Unified MPLS

« Multicast VPNs - mLDP Core Tree (/workbook/view/service-provider-v4/task/multicast-vpns-mldp-core-tree-Mjk0MA%3D%3D) | undefined »

Last updated: April 22, 2016

Note:

Initial Configuration & Diagrams: Load the initial configuration files for the section named **Unified MPLS**, 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 **Unified MPLS Diagram** in order to complete this task.

Task

- Configure a Unified MPLS architecture allowing PE routers R3, R8, R9, and XR3 to provide services to CE routers R7 and XR4.
- IGP Zones have been configured according to the diagram.
- Enable LDP label exchange in each IGP area. Use the Loopback0 of each device as the LDP router-id.
- Configure the following redistributions to facilitate BGP session establishment between ASBR routers in adjacent zones. Ensure that redistributed routes are not leaked beyond the IGP zone stated below.
  - Redistribute the Loopback0 of R6 and R10 into OSPF Process 2.
  - Redistribute the Loopback0 of R2 and XR2 into OSPF Process 2.
- Configure BGP AS1 as follows:
  - Use the Loopback0 of all devices as the source for the BGP sessions.
  - Configure R8 and R9 as Route Reflector Clients of R6 and R10.
  - Configure R6 and R10 as Route Reflector Clients of each other.
  - Configure R3 and XR3 as Route Reflector Clients of R2 and XR2.
  - Configure R2 and XR2 as Route Reflector Clients of each other.
  - Configure ASBR routers R6 and R10 to peer with ASBR routers R2 and XR2.
  - Enable label exchange using RFC-3107 on all peerings.
  - Advertise the Loopbacks within OSPF Process 1 into BGP on R6 and R10.
  - Advertise the Loopbacks within OSPF Process 3 into BGP on R2 and XR2.
  - Ensure all BGP routers involved in reflecting routes set themselves as the next-hop for routes being reflected.
- Configure the following services over the Unified MPLS network:
  - Configure PE Routers R3 and R9 to provide a point-to-point EoMPLS service between R7's Gig1.34 and XR4's Gig0/0/0/0.914.
    - R7 and XR4 should form and OSPFv2 adjacency over the EoMPLS service.
    - Advertise the Loopback1 of R7 and XR14 into OSPFv2 and ensure they have reachability between the Loopback1 interfaces using the EoMPLS service.

- Configure PE Routers R8 and XR3 to provide a L3VPN service to R7 and XR2.
    - Run eBGP as the PE/CE protocol.
    - Configure R7 and XR4 in BGP AS 100.
    - Ensure that R7 and XR4 have reachability between their Loopback0 over the L3VPN service.

Configuration **Click to *collapse***

CONTENTS

```
R1:
mpls ldp router-id loopback 0 force
!
router ospf 2
 mpls ldp autoconfig area 0
!



R2:
mpls ldp router-id loopback 0 force
!
route-map REDISTRIBUTE_LOOPBACK permit 10
 match interface lo0
!
router ospf 2
 mpls ldp autoconfig area 0
 redistribute connected subnets route-map REDISTRIBUTE_LOOPBACK
!
router ospf 3
 mpls ldp autoconfig area 0
!
ip prefix-list LOOPBACKS permit 0.0.0.0/0 ge 32 le 32
!
route-map REDISTRIBUTE_LOOPBACKS_BGP permit 10
 match ip add prefix-list LOOPBACKS
!
!
router bgp 1
 template peer-policy iBGP_RR_Client
  route-reflector-client
  next-hop-self all
  send-label
 exit-peer-policy
 !
 template peer-policy iBGP_Non_RR_Client
  next-hop-self all
  send-label
 exit-peer-policy
 !
 template peer-session iBGP_SESS
  remote-as 1
  update-source Loopback0
 exit-peer-session
 !
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 6.6.6.6 inherit peer-session iBGP_SESS
 neighbor 3.3.3.3 inherit peer-session iBGP_SESS
 neighbor 10.10.10.10 inherit peer-session iBGP_SESS
 neighbor 12.12.12.12 inherit peer-session iBGP_SESS
 neighbor 13.13.13.13 inherit peer-session iBGP_SESS
 !
 address-family ipv4
```

CONTENTS

```
  redistribute ospf 3 route-map REDISTRIBUTE_LOOPBACKS_BGP
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 inherit peer-policy iBGP_Non_RR_Client
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 inherit peer-policy iBGP_RR_Client
  neighbor 10.10.10.10 activate
  neighbor 10.10.10.10 inherit peer-policy iBGP_Non_RR_Client
  neighbor 12.12.12.12 activate
  neighbor 12.12.12.12 inherit peer-policy iBGP_RR_Client
  neighbor 13.13.13.13 activate
  neighbor 13.13.13.13 inherit peer-policy iBGP_RR_Client
 exit-address-family



R3:
mpls ldp router-id loopback 0 force
!
router ospf 3
 mpls ldp autoconfig area 0
!
router bgp 1
 no bgp default ipv4-unicast
 neighbor 2.2.2.2 remote-as 1
 neighbor 2.2.2.2 update-source lo0
 neighbor 12.12.12.12 remote-as 1
 neighbor 12.12.12.12 update-source lo0
 !
 address-family ipv4 unicast
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.2 send-label
 neighbor 12.12.12.12 activate
 neighbor 12.12.12.12 send-label
!
!
!
! Customer Services
!
!
interface GigabitEthernet1.37
 xconnect 9.9.9.9 39 encapsulation mpls
!
!



R4:
mpls ldp router-id loopback 0 force
!
router ospf 2
 mpls ldp autoconfig area 0



R5:
mpls ldp router-id loopback 0 force
!
```

```
router ospf 2
 mpls ldp autoconfig area 0



R6:
mpls ldp router-id loopback 0 force
!
route-map REDISTRIBUTE_LOOPBACK permit 10
 match interface lo0
!
router ospf 1
 mpls ldp autoconfig area 0
!
router ospf 2
 mpls ldp autoconfig area 0
 redistribute connected subnets route-map REDISTRIBUTE_LOOPBACK
!
ip prefix-list LOOPBACKS permit 0.0.0.0/0 ge 32 le 32
!
route-map REDISTRIBUTE_LOOPBACKS_BGP permit 10
 match ip add prefix-list LOOPBACKS
!
router bgp 1
 template peer-policy iBGP_RR_Client
  route-reflector-client
  next-hop-self all
  send-label
 exit-peer-policy
 !
 template peer-policy iBGP_Non_RR_Client
  next-hop-self all
  send-label
 exit-peer-policy
 !
 template peer-session iBGP_SESS
  remote-as 1
  update-source Loopback0
 exit-peer-session
 !
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2.2.2.2 inherit peer-session iBGP_SESS
 neighbor 12.12.12.12 inherit peer-session iBGP_SESS
 neighbor 8.8.8.8 inherit peer-session iBGP_SESS
 neighbor 9.9.9.9 inherit peer-session iBGP_SESS
 neighbor 10.10.10.10 inherit peer-session iBGP_SESS
 !
 address-family ipv4
  redistribute ospf 1 route-map REDISTRIBUTE_LOOPBACKS_BGP
  neighbor 2.2.2.2 activate
  neighbor 2.2.2.2 inherit peer-policy iBGP_Non_RR_Client
  neighbor 12.12.12.12 activate
  neighbor 12.12.12.12 inherit peer-policy iBGP_Non_RR_Client
  neighbor 8.8.8.8 activate
```

```
  neighbor 8.8.8.8 inherit peer-policy iBGP_RR_Client
  neighbor 9.9.9.9 activate
  neighbor 9.9.9.9 inherit peer-policy iBGP_RR_Client
  neighbor 10.10.10.10 activate
  neighbor 10.10.10.10 inherit peer-policy iBGP_RR_Client
 exit-address-family

R7:
interface GigabitEthernet1.37
 ip ospf 100 area 0
!
interface Loopback1
 ip ospf 100 area 0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 10.7.13.13 remote-as 1
 !
 address-family ipv4
  network 7.7.7.7 mask 255.255.255.255
  neighbor 10.7.13.13 activate
 exit-address-family


R8:
mpls ldp router-id loopback 0 force
!
router ospf 1
 mpls ldp autoconfig area 0
!
router bgp 1
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 6.6.6.6 remote-as 1
 neighbor 6.6.6.6 update-source Loopback0
 neighbor 10.10.10.10 remote-as 1
 neighbor 10.10.10.10 update-source Loopback0
 !
 address-family ipv4
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 send-label
  neighbor 10.10.10.10 activate
  neighbor 10.10.10.10 send-label
 exit-address-family
!
!
!
! Customer Services
!
!
router bgp 1
 neighbor 13.13.13.13 remote-as 1
 neighbor 13.13.13.13 update-source Loopback0
 !
```

