MPLS Lab 02

MPLS Layer 2 and 3 Virtual Private Networks

Topology:

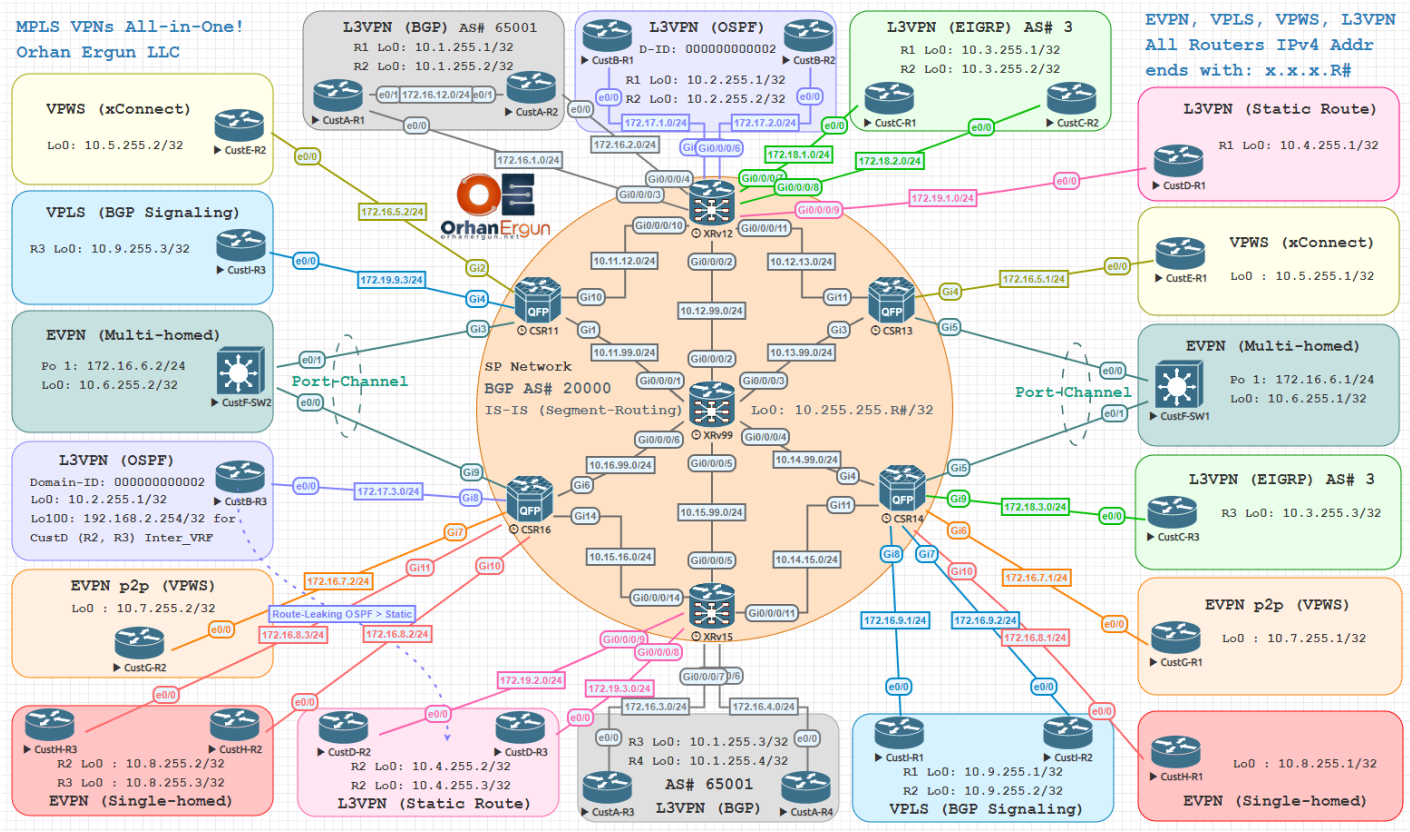

MP-BGP: VPNv4 Unicast / L2VPN VPLS / L2VPN EVPN

PE-CE Routing Protocols:

- Static Routing (Including VRF Receive command)
- EIGRP
- OSPF
- BGP

EVPN:

- VPWS
- Single-Homed
- Multi-Homed (Active-Active)

Task 01:

- Configure IS-IS in the Service Provider Network
- Enable Prefix-Suppression for IS-IS
- Configure Basic Segment-Routing (Prefix SID-MAP using Loopback 0 IPv4 addresses)
- Enable Segment-Routing MPLS in the IS-IS process

Solution:

Instead of LDP we are going to use segment routing for Label distribution, when we enable Segment-Routing MPLS, IS-IS advertises Prefix to SID MAP using LSPs, there is no need for LDP anymore:

```
XRv99:
router isis 20000
 is-type level-2-only
 net 49.0020.0000.0000.0099.00
 address-family ipv4 unicast
  metric-style wide
  advertise passive-only
  segment-routing mpls
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
   prefix-sid index 99
  !
 !
  interface GigabitEthernet0/0/0/1
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/2
  address-family ipv4 unicast
  !
 !
```

```
 interface GigabitEthernet0/0/0/3
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/4
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/5
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/6
  address-family ipv4 unicast
  !
 !
!
segment-routing
```

Enabling basic Segment-Routing has an easy procedure. First we enable segment-routing globally, then in IS-IS change the metric-style to wide, then IS-IS can carry SR mapping information, the next step is enabling segment-routing mpls, as the last step we need to map a prefix (/32 host network of loopback interface) to a Segment ID.

Let's configure the other SP routers:

```
XRv12:
router isis 20000
 is-type level-2-only
 net 49.0020.0000.0000.0012.00
 address-family ipv4 unicast
  metric-style wide
  advertise passive-only
  segment-routing mpls
 !
 interface Loopback0
```

```
  passive
  address-family ipv4 unicast
   prefix-sid index 12
  !
 !
 interface GigabitEthernet0/0/0/2
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/10
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/11
  address-family ipv4 unicast
  !
 !
!
segment-routing
```

XRv15:

```
router isis 20000
 is-type level-2-only
 net 47.0020.0000.0000.0015.00
 address-family ipv4 unicast
  metric-style wide
  advertise passive-only
  segment-routing mpls
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
   prefix-sid index 15
  !
```

```
 !
 interface GigabitEthernet0/0/0/5
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/11
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/14
  address-family ipv4 unicast
  !
 !
!
segment-routing
```

Configuring Segment-Routing in IOS-XE is a little bit different than IOS-XR, but the concept is the same:

```
CSR11:
segment-routing mpls
 !
 connected-prefix-sid-map
  address-family ipv4
   10.255.255.11/32 index 11
  exit-address-family
 !
!
router isis 20000
 net 49.0020.0000.0000.0011.00
 is-type level-2-only
 advertise passive-only
 metric-style wide
 segment-routing mpls
 passive-interface Loopback0
!
```

```
interface GigabitEthernet1
 ip router isis 20000
!
interface GigabitEthernet10
 ip router isis 20000
!
```

CSR13:

```
segment-routing mpls
 !
 connected-prefix-sid-map
  address-family ipv4
   10.255.255.13/32 index 13
  exit-address-family
 !
!
router isis 20000
 net 49.0020.0000.0000.0013.00
 is-type level-2-only
 advertise passive-only
 metric-style wide
 segment-routing mpls
 passive-interface Loopback0
!
interface GigabitEthernet3
 ip router isis 20000
!
interface GigabitEthernet11
 ip router isis 20000
!
```

