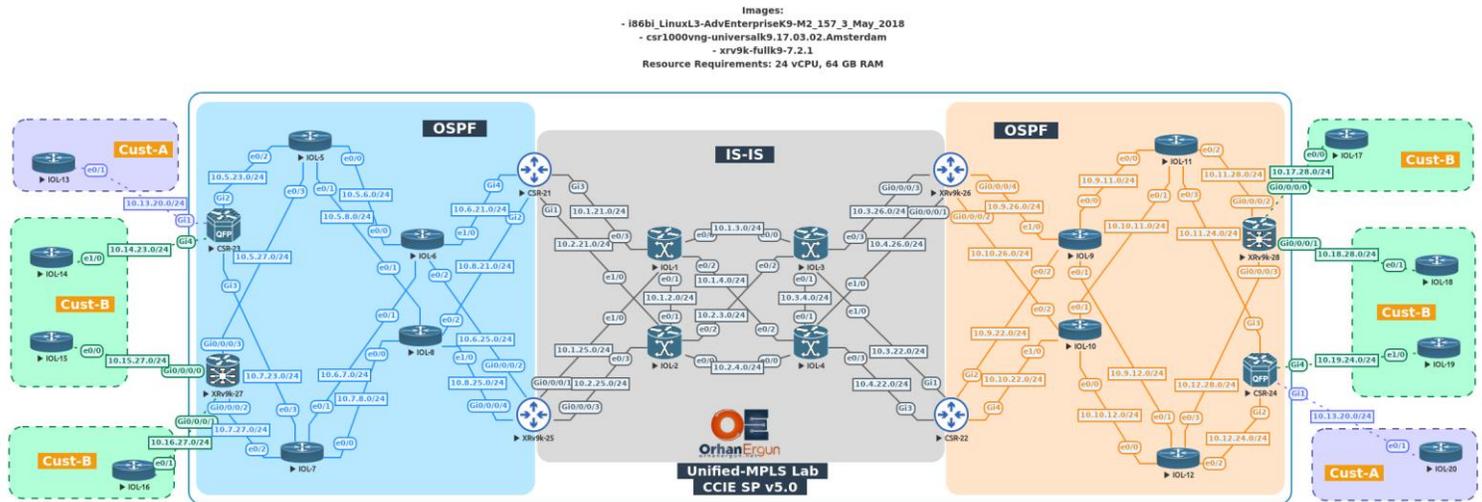


Unified/Seamless MPLS Lab

Topology:



BGP AF/SAF: IPv4 Unicast, VPNv4 Unicast, L2VPN VPLS

PE-CE Routing Protocol:

- BGP

VPLS: Martini with BGP Auto-Discovery



Task 01:

- Configure each routing domain as per the diagram
- Enable MPLS LDP for each IGP domain
- DO NOT do any redistribution between routing domains



Solution:

What is the idea behind Unified or Seamless MPLS framework?

Please take a look at the topology, You can see three different routing domains (IS-IS in the core, and two OSPFs in the Aggregation-Access network). The real Mobile Operator is not being designed this way (to have different IGPs in the entire Autonomous System), but in order to make Unified MPLS understandable (Because it is a kind of complex framework) we have designed the lab this way.

Unified MPLS framework is common in mobile operators. They have different parts in their network hierarchy and for each part they have a separate management team. For example a team manages the CORE, another team manages Aggregation Layer in some City or Province, etc... and also they have a team managing Access layer such as RAN (Radio Access Network).

As you realized from above words, A large mobile operator have so many parts to be managed.

In this lab we will call different parts: a routing domain or an Autonomous System.

The entire SP network in a single Autonomous System, but in Unified MPLS framework we divide it to multiple Autonomous System that has a different team managing that AS.

Take a look at the topology one more time: We have a routing domain in the Core (IGP: IS-IS), in the left side we have another routing domain (IGP: OSPF) and finally in the right side we have another routing domain (IGP: OSPF).

So we can device this Service Provider AS to 3 different Autonomous Systems:

- Middle AS (IS-IS)
- Left AS (OSPF)
- Right AS (OSPF)

NOTE: This is just a lab, we are trying to learn and understand how we can configure Seamless or Unified MPLS framework, so it is not designed based on the real Mobile Operator network.

Most service Providers (let's say almost all of them!) are using MPLS in their network, they can provide different services to the customers, as an Internet Access, MPLS L3VPNs or L2VPNs, etc... .

They need to have end-to-end LSPs (Label Switched Paths) between their access devices (or let's call them PE devices). In this topology our PE devices are:

- CSR-23
- XRv9k-27
- CSR-24
- XRv9k-28

In order to provide Loopback to Loopback LSP between these PE devices, one way is to do redistribution between these different routing domains:

- Two way redistribution between left OSPF and middle IS-IS
- Two way redistribution between right OSPF and middle IS-IS

So this way we can leak host routes (Loopback 0 interfaces /32 prefixes) and PE devices can form IPv4 Unicast, VPNv4 Unicast, L2VPN EVPN, L2VPN VPLS and so on... .

But there is something wrong with this kind of implementation. The managing of network becomes very difficult and it becomes very complex (Redistribution has its own pros and cons etc...), beside that, we need to deal with different routing protocols in our network.

The idea of Unified MPLS is there to help us to reduce this overcomplication.

Unified MPLS or in other words Seamless MPLS says:

Your different Islands/Routing Domains (Or whatever you call it) can be isolated from each other and you can use a unified method to build the PE to PE LSPs.

How it is possible? Thanks to the BGP we can achieve this goal.

Each Routing domain has the responsibility of providing LSPs between the devices inside that routing domain.

As an example from the topology:

Core devices (IGP: IS-IS) need to only provide LSPs inside that routing domain, no where else!

IOL-1, IOL-2, IOL-3 and IOL-4 DO NOT need to know about CSR-23, XRv9k-27 loopback IP addresses (They simply do not care).

This was a brief introduction to the seamless MPLS framework. As we continue doing the tasks you will totally understand what it is and how it works.

Let's start configuring the first task.

In this task we need to configure our three different routing domains (or let's call it Autonomous Systems). This is only a simple routing configuration on each router and enabling MPLS LDP on them:

IOL-1:

```
hostname IOL-1
!
clock timezone +0330 3 30
!
no ip icmp rate-limit unreachable
!
no ip domain lookup
ip cef
ip cef load-sharing algorithm include-ports source destination
no ipv6 cef
!
mpls label range 10000 10999
!
interface Loopback0
 ip address 10.255.255.1 255.255.255.255
 isis tag 1
!
interface Ethernet0/0
 ip address 10.1.3.1 255.255.255.0
 ip router isis CORE
 duplex auto
 isis network point-to-point
 isis metric 100
 isis three-way-handshake ietf
!
interface Ethernet0/1
 ip address 10.1.2.1 255.255.255.0
 ip router isis CORE
 duplex auto
 isis network point-to-point
 isis metric 100
 isis three-way-handshake ietf
!
interface Ethernet0/2
```

```
ip address 10.1.4.1 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 100
isis three-way-handshake ietf
!
interface Ethernet0/3
ip address 10.1.21.1 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 1000
isis three-way-handshake ietf
!
interface Ethernet1/0
ip address 10.1.25.1 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 1000
isis three-way-handshake ietf
!
router isis CORE
mpls ldp autoconfig level-2
net 49.0000.0000.0000.0001.00
is-type level-2-only
advertise passive-only
ispf level-2
metric-style wide
log-adjacency-changes all
passive-interface Loopback0
!
ip forward-protocol nd
```

```
!  
mpls ldp router-id Loopback0  
!  
control-plane  
!  
  
IOL-2:  
hostname IOL-2  
!  
boot-start-marker  
boot-end-marker  
!  
clock timezone +0330 3 30  
!  
no ip icmp rate-limit unreachable  
!  
no ip domain lookup  
ip cef  
ip cef load-sharing algorithm include-ports source destination  
no ipv6 cef  
!  
mpls label range 20000 20999  
!  
interface Loopback0  
 ip address 10.255.255.2 255.255.255.255  
 isis tag 1  
!  
interface Ethernet0/0  
 ip address 10.2.4.2 255.255.255.0  
 ip router isis CORE  
 duplex auto  
 isis network point-to-point  
 isis metric 100  
 isis three-way-handshake ietf  
!
```

```
interface Ethernet0/1
 ip address 10.1.2.2 255.255.255.0
 ip router isis CORE
 duplex auto
 isis network point-to-point
 isis metric 100
 isis three-way-handshake ietf
!
interface Ethernet0/2
 ip address 10.2.3.2 255.255.255.0
 ip router isis CORE
 duplex auto
 isis network point-to-point
 isis metric 100
 isis three-way-handshake ietf
!
interface Ethernet0/3
 ip address 10.2.25.2 255.255.255.0
 ip router isis CORE
 duplex auto
 isis network point-to-point
 isis metric 1000
 isis three-way-handshake ietf
!
interface Ethernet1/0
 ip address 10.2.21.2 255.255.255.0
 ip router isis CORE
 duplex auto
 isis network point-to-point
 isis metric 1000
 isis three-way-handshake ietf
!
router isis CORE
 mpls ldp autoconfig level-2
```

```
net 49.0000.0000.0000.0002.00
is-type level-2-only
advertise passive-only
ispf level-2
metric-style wide
log-adjacency-changes all
passive-interface Loopback0
!
ip forward-protocol nd
!
mpls ldp router-id Loopback0
!

IOL-3:
hostname IOL-3
!
clock timezone +0330 3 30
!
no ip icmp rate-limit unreachable
!
no ip domain lookup
ip cef
ip cef load-sharing algorithm include-ports source destination
no ipv6 cef
!
mpls label range 30000 30999
!
interface Loopback0
 ip address 10.255.255.3 255.255.255.255
 isis tag 1
!
interface Ethernet0/0
 ip address 10.1.3.3 255.255.255.0
 ip router isis CORE
```

```
duplex auto
isis network point-to-point
isis metric 100
isis three-way-handshake ietf
!
interface Ethernet0/1
ip address 10.3.4.3 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 100
isis three-way-handshake ietf
!
interface Ethernet0/2
ip address 10.2.3.3 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 100
isis three-way-handshake ietf
!
interface Ethernet0/3
ip address 10.3.26.3 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 1000
isis three-way-handshake ietf
!
interface Ethernet1/0
ip address 10.3.22.3 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
```

```
isis metric 1000
isis three-way-handshake ietf
!
router isis CORE
mpls ldp autoconfig level-2
net 49.0000.0000.0000.0003.00
is-type level-2-only
advertise passive-only
ispf level-2
metric-style wide
log-adjacency-changes all
passive-interface Loopback0
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
ipv6 ioam timestamp
!
!
mpls ldp router-id Loopback0
!

IOL-4:
hostname IOL-4
!
clock timezone +0330 3 30
no ip icmp rate-limit unreachable
!
no ip domain lookup
ip cef
ip cef load-sharing algorithm include-ports source destination
no ipv6 cef
```

```
!  
multilink bundle-name authenticated  
mpls label range 40000 40999  
!  
interface Loopback0  
 ip address 10.255.255.4 255.255.255.255  
 isis tag 1  
!  
interface Ethernet0/0  
 ip address 10.2.4.4 255.255.255.0  
 ip router isis CORE  
 duplex auto  
 isis network point-to-point  
 isis metric 100  
 isis three-way-handshake ietf  
!  
interface Ethernet0/1  
 ip address 10.3.4.4 255.255.255.0  
 ip router isis CORE  
 duplex auto  
 isis network point-to-point  
 isis metric 100  
 isis three-way-handshake ietf  
!  
interface Ethernet0/2  
 ip address 10.1.4.4 255.255.255.0  
 ip router isis CORE  
 duplex auto  
 isis network point-to-point  
 isis metric 100  
 isis three-way-handshake ietf  
!  
interface Ethernet0/3  
 ip address 10.4.22.4 255.255.255.0
```

```
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 1000
isis three-way-handshake ietf
!
interface Ethernet1/0
ip address 10.4.26.4 255.255.255.0
ip router isis CORE
duplex auto
isis network point-to-point
isis metric 1000
isis three-way-handshake ietf
!
router isis CORE
mpls ldp autoconfig level-2
net 49.0000.0000.0000.0004.00
is-type level-2-only
advertise passive-only
ispf level-2
metric-style wide
log-adjacency-changes all
passive-interface Loopback0
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
ipv6 ioam timestamp
!
!
mpls ldp router-id Loopback0
```

CSR-21:

```
interface Loopback0
 ip address 10.255.255.21 255.255.255.255
 ip ospf 1000 area 0
 isis tag 1
!
interface GigabitEthernet1
 ip address 10.2.21.21 255.255.255.0
 ip router isis CORE
 negotiation auto
 no mop enabled
 no mop sysid
 isis network point-to-point
 isis metric 1000 level-1
 isis metric 1000 level-2
!
interface GigabitEthernet2
 ip address 10.8.21.21 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 negotiation auto
 no mop enabled
 no mop sysid
!
interface GigabitEthernet3
 ip address 10.1.21.21 255.255.255.0
 ip router isis CORE
 negotiation auto
 no mop enabled
 no mop sysid
 isis network point-to-point
 isis metric 1000 level-1
 isis metric 1000 level-2
!
interface GigabitEthernet4
```

```
ip address 10.6.21.21 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
negotiation auto
no mop enabled
no mop sysid
!
router ospf 1000
router-id 10.255.255.21
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
mpls ldp autoconfig
!
router isis CORE
net 49.0000.0000.0000.0021.00
is-type level-2-only
advertise passive-only
metric-style wide
log-adjacency-changes all
passive-interface Loopback0
mpls ldp autoconfig level-2
!
mpls ldp router-id Loopback0
```

As you realized, CSR-21 is participating in both routing domain (Middle one: IS-IS as the IGP and Left one: OSPF as the IGP) and there is no redistribution between these two routing domain (IS-IS <-> OSPF). Later on we will call these devices: **ASBRs**.