CONTENTS ›

```
 address-family vpnv4
  neighbor 13.13.13.13 activate
  neighbor 13.13.13.13 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf A
  neighbor 10.8.14.14 remote-as 100
  neighbor 10.8.14.14 activate
  neighbor 10.8.14.14 as-override
 exit-address-family


R9:
mpls ldp router-id loopback 0 force
!
router ospf 1
 mpls ldp autoconfig area 0
!
router bgp 1
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 6.6.6.6 remote-as 1
 neighbor 6.6.6.6 update-source Loopback0
 neighbor 10.10.10.10 remote-as 1
 neighbor 10.10.10.10 update-source Loopback0
 !
 address-family ipv4
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 send-label
  neighbor 10.10.10.10 activate
  neighbor 10.10.10.10 send-label
 exit-address-family
!
!
! Customer Services
!
!
interface GigabitEthernet1.914
 xconnect 3.3.3.3 39 encapsulation mpls


R10:
mpls ldp router-id loopback 0 force
!
route-map REDISTRIBUTE_LOOPBACK permit 10
 match interface lo0
!
router ospf 1
 mpls ldp autoconfig area 0
!
router ospf 2
 mpls ldp autoconfig area 0
 redistribute connected subnets route-map REDISTRIBUTE_LOOPBACK
!
```

```
ip prefix-list LOOPBACKS permit 0.0.0.0/0 ge 32 le 32
!
route-map REDISTRIBUTE_LOOPBACKS_BGP permit 10
 match ip add prefix-list LOOPBACKS
!
router bgp 1
 template peer-policy iBGP_RR_Client
  route-reflector-client
  next-hop-self all
  send-label
 exit-peer-policy
 !
 template peer-policy iBGP_Non_RR_Client
  next-hop-self all
  send-label
 exit-peer-policy
 !
 template peer-session iBGP_SESS
  remote-as 1
  update-source Loopback0
 exit-peer-session
 !
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2.2.2.2 inherit peer-session iBGP_SESS
 neighbor 12.12.12.12 inherit peer-session iBGP_SESS
 neighbor 8.8.8.8 inherit peer-session iBGP_SESS
 neighbor 9.9.9.9 inherit peer-session iBGP_SESS
 neighbor 6.6.6.6 inherit peer-session iBGP_SESS
 !
 address-family ipv4
  redistribute ospf 1 route-map REDISTRIBUTE_LOOPBACKS_BGP
  neighbor 2.2.2.2 activate
  neighbor 2.2.2.2 inherit peer-policy iBGP_Non_RR_Client
  neighbor 12.12.12.12 activate
  neighbor 12.12.12.12 inherit peer-policy iBGP_Non_RR_Client
  neighbor 8.8.8.8 activate
  neighbor 8.8.8.8 inherit peer-policy iBGP_RR_Client
  neighbor 9.9.9.9 activate
  neighbor 9.9.9.9 inherit peer-policy iBGP_RR_Client
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 inherit peer-policy iBGP_RR_Client
 exit-address-family



XR1:
router ospf 2
 mpls ldp auto-config
!
mpls ldp
 router-id 11.11.11.11
```

CONTENTS ›

```
XR2:
router ospf 2
 mpls ldp auto-config
 redistribute connected route-policy REDISTRIBUTE_CONNECTED
!
router ospf 3
 area 0
 mpls ldp auto-config
!
mpls ldp
 router-id 12.12.12.12
!
route-policy REDISTRIBUTE_CONNECTED
  if destination in (12.12.12.12/32) then
    pass
  endif
end-policy
!
prefix-set LOOPBACKS
  0.0.0.0/0 ge 32 le 32
end-set
!
route-policy REDISTRIBUTE_LOOPBACKS_BGP
  if destination in LOOPBACKS then
    pass
  endif
end-policy
!
router bgp 1
 ibgp policy out enforce-modifications
 address-family ipv4 unicast
  redistribute ospf 3 route-policy REDISTRIBUTE_LOOPBACKS_BGP
  allocate-label all
 !
 neighbor-group iBGP_RR_Client
  remote-as 1
  update-source Loopback0
  address-family ipv4 labeled-unicast
   route-reflector-client
   next-hop-self
  !
 !
 neighbor-group iBGP_Non_RR_Client
  remote-as 1
  update-source Loopback0
  address-family ipv4 labeled-unicast
   next-hop-self
  !
 !
 neighbor 10.10.10.10
  use neighbor-group iBGP_Non_RR_Client
 !
 neighbor 2.2.2.2
  use neighbor-group iBGP_RR_Client
```

```
 !
 neighbor 3.3.3.3
  use neighbor-group iBGP_RR_Client
 !
 neighbor 6.6.6.6
  use neighbor-group iBGP_Non_RR_Client
 !
 neighbor 13.13.13.13
  use neighbor-group iBGP_RR_Client
 !
!


XR3:
router ospf 3
 mpls ldp auto-config
!
mpls ldp
 router-id 13.13.13.13
!
router bgp 1
 address-family ipv4 unicast
  allocate-label all
 !
 neighbor-group iBGP
  remote-as 1
  update-source Loopback0
  address-family ipv4 labeled-unicast
   !
 neighbor 2.2.2.2
  use neighbor-group iBGP
 !
 neighbor 12.12.12.12
  use neighbor-group iBGP
 !
!
!
!
! Customer Services
!
!
route-policy PASS
  pass
end-policy
!
router bgp 1
 address-family vpnv4 unicast
 !
 neighbor 8.8.8.8
  remote-as 1
  update-source Loopback0
  address-family vpnv4 unicast
   !
 !
```

CONTENTS

```
 vrf A
  rd 1300:1300
  address-family ipv4 unicast
  !
  neighbor 10.7.13.7
   remote-as 100
   address-family ipv4 unicast
    route-policy PASS in
    route-policy PASS out
    as-override
   !
  !
 !
!

XR4:
route-policy PASS
  pass
end-policy
!
router ospf 100
 area 0
  interface Loopback1
  !
  interface GigabitEthernet0/0/0.914
  !
 !
!
router bgp 100
 address-family ipv4 unicast
  network 14.14.14.14/32
 !
 neighbor 10.8.14.8
  remote-as 1
  address-family ipv4 unicast
   route-policy PASS in
   route-policy PASS out
  !
 !
!
```

CONTENTS

Verification

The idea behind the Unified MPLS architecture, also referred to as Seamless MPLS in the IETF draft draft-ietf-mpls-seamless-mpls-07 (https://tools.ietf.org/html/draft-ietf-mpls-seamless-mpls-07), is to allow the Service Provider network to scale to hundreds of thousand nodes, while still being able to provide end-to-end MPLS services to customers and end devices. One of the biggest scaling limitations of an MPLS network is the size of the IGP area/level/process. Additionally, dealing with the operational challenges introduced by the protocol interactions at such a large scale is another major challenge.

ISIS and OSPF can only scale up to a few thousand nodes. However, in order to provide end-to-end MPLS services throughout a large scale network, some sort of "stitching" is needed in order to advertise transport labels (also referred to as IGP labels throughout this workbook) in a network scaling past the bounds of an IGP. Several enhancements and best practices have been proposed to enhance

IGP scale over the years, such as configuring point-to-point Ethernet links as point-to-point, using prefix suppression, only advertising customer routes in BGP instead of the IGP, or enabling iSPF, but even with such enhancements, it is still not possible to dump 100,000 routes into an IGP today.

Unified MPLS is a design that breaks the scaling barriers imposed by the IGPs, by using BGP - the workhorse of the internet and other large networks - to carry information previously relayed by the IGP + LDP. In a Unified MPLS design, BGP is used to carry the PE or "service" router endpoint Loopbacks, plus a label for each loopback, using RFC-3107. Note that the Unified MPLS design does not introduce any new technologies. It is simply a re-architecture of how existing technologies and protocols interact. The goal is to extend MPLS throughout all tiers of the network, including the Access tiers, and allow PE routers at any Edge/Access tier to instantiate services over MPLS seamlessly.