CSR14:

```
segment-routing mpls
 !
 connected-prefix-sid-map
  address-family ipv4
```

```
   10.255.255.14/32 index 14
  exit-address-family
 !
!
router isis 20000
 net 49.0020.0000.0000.0014.00
 is-type level-2-only
 advertise passive-only
 metric-style wide
 segment-routing mpls
 passive-interface Loopback0
!
interface GigabitEthernet4
 ip router isis 20000
!
interface GigabitEthernet11
 ip router isis 20000
!
```

CSR16:

```
segment-routing mpls
 !
 connected-prefix-sid-map
  address-family ipv4
   10.255.255.16/32 index 16
  exit-address-family
 !
!
router isis 20000
 net 49.0020.0000.0000.0016.00
 is-type level-2-only
 advertise passive-only
 metric-style wide
 segment-routing mpls
```

```
 passive-interface Loopback0
!
interface GigabitEthernet6
 ip router isis 20000
!
interface GigabitEthernet14
 ip router isis 20000
!
```

Verification:

```
RP/0/0/CPU0:XRv99(config)#do sh isis neighbors
Thu Apr 23 08:16:54.964 UTC

IS-IS 20000 neighbors:
System Id      Interface        SNPA           State Holdtime Type IETF-NSF
XRv12          Gi0/0/0/2        5000.0008.0003 Up    27       L2   Capable
CSR14          Gi0/0/0/4        5000.0003.0003 Up    26       L2   Capable
CSR16          Gi0/0/0/6        5000.0004.0005 Up    25       L2   Capable
CSR11          Gi0/0/0/1        5000.0001.0000 Up    27       L2   Capable
CSR13          Gi0/0/0/3        5000.0002.0002 Up    22       L2   Capable
XRv15          Gi0/0/0/5        5000.0007.0006 Up    29       L2   Capable


RP/0/0/CPU0:XRv99(config)#do sh  isis database CSR11.00-00 verbose detail
Thu Apr 23 08:23:59.595 UTC

IS-IS 20000 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
CSR11.00-00           0x0000000f   0x2834        1027            0/0/0
  Area Address:   49.0020
  NLPID:          0xcc
  Router Cap:     10.255.255.11, D:0, S:0
    Segment Routing: I:1 V:0, SRGB Base: 16000 Range: 8000
    SubTLV 22Length: 9
    SubTLV 19Length: 2
```

```
    SubTLV 23Length: 2
  Hostname:       CSR11
  Metric: 10         IS-Extended XRv99.17
    Admin. Weight: 10
    Physical BW: 1000000 kbits/sec
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid: 110013 System ID:XRv99
  Metric: 10         IS-Extended XRv12.0b
    Admin. Weight: 10
    Physical BW: 1000000 kbits/sec
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid: 110014 System ID:XRv12
  IP Address:     10.255.255.11
  Metric: 0          IP-Extended 10.255.255.11/32
    SubTLV 4Length: 1
    Prefix-SID Index: 11, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
```

LSP of CSR11 includes Prefix-SID. Let's traceroute to that prefix on CSR14:

```
CSR14(config)#do trace 10.255.255.11 source lo 0
Type escape sequence to abort.
Tracing the route to 10.255.255.11
VRF info: (vrf in name/id, vrf out name/id)
  1 10.14.99.99 [MPLS: Label 16011 Exp 0] 62 msec 31 msec 14 msec
  2 10.11.99.11 15 msec 13 msec 9 msec


CSR14(config)#do sh mpls forwarding-table 10.255.255.11
Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
16011      16011      10.255.255.11/32 0             Gi4         10.14.99.99


RP/0/0/CPU0:XRv15(config)#do sh cef ipv4 10.255.255.11/32
Thu Apr 23 08:31:00.237 UTC
10.255.255.11/32, version 152, internal 0x1000001 0x81 (ptr 0xa141b174) [1], 0x0 (0xa13e5e18), 0xa28
(0xa1583410)
 Updated Apr 23 07:54:33.787
 local adjacency 10.15.99.99
 Prefix Len 32, traffic index 0, precedence n/a, priority 1
```

```
   via 10.15.99.99/32, GigabitEthernet0/0/0/5, 7 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa10f32a4 0x0]
    next hop 10.15.99.99/32
    local adjacency
     local label 16011      labels imposed {16011}  // SRGB + Index Value
```

CSR14 and every other router is aware of label 16011 to reach that prefix, in fact every router receives an SID for /32 prefixes and create label using SRGB + Index (The default SRGB on IOS and IOS-XR is 16000 – 23999).

```
CSR11(config)#do sh segment-routing mpls connected-prefix-sid-map ipv4


               PREFIX_SID_CONN_MAP ALGO_0


    Prefix/masklen   SID Type Range Flags SRGB
  10.255.255.11/32    11 Indx     1         Y


               PREFIX_SID_PROTOCOL_ADV_MAP ALGO_0


    Prefix/masklen   SID Type Range Flags SRGB Source
  10.255.255.11/32    11 Indx     1         Y  IS-IS Level 2 0000.0000.0011
  10.255.255.12/32    12 Indx     1         Y  IS-IS Level 2 0000.0000.0012
  10.255.255.13/32    13 Indx     1         Y  IS-IS Level 2 0000.0000.0013
  10.255.255.14/32    14 Indx     1         Y  IS-IS Level 2 0000.0000.0014
  10.255.255.15/32    15 Indx     1         Y  IS-IS Level 2 0000.0000.0015
  10.255.255.16/32    16 Indx     1         Y  IS-IS Level 2 0000.0000.0016
  10.255.255.99/32    99 Indx     1         Y  IS-IS Level 2 0000.0000.0099
               PREFIX_SID_CONN_MAP ALGO_1
    Prefix/masklen   SID Type Range Flags SRGB
               PREFIX_SID_PROTOCOL_ADV_MAP ALGO_1
    Prefix/masklen   SID Type Range Flags SRGB Source
CSR11(config)#do sh mpls fo | include 10.255.255
16012 [M]  Pop Label  10.255.255.12/32 0            Gi10       10.11.12.12
16013 [M]  16013      10.255.255.13/32 0            Gi10       10.11.12.12
      [M]  16013      10.255.255.13/32 0            Gi1        10.11.99.99
16014 [M]  16014      10.255.255.14/32 0            Gi1        10.11.99.99
```

```
16015 [M]  16015      10.255.255.15/32 0            Gi1        10.11.99.99
16016 [M]  16016      10.255.255.16/32 0            Gi1        10.11.99.99
16099 [M]  Pop Label  10.255.255.99/32 0            Gi1        10.11.99.99
110015[M]  Pop Label  10.255.255.12/32 0            Gi10       10.11.12.12
110016[M]  16013      10.255.255.13/32 0            Gi10       10.11.12.12
      [M]  16013      10.255.255.13/32 0            Gi1        10.11.99.99
110017[M]  Pop Label  10.255.255.99/32 0            Gi1        10.11.99.99
110018[M]  16014      10.255.255.14/32 0            Gi1        10.11.99.99
110019[M]  16015      10.255.255.15/32 0            Gi1        10.11.99.99
110020[M]  16016      10.255.255.16/32 0            Gi1        10.11.99.99
```

Task 02:

- Configure BGP on all Edge routers (XRv99 is a Core/P router).
- **ONLY** Enable VPNv4 Unicast Address-family.
- CSR11 and CSR13 must also act as Route-Reflectors.
- Configure Router-ID on all Edge routers and Cluster-ID on RRs statically (Use Lo0 IPv4 Address).