Let's also configure the other three ASBRs:

CSR-22:

```
interface Loopback0
ip address 10.255.255.22 255.255.255.255
ip ospf 1000 area 0
isis tag 1
!
```

```
interface GigabitEthernet1
 ip address 10.3.22.22 255.255.255.0
 ip router isis CORE
 negotiation auto
 no mop enabled
 no mop sysid
 isis network point-to-point
 isis metric 1000 level-1
 isis metric 1000 level-2
!
interface GigabitEthernet2
 ip address 10.9.22.22 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 negotiation auto
 no mop enabled
 no mop sysid
!
interface GigabitEthernet3
 ip address 10.4.22.22 255.255.255.0
 ip router isis CORE
 negotiation auto
 no mop enabled
 no mop sysid
 isis network point-to-point
 isis metric 1000 level-1
 isis metric 1000 level-2
!
interface GigabitEthernet4
 ip address 10.10.22.22 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 negotiation auto
 no mop enabled
```

```
no mop sysid
!
router ospf 1000
  router-id 10.255.255.22
  log-adjacency-changes detail
  auto-cost reference-bandwidth 1000
  prefix-suppression
  passive-interface Loopback0
  mpls ldp autoconfig
!
router isis CORE
  net 49.0000.0000.0000.0022.00
  is-type level-2-only
  advertise passive-only
  metric-style wide
  log-adjacency-changes all
  passive-interface Loopback0
  mpls ldp autoconfig level-2
!
mpls ldp router-id Loopback0
!
```

XRv9k-25:

```
router ospf 1000
  log adjacency changes detail
  router-id 10.255.255.25
  mpls ldp auto-config
  prefix-suppression
  auto-cost reference-bandwidth 1000
  address-family ipv4
  area 0
    interface Loopback0
      passive enable
  !
  interface GigabitEthernet0/0/0/2
```

```
network point-to-point
!
interface GigabitEthernet0/0/0/4
network point-to-point
!
!
!
router isis CORE
is-type level-2-only
net 49.0000.0000.0000.0025.00
log adjacency changes
address-family ipv4 unicast
metric-style wide
advertise passive-only
mpls ldp auto-config
!
interface Loopback0
passive
address-family ipv4 unicast
tag 1
!
!
interface GigabitEthernet0/0/0/1
point-to-point
address-family ipv4 unicast
metric 1000
!
!
interface GigabitEthernet0/0/0/3
point-to-point
address-family ipv4 unicast
metric 1000
!
!
!
```

```
mpls ldp
  log
  neighbor
!
  router-id 10.255.255.25
!
mpls label range table 0 25000 25999
!
commit
!
```

XRv9k-26:

```
router ospf 1000
  log adjacency changes detail
  router-id 10.255.255.26
  mpls ldp auto-config
  prefix-suppression
  auto-cost reference-bandwidth 1000
  address-family ipv4
  area 0
    interface Loopback0
      passive enable
    !
    interface GigabitEthernet0/0/0/2
      network point-to-point
    !
    interface GigabitEthernet0/0/0/4
      network point-to-point
    !
  !
!
router isis CORE
  is-type level-2-only
  net 49.0000.0000.0000.0026.00
  log adjacency changes
```

```
address-family ipv4 unicast
  metric-style wide
  advertise passive-only
  mpls ldp auto-config
!
interface Loopback0
  passive
  address-family ipv4 unicast
    tag 1
  !
!
interface GigabitEthernet0/0/0/1
  point-to-point
  address-family ipv4 unicast
    metric 1000
  !
!
interface GigabitEthernet0/0/0/3
  point-to-point
  address-family ipv4 unicast
    metric 1000
  !
!
!
mpls ldp
  log
  neighbor
!
router-id 10.255.255.26
!
mpls label range table 0 26000 26999
!
commit
!
```

IOL-6:

```
interface Loopback0
  ip address 10.255.255.6 255.255.255.255
  ip ospf 1000 area 0
!
interface Ethernet0/0
  bandwidth 100000
  ip address 10.5.6.6 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/1
  bandwidth 100000
  ip address 10.6.7.6 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/2
  bandwidth 1000000
  ip address 10.6.25.6 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/3
  no ip address
  shutdown
  duplex auto
!
interface Ethernet1/0
  bandwidth 1000000
  ip address 10.6.21.6 255.255.255.0
  ip ospf network point-to-point
```

```
ip ospf 1000 area 0
duplex auto
!
router ospf 1000
mpls ldp autoconfig
router-id 10.255.255.6
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
!
mpls ldp router-id Loopback0
!
```

IOL-8:

```
interface Loopback0
ip address 10.255.255.8 255.255.255.255
ip ospf 1000 area 0
!
interface Ethernet0/0
bandwidth 100000
ip address 10.7.8.8 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/1
bandwidth 100000
ip address 10.5.8.8 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/2
bandwidth 1000000
```

```
ip address 10.8.21.8 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/3
no ip address
shutdown
duplex auto
!
interface Ethernet1/0
bandwidth 1000000
ip address 10.8.25.8 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
router ospf 1000
mpls ldp autoconfig
router-id 10.255.255.8
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
!
mpls ldp router-id Loopback0
!

IOL-5:
interface Loopback0
ip address 10.255.255.5 255.255.255.255
ip ospf 1000 area 0
!
interface Ethernet0/0
```

```
bandwidth 100000
ip address 10.5.6.5 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/1
bandwidth 100000
ip address 10.5.8.5 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/2
bandwidth 10000
ip address 10.5.23.5 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/3
bandwidth 10000
ip address 10.5.27.5 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
router ospf 1000
mpls ldp autoconfig
router-id 10.255.255.5
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
```

```
!  
mpls ldp router-id Loopback0  
!  
  
IOL-7:  
  
interface Loopback0  
  ip address 10.255.255.7 255.255.255.255  
  ip ospf 1000 area 0  
!  
interface Ethernet0/0  
  bandwidth 100000  
  ip address 10.7.8.7 255.255.255.0  
  ip ospf network point-to-point  
  ip ospf 1000 area 0  
  duplex auto  
!  
interface Ethernet0/1  
  bandwidth 100000  
  ip address 10.6.7.7 255.255.255.0  
  ip ospf network point-to-point  
  ip ospf 1000 area 0  
  duplex auto  
!  
interface Ethernet0/2  
  bandwidth 10000  
  ip address 10.7.27.7 255.255.255.0  
  ip ospf network point-to-point  
  ip ospf 1000 area 0  
  duplex auto  
!  
interface Ethernet0/3  
  bandwidth 10000  
  ip address 10.7.23.7 255.255.255.0  
  ip ospf network point-to-point  
  ip ospf 1000 area 0
```

```
duplex auto
!
router ospf 1000
 mpls ldp autoconfig
 router-id 10.255.255.7
 log-adjacency-changes detail
 auto-cost reference-bandwidth 1000
 prefix-suppression
 passive-interface Loopback0
!
mpls ldp router-id Loopback0
!
```

CSR-23:

```
interface Loopback0
 ip address 10.255.255.23 255.255.255.255
 ip ospf 1000 area 0
!
interface GigabitEthernet2
 bandwidth 10000
 ip address 10.5.23.23 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 negotiation auto
 no mop enabled
 no mop sysid
!
interface GigabitEthernet3
 bandwidth 10000
 ip address 10.7.23.23 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 negotiation auto
 no mop enabled
 no mop sysid
```

```
!  
router ospf 1000  
  router-id 10.255.255.23  
  log-adjacency-changes detail  
  auto-cost reference-bandwidth 1000  
  prefix-suppression  
  passive-interface Loopback0  
  mpls ldp autoconfig  
!  
mpls ldp router-id Loopback0  
!
```