The core of the Unified MPLS design consists of breaking up the Service Provider network into isolated IGP zones. These would typically consists of Access, Edge, Aggregation, and Core tiers within a Service Provider network. A separate IGP process (not Area/Level) is used in each zone, such that no topology information is leaked between the zones. LDP is used within each zone to exchange labels for the local loopbacks of routers in the zone. In our network this consists of two "Access/Edge" zones, using OSPF Process 1 and 3 respectively, and the Core zone, using OSPF Process 2. BGP is then used to merge the IGP zones together, by transporting each zone's Loopbacks plus their corresponding labels to all other zones within the SP network. Carrying this information in BGP allows each zone's IGP to stay slim and efficient, yet still be able to construct an LSP between any Access/Edge zones in order to provide a service to a customer. All zones are "Unified" into a single "Seamless" MPLS network. From the a PE router's perspective, the end goal is to mimic what a network would look like if it ran a flat IGP area throughout the entire network - any PE router would be able to establish an LSP to any other PE router. The problem as mentioned previously however is that no IGP can scale as large as is needed by today's large SP networks.

The BGP peerings are established between PE routers in each Access/Edge network to the ASBRs of that zone. In our network, PE routers R8 and R9 establish iBGP sessions to their local ASBR routers R6 and R10. Similarly, PE routers R3 and XR3 establish iBGP sessions with their zone's ASBR routers R2 and XR2. Note that the ASBR routers of each zone are configured as route-reflectors towards the PE routers they service. The zone's are then merged together to provide end to end route+label exchange by establishing peerings between each zone's ASBRs. The design in our current network only has three tiers, however it can scale to many more tiers following the same architecture: The goal of the design is to enable end-to-end route+label exchange between all zones, using the ASBRs of each zone as the inter-area peering points. The ASBRs of each zone serve as route reflectors for the PE routers of that zone, or any downstream ASBRs from other zones.

CONTENTS ›

In order for R6 and R10 to establish an iBGP session with R2 and XR2, the loopbacks of these devices need to be advertised into the "Core" IGP zone. We accomplish this by using tightly controlled redistribution of the Loopback of each ASBRs into the Core zone. Note that this is the only leaking of routing information that is needed between the IGP zones. Topology information is not leaked, as these are separate processes, not separate areas/levels. Other than this redistribution, each zone is self contained and as slim as possible. Additionally, LDP labels are exchanged for these redistributed routes, since each zone is running LDP and generates labels for all IGP routes. The redistribution of the ASBR loopbacks between adjacent zones would also be needed in larger designs with more zones in order to facilitate BGP peerings between ASBRs.

In order to exchange labels over BGP for the endpoint loopbacks, RFC-3107 (https://tools.ietf.org/html/rfc3107) is enabled on all active BGP sessions through the SP network. There are different methodologies to advertising endpoint Loopbacks into BGP in this type of design. The one used in this example was redistributing all Loopbacks within each Edge/Access zone (prefixes with a /32) into BGP by the ASBRs. This has the benefit of centralizing the route advertisement at the "edge" of each zone. Another way of accomplishing the the same thing would have been to use a network statement on each PE router for its loopback. Keep in mind that the routes we are interested in carrying in BGP are the Loopbacks of the routers who will be providing a service for a customer. In our network these are PE routers R3, R8, R9, and XR3.

One subtlety about this design is the way BGP next-hops are handled. This is a large scale iBGP architecture, where the route-reflectors have to insert themselves in the forwarding path. Normally route-reflection is done by off-line boxes hosted on virtual routers, or other software based routers without much hardware forwarding performance. Generally these devices are not in the forwarding path, as they simply reflect routes between the iBGP peers in the network. In a Unified MPLS design, the route-reflectors insert themselves in the forwarding path by setting themselves as the next-hop for all reflected routes (i,e `next-hop-self`). What is particular about this scenario however, is that normally a route-reflector is not supposed to change any attribute of routes it is reflecting. Prior to the BGP "knob" added in recent code versions, a route-reflector was not able to change the next-hop of an iBGP reflected route when it was being advertised to another iBGP peer - only routes received from an eBGP peer could have their next-hops changed. The new knob added, `next-hop-self all` in IOS and `ibgp policy out enforce-modifications` in IOS-XR, allows a route-reflector to change the next-hop of iBGP reflected routes, and in doing so the route-reflector inserts itself in the forwarding path for all routes it reflects. This is paramount to the Unified MPLS design, as it dictates where the underlying transport LSPs are terminated.

Lets begin our verification by validating the end services being overlaid on top of the Unified MPLS network - the EoMPLS xconnect between R3 and R9, and the L3VPN between R8 and XR3:

R7 and XR4 have established an OSPFv2 adjacency over the point-to-point EoMPLS service. This validates that MPLS services have been extended end-to-end - from Access tier to Access tier.

```
RP/0/0/CPU0:XR4#show ospf neighbor
Sun Jul 12 21:33:51.898 UTC

* Indicates MADJ interface

Neighbors for OSPF 100

Neighbor ID     Pri   State           Dead Time   Address         Interface
 77.77.77.77     1     FULL/BDR        00:00:35    10.0.0.7        GigabitEthernet0/0/0/0.914
    Neighbor is up for 00:11:42

Total neighbor count: 1



RP/0/0/CPU0:XR4#show route ospf
Sun Jul 12 21:34:08.587 UTC

 O    77.77.77.77/32 [110/2] via 10.0.0.7, 00:11:54, GigabitEthernet0/0/0/0.914

XR4 has learned about R7's Loopback1 as is able to reach it over the EoMPLS service.



RP/0/0/CPU0:XR4#ping 77.77.77.77 source 114.114.114.114
Sun Jul 12 21:36:02.269 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 77.77.77.77, timeout is 2 seconds:
 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/9/29 ms



RP/0/0/CPU0:XR4#traceroute 77.77.77.77 source 114.114.114.114
Sun Jul 12 21:36:10.639 UTC

Type escape sequence to abort.
Tracing the route to 77.77.77.77

  1  10.0.0.7 9 msec  *  0 msec



RP/0/0/CPU0:XR4#show arp GigabitEthernet0/0/0/0.914
Sun Jul 12 21:36:20.888 UTC

-------------------------------------------------------------------------------
0/0/CPU0
-------------------------------------------------------------------------------
Address         Age        Hardware Addr   State      Type  Interface
 10.0.0.7        01:14:04   0050.569e.5c4d  Dynamic    ARPA  GigabitEthernet0/0/0/0.914
10.0.0.14       -          0050.569e.762b  Interface  ARPA  GigabitEthernet0/0/0/0.914



R7#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
```

CONTENTS

```
114.114.114.114   1   FULL/DR         00:00:35    10.0.0.14       GigabitEthernet1.37



R7#show ip route ospf
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

      114.0.0.0/32 is subnetted, 1 subnets
O        114.114.114.114 [110/2] via 10.0.0.14, 00:16:11, GigabitEthernet1.37



R7#traceroute 114.114.114.114 source loopback1
Type escape sequence to abort.
Tracing the route to 114.114.114.114
VRF info: (vrf in name/id, vrf out name/id)
  1 10.0.0.14 8 msec *  5 msec
```

The L3VPN service is also working as expected. XR4 and R7 have routes and reachability to each others Loopback0.

```
RP/0/0/CPU0:XR4#show bgp ipv4 unicast
Sun Jul 12 21:39:52.354 UTC
BGP router identifier 14.14.14.14, local AS number 100
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000   RD version: 4
BGP main routing table version 4
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
 *> 7.7.7.7/32         10.8.14.8                             0 1 1 i
*> 14.14.14.14/32     0.0.0.0                  0         32768 i

Processed 2 prefixes, 2 paths

RP/0/0/CPU0:XR4#traceroute 7.7.7.7 source 14.14.14.14
Sun Jul 12 21:40:25.281 UTC

Type escape sequence to abort.
Tracing the route to 7.7.7.7

 1  10.8.14.8 0 msec  0 msec  0 msec
 2  10.6.8.6 [MPLS: Labels 41/16011 Exp 0] 19 msec  29 msec  29 msec
 3  10.4.6.4 [MPLS: Labels 17/37/16011 Exp 0] 29 msec  39 msec  19 msec
 4  10.4.11.11 [MPLS: Labels 16012/37/16011 Exp 0] 19 msec  19 msec  19 msec
 5  10.2.11.2 [MPLS: Labels 37/16011 Exp 0] 19 msec  19 msec  29 msec
 6  10.2.13.13 [MPLS: Label 16011 Exp 0] 29 msec  19 msec  19 msec
 7  10.7.13.7 19 msec  *  9 msec

R7#show bgp ipv4 unicast
BGP table version is 8, local router ID is 7.7.7.7
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
 *>  7.7.7.7/32       0.0.0.0                  0         32768 i
  *>  14.14.14.14/32   10.7.13.13                            0 1 1 i

R7#traceroute 14.14.14.14 source loopback 0
Type escape sequence to abort.
Tracing the route to 14.14.14.14
VRF info: (vrf in name/id, vrf out name/id)
  1 10.7.13.13 4 msec 2 msec 1 msec
  2 10.2.13.2 [MPLS: Labels 39/22 Exp 0] 20 msec 15 msec 16 msec
  3 10.2.11.11 [MPLS: Labels 16013/20/22 Exp 0] 16 msec 32 msec 32 msec
  4 10.4.11.4 [MPLS: Labels 29/20/22 Exp 0] 36 msec 28 msec 40 msec
  5 10.4.6.6 [MPLS: Labels 20/22 Exp 0] 36 msec 32 msec 55 msec
```

```
 6 10.8.14.8 [MPLS: Label 22 Exp 0] 32 msec 20 msec 20 msec
 7 10.8.14.14 20 msec *  22 msec
```