Solution:

CSR11:

```
router bgp 20000
 bgp router-id 10.255.255.11
 bgp cluster-id 10.255.255.11
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.255.255.12 remote-as 20000
 neighbor 10.255.255.12 update-source Loopback0
 neighbor 10.255.255.13 remote-as 20000
 neighbor 10.255.255.13 update-source Loopback0
 neighbor 10.255.255.14 remote-as 20000
 neighbor 10.255.255.14 update-source Loopback0
 neighbor 10.255.255.15 remote-as 20000
 neighbor 10.255.255.15 update-source Loopback0
 neighbor 10.255.255.16 remote-as 20000
 neighbor 10.255.255.16 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 10.255.255.12 activate
  neighbor 10.255.255.12 send-community extended
  neighbor 10.255.255.12 route-reflector-client
```

```
neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community extended
  neighbor 10.255.255.14 activate
  neighbor 10.255.255.14 send-community extended
  neighbor 10.255.255.14 route-reflector-client
  neighbor 10.255.255.15 activate
  neighbor 10.255.255.15 send-community both
  neighbor 10.255.255.15 route-reflector-client
  neighbor 10.255.255.16 activate
  neighbor 10.255.255.16 send-community extended
  neighbor 10.255.255.16 route-reflector-client
 exit-address-family
 !
```

The iBGP configuration Steps are the same as MPLS-Lab-1. If you need any detailed explanation refer to that Lab.

IOS-XE automatically enables send-community extended for the VPNv4 address family as soon as the peer is activated using neighbor x.x.x.x activate command.

Note: IOS-XR (XRv 6.1.3) only supports L3VPNs, the Data Plane is not forwarding any packets for L2VPN services, we will use IOX-XR for only L3VPNs part of this Lab.

CSR13:

```
router bgp 20000
 bgp router-id 10.255.255.13
 bgp cluster-id 10.255.255.13
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.255.255.11 remote-as 20000
 neighbor 10.255.255.11 update-source Loopback0
 neighbor 10.255.255.12 remote-as 20000
 neighbor 10.255.255.12 update-source Loopback0
 neighbor 10.255.255.14 remote-as 20000
 neighbor 10.255.255.14 update-source Loopback0
 neighbor 10.255.255.15 remote-as 20000
 neighbor 10.255.255.15 update-source Loopback0
 neighbor 10.255.255.16 remote-as 20000
```

```
 neighbor 10.255.255.16 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community extended
  neighbor 10.255.255.12 activate
  neighbor 10.255.255.12 send-community extended
  neighbor 10.255.255.12 route-reflector-client
  neighbor 10.255.255.14 activate
  neighbor 10.255.255.14 send-community extended
  neighbor 10.255.255.14 route-reflector-client
  neighbor 10.255.255.15 activate
  neighbor 10.255.255.15 send-community both
  neighbor 10.255.255.15 route-reflector-client
  neighbor 10.255.255.16 activate
  neighbor 10.255.255.16 send-community extended
  neighbor 10.255.255.16 route-reflector-client
 exit-address-family
 !
```

XRv12:

```
router bgp 20000
 bgp router-id 10.255.255.12
 address-family vpnv4 unicast
 !
 neighbor 10.255.255.11
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
```

```
neighbor 10.255.255.13
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
!
```

XRv15:

```
router bgp 20000
 bgp router-id 10.255.255.15
 address-family vpnv4 unicast
 !
 neighbor 10.255.255.11
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
 neighbor 10.255.255.13
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
!
```

CSR14:

```
router bgp 20000
 bgp router-id 10.255.255.14
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.255.255.11 remote-as 20000
 neighbor 10.255.255.11 update-source Loopback0
 neighbor 10.255.255.13 remote-as 20000
 neighbor 10.255.255.13 update-source Loopback0
 .
```

```
neighbor 10.255.255.13
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
!
```

XRv15:

```
router bgp 20000
 bgp router-id 10.255.255.15
 address-family vpnv4 unicast
 !
 neighbor 10.255.255.11
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
 neighbor 10.255.255.13
  remote-as 20000
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
!
```

CSR14:

```
router bgp 20000
 bgp router-id 10.255.255.14
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.255.255.11 remote-as 20000
 neighbor 10.255.255.11 update-source Loopback0
 neighbor 10.255.255.13 remote-as 20000
 neighbor 10.255.255.13 update-source Loopback0
```

```
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community extended
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community extended
 exit-address-family
 !
```

CSR16:

```
router bgp 20000
 bgp router-id 10.255.255.16
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.255.255.11 remote-as 20000
 neighbor 10.255.255.11 update-source Loopback0
 neighbor 10.255.255.13 remote-as 20000
 neighbor 10.255.255.13 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community extended
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community extended
 exit-address-family
```

Verification:

```
CSR11(config)#do sh bgp vpnv4 uni all summ | begin Neighbor
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.12   4        20000      78      91       28    0    0 01:07:39        0
10.255.255.13   4        20000     288     300       28    0    0 01:07:35        0
10.255.255.14   4        20000     154     260       28    0    0 01:07:27        0
10.255.255.15   4        20000      73      89       28    0    0 01:07:30        0
10.255.255.16   4        20000      92     287       28    0    0 00:45:28        0

CSR13(config)#do sh bgp vpnv4 uni all summ | begin Neighbor
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.11   4        20000     305     293       28    0    0 01:08:31        0
10.255.255.12   4        20000      79      93       28    0    0 01:08:36        0
10.255.255.14   4        20000     159     266       28    0    0 01:08:24        0
10.255.255.15   4        20000      74      89       28    0    0 01:08:22        0
10.255.255.16   4        20000      90     292       28    0    0 00:46:27        0

CSR13(config)#do sh ip protocols | begin bgp
Routing Protocol is "bgp 20000"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Route Reflector for address family VPNv4 Unicast with the cluster-id 10.255.255.13, 4 clients
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: external 20 internal 200 local 200
CSR11(config)#do sh ip protocols | begin bgp
Routing Protocol is "bgp 20000"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Route Reflector for address family VPNv4 Unicast with the cluster-id 10.255.255.11, 4 clients
  !!!!!###
```

All BGP VPNv4 peerings are UP and IPv4 unicast Address-Family is disabled using **no bgp default ipv4-unicast** command on IOS-XE nodes.

Task 03:

- Provide MPLS L3VPN Service to Customer A (CustA)
- Use BGP as the PE-CE Routing Protocol
- Use the same AS# 65001 on all CE devices for BGP
- CustA-R1 and CustA-R2 should form iBGP sessions with each other, Configure SoO extended community value on PE devices for the sites **if it is necessary**.