XRv9k-27:

```
router ospf 1000  
  log adjacency changes detail  
  router-id 10.255.255.27  
  mpls ldp auto-config  
  prefix-suppression  
  auto-cost reference-bandwidth 1000000  
  area 0  
    interface Loopback0  
      passive enable  
    !  
    interface GigabitEthernet0/0/0/2  
      network point-to-point  
    !  
    interface GigabitEthernet0/0/0/3  
      network point-to-point  
    !  
  !  
  !  
  !  
mpls ldp  
  log  
  neighbor  
!
```

```
router-id 10.255.255.27
!
mpls label range table 0 27000 27999

IOL-9:
interface Loopback0
 ip address 10.255.255.9 255.255.255.255
 ip ospf 1000 area 0
!
interface Ethernet0/0
 bandwidth 100000
 ip address 10.9.11.9 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 duplex auto
!
interface Ethernet0/1
 bandwidth 100000
 ip address 10.9.12.9 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 duplex auto
!
interface Ethernet0/2
 bandwidth 1000000
 ip address 10.9.22.9 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 duplex auto
!
interface Ethernet0/3
 no ip address
 shutdown
 duplex auto
!
```

```
interface Ethernet1/0
  bandwidth 1000000
  ip address 10.9.26.9 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
router ospf 1000
  mpls ldp autoconfig
  router-id 10.255.255.9
  log-adjacency-changes detail
  auto-cost reference-bandwidth 1000
  prefix-suppression
  passive-interface Loopback0
!
mpls ldp router-id Loopback0
!
```

IOL-10:

```
interface Loopback0
  ip address 10.255.255.10 255.255.255.255
  ip ospf 1000 area 0
!
interface Ethernet0/0
  bandwidth 100000
  ip address 10.10.12.10 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/1
  bandwidth 100000
  ip address 10.10.11.10 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
```

```
duplex auto
!
interface Ethernet0/2
bandwidth 1000000
ip address 10.10.26.10 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
interface Ethernet0/3
no ip address
shutdown
duplex auto
!
!
interface Ethernet1/0
bandwidth 1000000
ip address 10.10.22.10 255.255.255.0
ip ospf network point-to-point
ip ospf 1000 area 0
duplex auto
!
!
router ospf 1000
mpls ldp autoconfig
router-id 10.255.255.10
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
!
!
!
mpls ldp router-id Loopback0
```

IOL-11:

```
interface Loopback0
  ip address 10.255.255.11 255.255.255.255
  ip ospf 1000 area 0
!
interface Ethernet0/0
  bandwidth 100000
  ip address 10.9.11.11 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/1
  bandwidth 100000
  ip address 10.10.11.11 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/2
  bandwidth 10000
  ip address 10.11.28.11 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
interface Ethernet0/3
  bandwidth 10000
  ip address 10.11.24.11 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
router ospf 1000
  mpls ldp autoconfig
```

```
router-id 10.255.255.11
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
!
mpls ldp router-id Loopback0
!
```

IOL-12:

```
interface Loopback0
 ip address 10.255.255.12 255.255.255.255
 ip ospf 1000 area 0
!
interface Ethernet0/0
 bandwidth 100000
 ip address 10.10.12.12 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 duplex auto
!
interface Ethernet0/1
 bandwidth 100000
 ip address 10.9.12.12 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 duplex auto
!
interface Ethernet0/2
 bandwidth 10000
 ip address 10.12.24.12 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1000 area 0
 duplex auto
!
```

```
interface Ethernet0/3
  bandwidth 10000
  ip address 10.12.28.12 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  duplex auto
!
router ospf 1000
  mpls ldp autoconfig
  router-id 10.255.255.12
  log-adjacency-changes detail
  auto-cost reference-bandwidth 1000
  prefix-suppression
  passive-interface Loopback0
!
mpls ldp router-id Loopback0
!
```

CSR-24:

```
interface Loopback0
  ip address 10.255.255.24 255.255.255.255
  ip ospf 1000 area 0
!
interface GigabitEthernet2
  bandwidth 10000
  ip address 10.12.24.24 255.255.255.0
  ip ospf network point-to-point
  ip ospf 1000 area 0
  negotiation auto
  no mop enabled
  no mop sysid
!
interface GigabitEthernet3
  bandwidth 10000
  ip address 10.11.24.24 255.255.255.0
```

```
ip ospf network point-to-point
ip ospf 1000 area 0
negotiation auto
no mop enabled
no mop sysid
!
router ospf 1000
router-id 10.255.255.24
log-adjacency-changes detail
auto-cost reference-bandwidth 1000
prefix-suppression
passive-interface Loopback0
mpls ldp autoconfig
!
mpls ldp router-id Loopback0
!
```

XRv9k-28:

```
router ospf 1000
log adjacency changes detail
router-id 10.255.255.28
mpls ldp auto-config
prefix-suppression
auto-cost reference-bandwidth 1000000
area 0
interface Loopback0
passive enable
!
interface GigabitEthernet0/0/0/2
network point-to-point
!
interface GigabitEthernet0/0/0/3
network point-to-point
!
```

```

!
!
mpls ldp
 log
  neighbor
!
router-id 10.255.255.28
!
mpls label range table 0 28000 28999
!
commit
!
    
```

Task 01 is done. It was very easy to configure this task, the only thing you needed to do was to configure IGP on each of the routers and DO NOT do any redistribution between those IGPs.

Verification:

```

RP/0/RP0/CPU0:XRv9k-28#show ospf neighbor
Tue Apr 13 16:46:21.984 UTC

* Indicates MADJ interface
# Indicates Neighbor awaiting BFD session up

Neighbors for OSPF 1000

Neighbor ID    Pri   State           Dead Time   Address        Interface
10.255.255.11  1     FULL/ -         00:00:37   10.11.28.11   GigabitEthernet0/0/0/2
    Neighbor is up for 03:30:53
10.255.255.12  1     FULL/ -         00:00:32   10.12.28.12   GigabitEthernet0/0/0/3
    Neighbor is up for 03:30:54

Total neighbor count: 2
    
```

We do not need to verify the configuration on every single router, on each routing domain I will pick a random router to verify it's IGP neighbors and also the routing table:

```

RP/0/RP0/CPU0:XRv9k-28#show route ipv4 ospf
Tue Apr 13 16:47:42.292 UTC

O   10.255.255.9/32 [110/65546] via 10.11.28.11, 03:32:12, GigabitEthernet0/0/0/2
      [110/65546] via 10.12.28.12, 03:32:12, GigabitEthernet0/0/0/3
O   10.255.255.10/32 [110/65546] via 10.11.28.11, 03:32:12, GigabitEthernet0/0/0/2
      [110/65546] via 10.12.28.12, 03:32:12, GigabitEthernet0/0/0/3
O   10.255.255.11/32 [110/65536] via 10.11.28.11, 03:32:12, GigabitEthernet0/0/0/2
O   10.255.255.12/32 [110/65536] via 10.12.28.12, 03:32:14, GigabitEthernet0/0/0/3
O   10.255.255.22/32 [110/65547] via 10.11.28.11, 03:32:12, GigabitEthernet0/0/0/2
      [110/65547] via 10.12.28.12, 03:32:12, GigabitEthernet0/0/0/3
O   10.255.255.24/32 [110/65636] via 10.11.28.11, 03:32:12, GigabitEthernet0/0/0/2
      [110/65636] via 10.12.28.12, 03:32:12, GigabitEthernet0/0/0/3
O   10.255.255.26/32 [110/65547] via 10.11.28.11, 03:32:12, GigabitEthernet0/0/0/2
      [110/65547] via 10.12.28.12, 03:32:12, GigabitEthernet0/0/0/3
    
```

There are only 8 routers in the right side OSPF domain, that's why it only learned 7 host routes in from OSPF.