Notice that both of these services, L3VPN and EoMPLS, which rely on an end-to-end LSP between the endpoint PE routers, are working over our Unified MPLS network. These MPLS services were used to demonstrate that any MPLS service can be levered on top of this type of design.

Lets begin by looking at the internals of how the Unified MPLS architecture works.

Notice that each zone is kept slim. PE router R9 in an Access/Edge zone only sees the internal routes and LDP labels for other zone members.

```
R9#show ip route ospf
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

      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/2] via 10.6.9.6, 1d20h, GigabitEthernet1.69
      8.0.0.0/32 is subnetted, 1 subnets
O        8.8.8.8 [110/2] via 10.8.9.8, 1d20h, GigabitEthernet1.89
      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
O        10.6.8.0/24 [110/2] via 10.8.9.8, 1d20h, GigabitEthernet1.89
                     [110/2] via 10.6.9.6, 1d20h, GigabitEthernet1.69
O        10.6.10.0/24 [110/2] via 10.9.10.10, 1d20h, GigabitEthernet1.910
                      [110/2] via 10.6.9.6, 1d20h, GigabitEthernet1.69
O        10.8.10.0/24 [110/2] via 10.9.10.10, 1d20h, GigabitEthernet1.910
                      [110/2] via 10.8.9.8, 1d20h, GigabitEthernet1.89
O        10.10.10.10/32 [110/2] via 10.9.10.10, 1d20h, GigabitEthernet1.910


R9#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  6.6.6.6/32       0             Gi1.69     10.6.9.6
17         Pop Label  8.8.8.8/32       0             Gi1.89     10.8.9.8
18         Pop Label  10.6.8.0/24      0             Gi1.69     10.6.9.6
           Pop Label  10.6.8.0/24      0             Gi1.89     10.8.9.8
19         Pop Label  10.6.10.0/24     0             Gi1.69     10.6.9.6
           Pop Label  10.6.10.0/24     0             Gi1.910    10.9.10.10
20         Pop Label  10.8.10.0/24     0             Gi1.89     10.8.9.8
           Pop Label  10.8.10.0/24     0             Gi1.910    10.9.10.10
21         Pop Label  10.10.10.10/32   0             Gi1.910    10.9.10.10
22         No Label   l2ckt(1)         72794         Gi1.914    point2point
```

R6 is an ASBR router who's loopback is natively advertised into OSPF Process 1.

CONTENTS ›

```
R6#show ip ospf int brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C

Gi1.56       2     0               10.5.6.6/24        1     DR    1/1

Gi1.46       2     0               10.4.6.6/24        1     DR    1/1

 Lo0          1     0               6.6.6.6/32         1     LOOP  0/0

Gi1.68       1     0               10.6.8.6/24        1     BDR   1/1

Gi1.69       1     0               10.6.9.6/24        1     BDR   1/1

Gi1.610      1     0               10.6.10.6/24       1     BDR   1/1
```

In order to allow the ASBR routers of the adjacent zone (R2 and XR2) to form a BGP peering with R6, its loopback must be redistributed into the adjacent zone. This is why external routes for the ASBR routers are seen in the Core zone.

```
R6#show ip route ospf 2
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

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                 [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
      2.0.0.0/32 is subnetted, 1 subnets
 O E2     2.2.2.2 [110/20] via 10.5.6.5, 1d19h, GigabitEthernet1.56
                 [110/20] via 10.4.6.4, 1d19h, GigabitEthernet1.46
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/2] via 10.4.6.4, 1d20h, GigabitEthernet1.46
      5.0.0.0/32 is subnetted, 1 subnets
O        5.5.5.5 [110/2] via 10.5.6.5, 1d20h, GigabitEthernet1.56
      10.0.0.0/8 is variably subnetted, 26 subnets, 2 masks
O        10.1.2.0/24 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                     [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.1.4.0/24 [110/2] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.1.5.0/24 [110/2] via 10.5.6.5, 1d20h, GigabitEthernet1.56
O        10.1.11.0/24 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                      [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.1.12.0/24 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                      [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.2.11.0/24 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                      [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.4.5.0/24 [110/2] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                     [110/2] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.4.10.0/24 [110/2] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.4.11.0/24 [110/2] via 10.4.6.4, 1d20h, GigabitEthernet1.46
O        10.5.10.0/24 [110/2] via 10.5.6.5, 1d20h, GigabitEthernet1.56
O        10.5.11.0/24 [110/2] via 10.5.6.5, 1d20h, GigabitEthernet1.56
O        10.11.12.0/24 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                       [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
      11.0.0.0/32 is subnetted, 1 subnets
O        11.11.11.11 [110/3] via 10.5.6.5, 1d20h, GigabitEthernet1.56
                     [110/3] via 10.4.6.4, 1d20h, GigabitEthernet1.46
      12.0.0.0/32 is subnetted, 1 subnets
 O E2     12.12.12.12 [110/20] via 10.5.6.5, 1d19h, GigabitEthernet1.56
                     [110/20] via 10.4.6.4, 1d19h, GigabitEthernet1.46
```