Solution:

MPLS L3VPN configuration concepts are the same as MPLS-Lab-1 tasks, the only difference is we configure MPLS L3VPNs here mostly on IOS-XR instead of IOS/IOS-XE. If you need any detailed explanation for MPLS L3VPNs please refer to the MPLS-Lab-1, **we recommend you to do MPLS-Lab-1 before doing this lab.**

Customer side configuration is straight forward, Let's configure that part first (CE Devices):

CustA-R1:

```
router bgp 65001
 bgp router-id 10.1.255.1
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.1.255.2 remote-as 65001
 neighbor 10.1.255.2 update-source Loopback0
 neighbor 172.16.1.12 remote-as 20000
 !
 address-family ipv4
  network 10.1.255.1 mask 255.255.255.255
  network 172.16.1.0 mask 255.255.255.0
  neighbor 10.1.255.2 activate
  neighbor 172.16.1.12 activate
 exit-address-family
!
ip route 10.1.255.2 255.255.255.255 Ethernet0/1 172.16.12.2
```

CustA-R2:

```
router bgp 65001
 bgp router-id 10.1.255.2
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.1.255.1 remote-as 65001
 neighbor 10.1.255.1 update-source Loopback0
 neighbor 172.16.2.12 remote-as 20000
 !
 address-family ipv4
  network 10.1.255.2 mask 255.255.255.255
  network 172.16.2.0 mask 255.255.255.0
  neighbor 10.1.255.1 activate
  neighbor 172.16.2.12 activate
 exit-address-family
!
ip route 10.1.255.1 255.255.255.255 Ethernet0/1 172.16.12.1
```

CustA-R3:

```
router bgp 65001
 bgp router-id 10.1.255.3
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 172.16.3.15 remote-as 20000
 !
 address-family ipv4
  network 10.1.255.3 mask 255.255.255.255
  network 172.16.3.0 mask 255.255.255.0
  neighbor 172.16.3.15 activate
 exit-address-family
```

CustA-R4:

```
router bgp 65001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 172.16.4.15 remote-as 20000
 !
 address-family ipv4
  network 10.1.255.4 mask 255.255.255.255
  network 172.16.4.0 mask 255.255.255.0
  neighbor 172.16.4.15 activate
 exit-address-family
```

Customer's CE devices are not aware of Service Provider's Underlay network, they just form IPv4 unicast BGP peering with PE devices and advertise connected networks to the BGP. There is no need to advertise PE-CE connections network (PE-CE physical network subnet) but to provide full connectivity we advertise those subnets as well.

Now, it is time to configure Service Provider side, the configuration steps only differs a little bit compared to IOS-XE. Two major differences are:

- On IOS-XR, Route-Distinguisher must be configured under BGP Process (for a particular VRF)
- eBGP session between PE and CE device requires route-policy (for simplicity we only use PASS ALL route-policy), by default IOS-XR does not accept nor send any updates from/to eBGP peer.

XRv15:

```
vrf CustA
 address-family ipv4 unicast
  import route-target
   10.255.255.12:1
  !
  export route-target
   10.255.255.15:1
  !
 !
```

```
!
interface GigabitEthernet0/0/0/6
 vrf CustA
 ipv4 address 172.16.4.15 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/7
 vrf CustA
 ipv4 address 172.16.3.15 255.255.255.0
 no shutdown
!
route-policy PASS_ALL_EXTERNAL
  pass
end-policy
!
router bgp 20000
 vrf CustA
  rd 20000:1
  address-family ipv4 unicast
  !
  neighbor 172.16.3.3
   remote-as 65001
   address-family ipv4 unicast
    route-policy PASS_ALL_EXTERNAL in
    route-policy PASS_ALL_EXTERNAL out
    as-override
   !
  !
  neighbor 172.16.4.4
   remote-as 65001
   address-family ipv4 unicast
    route-policy PASS_ALL_EXTERNAL in
    route-policy PASS_ALL_EXTERNAL out
    as-override
   !
```

We used different route-target values for import and export. The schema is:

PE_Loopback0_IPv4_Address:Customer_Number.

RD is configured under BGP Customer VRF.

As-override is required for this scenario, task expilicitly asks us to use the same AS numbers on Customer Edge devices, by default BGP peer ignores the updates if it sees its own AS number in the AS_Path attribute. There is another option that we could use on CE devices: allow-as. But we prefer provider takes care of this situation.

Route-Policy contains only PASS statement, means if we apply it inbound or outbound, it is going to pass all the updates to/from neighbor.

It is a good idea that RD configuration is under BGP process for a particular VRF in IOS-XR. Configuring it under the global VRF definition (The IOS-XE way) made confusions. RD is a 32 bit value that will be added to the Customer Routes to make them distinguishable from other customer routes, though different customers can use overlapping subnets). RT is an extended community value that will be added to the route community attribute. We can use that value in order to import/export routes into/from VRF table. A VRF can have multiple RT import/export values.

```
XRv12:
vrf CustA
 address-family ipv4 unicast
  import route-target
   10.255.255.15:1
  !
  export route-target
   10.255.255.12:1
  !
 !
!
interface GigabitEthernet0/0/0/3
 vrf CustA
 ipv4 address 172.16.1.12 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/4
 vrf CustA
 ipv4 address 172.16.2.12 255.255.255.0
```

```
 no shutdown
!
route-policy PASS_ALL_EXTERNAL
  pass
end-policy
!
router bgp 20000
 vrf CustA
  rd 20000:1
  address-family ipv4 unicast
  !
  neighbor 172.16.1.1
   remote-as 65001
   address-family ipv4 unicast
    route-policy PASS_ALL_EXTERNAL in
    route-policy PASS_ALL_EXTERNAL out
    as-override
    site-of-origin 65001:12
   !
  !
  neighbor 172.16.2.2
   remote-as 65001
   address-family ipv4 unicast
    route-policy PASS_ALL_EXTERNAL in
    route-policy PASS_ALL_EXTERNAL out
    as-override
    site-of-origin 65001:12
   !
  !
 !
```

The task asks us to configure SoO (Site-Of-Origin) value on PE devices if it is needed. But in this scenario it seems there is no point for using SoO value! CustA-R1 and CustA-R2 have iBGP peering but they are both connected to XRv12.

Let's Try some tests to see what happens:

When SoO is configured and we shutdown the e0/1 of R1 and R2, XRv12 does not advertise R2 learned prefixes to R1:

```
CustA-R1(config-if)#do sh ip int br | ex unas
Interface                  IP-Address      OK? Method Status                Protocol
Ethernet0/0                172.16.1.1      YES TFTP   up                    up
Ethernet0/1                172.16.12.1     YES TFTP   administratively down down
Loopback0                  10.1.255.1      YES TFTP   up                    up


CustA-R1(config-if)#do sh ip bgp summ | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.255.2      4        65001       0       0        1    0    0 00:04:23 Idle
172.16.1.12     4        20000      29      25       55    0    0 00:11:56        4


 CustA-R1(config-if)#do sh ip bgp | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
 *>   10.1.255.1/32    0.0.0.0                  0         32768 i
 *>   10.1.255.3/32    172.16.1.12                            0 20000 20000 i
 *>   10.1.255.4/32    172.16.1.12                            0 20000 20000 i
 *>   172.16.1.0/24    0.0.0.0                  0         32768 i
 *>   172.16.3.0/24    172.16.1.12                            0 20000 20000 i
 *>   172.16.4.0/24    172.16.1.12                            0 20000 20000 i
```

This time, we remove the Site-Of-Origin value to see what happens:

XRv12:

```
router bgp 20000
 vrf CustA
  address-family ipv4 unicast
  !
  neighbor 172.16.1.1
   address-family ipv4 unicast
    no site-of-origin 65001:12
   !
  !
  neighbor 172.16.2.2
```

```
   address-family ipv4 unicast
    no site-of-origin 65001:12
   !
  !
 !
commit
```

```
CustA-R1(config-if)#
*Apr 23 18:38:32.303: %BGP-5-ADJCHANGE: neighbor 10.1.255.2 Up


CustA-R1(config-if)#do sh ip bgp | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
 * i  10.1.255.1/32    172.16.2.12              0    100      0 20000 20000 i
 *>                    0.0.0.0                  0         32768 i
 * i  10.1.255.2/32    10.1.255.2               0    100      0 i
 * i  10.1.255.3/32    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
 * i  10.1.255.4/32    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
 * i  172.16.1.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                    0.0.0.0                  0         32768 i
 * i  172.16.2.0/24    10.1.255.2               0    100      0 i
 * i  172.16.3.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
 * i  172.16.4.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
```

Interesting! iBGP session between R1 and R2 has been established using the XRv12 links as the transit network! And R1 learns 10.1.255.2 (The loopback 0 interface of R2) from R2! Let's continue watching the logs:

```
CustA-R1(config-if)#do sh ip bgp | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
 * i  10.1.255.1/32    172.16.2.12              0    100      0 20000 20000 i
 *>                    0.0.0.0                  0         32768 i
 *    10.1.255.2/32    172.16.1.12                           0 20000 20000 i
 *>i                   10.1.255.2               0    100      0 i
```

```
* i  10.1.255.3/32    172.16.2.12              0    100      0 20000 20000 i
 *>                   172.16.1.12                            0 20000 20000 i
 * i  10.1.255.4/32    172.16.2.12              0    100      0 20000 20000 i
 *>                   172.16.1.12                            0 20000 20000 i
 * i  172.16.1.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                   0.0.0.0                  0         32768 i
 *    172.16.2.0/24    172.16.1.12                            0 20000 20000 i
 *>i                  10.1.255.2               0    100      0 i
 * i  172.16.3.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                   172.16.1.12                            0 20000 20000 i
 * i  172.16.4.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                   172.16.1.12                            0 20000 20000 i
CustA-R1(config-if)#
*Apr 23 18:41:32.371: %BGP-5-NBR_RESET: Neighbor 10.1.255.2 reset (Peer closed the session)
*Apr 23 18:41:32.371: %BGP-5-ADJCHANGE: neighbor 10.1.255.2 Down Peer closed the session
*Apr 23 18:41:32.371: %BGP_SESSION-5-ADJCHANGE: neighbor 10.1.255.2 IPv4 Unicast topology base removed from
session  Peer closed the session
CustA-R1(config-if)#
*Apr 23 18:41:42.329: %BGP-5-ADJCHANGE: neighbor 10.1.255.2 Up
*Apr 23 18:44:42.929: %BGP-3-NOTIFICATION: sent to neighbor 10.1.255.2 4/0 (hold time expired) 0 bytes
CustA-R1(config-if)#
*Apr 23 18:44:42.930: %BGP-5-NBR_RESET: Neighbor 10.1.255.2 reset (BGP Notification sent)
*Apr 23 18:44:42.930: %BGP-5-ADJCHANGE: neighbor 10.1.255.2 Down BGP Notification sent
*Apr 23 18:44:42.930: %BGP_SESSION-5-ADJCHANGE: neighbor 10.1.255.2 IPv4 Unicast topology base removed from
session  BGP Notification sent
```

R1 learns 10.1.255.2 from both XRv12 and R2 (using XRv12 as a transit link) and because of Shorter AS_Path value, R1 chooses 10.1.255.2 as the next hop.

This is called recursive routing error. Just like the recursive routing error we face using tunnel interfaces.

By enabling SoO again, we solve this problem:

XRv12:

```
router bgp 20000
 vrf CustA
  address-family ipv4 unicast
  !
```

```
  neighbor 172.16.1.1
   address-family ipv4 unicast
    site-of-origin 65001:12
   !
  !
  neighbor 172.16.2.2
   address-family ipv4 unicast
    site-of-origin 65001:12
   !
  !
 !
commit
```

This time XRv12 is not going to advertise R2 learned routes to R1, though R1 cannot form iBGP session with R2 using Loopback 0 interfaces.

Verification:

```
CustA-R1(config-if)#do sh ip bgp | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
 *>   10.1.255.1/32    0.0.0.0                  0         32768 i
 *>   10.1.255.3/32    172.16.1.12                            0 20000 20000 i
 *>   10.1.255.4/32    172.16.1.12                            0 20000 20000 i
 *>   172.16.1.0/24    0.0.0.0                  0         32768 i
 *>   172.16.3.0/24    172.16.1.12                            0 20000 20000 i
 *>   172.16.4.0/24    172.16.1.12                            0 20000 20000 i



CustA-R1(config-if)#int e0/1
CustA-R1(config-if)#no shutdown
CustA-R1(config-if)#
*Apr 23 18:54:13.421: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to up
*Apr 23 18:54:14.421: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/1, changed state to up
CustA-R1(config-if)#
*Apr 23 18:54:14.732: %BGP-5-ADJCHANGE: neighbor 10.1.255.2 Up


CustA-R1(config-if)#do sh ip bgp | begin Net
```

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>   10.1.255.1/32    0.0.0.0                  0         32768 i
 r>i  10.1.255.2/32    10.1.255.2               0    100      0 i
 * i  10.1.255.3/32    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
 * i  10.1.255.4/32    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
 *>   172.16.1.0/24    0.0.0.0                  0         32768 i
 *>i  172.16.2.0/24    10.1.255.2               0    100      0 i
 * i  172.16.3.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i
 * i  172.16.4.0/24    172.16.2.12              0    100      0 20000 20000 i
 *>                    172.16.1.12                           0 20000 20000 i


CustA-R1(config-if)#do trace 10.1.255.3 source lo 0 probe 2
Type escape sequence to abort.
Tracing the route to 10.1.255.3
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.1.12 10 msec 3 msec
  2 10.12.99.99 [MPLS: Labels 16015/150002 Exp 0] 173 msec 24 msec
  3 10.15.99.15 [MPLS: Label 150002 Exp 0] 28 msec 19 msec
  4 172.16.3.3 [AS 20000] 19 msec 21 msec


CustA-R1(config-if)#do trace 10.1.255.4 source lo 0 probe 2
Type escape sequence to abort.
Tracing the route to 10.1.255.4
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.1.12 4 msec 2 msec
  2 10.12.99.99 [MPLS: Labels 16015/150003 Exp 0] 26 msec 24 msec
  3 10.15.99.15 [MPLS: Label 150003 Exp 0] 19 msec 19 msec
  4 172.16.4.4 [AS 20000] 23 msec 18 msec
```

Task 04:

- Provide MPLS L3VPN Service to Customer B (CustB).
- Use OSPF as the PE-CE Routing Protocol.
- Customer Edge devices sould receive PE learned routes as Inter-Area.

Solution:

Configuration steps are the same as OSPF configuration in MPLS-Lab-1, but instead we use IOS-XR here.

By default IOS/IOS-XE sets the domain-id based on the OSPF Process ID value, but in IOS-XR we smust configure Domain-ID value manually.

First, we configure CE devices, the steps are straight forward:

```
CustB-R1:
router ospf 2
 router-id 10.2.255.1
 passive-interface default
 no passive-interface Ethernet0/0
 network 10.2.255.1 0.0.0.0 area 0
 network 172.17.1.0 0.0.0.255 area 0
!