NOTE: there is no information about other IGP domains host routes learned from OSPF.

```

IOL-1#show ip route isis | begin Gate
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
i L2   10.255.255.2/32 [115/100] via 10.1.2.2, 03:46:13, Ethernet0/1
i L2   10.255.255.3/32 [115/100] via 10.1.3.3, 03:46:13, Ethernet0/0
i L2   10.255.255.4/32 [115/100] via 10.1.4.4, 03:46:03, Ethernet0/2
i L2   10.255.255.21/32 [115/1000] via 10.1.21.21, 03:44:55, Ethernet0/3
i L2   10.255.255.22/32 [115/1100] via 10.1.4.4, 03:44:55, Ethernet0/2
      [115/1100] via 10.1.3.3, 03:44:55, Ethernet0/0
i L2   10.255.255.25/32 [115/1000] via 10.1.25.25, 03:41:39, Ethernet1/0
i L2   10.255.255.26/32 [115/1100] via 10.1.4.4, 03:41:39, Ethernet0/2
      [115/1100] via 10.1.3.3, 03:41:39, Ethernet0/0
    
```

The core devices (IOL-1 to IOL-4) only learned the /32 prefixes inside IS-IS domain nothing else.

But what about ASBRs (in the future we will configure them as being the ASBRs)?

The ASBRs need to learn routes from both IS-IS and OSPF. Because they are part of two different routing domains:

```
RP/0/RP0/CPU0:XRv9k-25#show route ipv4 ospf
Tue Apr 13 16:52:26.640 UTC

O   10.255.255.5/32 [110/12] via 10.6.25.6, 03:43:42, GigabitEthernet0/0/0/2
      [110/12] via 10.8.25.8, 03:43:42, GigabitEthernet0/0/0/4
O   10.255.255.6/32 [110/2] via 10.6.25.6, 03:43:45, GigabitEthernet0/0/0/2
O   10.255.255.7/32 [110/12] via 10.6.25.6, 03:43:42, GigabitEthernet0/0/0/2
      [110/12] via 10.8.25.8, 03:43:42, GigabitEthernet0/0/0/4
O   10.255.255.8/32 [110/2] via 10.8.25.8, 03:43:42, GigabitEthernet0/0/0/4
O   10.255.255.21/32 [110/3] via 10.6.25.6, 03:43:42, GigabitEthernet0/0/0/2
      [110/3] via 10.8.25.8, 03:43:42, GigabitEthernet0/0/0/4
O   10.255.255.23/32 [110/112] via 10.6.25.6, 03:43:42, GigabitEthernet0/0/0/2
      [110/112] via 10.8.25.8, 03:43:42, GigabitEthernet0/0/0/4
O   10.255.255.27/32 [110/112] via 10.6.25.6, 03:36:40, GigabitEthernet0/0/0/2
      [110/112] via 10.8.25.8, 03:36:40, GigabitEthernet0/0/0/4

RP/0/RP0/CPU0:XRv9k-25#show route ipv4 isis
Tue Apr 13 16:52:35.602 UTC

i L2 10.255.255.1/32 [115/1000] via 10.1.25.1, 03:44:00, GigabitEthernet0/0/0/1
i L2 10.255.255.2/32 [115/1000] via 10.2.25.2, 03:44:00, GigabitEthernet0/0/0/3
i L2 10.255.255.3/32 [115/1100] via 10.1.25.1, 03:43:59, GigabitEthernet0/0/0/1
      [115/1100] via 10.2.25.2, 03:43:59, GigabitEthernet0/0/0/3
i L2 10.255.255.4/32 [115/1100] via 10.1.25.1, 03:43:59, GigabitEthernet0/0/0/1
      [115/1100] via 10.2.25.2, 03:43:59, GigabitEthernet0/0/0/3
i L2 10.255.255.22/32 [115/2100] via 10.1.25.1, 03:43:59, GigabitEthernet0/0/0/1
      [115/2100] via 10.2.25.2, 03:43:59, GigabitEthernet0/0/0/3
i L2 10.255.255.26/32 [115/2100] via 10.1.25.1, 03:43:59, GigabitEthernet0/0/0/1
      [115/2100] via 10.2.25.2, 03:43:59, GigabitEthernet0/0/0/3
```

Everything seems to be working fine. Let's continue with the next task.



Task 02:

- Provide MPLS L3VPN service to the Green Customer (PE-CE routing protocol: BGP)
- Provide VPLS Service to the Purple Customer (Martini with BGP A-D)



Solution:

You already know how to configure this task. The only thing you don't already know is:

How to provide connectivity between all the customer sites, because there is no LSP between all PE devices (Left side PEs and Right side PEs).

We will walk through how to provide this end-to-end LSPs with the help of BGP Labeled unicast in task 03.

Let's configure the necessary configuration on the PE devices for this task:

CSR-23:

```
vrf definition Cust-B
 rd 1000:2
 route-target export 1000:2
 route-target import 1000:2
 !
 address-family ipv4
 exit-address-family
 !
 !
 interface GigabitEthernet4
 vrf forwarding Cust-B
 ip address 10.14.23.23 255.255.255.0
 negotiation auto
 no mop enabled
 no mop sysid
 !
 router bgp 1000
 !
 address-family ipv4 vrf Cust-B
 redistribute connected
```

```
neighbor 10.14.23.14 remote-as 65002
neighbor 10.14.23.14 activate
neighbor 10.14.23.14 as-override
exit-address-family
!
!
interface GigabitEthernet1
no ip address
negotiation auto
no mop enabled
no mop sysid
service instance 1 ethernet
encapsulation default
l2protocol forward cdp
!
l2vpn vfi context VPLS
vpn id 1
autodiscovery bgp signaling ldp
vpls-id 1000:1
!
bridge-domain 1
member GigabitEthernet1 service-instance 1
member vfi VPLS
!

CSR-24:
vrf definition Cust-B
rd 1000:2
route-target export 1000:2
route-target import 1000:2
!
address-family ipv4
exit-address-family
!
```

```
!  
interface GigabitEthernet4  
  vrf forwarding Cust-B  
  ip address 10.19.24.24 255.255.255.0  
  negotiation auto  
  no mop enabled  
  no mop sysid  
!  
router bgp 1000  
!  
  address-family ipv4 vrf Cust-B  
    redistribute connected  
    neighbor 10.19.24.19 remote-as 65002  
    neighbor 10.19.24.19 activate  
    neighbor 10.19.24.19 as-override  
  exit-address-family  
!  
interface GigabitEthernet1  
  service instance 1 ethernet  
  encapsulation default  
  l2protocol forward cdp  
!  
!  
l2vpn vfi context VPLS  
  vpn id 1  
  autodiscovery bgp signaling ldp  
  vpls-id 1000:1  
!  
bridge-domain 1  
  member GigabitEthernet1 service-instance 1  
  member vfi VPLS  
!
```

all IOS-XE PE devices are configured and they are ready to serve the customer devices.

Let's configure IOS-XR devices:

XRv9k-27:

```
vrf Cust-B
  address-family ipv4 unicast
    import route-target
      1000:2
    !
  export route-target
    1000:2
  !
!
!
!
interface GigabitEthernet0/0/0/0
  vrf Cust-B
  ipv4 address 10.15.27.27 255.255.255.0
!
interface GigabitEthernet0/0/0/1
  vrf Cust-B
  ipv4 address 10.16.27.27 255.255.255.0
!
router bgp 1000
!
vrf Cust-B
  rd 1000:1
  address-family ipv4 unicast
    redistribute connected
  !
  neighbor 10.15.27.15
    remote-as 65002
    address-family ipv4 unicast
      route-policy PASS in
      route-policy PASS out
      as-override
  !
!
neighbor 10.16.27.16
```

```
remote-as 65002
address-family ipv4 unicast
  route-policy PASS in
  route-policy PASS out
as-override
!
!
!
!
```

XRv9k-28:

```
vrf Cust-B
address-family ipv4 unicast
import route-target
  1000:2
!
export route-target
  1000:2
!
!
!
interface GigabitEthernet0/0/0/0
vrf Cust-B
ipv4 address 10.17.28.28 255.255.255.0
!
interface GigabitEthernet0/0/0/1
vrf Cust-B
ipv4 address 10.18.28.28 255.255.255.0
!
router bgp 1000
!
vrf Cust-B
rd 1000:1
address-family ipv4 unicast
```

```

redistribute connected
!
neighbor 10.17.28.17
  remote-as 65002
  address-family ipv4 unicast
    route-policy PASS in
    route-policy PASS out
  as-override
!
!
neighbor 10.18.28.18
  remote-as 65002
  address-family ipv4 unicast
    route-policy PASS in
    route-policy PASS out
  as-override
!
!
!
!

