00:00:53
      10.0.0.0/32 is subnetted, 1 subnets
R        10.10.10.10 [120/1] via 30.1.19.19, 00:00:25, GigabitEthernet1.30
      30.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        30.1.19.0/24 is directly connected, GigabitEthernet1.30
L        30.1.19.1/32 is directly connected, GigabitEthernet1.30
B        30.2.9.0/24 [200/0] via 2.2.2.2, 00:00:53
R        30.10.20.0/24 [120/1] via 30.1.19.19, 00:00:25, GigabitEthernet1.30

R2#show bgp vpnv4 unicast vrf VPN_A
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: 100:1 (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                  1    100      0 ?
*>i 30.1.19.0/24     1.1.1.1                  0    100      0 ?
*>  30.2.9.0/24      0.0.0.0                  0         32768 ?
*>i 30.10.20.0/24    1.1.1.1                  1    100      0 ?


R1#show bgp vpnv4  unicast vrf VPN_A
BGP table version is 15, 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: 100:1 (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               1         32768 ?
 *>  30.1.19.0/24     0.0.0.0                  0         32768 ?
 *>i 30.2.9.0/24      2.2.2.2                  0    100      0 ?
 *>  30.10.20.0/24    30.1.19.19               1         32768 ?
```

When the routes are advertised to the other Service Provider, AS 200 in this case, they treat AS 100 just like any other customer site.

```
RP/0/0/CPU0:XR2#show route vrf VPN_A

Tue May 12 23:59:52.338 UTC


Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path

       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, E - EGP

       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2

       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default

       U - per-user static route, o - ODR, L - local, G  - DAGR, l - LISP

       A - access/subscriber, a - Application route, (!) - FRR Backup path


Gateway of last resort is not set


B    9.9.9.9/32 [200/1] via 19.19.19.19 (nexthop in vrf default), 00:02:56

R    10.10.10.10/32 [120/1] via 30.10.20.10, 00:09:25, GigabitEthernet0/0/0/0.1020

B    30.1.19.0/24 [200/0] via 19.19.19.19 (nexthop in vrf default), 00:08:07

B    30.2.9.0/24 [200/1] via 19.19.19.19 (nexthop in vrf default), 00:02:56

C    30.10.20.0/24 is directly connected, 00:10:09, GigabitEthernet0/0/0/0.1020

L    30.10.20.20/32 is directly connected, 00:10:09, GigabitEthernet0/0/0/0.1020



RP/0/0/CPU0:XR1#show route vrf VPN_A

Wed May 13 00:00:05.186 UTC


Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path

       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, E - EGP

       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2

       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default

       U - per-user static route, o - ODR, L - local, G  - DAGR, l - LISP

       A - access/subscriber, a - Application route, (!) - FRR Backup path


Gateway of last resort is not set


R    9.9.9.9/32 [120/1] via 30.1.19.1, 00:03:09, GigabitEthernet0/0/0/0.30

B    10.10.10.10/32 [200/1] via 20.20.20.20 (nexthop in vrf default), 00:08:20

C    30.1.19.0/24 is directly connected, 00:09:55, GigabitEthernet0/0/0/0.30

L    30.1.19.19/32 is directly connected, 00:09:55, GigabitEthernet0/0/0/0.30

R    30.2.9.0/24 [120/1] via 30.1.19.1, 00:03:09, GigabitEthernet0/0/0/0.30

B    30.10.20.0/24 [200/0] via 20.20.20.20 (nexthop in vrf default), 00:08:20
```

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

Wed May 13 00:00:18.366 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 17

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: 200:1 (default for vrf VPN_A)
*>i9.9.9.9/32         19.19.19.19              1    100      0 ?
*> 10.10.10.10/32     30.10.20.10              1         32768 ?
*>i30.1.19.0/24       19.19.19.19              0    100      0 ?
*>i30.2.9.0/24        19.19.19.19              1    100      0 ?
*> 30.10.20.0/24      0.0.0.0                  0         32768 ?

Processed 5 prefixes, 5 paths



RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A
Wed May 13 00:00:34.384 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 18
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: 200:1 (default for vrf VPN_A)
*> 9.9.9.9/32         30.1.19.1                1         32768 ?
*>i10.10.10.10/32     20.20.20.20              1    100      0 ?
*> 30.1.19.0/24       0.0.0.0                  0         32768 ?
*> 30.2.9.0/24        30.1.19.1                1         32768 ?
*>i30.10.20.0/24      20.20.20.20              0    100      0 ?

Processed 5 prefixes, 5 paths
```

The final result of this is that although the customer sites do have end-to-end reachability to each other, they do not use the same MPLS Label Switch Path to exchange their traffic. Instead there is one LSP from PEs R1 to R2, and second LSP from PEs XR1 to XR2, and vice versa. The Inter-AS traffic exchanged on the R1 to XR1 link is sent as normal *unlabeled* IPv4 traffic, resulting in a path that is not end to end MPLS encapsulated. Verification of this can be seen through a traceroute.

```
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 2 msec 1 msec 1 msec

 2 10.2.3.3 [MPLS: Labels 17/23 Exp 0] 3 msec 2 msec 10 msec

 3 40.1.19.1 [MPLS: Label 23 Exp 0] 19 msec 18 msec 19 msec

 4 40.1.19.19 19 msec 14 msec 14 msec

 5 20.4.19.4 [MPLS: Labels 17/16005 Exp 0] 31 msec 23 msec 20 msec

 6 20.4.20.20 [MPLS: Label 16005 Exp 0] 24 msec 24 msec 20 msec

 7 40.8.20.8 20 msec *  25 msec


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 1 msec 2 msec 1 msec

  2 20.4.20.4 [MPLS: Labels 16/16009 Exp 0] 26 msec 19 msec 19 msec

  3 20.4.19.19 [MPLS: Label 16009 Exp 0] 20 msec 20 msec 20 msec

  4 30.1.19.1 20 msec 20 msec 20 msec

  5 10.1.3.3 [MPLS: Labels 16/16 Exp 0] 20 msec 20 msec 20 msec

  6 30.2.9.2 [MPLS: Label 16 Exp 0] 25 msec 23 msec 20 msec

  7 30.2.9.9 20 msec *  28 msec
```

The above output indicates that when traffic from R7's site in VPN_B goes to 8.8.8.8 – the Loopback0 of R8 – it uses the VPN label 23 to reach R1, traffic is unlabeled IPv4 traffic between R1 and XR1, then XR1 uses the VPN label 16005 to reach XR2, where the traffic is sent towards the final customer site. However, from the customer's point of view, end-to-end IP transport is provided, and they have no way of knowing that their traffic is transiting more than one service provider.

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

« Multi-VRF CE (VRF Lite) (/workbook/view/service-provider-v4/task/multi-vrf-ce-vrf-lite-Mjg2Ng%3D%3D) | MPLS L3VPN Inter-AS Option B - VPNv4 EBGP Exchange (/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-b-vpnv4-ebgp-exchange-Mjg2OA%3D%3D) »