CustB-R2:
router ospf 2
 router-id 10.2.255.2
 passive-interface default
 no passive-interface Ethernet0/0
 network 10.2.255.2 0.0.0.0 area 0
 network 172.17.2.0 0.0.0.255 area 0
!

CustB-R3:
router ospf 2
 router-id 10.2.255.3
 passive-interface default
```

```
 no passive-interface Ethernet0/0
 network 10.2.255.3 0.0.0.0 area 0
 network 172.17.3.0 0.0.0.255 area 0
 network 192.168.2.254 0.0.0.0 area 0
```

Now, time to configure PE devices. If we don't configure domain-id on PE devices, CE devices receive routes as External. The task expilicitly asks us to make them Inter-Area.

IOS/IOS-XE by default create domain-id type 0005 based on the ospf process id, but IOS-XR needs to be manually configured. Because of the IOS-XE default behavior compatibility, we use Type 0005 domain-id.

```
XRv12:
vrf CustB
 address-family ipv4 unicast
  import route-target
   10.255.255.16:2
  !
  export route-target
   10.255.255.12:2
  !
 !
!
interface GigabitEthernet0/0/0/5
 vrf CustB
 ipv4 address 172.17.1.12 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/6
 vrf CustB
 ipv4 address 172.17.2.12 255.255.255.0
 no shutdown
!
router ospf 2
 vrf CustB
  router-id 10.2.255.12
  domain-id type 0005 value 000000000002
```

```
  redistribute bgp 20000
  address-family ipv4 unicast
  area 0
   interface GigabitEthernet0/0/0/5
   !
   interface GigabitEthernet0/0/0/6
   !
  !
 !
!
router bgp 20000
 vrf CustB
  rd 20000:2
  address-family ipv4 unicast
   redistribute ospf 2 match internal external
  !
 !
!
```

By default, IOS-XR does redistribute both OSPF internal and external prefixes, but IOS-XE does not redistribute OSPF external type routes.

The mutual redistribution is needed in order to import and export routes to/from BGP VPNv4.

This step is not required when we use BGP as a PE-CE routing protocol.

CSR16:

```
vrf definition CustB
 rd 20000:2
 !
 address-family ipv4
  route-target import 10.255.255.12:2
  route-target export 10.255.255.16:2
 exit-address-family
!
!
interface GigabitEthernet8
```

```
 vrf forwarding CustB
 ip address 172.17.3.16 255.255.255.0
 negotiation auto
 no shutdown
!
router ospf 2 vrf CustB
 router-id 10.2.255.16
 redistribute bgp 20000
 passive-interface default
 no passive-interface GigabitEthernet8
 network 172.17.3.0 0.0.0.255 area 0
!
router bgp 20000
 !
 address-family ipv4 vrf CustB
  redistribute ospf 2 match internal external 1 external 2
 exit-address-family
!
```

Verification:

```
CustB-R3(config-router)#do sh ip route ospf | include 10.2.255
O E2     10.2.255.1 [110/2] via 172.17.3.16, 00:08:59, Ethernet0/0
O E2     10.2.255.2 [110/2] via 172.17.3.16, 00:08:59, Ethernet0/0

CustB-R1(config-router)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
O        10.2.255.2 [110/12] via 172.17.1.12, 12:50:54, Ethernet0/0
O E2     10.2.255.3 [110/2] via 172.17.1.12, 00:11:17, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.2.0/24 [110/11] via 172.17.1.12, 12:50:54, Ethernet0/0
O E2     172.17.3.0/24 [110/1] via 172.17.1.12, 00:11:49, Ethernet0/0

CustB-R2(config-router)#do sh ip route ospf | begin Gate
Gateway of last resort is not set
```

```
      10.0.0.0/32 is subnetted, 3 subnets
O        10.2.255.1 [110/12] via 172.17.2.12, 12:51:26, Ethernet0/0
O E2     10.2.255.3 [110/2] via 172.17.2.12, 00:11:41, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.1.0/24 [110/11] via 172.17.2.12, 12:51:26, Ethernet0/0
O E2     172.17.3.0/24 [110/1] via 172.17.2.12, 00:12:13, Ethernet0/0
```

I have already mentioned that IOS-XE sets the domain-id value based on OSPF Process ID automatically. But Routes are still External Type 2, Let's check BGP VPNv4 route details:

```
CSR16(config)#do sh bgp vpnv4 uni rd 20000:2 10.2.255.2/32
BGP routing table entry for 20000:2:10.2.255.2/32, version 12
Paths: (2 available, best #1, table CustB)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    10.255.255.12 (metric 20) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 2, localpref 100, valid, internal, best
      Extended Community: OSPF DOMAIN ID:0x0005:0x000000000002
        RT:10.255.255.12:2 OSPF RT:0.0.0.0:1:0 OSPF ROUTER ID:10.2.255.12:0
      Originator: 10.255.255.12, Cluster list: 10.255.255.11
      mpls labels in/out nolabel/24007
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 08:16:49 UTC
  Refresh Epoch 2
  Local
    10.255.255.12 (metric 20) (via default) from 10.255.255.13 (10.255.255.13)
      Origin incomplete, metric 2, localpref 100, valid, internal
      Extended Community: OSPF DOMAIN ID:0x0005:0x000000000002
        RT:10.255.255.12:2 OSPF RT:0.0.0.0:1:0 OSPF ROUTER ID:10.2.255.12:0
      Originator: 10.255.255.12, Cluster list: 10.255.255.13
      mpls labels in/out nolabel/24007
      rx pathid: 0, tx pathid: 0
      Updated on Apr 23 2020 08:16:47 UTC


CSR16(config)#do sh bgp vpnv4 uni rd 20000:2 10.2.255.3/32
```

```
BGP routing table entry for 20000:2:10.2.255.3/32, version 27
Paths: (1 available, best #1, table CustB)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    172.17.3.3 (via vrf CustB) from 0.0.0.0 (10.255.255.16)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid, sourced, best
      Extended Community: OSPF DOMAIN ID:0x0005:0x000000020200
        RT:10.255.255.16:2 OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.2.255.16:0
      mpls labels in/out 160007/nolabel
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 19:25:18 UTC
```

IOS-XE automatically generated domain-id is different than our manual domain-id on IOS-XR, if both PE devices were both IOS-XE and the OSPF Process ID were the same, then the CE learned routes would be Inter-Area. Let's configure domain-id on IOS-XE manually too:

```
CSR16:
router ospf 2 vrf CustB
 domain-id type 0005 value 000000000002
!
```

Verification:

```
CustB-R3(config-router)#do sh ip route ospf | include 10.2.255
O IA     10.2.255.1 [110/12] via 172.17.3.16, 00:01:27, Ethernet0/0
O IA     10.2.255.2 [110/12] via 172.17.3.16, 00:01:27, Ethernet0/0

CustB-R1(config-router)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
O        10.2.255.2 [110/12] via 172.17.1.12, 13:00:57, Ethernet0/0
O IA     10.2.255.3 [110/12] via 172.17.1.12, 00:00:52, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.2.0/24 [110/11] via 172.17.1.12, 13:00:57, Ethernet0/0
```