```

NOTE: We have not configured VPLS on the IOS-XR devices, because: there is no any virtual IOS-XR device (At the time of writing this document) that supports any P2MP L2 Forwarding in the dataplane, so we cannot have VPLS on the XRv9k device so far.

Let's do a quick verification on one of the PE devices:

Verification:

```

RP/0/RP0/CPU0:XRv9k-28#show bgp vrf Cust-B ipv4 unicast summary | begin Neigh
Tue Apr 13 17:12:20.453 UTC
Neighbor      Spk    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  St/PfxRcd
10.17.28.17   0 65002   266    246     30    0    0 03:56:50    1
10.18.28.18   0 65002   266    246     30    0    0 03:56:51    1

```



Task 03:

- PE devices should form BGP peering with each other (Full Mesh)
- End to End LSPs from one PE device to the other PE devices should be provided by BGP IPv4 Labeled unicast
- R21, R22, R25 and R26 must be the ASBRs
- Use traceroute on customer devices to make sure they have end to end reachability



Solution:

So far there is ONLY LSP between Loopback interfaces of the routers inside a routing domain.

PE devices have no idea how to reach to the other PE devices in another routing domain (They even don't have any IP reachability to those devices) because: There is no redistribution between the routing domains.

Our Unified MPLS journey begins at this point!

Four devices are going to be configured as the ASBRs:

- CSR-21
- CSR-22
- XRv9k-25
- XRv9k-26

These devices will form IPv4 Labeled unicast neighborship with each other. Besides that, they also need to form the same neighborship with the PE devices in their routing domain.

As an example:

- CSR-21 should form MP-BGP IPv4 Labeled Unicast with XRv9k-26, XRv9k-25 and CSR-22
- CSR-21 should also form MP-BGP IPv4 Labeled Unicast with CSR-23 and XRv9k-27

NOTE: These ASBRs should be configured as the Route Reflectors, and PE devices should be their Route-Reflector Clients, why?!

ASBRs will get the routes from other ASBRs and advertise them to their clients.

Why we even need Labeled Unicast? Can't we use IPv4 Unicast instead?

In the Seamless MPLS Framework, our final goal is to provide PE to PE LSPs, but without the redistribution and MPLS LDP or Segment-Routing, how can we achieve that goal?!

That is why there must be some kind of technology to advertise the labels for the PE devices prefixes to the other PE devices at the other part of the network (inside other routing domain).

Let's begin the configuration, as I mentioned before, I'm going to start configuring the ASBRs first:

```
CSR-21:
router bgp 1000
  bgp cluster-id 1.1.1.1
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 10.255.255.22 remote-as 1000
  neighbor 10.255.255.22 update-source Loopback0
  neighbor 10.255.255.23 remote-as 1000
  neighbor 10.255.255.23 update-source Loopback0
  neighbor 10.255.255.26 remote-as 1000
  neighbor 10.255.255.26 update-source Loopback0
  neighbor 10.255.255.27 remote-as 1000
  neighbor 10.255.255.27 update-source Loopback0
!
address-family ipv4
  neighbor 10.255.255.22 activate
  neighbor 10.255.255.22 next-hop-self all
  neighbor 10.255.255.22 send-label
  neighbor 10.255.255.23 activate
  neighbor 10.255.255.23 route-reflector-client
  neighbor 10.255.255.23 next-hop-self all
  neighbor 10.255.255.23 send-label
  neighbor 10.255.255.26 activate
  neighbor 10.255.255.26 next-hop-self all
  neighbor 10.255.255.26 send-label
  neighbor 10.255.255.27 activate
  neighbor 10.255.255.27 route-reflector-client
  neighbor 10.255.255.27 next-hop-self all
  neighbor 10.255.255.27 send-label
exit-address-family
!
```

Another important thing about the ASBR configuration is:

We need to use `next-hop-self all` command for each of our neighbors.

Imagin:

- CSR-21 receives a prefix from CSR-23 (CSR-23 Loopback 0 prefix)
- It advertises that prefix to the XRv9k-26
- It will be an iBGP route and by default a route-reflector cannot make any changes to the next-hop of the iBGP routes
- In addition to that, If CSR-21 does not change the next hop, the next hop will be the loopback 0 IP address of the CSR-23, we can conclude that: XRv9k-26 will never have any idea about how to reach to that next hop (there is no redistribution between IGP, there is no LSP for that prefix).

That is why we need to force the ASBR to change the next-hop to it's own Update-Source interface IP address (Loopback 0) even if it is an iBGP route (that is why at the end of the next-hop-self command we need to use "all" option as well).

Time to configure the other IOS-XE ASBR:

CSR-22:

```
router bgp 1000
  bgp cluster-id 2.2.2.2
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 10.255.255.21 remote-as 1000
  neighbor 10.255.255.21 update-source Loopback0
  neighbor 10.255.255.24 remote-as 1000
  neighbor 10.255.255.24 update-source Loopback0
  neighbor 10.255.255.25 remote-as 1000
  neighbor 10.255.255.25 update-source Loopback0
  neighbor 10.255.255.28 remote-as 1000
  neighbor 10.255.255.28 update-source Loopback0
!
address-family ipv4
  neighbor 10.255.255.21 activate
  neighbor 10.255.255.21 route-reflector-client
  neighbor 10.255.255.21 next-hop-self all
  neighbor 10.255.255.21 send-label
  neighbor 10.255.255.24 activate
  neighbor 10.255.255.24 route-reflector-client
```

```
neighbor 10.255.255.24 next-hop-self all
neighbor 10.255.255.24 send-label
neighbor 10.255.255.25 activate
neighbor 10.255.255.25 route-reflector-client
neighbor 10.255.255.25 next-hop-self all
neighbor 10.255.255.25 send-label
neighbor 10.255.255.28 activate
neighbor 10.255.255.28 route-reflector-client
neighbor 10.255.255.28 next-hop-self all
neighbor 10.255.255.28 send-label
exit-address-family
```

Let's configure The IOS-XR ASBR devices the same way:

XRv9k-25:

```
router bgp 1000
  bgp cluster-id 1.1.1.1
  bgp log neighbor changes detail
  ibgp policy out enforce-modifications
  address-family ipv4 unicast
    allocate-label all
  !
  neighbor 10.255.255.22
    remote-as 1000
    update-source Loopback0
    address-family ipv4 labeled-unicast
      next-hop-self
  !
  !
  neighbor 10.255.255.23
    remote-as 1000
    update-source Loopback0
    address-family ipv4 labeled-unicast
      route-reflector-client
      next-hop-self
  !
  !
```

```
neighbor 10.255.255.26
  remote-as 1000
  update-source Loopback0
  address-family ipv4 labeled-unicast
    next-hop-self
  !
!
```

```
neighbor 10.255.255.27
  remote-as 1000
  update-source Loopback0
  address-family ipv4 labeled-unicast
    route-reflector-client
    next-hop-self
  !
!
```