Notice that a P router in the Core zone sees the redistributed loopbacks of all four ASBR routers in our network.

```
RP/0/0/CPU0:XR1#show route ospf
Sun Jul 12 22:34:40.448 UTC

O    1.1.1.1/32 [110/2] via 10.1.11.1, 1d20h, GigabitEthernet0/0/0/0.111
 O E2 2.2.2.2/32 [110/20] via 10.2.11.2, 1d19h, GigabitEthernet0/0/0/0.211
O    4.4.4.4/32 [110/2] via 10.4.11.4, 1d20h, GigabitEthernet0/0/0/0.411
O    5.5.5.5/32 [110/2] via 10.5.11.5, 1d20h, GigabitEthernet0/0/0/0.511
 O E2 6.6.6.6/32 [110/20] via 10.5.11.5, 1d19h, GigabitEthernet0/0/0/0.511
                [110/20] via 10.4.11.4, 1d19h, GigabitEthernet0/0/0/0.411
O    10.1.2.0/24 [110/2] via 10.2.11.2, 1d20h, GigabitEthernet0/0/0/0.211
                 [110/2] via 10.1.11.1, 1d20h, GigabitEthernet0/0/0/0.111
O    10.1.4.0/24 [110/2] via 10.4.11.4, 1d20h, GigabitEthernet0/0/0/0.411
                 [110/2] via 10.1.11.1, 1d20h, GigabitEthernet0/0/0/0.111
O    10.1.5.0/24 [110/2] via 10.5.11.5, 1d20h, GigabitEthernet0/0/0/0.511
                 [110/2] via 10.1.11.1, 1d20h, GigabitEthernet0/0/0/0.111
O    10.1.12.0/24 [110/2] via 10.11.12.12, 1d20h, GigabitEthernet0/0/0/0.1112
                  [110/2] via 10.1.11.1, 1d20h, GigabitEthernet0/0/0/0.111
O    10.4.5.0/24 [110/2] via 10.5.11.5, 1d20h, GigabitEthernet0/0/0/0.511
                 [110/2] via 10.4.11.4, 1d20h, GigabitEthernet0/0/0/0.411
O    10.4.6.0/24 [110/2] via 10.4.11.4, 1d20h, GigabitEthernet0/0/0/0.411
O    10.4.10.0/24 [110/2] via 10.4.11.4, 1d20h, GigabitEthernet0/0/0/0.411
O    10.5.6.0/24 [110/2] via 10.5.11.5, 1d20h, GigabitEthernet0/0/0/0.511
O    10.5.10.0/24 [110/2] via 10.5.11.5, 1d20h, GigabitEthernet0/0/0/0.511
 O E2 10.10.10.10/32 [110/20] via 10.5.11.5, 1d19h, GigabitEthernet0/0/0/0.511
                    [110/20] via 10.4.11.4, 1d19h, GigabitEthernet0/0/0/0.411
 O E2 12.12.12.12/32 [110/20] via 10.11.12.12, 1d19h, GigabitEthernet0/0/0/0.1112


RP/0/0/CPU0:XR1#show mpls forwarding
Sun Jul 12 22:39:41.297 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- -----------
16000  Pop         5.5.5.5/32         Gi0/0/0/0.511 10.5.11.5       291144
16001  Pop         10.5.6.0/24        Gi0/0/0/0.511 10.5.11.5       0
16002  Pop         1.1.1.1/32         Gi0/0/0/0.111 10.1.11.1       290248
16003  Pop         4.4.4.4/32         Gi0/0/0/0.411 10.4.11.4       290284
16004  Pop         10.1.5.0/24        Gi0/0/0/0.111 10.1.11.1       0
       Pop         10.1.5.0/24        Gi0/0/0/0.511 10.5.11.5       0
16005  Pop         10.4.6.0/24        Gi0/0/0/0.411 10.4.11.4       0
16006  Pop         10.1.4.0/24        Gi0/0/0/0.111 10.1.11.1       0
       Pop         10.1.4.0/24        Gi0/0/0/0.411 10.4.11.4       0
16007  Pop         10.4.5.0/24        Gi0/0/0/0.411 10.4.11.4       0
       Pop         10.4.5.0/24        Gi0/0/0/0.511 10.5.11.5       0
16008  Pop         10.1.2.0/24        Gi0/0/0/0.111 10.1.11.1       0
       Pop         10.1.2.0/24        Gi0/0/0/0.211 10.2.11.2       0
16009  Pop         10.5.10.0/24       Gi0/0/0/0.511 10.5.11.5       0
16010  Pop         10.1.12.0/24       Gi0/0/0/0.111 10.1.11.1       0
       Pop         10.1.12.0/24       Gi0/0/0/0.1112 10.11.12.12     0
16011  Pop         10.4.10.0/24       Gi0/0/0/0.411 10.4.11.4       0
16012  Pop         2.2.2.2/32         Gi0/0/0/0.211 10.2.11.2       313525
16013  29          6.6.6.6/32         Gi0/0/0/0.411 10.4.11.4       22633
       29          6.6.6.6/32         Gi0/0/0/0.511 10.5.11.5       690198
```

```
16014  30         10.10.10.10/32     Gi0/0/0/0.411 10.4.11.4       291401
       30         10.10.10.10/32     Gi0/0/0/0.511 10.5.11.5       0
16015  Pop        12.12.12.12/32     Gi0/0/0/0.1112 10.11.12.12     291577
```

PE router R3 in an Access/Edge zone only sees the internal routes for other zone members.

```
R3#show ip route ospf
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

      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/2] via 10.2.3.2, 1d20h, GigabitEthernet1.23
      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
O        10.2.12.0/24 [110/2] via 10.3.12.12, 1d20h, GigabitEthernet1.312
                      [110/2] via 10.2.3.2, 1d20h, GigabitEthernet1.23
O        10.2.13.0/24 [110/2] via 10.3.13.13, 1d20h, GigabitEthernet1.313
                      [110/2] via 10.2.3.2, 1d20h, GigabitEthernet1.23
O        10.12.13.0/24 [110/2] via 10.3.13.13, 1d20h, GigabitEthernet1.313
                       [110/2] via 10.3.12.12, 1d20h, GigabitEthernet1.312
      12.0.0.0/32 is subnetted, 1 subnets
O        12.12.12.12 [110/2] via 10.3.12.12, 1d20h, GigabitEthernet1.312
      13.0.0.0/32 is subnetted, 1 subnets
O        13.13.13.13 [110/2] via 10.3.13.13, 1d20h, GigabitEthernet1.313


R3#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  2.2.2.2/32       0             Gi1.23     10.2.3.2
17         Pop Label  10.2.12.0/24     0             Gi1.23     10.2.3.2
           Pop Label  10.2.12.0/24     0             Gi1.312    10.3.12.12
18         Pop Label  10.2.13.0/24     0             Gi1.23     10.2.3.2
           Pop Label  10.2.13.0/24     0             Gi1.313    10.3.13.13
19         Pop Label  10.12.13.0/24    0             Gi1.312    10.3.12.12
           Pop Label  10.12.13.0/24    0             Gi1.313    10.3.13.13
20         Pop Label  12.12.12.12/32   0             Gi1.312    10.3.12.12
21         Pop Label  13.13.13.13/32   0             Gi1.313    10.3.13.13
22         No Label   l2ckt(7)         64956         Gi1.37     point2point
```

Now that we covered how each IGP/LDP zone is structured, lets look at how BGP is merging these zones together.

PE router R9 established a Targeted-LDP session with R3, which is being used to exchange the EoMPLS VC label for the point-to-point service.

```
R9#show mpls l2transport vc 39

Local intf     Local circuit              Dest address    VC ID      Status
-------------  -------------------------- --------------- ---------- ----------
 Gi1.914        Eth VLAN 914               3.3.3.3         39         UP
```

PE router R9 has a BGP peering with R6 and R10 - the zone's ASBR routers who are serving as route-reflectors.

```
R9#show bgp ipv4 unicast summary
BGP router identifier 9.9.9.9, local AS number 1
BGP table version is 14, main routing table version 14
8 network entries using 1984 bytes of memory
16 path entries using 1920 bytes of memory
2/2 BGP path/bestpath attribute entries using 496 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4560 total bytes of memory
BGP activity 8/0 prefixes, 16/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
 6.6.6.6         4            1    2911    2871       14    0    0 1d19h           8
 10.10.10.10     4            1    2898    2875       14    0    0 1d19h           8
```

In order to establish an LSP from R9 to R3 to provide the EoMPLS service, R9 first needs a route to R3's endpoint loopback. This route is being received from the ASBRs. Notice that 6.6.6.6 and 10.10.10.10 appear in the `cluster-list` attribute of the route - the route-reflectors "stamping" the route as it is reflected with their router-id. The route from R6 is selected as the best path as it has a shorter cluster-list length.

CONTENTS

```
R9#show bgp ipv4 unicast 3.3.3.3/32
BGP routing table entry for 3.3.3.3/32, version 3
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.10.10.10 (metric 2) from 10.10.10.10 (10.10.10.10)
      Origin incomplete, metric 2, localpref 100, valid, internal
       Originator: 2.2.2.2, Cluster list: 10.10.10.10, 6.6.6.6
       mpls labels in/out nolabel/39
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 2
  Local
    6.6.6.6 (metric 2) from 6.6.6.6 (6.6.6.6)
      Origin incomplete, metric 2, localpref 100, valid, internal,  best
       Originator: 2.2.2.2, Cluster list: 6.6.6.6
       mpls labels in/out nolabel/39
      rx pathid: 0, tx pathid: 0x0
```

Notice too that this IPv4 BGP route is also carrying a label. Both R6 and R10 are allocating and advertising label 39 for 3.3.3.3/32. Lastly, notice that the next-hop for this route points to 6.6.6.6 and 10.10.10.10 - the loopbacks of R6 and R10 (the iBGP session's source-interface). R9 sees this as the next-hop because R6 and R10 have `next-hop-self all` configured, and thus are able to change the next-hop of an iBGP route that is being reflected.