```
O IA     172.17.3.0/24 [110/11] via 172.17.1.12, 00:00:52, Ethernet0/0



CustB-R2(config-router)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
O        10.2.255.1 [110/12] via 172.17.2.12, 13:01:15, Ethernet0/0
O IA     10.2.255.3 [110/12] via 172.17.2.12, 00:01:02, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.1.0/24 [110/11] via 172.17.2.12, 13:01:15, Ethernet0/0
O IA     172.17.3.0/24 [110/11] via 172.17.2.12, 00:01:02, Ethernet0/0
```

The other default behavior of IOS-XE is redistribution of OSPF into BGP:

```
CSR16(config-router)#do sh run vrf CustB | begin router bgp
router bgp 20000
 !
 address-family ipv4 vrf CustB
  redistribute ospf 2 match internal external 1 external 2
 exit-address-family
!
CSR16(config)#router bgp 20000
CSR16(config-router)#address-family ipv4 vrf CustB
CSR16(config-router-af)#no redistribute ospf 2 match internal external 1 external 2
CSR16(config-router-af)#do sh run vrf CustB | begin router bgp
router bgp 20000
 !
 address-family ipv4 vrf CustB
  redistribute ospf 2
 exit-address-family
!
```

Match internal external is now removed and only **redistribute ospf 2** command is used.

Let's create a new loopback interface on CustB-R1 and redistribute it into OSPF to see if CSR16 by default injects it into BGP or not:

```
CustB-R3:
interface Loopback130
 ip address 10.130.255.3 255.255.255.255
!
router ospf 2
 redistribute connected subnets
!
```

Verification:

```
CustB-R1(config)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
O        10.2.255.2 [110/12] via 172.17.1.12, 13:21:12, Ethernet0/0
O IA     10.2.255.3 [110/12] via 172.17.1.12, 00:21:07, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.2.0/24 [110/11] via 172.17.1.12, 13:21:12, Ethernet0/0
O IA     172.17.3.0/24 [110/11] via 172.17.1.12, 00:21:07, Ethernet0/0

CustB-R2(config-router)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
O        10.2.255.1 [110/12] via 172.17.2.12, 13:21:41, Ethernet0/0
O IA     10.2.255.3 [110/12] via 172.17.2.12, 00:21:28, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.1.0/24 [110/11] via 172.17.2.12, 13:21:41, Ethernet0/0
O IA     172.17.3.0/24 [110/11] via 172.17.2.12, 00:21:28, Ethernet0/0
```

There is no External Type 5 routes on R1 and R2, This time we expilicitly tell IOS-XE to redistribute both Internal and External routes:

```
CSR16:
router bgp 20000
 !
 address-family ipv4 vrf CustB
  redistribute ospf 2 match internal external 1 external 2
```

```
 exit-address-family
!
```

Verification:

```
CustB-R1(config)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 4 subnets
O        10.2.255.2 [110/12] via 172.17.1.12, 13:23:20, Ethernet0/0
O IA     10.2.255.3 [110/12] via 172.17.1.12, 00:23:15, Ethernet0/0
O E2     10.130.255.3 [110/20] via 172.17.1.12, 00:01:15, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.2.0/24 [110/11] via 172.17.1.12, 13:23:20, Ethernet0/0
O IA     172.17.3.0/24 [110/11] via 172.17.1.12, 00:23:15, Ethernet0/0

CustB-R2(config-router)#do sh ip route ospf | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 4 subnets
O        10.2.255.1 [110/12] via 172.17.2.12, 13:23:41, Ethernet0/0
O IA     10.2.255.3 [110/12] via 172.17.2.12, 00:23:28, Ethernet0/0
O E2     10.130.255.3 [110/20] via 172.17.2.12, 00:01:29, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.17.1.0/24 [110/11] via 172.17.2.12, 13:23:41, Ethernet0/0
O IA     172.17.3.0/24 [110/11] via 172.17.2.12, 00:23:28, Ethernet0/0

CustB-R2(config-router)#do trace 10.130.255.3 source lo 0 probe 1
Type escape sequence to abort.
Tracing the route to 10.130.255.3
VRF info: (vrf in name/id, vrf out name/id)
  1 172.17.2.12 14 msec
  2 10.12.99.99 [MPLS: Labels 16016/160011 Exp 0] 23 msec
  3 172.17.3.16 [MPLS: Label 160011 Exp 0] 16 msec
  4 172.17.3.3 14 msec
```

This time CSR16 redistributes both internal and external OSPF routes and sends them into BGP VPNv4.

Task 05:

- Provide MPLS L3VPN Service to Customer C (CustC).
- Use EIGRP as the PE-CE Routing Protocol.

Solution:

Configuring EIGRP as the PE-CE routing protocol has less steps than OSPF, and it is also straight forward, we will use EIGRP named mode because IOS-XR only supports EIGRP named mode, of course EIGRP Classic and Named Mode are compatible with each other but anyways it is better to use Named Mode on all devices:

CustC-R1:

```
router eigrp CustC
 !
 address-family ipv4 unicast autonomous-system 3
  !
  af-interface default
   passive-interface
  exit-af-interface
  !
  af-interface Ethernet0/0
   no passive-interface
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 10.3.255.1 0.0.0.0
  network 172.18.1.0 0.0.0.255
 exit-address-family
```

CustC-R2:

```
router eigrp CustC
 !
 address-family ipv4 unicast autonomous-system 3
```

```
  !
  af-interface default
   passive-interface
  exit-af-interface
  !
  af-interface Ethernet0/0
   no passive-interface
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 10.3.255.2 0.0.0.0
  network 172.18.2.0 0.0.0.255
 exit-address-family
```

CustC-R3:

```
router eigrp CustC
 !
 address-family ipv4 unicast autonomous-system 3
  !
  af-interface default
   passive-interface
  exit-af-interface
  !
  af-interface Ethernet0/0
   no passive-interface
  exit-af-interface
  !
  topology base
  exit-af-topology
  network 10.3.255.3 0.0.0.0
  network 172.18.3.0 0.0.0.255
 exit-address-family
```

XRv12:

```
vrf CustC
 address-family ipv4 unicast
  import route-target
   10.255.255.14:3
  !
  export route-target
   10.255.255.12:3
  !
 !
!
interface GigabitEthernet0/0/0/7
 vrf CustC
 ipv4 address 172.18.1.12 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/8
 vrf CustC
 ipv4 address 172.18.2.12 255.255.255.0
 no shutdown
!
router eigrp CustC
 vrf CustC
  address-family ipv4
   autonomous-system 3
   redistribute bgp 20000
   interface GigabitEthernet0/0/0/7
   !
   interface GigabitEthernet0/0/0/8
   !
  !
 !
!
router bgp 20000
 vrf CustC
```

```
  rd 20000:3
  address-family ipv4 unicast
   redistribute eigrp 3
  !
 !
!
commit

CSR14:
vrf definition CustC
 rd 20000:3
 !
 address-family ipv4
  route-target export 10.255.255.14:3
  route-target import 10.255.255.12:3
 exit-address-family
!
!
interface GigabitEthernet9
 vrf forwarding CustC
 ip address 172.18.3.14 255.255.255.0
 negotiation auto
 no shutdown
!
router eigrp CustC
!
address-family ipv4 unicast vrf CustC autonomous-system 3
  !
  af-interface default
   passive-interface
  exit-af-interface
  !
  af-interface GigabitEthernet9
   no passive-interface
  exit-af-interface
```