XRv9k-26:

```
router bgp 1000
  bgp cluster-id 2.2.2.2
  bgp log neighbor changes detail
  ibgp policy out enforce-modifications
  address-family ipv4 unicast
    allocate-label all
  !
  neighbor 10.255.255.21
    remote-as 1000
    update-source Loopback0
    address-family ipv4 labeled-unicast
      next-hop-self
  !
!
```

```
neighbor 10.255.255.24
  remote-as 1000
```

```
update-source Loopback0
address-family ipv4 labeled-unicast
  route-reflector-client
  next-hop-self
!
!
neighbor 10.255.255.25
  remote-as 1000
  update-source Loopback0
  address-family ipv4 labeled-unicast
    next-hop-self
!
!
neighbor 10.255.255.28
  remote-as 1000
  update-source Loopback0
  address-family ipv4 labeled-unicast
    route-reflector-client
    next-hop-self
!
!
!
```

On IOS-XR there is no “all” option at the end of the next-hop-self command.

Instead we have this command:

```
ibgp policy out enforce-modifications
```

This command at the Router BGP level, forces the route-reflector to change the next-hop to it's own update source interface IP address (Loopback 0) even if the route is iBGP.

Now, our ASBRs are fully configured and ready to make neighborship with PE devices in their routing domain.

We will configure the PE devices the same way, this time they are not Route-Reflectors or ASBRs etc... .

They are just simple iBGP neighbors with ASBRs:

CSR-23:

```

router bgp 1000
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 10.255.255.21 remote-as 1000
  neighbor 10.255.255.21 update-source Loopback0
  neighbor 10.255.255.25 remote-as 1000
  neighbor 10.255.255.25 update-source Loopback0
  !
  address-family ipv4
    network 10.255.255.23 mask 255.255.255.255
  neighbor 10.255.255.21 activate
  neighbor 10.255.255.21 send-label
  neighbor 10.255.255.25 activate
  neighbor 10.255.255.25 send-label
  exit-address-family
    
```

What an easy step!

There is something happening under the IPv4 Address-Family:

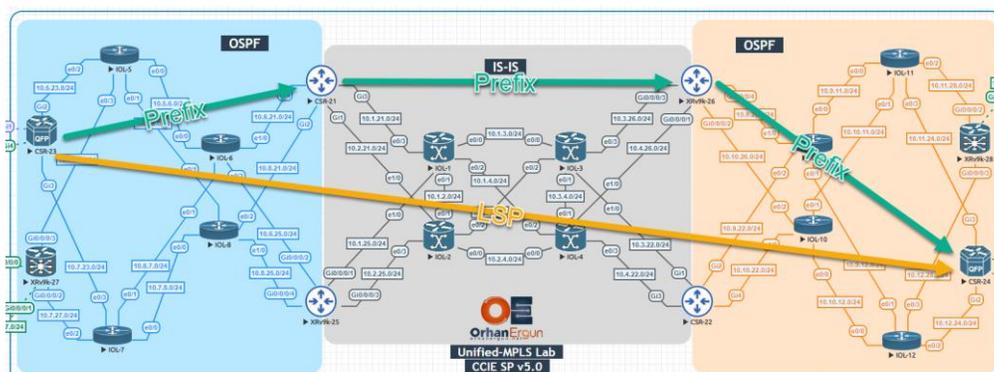
```
network 10.255.255.23 mask 255.255.255.255
```

The PE device needs to advertise it's own Loopback 0 IP address to the ASBRs.

The problem is solved!

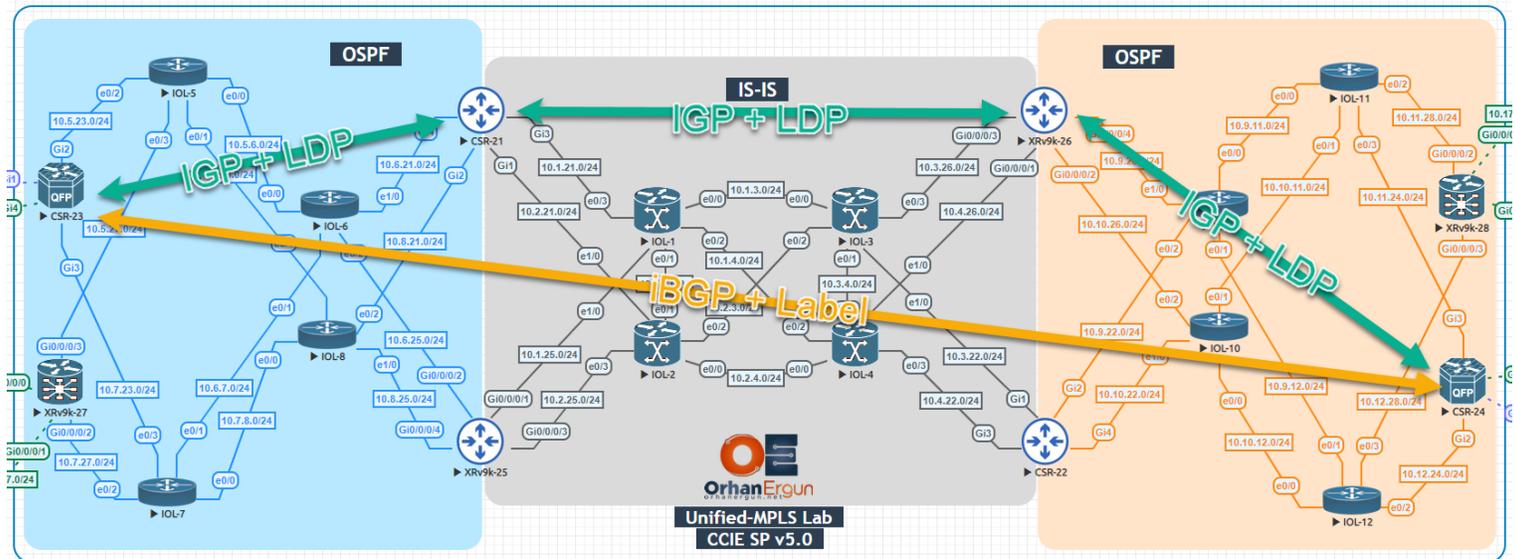
- CSR-23 (Left side PE device) advertises it's Loopback 0 prefix to the ASBRs
- ASBRs will advertise it to their neighboring ASBRs (for example XRv9k-26)
- XRv9k-26 advertises it to it's client (for example CSR-24)

Now the two PE devices (CSR-26 and CSR-24) will have LSPs between their loopback interfaces:



This is how Unified MPLS works.

- CSR-24 can reach to the XRv9k-26 (IGP LDP provides the LSP)
- XRv9k-26 can reach CSR-21 (IGP LDP Provides the LSP)
- CSR-21 can reach CSR-23 (IGP LDP Provides the LSP)
- Finally: CSR-24 can reach the CSR-23 (iBGP Labeled Unicast provides the LSP)
- In fact iBGP + Label is between **PE-ASBR-ASBR-PE**, for simplicity we draw it in a shape of PE-PE:



Now we have Label switched path between all PE devices, we can form any MP-BGP neighborhood between them (Like VPNv4 Unicast, L2VPN VPLS, etc...).