R9 knows about 6.6.6.6/32 and 10.10.10.10/32, the next-hops of the BGP routes, via the local IGP.

```
R9#show ip route 6.6.6.6
 Routing entry for 6.6.6.6/32
   Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.6.9.6 on GigabitEthernet1.69, 1d20h ago
  Routing Descriptor Blocks:
  * 10.6.9.6, from 6.6.6.6, 1d20h ago, via GigabitEthernet1.69
      Route metric is 2, traffic share count is 1

R9#show ip route 10.10.10.10
 Routing entry for 10.10.10.10/32
   Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.9.10.10 on GigabitEthernet1.910, 1d20h ago
  Routing Descriptor Blocks:
  * 10.9.10.10, from 10.10.10.10, 1d20h ago, via GigabitEthernet1.910
      Route metric is 2, traffic share count is 1
```

Since R9 is directly connected to R6 and R10, the ASBRs advertise Label-3, or "Pop-Label", via LDP to R9.

```
R9#show mpls forwarding-table 6.6.6.6
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16          Pop Label  6.6.6.6/32       0             Gi1.69     10.6.9.6

R9#show mpls forwarding-table 10.10.10.10
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
21          Pop Label  10.10.10.10/32   0             Gi1.910    10.9.10.10
```

CONTENTS ›

Tying all of these recursive steps together, the CEF output shows us that R9 is routing towards 6.6.6.6, and using label 39, to reach 3.3.3.3/32.

```
R9#show ip cef 3.3.3.3/32 detail
3.3.3.3/32, epoch 2, flags [rib defined all labels]
  1 RR source [no flags]
   recursive via 6.6.6.6 label 39
     nexthop 10.6.9.6 GigabitEthernet1.69
```

Moving to ASBR router R6, we validate that it has active BGP peering with PE routers local to its zone, and with the ASBR routers in the adjacent zone. Additionally R6 has an active peering with the other ASBR servicing the zone, R10, but this peering is not necessary to the Unified MPLS architecture.

```
R6#show bgp ipv4 unicast summary
BGP router identifier 6.6.6.6, local AS number 1
BGP table version is 68, main routing table version 68
8 network entries using 1984 bytes of memory
18 path entries using 2160 bytes of memory
4/4 BGP path/bestpath attribute entries using 992 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5216 total bytes of memory
BGP activity 8/0 prefixes, 18/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
 2.2.2.2         4            1    2946    2941       68    0    0 1d20h           4
8.8.8.8         4            1    2886    2921       68    0    0 1d19h           0
9.9.9.9         4            1    2885    2925       68    0    0 1d19h           0
10.10.10.10     4            1    2913    2912       68    0    0 1d19h           5
 12.12.12.12     4            1    2618    2908       68    0    0 1d19h           5
```

R6 learns about 3.3.3.3/32 from the ASBRs in the adjacent zone, R2 and XR2. R2 advertised a label of 17 and XR2 a label of 16001 for this prefix. Notice both BGP routes seen in the output below have a next-hop of the advertising router - this is caused by the use of `next-hop-self all` on R2 and XR2, and is what allows this entire design to work. The path via ASBR router R2 is selected as the best path.

```
R6#show bgp ipv4 unicast 3.3.3.3/32
BGP routing table entry for 3.3.3.3/32, version 3
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    12.12.12.12 (metric 20) from 12.12.12.12 (12.12.12.12)
      Origin incomplete, metric 2, localpref 100, valid, internal
       mpls labels in/out 39/16001
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  Local
    2.2.2.2 (metric 20) from 2.2.2.2 (2.2.2.2)
      Origin incomplete, metric 2, localpref 100, valid, internal,  best
       mpls labels in/out 39/17
      rx pathid: 0, tx pathid: 0x0
```

Notice that R6 shows an "in" label of 39 in each BGP route. This is the same label that R6 is advertising to R9 and R8. R6 allocates and advertises a new label, instead of simply passing along the label received from R2 and XR2, because it too is setting itself as the next-hop towards the PE routers, causing this label re-generation to take place.

```
R6#show ip bgp labels
   Network          Next Hop      In label/Out label
   2.2.2.2/32       12.12.12.12     17/16000
                    2.2.2.2         17/imp-null
   3.3.3.3/32        12.12.12.12     39/16001
                     2.2.2.2         39/17
   6.6.6.6/32       10.10.10.10     imp-null/20
                    0.0.0.0         imp-null/nolabel
   8.8.8.8/32       10.10.10.10     20/21
                    10.6.8.8        20/nolabel
   9.9.9.9/32       10.10.10.10     21/22
                    10.6.9.9        21/nolabel
   10.10.10.10/32   12.12.12.12     36/16023
                    10.10.10.10     36/imp-null
                    10.6.10.10      36/nolabel
   12.12.12.12/32   12.12.12.12     40/imp-null
                    10.10.10.10     40/40
                    2.2.2.2         40/36
   13.13.13.13/32   12.12.12.12     41/16003
                    2.2.2.2         41/37
```

R6 knows how to reach the next-hop of the BGP route, 2.2.2.2 and 12.12.12.12, as these routes exist in the Core zone IGP from the redistribution done earlier. Without these routes, R6 would not be able to establish a BGP session to R2 and XR2.

```
R6#show ip route 2.2.2.2
 Routing entry for 2.2.2.2/32
   Known via "ospf 2" , distance 110, metric 20, type extern 2, forward metric 3
  Last update from 10.4.6.4 on GigabitEthernet1.46, 1d20h ago
  Routing Descriptor Blocks:
  * 10.5.6.5, from 10.2.11.2, 1d20h ago, via GigabitEthernet1.56
      Route metric is 20, traffic share count is 1
    10.4.6.4, from 10.2.11.2, 1d20h ago, via GigabitEthernet1.46
      Route metric is 20, traffic share count is 1

R6#show ip route 12.12.12.12
 Routing entry for 12.12.12.12/32
   Known via "ospf 2" , distance 110, metric 20, type extern 2, forward metric 3
  Last update from 10.4.6.4 on GigabitEthernet1.46, 1d19h ago
  Routing Descriptor Blocks:
  * 10.5.6.5, from 10.1.12.12, 1d19h ago, via GigabitEthernet1.56
      Route metric is 20, traffic share count is 1
    10.4.6.4, from 10.1.12.12, 1d19h ago, via GigabitEthernet1.46
      Route metric is 20, traffic share count is 1
```

R6 has LDP labels received from R4 and R5 for the next-hops.

```
R6#show mpls forwarding-table 2.2.2.2
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
17          17         2.2.2.2/32       0             Gi1.46     10.4.6.4
            17         2.2.2.2/32       0             Gi1.56     10.5.6.5
R6#show mpls forwarding-table 12.12.12.12
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
40          31         12.12.12.12/32   0             Gi1.46     10.4.6.4
            31         12.12.12.12/32   0             Gi1.56     10.5.6.5
```

Tying it all together, R6 is routing towards ASBR router R2, using BGP label of 17, and LDP label of 17.

```
R6#show ip cef 3.3.3.3/32 detail
3.3.3.3/32, epoch 2, flags [rib defined all labels]
  local label info: global/39
  recursive via 2.2.2.2 label 17
    nexthop 10.4.6.4  GigabitEthernet1.46 label 17
    nexthop 10.5.6.5  GigabitEthernet1.56 label 17
```

R2 has a similar peering scheme as R6:

```
R2#show bgp ipv4 unicast summary
BGP router identifier 2.2.2.2, local AS number 1
BGP table version is 75, main routing table version 75
8 network entries using 1984 bytes of memory
17 path entries using 2040 bytes of memory
4/4 BGP path/bestpath attribute entries using 992 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5056 total bytes of memory
BGP activity 8/0 prefixes, 17/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4            1    2933    2978       75    0    0 1d20h           0
 6.6.6.6         4            1    2967    2973       75    0    0 1d20h           5
 10.10.10.10     4            1    2951    2948       75    0    0 1d20h           4
12.12.12.12     4            1    2642    2933       75    0    0 1d19h           4
13.13.13.13     4            1    2637    2935       75    0    0 1d19h           0
```

R2 is locally sourcing the route to 3.3.3.3/32. It is allocating and advertising a label of 17 for this prefix.

```
R2#show bgp ipv4 unicast 3.3.3.3/32
BGP routing table entry for 3.3.3.3/32, version 3
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     1          2
  Refresh Epoch 1
  Local, (Received from a RR-client)
    12.12.12.12 (metric 2) from 12.12.12.12 (12.12.12.12)
      Origin incomplete, metric 2, localpref 100, valid, internal
      mpls labels in/out 17/16001
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  Local
    10.2.3.3 from 0.0.0.0 (2.2.2.2)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid,  sourced, best
       mpls labels in/out 17/nolabel
      rx pathid: 0, tx pathid: 0x0


R2#show ip bgp labels
   Network          Next Hop      In label/Out label
   2.2.2.2/32       12.12.12.12     imp-null/16000
                    0.0.0.0         imp-null/nolabel
    3.3.3.3/32       12.12.12.12     17/16001
                     10.2.3.3        17/nolabel
   6.6.6.6/32       10.10.10.10     38/20
                    6.6.6.6         38/imp-null
   8.8.8.8/32       10.10.10.10     39/21
                    6.6.6.6         39/20
   9.9.9.9/32       10.10.10.10     40/22
                    6.6.6.6         40/21
   10.10.10.10/32   10.10.10.10     41/imp-null
                    6.6.6.6         41/36
   12.12.12.12/32   12.12.12.12     36/imp-null
                    6.6.6.6         36/40
                    10.2.12.12      36/nolabel
   13.13.13.13/32   12.12.12.12     37/16003
                    10.2.13.13      37/nolabel
```