Let's configure the final step (Full mesh BGP between PE devices):

CSR-23:

```

router bgp 1000
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 10.255.255.24 remote-as 1000
  neighbor 10.255.255.24 update-source Loopback0
  neighbor 10.255.255.27 remote-as 1000
  neighbor 10.255.255.27 update-source Loopback0
  neighbor 10.255.255.28 remote-as 1000
  neighbor 10.255.255.28 update-source Loopback0
  !
  address-family vpnv4
    neighbor 10.255.255.24 activate
    
```

```
neighbor 10.255.255.24 send-community extended
neighbor 10.255.255.27 activate
neighbor 10.255.255.27 send-community extended
neighbor 10.255.255.28 activate
neighbor 10.255.255.28 send-community extended
exit-address-family
!
address-family l2vpn vpls
neighbor 10.255.255.24 activate
exit-address-family
!
```

XRv9k-27:

```
router bgp 1000
bgp log neighbor changes detail
address-family ipv4 unicast
network 10.255.255.27/32
allocate-label all
!
address-family vpnv4 unicast
!
neighbor 10.255.255.23
remote-as 1000
update-source Loopback0
address-family vpnv4 unicast
!
!
neighbor 10.255.255.24
remote-as 1000
update-source Loopback0
address-family vpnv4 unicast
!
!
neighbor 10.255.255.28
remote-as 1000
```

```
update-source Loopback0
address-family vpnv4 unicast
!
!
!
```

XRv9k-28:

```
router bgp 1000
  bgp log neighbor changes detail
  address-family ipv4 unicast
    network 10.255.255.28/32
    allocate-label all
  !
  address-family vpnv4 unicast
  !
  neighbor 10.255.255.23
    remote-as 1000
    update-source Loopback0
    address-family vpnv4 unicast
  !
  !
  neighbor 10.255.255.24
    remote-as 1000
    update-source Loopback0
    address-family vpnv4 unicast
  !
  !
  neighbor 10.255.255.27
    remote-as 1000
    update-source Loopback0
    address-family vpnv4 unicast
  !
  !
  !
```

CSR-24:

```
router bgp 1000
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 10.255.255.23 remote-as 1000
  neighbor 10.255.255.23 update-source Loopback0
  neighbor 10.255.255.27 remote-as 1000
  neighbor 10.255.255.27 update-source Loopback0
  neighbor 10.255.255.28 remote-as 1000
!
address-family vpnv4
  neighbor 10.255.255.23 activate
  neighbor 10.255.255.23 send-community extended
  neighbor 10.255.255.27 activate
  neighbor 10.255.255.27 send-community extended
  neighbor 10.255.255.28 activate
  neighbor 10.255.255.28 send-community extended
exit-address-family
!
address-family l2vpn vpls
  neighbor 10.255.255.23 activate
exit-address-family
!
```

Congratulations 😊 all tasks are done!

It's time to do some verification:

```
CSR-23#traceroute 10.255.255.28 source lo 0 probe 1
Type escape sequence to abort.
Tracing the route to 10.255.255.28
VRF info: (vrf in name/id, vrf out name/id)
 1 10.5.23.5 [MPLS: Labels 50003/21014 Exp 0] 10 msec
 2 10.6.7.6 [MPLS: Labels 60003/21014 Exp 0] 3 msec
 3 10.8.21.21 [MPLS: Label 21014 Exp 0] 4 msec
 4 10.2.21.2 [MPLS: Labels 20004/22013 Exp 0] 2 msec
 5 10.2.4.4 [MPLS: Labels 40003/22013 Exp 0] 3 msec
```

```

6 10.3.22.22 [MPLS: Label 22013 Exp 0] 4 msec
7 10.10.22.10 [MPLS: Label 10006 Exp 0] 3 msec
8 *
9 10.12.28.28 14 msec

```

```
CSR-23#show ip cef 10.255.255.28/32
```

```
10.255.255.28/32
```

```

nexthop 10.5.23.5 GigabitEthernet2 label 50003-(local:23000) 21014
nexthop 10.7.23.7 GigabitEthernet3 label 70003-(local:23000) 21014

```

```
CSR-23#show bgp ipv4 unicast labels
```

Network	Next Hop	In label/Out label
10.255.255.23/32	0.0.0.0	imp-null/nolabel
10.255.255.24/32	10.255.255.25	nolabel/25012
	10.255.255.21	nolabel/21010
10.255.255.27/32	10.255.255.25	nolabel/25013
	10.255.255.21	nolabel/21013
10.255.255.28/32	10.255.255.25	nolabel/25014
	10.255.255.21	nolabel/21014

```
CSR-23#show bgp ipv4 unicast | begin Net
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.255.255.23/32	0.0.0.0	0		32768	i
* i 10.255.255.24/32	10.255.255.25	0	100		0 i
*>i	10.255.255.21	0	100		0 i
r i 10.255.255.27/32	10.255.255.25	0	100		0 i
r>i	10.255.255.21	0	100		0 i
* i 10.255.255.28/32	10.255.255.25	0	100		0 i
*>i	10.255.255.21	0	100		0 i

```
CSR-23#show ip cef 10.255.255.21/32
```

```
10.255.255.21/32
```

```

nexthop 10.5.23.5 GigabitEthernet2 label 50003-(local:23000)
nexthop 10.7.23.7 GigabitEthernet3 label 70003-(local:23000)

```

```
IOL-14#show ip bgp | begin Net
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
r>	10.14.23.0/24	10.14.23.23	0		0	1000 ?
*>	10.15.27.0/24	10.14.23.23			0	1000 ?
*>	10.16.27.0/24	10.14.23.23			0	1000 ?
*>	10.17.28.0/24	10.14.23.23			0	1000 ?
*>	10.18.28.0/24	10.14.23.23			0	1000 ?
*>	10.19.24.0/24	10.14.23.23			0	1000 ?
*>	10.255.255.14/32	0.0.0.0	0		32768	i
*>	10.255.255.15/32	10.14.23.23			0	1000 1000 i
*>	10.255.255.16/32	10.14.23.23			0	1000 1000 i
*>	10.255.255.17/32	10.14.23.23			0	1000 1000 i
*>	10.255.255.18/32	10.14.23.23			0	1000 1000 i
*>	10.255.255.19/32	10.14.23.23			0	1000 1000 i

```
IOL-14#traceroute 10.255.255.18 source lo 0 probe 1
```

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

```
Tracing the route to 10.255.255.18
```

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

```

 1 10.14.23.23 [AS 1000] 0 msec
 2 10.7.23.7 [MPLS: Labels 70003/21014/28011 Exp 0] 4 msec
 3 10.6.7.6 [MPLS: Labels 60003/21014/28011 Exp 0] 3 msec
 4 10.6.21.21 [MPLS: Labels 21014/28011 Exp 0] 4 msec
 5 10.1.21.1 [MPLS: Labels 10004/22013/28011 Exp 0] 2 msec
 6 10.1.4.4 [MPLS: Labels 40003/22013/28011 Exp 0] 3 msec
 7 10.4.22.22 [MPLS: Labels 22013/28011 Exp 0] 4 msec
 8 10.10.22.10 [MPLS: Labels 10006/28011 Exp 0] 3 msec
 9 10.10.12.12 [MPLS: Labels 12006/28011 Exp 0] 3 msec
10 10.12.28.28 [MPLS: Label 28011 Exp 0] 4 msec
11 10.18.28.18 [AS 1000] 3 msec
    
```