R2 knows about 3.3.3.3/32 via the local zone's IGP:

```
R2#show ip route 3.3.3.3
Routing entry for 3.3.3.3/32
   Known via "ospf 3" , distance 110, metric 2, type intra area
  Redistributing via bgp 1
   Advertised by bgp 1 route-map REDISTRIBUTE_LOOPBACKS_BGP
  Last update from 10.2.3.3 on GigabitEthernet1.23, 1d20h ago
  Routing Descriptor Blocks:
  * 10.2.3.3, from 3.3.3.3, 1d20h ago,  via GigabitEthernet1.23
      Route metric is 2, traffic share count is 1
```

And is using LDP label of 3, or Pop-Label, as it is directly connected to the final destination:

CONTENTS >

```
R2#show mpls forwarding-table 3.3.3.3

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop

Label      Label      or Tunnel Id     Switched      interface

 17         Pop Label  3.3.3.3/32       244803        Gi1.23     10.2.3.3
```

Recall that in an MPLS network, the BGP next-hop determines where the LSP terminates. In this network, the underlying transport LSPs consists of a series of LSPs which terminate at each zone's ASBRs. For the path between R9 and R3, R9 has an LSP to R6 - using label 39, R6 has an LSP to R2 - using label 17, and R2 is directly connected to R3 so it simply uses PHP. What this means for the services overlayed on top is that they can run their service, whether it is L2VPN, L3VPN, or any other service requiring an end to end LSP, over the Unified MPLS network.

Lets look at the path from R8 to XR3, which is being used for the L3VPN service.

Notice that R8 established a VPNv4 session with XR3 - 13.13.13.13 - over which the L3VPN routes have been exchanged.

```
R8#show bgp vpnv4 unicast all summary

BGP router identifier 8.8.8.8, local AS number 1

BGP table version is 8, main routing table version 8

3 network entries using 768 bytes of memory

3 path entries using 360 bytes of memory

3/2 BGP path/bestpath attribute entries using 792 bytes of memory

4 BGP rrinfo entries using 160 bytes of memory

1 BGP AS-PATH entries using 24 bytes of memory

1 BGP extended community entries using 24 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 2128 total bytes of memory

BGP activity 11/0 prefixes, 19/0 paths, scan interval 60 secs


Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

10.8.14.14      4          100     175     193        8    0    0 02:49:19        1

 13.13.13.13     4            1     209     234        8    0    0 03:26:11        1
```

R8 receives VPNv4 route 7.7.7.7/32, with VPN label 16011, from XR3.

```
R8#show bgp vpnv4 unicast all 7.7.7.7/32
BGP routing table entry for 800:800:7.7.7.7/32, version 7
Paths: (1 available, best #1, table A)
  Advertised to update-groups:
     3
  Refresh Epoch 1
  100, imported path from 1300:1300:7.7.7.7/32 (global)
    13.13.13.13 (metric 2) (via default) from 13.13.13.13 (13.13.13.13)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:83:83
      mpls labels in/out nolabel/16011
      rx pathid: 0, tx pathid: 0x0
BGP routing table entry for 1300:1300:7.7.7.7/32, version 8
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  Refresh Epoch 1
  100
    13.13.13.13 (metric 2) (via default) from 13.13.13.13 (13.13.13.13)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:83:83
      mpls labels in/out nolabel/16011
      rx pathid: 0, tx pathid: 0x0
```

R8 installs the BGP route in the VRF:

```
R8#show ip route vrf A 7.7.7.7

 Routing Table: A
Routing entry for 7.7.7.7/32
  Known via "bgp 1", distance 200, metric 0
  Tag 100, type internal
  Last update from 13.13.13.13 02:44:45 ago
  Routing Descriptor Blocks:
  * 13.13.13.13 (default), from 13.13.13.13, 02:44:45 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: 16011
      MPLS Flags: MPLS Required
```

The next-hop for the VPNv4 route is 13.13.13.13, which R8 is learning via BGP from the ASBRs. R8 selects the path via R6 as the best path, again due to shorter cluster-list. Note that R6 is advertising Label 41 for 13.13.13.13/32.

CONTENTS ›

```
R8#show ip route 13.13.13.13
Routing entry for 13.13.13.13/32
  Known via "bgp 1", distance 200, metric 2, type internal
  Last update from 6.6.6.6 1d20h ago
  Routing Descriptor Blocks:
  * 6.6.6.6, from 6.6.6.6, 1d20h ago
      Route metric is 2, traffic share count is 1
      AS Hops 0
      MPLS label: 41

R8#show bgp ipv4 unicast 13.13.13.13/32
BGP routing table entry for 13.13.13.13/32, version 9
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.10.10.10 (metric 2) from 10.10.10.10 (10.10.10.10)
      Origin incomplete, metric 2, localpref 100, valid, internal
      Originator: 2.2.2.2, Cluster list: 10.10.10.10, 6.6.6.6
      mpls labels in/out nolabel/41
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  Local
    6.6.6.6 (metric 2) from 6.6.6.6 (6.6.6.6)
      Origin incomplete, metric 2, localpref 100, valid, internal, best
      Originator: 2.2.2.2, Cluster list: 6.6.6.6
      mpls labels in/out nolabel/41
      rx pathid: 0, tx pathid: 0x0
```

Looking at CEF, we can see that R8 has a label to get to the VPNv4 next-hop - the service endpoint, PE router XR3.

```
R8#show ip cef 13.13.13.13/32 detail
13.13.13.13/32, epoch 2, flags [rib defined all labels]
  1 RR source [no flags]
  recursive via 6.6.6.6 label 41
    nexthop 10.6.8.6 GigabitEthernet1.68
```

And we can validate that we have an underlying "stitched" LSP to get to XR3, consisting of the LSP from R8 to R6, from R6 to R2, and from R2 to XR3.

```
R8#traceroute 13.13.13.13 source 8.8.8.8
Type escape sequence to abort.
Tracing the route to 13.13.13.13
VRF info: (vrf in name/id, vrf out name/id)
  1 10.6.8.6 [MPLS: Label 41 Exp 0] 14 msec 16 msec 16 msec
  2 10.4.6.4 [MPLS: Labels 17/37 Exp 0] 20 msec 18 msec 30 msec
  3 10.4.11.11 [MPLS: Labels 16012/37 Exp 0] 25 msec 30 msec 31 msec
  4 10.2.11.2 [MPLS: Label 37 Exp 0] 15 msec 16 msec 28 msec
  5 10.2.13.13 29 msec *  15 msec
```

By tracing the path from the customers point of view, we can see that label 16011, the VPNv4 label advertised by XR3 for 7.7.7.7/32 is maintained end-to-end.

CONTENTS

```
RP/0/0/CPU0:XR4#traceroute 7.7.7.7 source 14.14.14.14

Mon Jul 13 00:07:51.425 UTC


Type escape sequence to abort.

Tracing the route to 7.7.7.7


 1  10.8.14.8 0 msec  0 msec  0 msec

 2  10.6.8.6 [MPLS: Labels 41/16011  Exp 0] 9 msec  19 msec  19 msec

 3  10.4.6.4 [MPLS: Labels 17/37/16011  Exp 0] 19 msec  29 msec  19 msec

 4  10.4.11.11 [MPLS: Labels 16012/37/16011  Exp 0] 9 msec  9 msec  19 msec

 5  10.2.11.2 [MPLS: Labels 37/16011  Exp 0] 29 msec  49 msec  19 msec

 6  10.2.13.13 [MPLS: Label  16011  Exp 0] 49 msec  29 msec  19 msec

 7  10.7.13.7 19 msec  *  19 msec
```

Within each IGP zone an LDP based LSP is established between the PE/ASBRs and the "exit" ASBRs acting as in-lined route-reflectors. Then an iBGP + Label based LSP is established between PE <-> PE, riding on top of the LDP based LSPs within each zone its traversing. A service LSP can then be established on top of the iBGP + Label LSP, to provide any MPLS based service towards the customers.

```
RP/0/0/CPU0:XR3#sh cef vrf A 14.14.14.14/32

Wed Jul 15 00:26:00.857 UTC

%Prefix not found or IP is not running. VRF **nVSatellite.


VRF: A

  _

14.14.14.14/32, version 6, internal 0x14004001 0x0 (ptr 0xa0ec2af4) [1], 0x0 (0x0), 0x208 (0xa1480168)

 Updated Jul 15 00:17:31.472

 Prefix Len 32, traffic index 0, precedence n/a, priority 3

   via 8.8.8.8, 3 dependencies, recursive [flags 0x6000]

    path-idx 0 NHID 0x0 [0xa14e8bf4 0x0]

    next hop VRF - 'default', table - 0xe0000000

     next hop 8.8.8.8 via 16008/0/21

      next hop 10.12.13.12/32 Gi0/0/0/0.1213 labels imposed { ImplNull 16023 27 }

%Prefix not found or IP is not running. VRF default.
```

CONTENTS ›

The following diagram illustrates that we can have an arbitrary amount of IGP zones within the Unified MPLS network, as long as each zone's upstream ASBRs set themselves as the next-hop for the BGP+Label prefixes receives from other zones. An additional zone is added to out original diagram, called "Core". A pictorial representation of the hierarchical LSPs is also depicted.

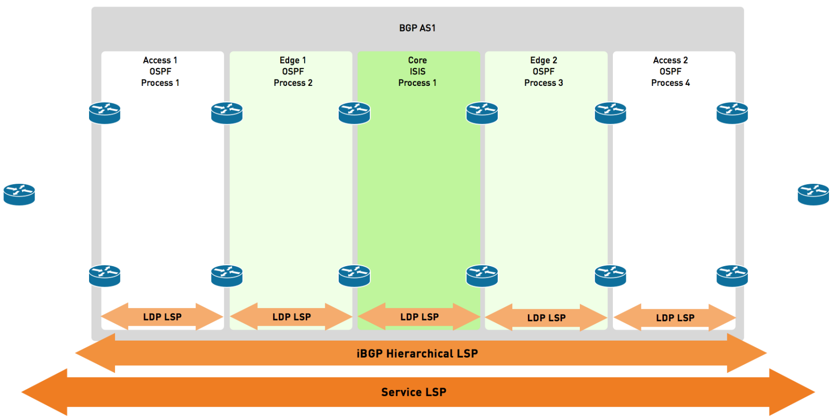

Lets investigate a break in the data-plane, by removing the "next-hop-self all" feature from one of the route-reflectors towards a PE. Shut down Gig1 on R10 to remove redundant paths out of the Access Zone where PE routers R8 and R9 are located.

```
R10(config)#interface Gig1
R10(config-if)#shut
R10(config-if)#end

R6(config)#router bgp 1
R6(config-router)# template peer-policy BREAK_LSP
R6(config-router-ptmp)#  send-label
R6(config-router-ptmp)#  route-reflector-client
R6(config-router-ptmp)# exit-peer-policy
R6(config-router)# address-family ipv4
R6(config-router-af)#no  neighbor 8.8.8.8 inherit peer-policy iBGP_RR_Client

%BGP-5-NBR_RESET: Neighbor 8.8.8.8 reset (RR client config change)
%BGP-5-ADJCHANGE: neighbor 8.8.8.8 Down RR client config change
%BGP_SESSION-5-ADJCHANGE: neighbor 8.8.8.8 IPv4 Unicast topology base removed from session  RR client config change
%BGP-5-ADJCHANGE: neighbor 8.8.8.8 Up

R6(config-router-af)#  neighbor 8.8.8.8 inherit peer-policy BREAK_LSP
R6(config-router-af)#
%BGP-5-ADJCHANGE: neighbor 8.8.8.8 Down Capability changed
%BGP_SESSION-5-ADJCHANGE: neighbor 8.8.8.8 IPv4 Unicast topology base removed from session  Capability changed
%BGP-5-ADJCHANGE: neighbor 8.8.8.8 Up

R6#clear ip bgp *
```

Notice that R8 receives the updates from R8, but the next-hop is still R2. There is no reachability in the current IGP zone for 2.2.2.2, thus the path is not usable. Changing the BGP next-hop on the route-reflectors is what faciliates keeping the IGP zones slim, as no unnecessary routes to reach other next-hops have to be leaked between the IGP zones.

```
R8#show ip bgp 13.13.13.13/32
BGP routing table entry for 13.13.13.13/32, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    2.2.2.2 (inaccessible) from 6.6.6.6 (6.6.6.6)
      Origin incomplete, metric 2, localpref 100, valid, internal
      Originator: 2.2.2.2, Cluster list: 6.6.6.6
      mpls labels in/out nolabel/40
      rx pathid: 0, tx pathid: 0
```

Also, try using the `next-hop-self` command without the `all` argument.

› CONTENTS

```
R6(config)#router bgp 1

R6(config-router)# template peer-policy BREAK_LSP

R6(config-router-ptmp)#next-hop-self

R6(config-router-ptmp)#do clear ip bgp * out

R6(config-router-ptmp)#
```

```
R8#debug ip bgp updates

BGP updates debugging is on for address family: IPv4 Unicast


R8#clear ip bgp 6.6.6.6 in

R8#

BGP: nbr_topo global 6.6.6.6 IPv4 Unicast:base (0x7FCC40220C80:1) rcvd Refresh Start-of-RIB

BGP: nbr_topo global 6.6.6.6 IPv4 Unicast:base (0x7FCC40220C80:1) refresh_epoch is 4

BGP: nbr_topo global 6.6.6.6 IPv4 Unicast:base (0x7FCC40220C80:1) rcvd Refresh Start-of-RIB

BGP: nbr_topo global 6.6.6.6 IPv4 Unicast:base (0x7FCC40220C80:1) refresh_epoch is 5

BGP(0): 6.6.6.6 rcvd UPDATE w/ attr:  nexthop 2.2.2.2 , origin ?, localpref 100, metric 2, originator 2.2.2.2, clusterlist 6.6.6.6

BGP(0): 6.6.6.6 rcvd

R8#3.3.3.3/32, label 36...duplicate ignored

BGP(0): 6.6.6.6 rcvd 13.13.13.13/32, label 40...duplicate ignored

BGP(0): 6.6.6.6 rcvd UPDATE w/ attr:  nexthop 12.12.12.12 , origin ?, localpref 100, metric 0, originator 12.12.12.12, clusterlist 6.6.6.6

BGP(0): 6.6.6.6 rcvd 12.12.12.12/32, label 3...duplicate ignored

BGP(0): 6.6.6.6 rcvd UPDATE w/ attr:  nexthop 2.2.2.2 , origin ?, localpref 100, metric 0, originator 2.2.2.2, clusterlist 6.6.6.6

BGP(0): 6.6.6.6 rcvd 2.2.2.2/32, label 3...duplicate ignored

BGP(0)

R8#: 6.6.6.6 rcvd UPDATE w/ attr: nexthop 6.6.6.6, origin ?, localpref 100, metric 2

BGP(0): 6.6.6.6 rcvd 8.8.8.8/32, label 46...duplicate ignored

BGP(0): 6.6.6.6 rcvd 9.9.9.9/32, label 42

BGP(0): 6.6.6.6 rcvd UPDATE w/ attr: nexthop 6.6.6.6, origin ?, localpref 100, metric 0

BGP(0): 6.6.6.6 rcvd 6.6.6.6/32, label 3...duplicate ignored

BGP: nbr_topo global 6.6.6.6 IPv4 Unicast:base (0x7FCC40220C80:1) rcvd Refresh End-of-RIB

BGP(0): Revise route installing 1 of 1 routes for 9.9.9.9/32 -> 6.6.6.6(global)

R8#to main IP table
```


Without the `all` argument, iBGP routes being reflected do not have their the next-hops changed. Note that the same is true in IOS-XR - with the `ibgp policy out enforce-modifications` command.

```
R8#show ip bgp

BGP table version is 12, local router ID is 8.8.8.8

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

 * i 2.2.2.2/32       2.2.2.2                  0    100      0 ?

 * i 3.3.3.3/32       2.2.2.2                  2    100      0 ?

 r>i 6.6.6.6/32       6.6.6.6                  0    100      0 ?

 r>i 8.8.8.8/32       6.6.6.6                  2    100      0 ?

 r>i 9.9.9.9/32       6.6.6.6                  2    100      0 ?

 * i 12.12.12.12/32   12.12.12.12              0    100      0 ?

 * i 13.13.13.13/32   2.2.2.2                  2    100      0 ?
```

CONTENTS