```
  !
  topology base
   redistribute bgp 20000
  exit-af-topology
  network 172.18.3.0 0.0.0.255
 exit-address-family
!
router bgp 20000
 !
 address-family ipv4 vrf CustC
  redistribute eigrp 3
 exit-address-family
!
```

Verification:

```
CSR14(config)#do sh bgp vpnv4 uni rd 20000:3 10.3.255.2/32
BGP routing table entry for 20000:3:10.3.255.2/32, version 13
Paths: (2 available, best #1, table CustC)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.255.255.12 (metric 20) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 10880, localpref 100, valid, internal, best
      Extended Community: RT:10.255.255.12:3 Cost:pre-bestpath:128:10880
        0x8800:32768:0 0x8801:3:288 0x8802:65281:2560 0x8803:1:1500
        0x8806:0:168034050
      Originator: 10.255.255.12, Cluster list: 10.255.255.11
      mpls labels in/out nolabel/24010
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 07:55:48 UTC
  Refresh Epoch 1
  Local
    10.255.255.12 (metric 20) (via default) from 10.255.255.13 (10.255.255.13)
      Origin incomplete, metric 10880, localpref 100, valid, internal
      Extended Community: RT:10.255.255.12:3 Cost:pre-bestpath:128:10880
```

```
        0x8800:32768:0 0x8801:3:288 0x8802:65281:2560 0x8803:1:1500
        0x8806:0:168034050
      Originator: 10.255.255.12, Cluster list: 10.255.255.13
      mpls labels in/out nolabel/24010
      rx pathid: 0, tx pathid: 0
      Updated on Apr 23 2020 07:55:48 UTC



CSR14(config)#do sh ip eigrp vrf CustC topology 10.3.255.2/32
EIGRP-IPv4 VR(CustC) Topology Entry for AS(3)/ID(172.18.3.14)
           Topology(base) TID(0) VRF(CustC)
EIGRP-IPv4(3): Topology base(0) entry for 10.3.255.2/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 1392640
  Descriptor Blocks:
  10.255.255.12, from VPNv4 Sourced, Send flag is 0x0
      Composite metric is (1392640/0), route is Internal (VPNv4 Sourced)
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
        Total delay is 11250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 10.3.255.2
```

EIGRP uses composite metric, by default metric is calculated based on Bandwidth and Delay. VPNv4 can transfer the composite metric parameters using Extended Community attribute to other PE devices. The only caveat is to do mutual redistribution without any seed/specific metric at the end of the redistribute command.

```
CustC-R1(config)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set


      10.0.0.0/32 is subnetted, 3 subnets
D        10.3.255.2 [90/1029760] via 172.18.1.12, 13:48:37, Ethernet0/0
D        10.3.255.3 [90/1029760] via 172.18.1.12, 12:38:23, Ethernet0/0
```

```
      172.18.0.0/16 is variably subnetted, 4 subnets, 2 masks
D        172.18.2.0/24 [90/1029120] via 172.18.1.12, 13:48:39, Ethernet0/0
D        172.18.3.0/24 [90/1029120] via 172.18.1.12, 12:38:23, Ethernet0/0



CustC-R3(config)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
D        10.3.255.1 [90/1029760] via 172.18.3.14, 12:39:25, Ethernet0/0
D        10.3.255.2 [90/1029760] via 172.18.3.14, 12:39:25, Ethernet0/0
      172.18.0.0/16 is variably subnetted, 4 subnets, 2 masks
D        172.18.1.0/24 [90/1029120] via 172.18.3.14, 12:39:25, Ethernet0/0
D        172.18.2.0/24 [90/1029120] via 172.18.3.14, 12:39:25, Ethernet0/0

CustC-R3(config)#do sh ip eigrp topol 10.3.255.2/32
EIGRP-IPv4 VR(CustC) Topology Entry for AS(3)/ID(10.3.255.3) for 10.3.255.2/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 131809280, RIB is 1029760
  Descriptor Blocks:
  172.18.3.14 (Ethernet0/0), from 172.18.3.14, Send flag is 0x0
      Composite metric is (131809280/1392640), route is Internal
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 1011250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 10.3.255.2

CustC-R3(config)#do sh ip eigrp topol 10.3.255.1/32
EIGRP-IPv4 VR(CustC) Topology Entry for AS(3)/ID(10.3.255.3) for 10.3.255.1/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 131809280, RIB is 1029760
```

```
Descriptor Blocks:
  172.18.3.14 (Ethernet0/0), from 172.18.3.14, Send flag is 0x0
      Composite metric is (131809280/1392640), route is Internal
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 1011250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 10.3.255.1



CustC-R3(config)#do trace 10.3.255.1 source lo 0 probe 1
Type escape sequence to abort.
Tracing the route to 10.3.255.1
VRF info: (vrf in name/id, vrf out name/id)
  1 172.18.3.14 3 msec
  2 10.14.99.99 [MPLS: Labels 16012/24009 Exp 0] 33 msec
  3 10.12.99.12 [MPLS: Label 24009 Exp 0] 27 msec
  4 172.18.1.1 14 msec



CustC-R3(config)#do trace 10.3.255.2 source lo 0 probe 1
Type escape sequence to abort.
Tracing the route to 10.3.255.2
VRF info: (vrf in name/id, vrf out name/id)
  1 172.18.3.14 1 msec
  2 10.14.99.99 [MPLS: Labels 16012/24010 Exp 0] 15 msec
  3 10.12.99.12 [MPLS: Label 24010 Exp 0] 13 msec
  4 172.18.2.2 11 msec
```

As you can see, the composite metric of the routes is not lost during redistribution.

Task 06:

- Provide MPLS L3VPN Service to Customer D (CustD).
- Use Static-Routes as the PE-CE Routing Protocol.

Solution:

Instead of using Dynamic Routing protocols like OSPF, BGP, EIGRP, RIP and even IS-IS as a PE-CE protocol, we can use Static-Routes and then redistribute them into BGP VPNv4 Unicast.

Remember that using Static Routes is not a scalable solution, and puts a lot of burden on Network Operator/Admin.

```
CustD-R1:
ip route 0.0.0.0 0.0.0.0 Ethernet0/0 172.19.1.12


CustD-R2:


ip route 0.0.0.0 0.0.0.0 Ethernet0/0 172.19.2.15


CustD-R3:


ip route 0.0.0.0 0.0.0.0 Ethernet0/0 172.19.3.15
```

Only a s imple step on Customer Edge devices! A default route with the next-hop of PE device.

```
XRv12:
vrf CustD
 address-family ipv4 unicast
  import route-target
   10.255.255.15:4
  !
  export route-target
   10.255.255.12:4
  !
 !
!
interface GigabitEthernet0/0/0/9
```

```
vrf CustD
 ipv4 address 172.19.1.12 255.255.255.0
 no shutdown
!
router static
 vrf CustD
  address-family ipv4 unicast
   10.4.255.1/32 GigabitEthernet0/0/0/9 172.19.1.1 metric 10
  !
 !
!
router bgp 20000
 vrf CustD
  rd 20000:4
  address-family ipv4 unicast
   redistribute connected metric 110
   redistribute static metric 110
  !
 !
commit
```

In IOS-XR, **router static** command is used to write static route statements instead of **ip route** command of IOS-XE.

```
RP/0/0/CPU0:XRv12(config)#do sh route vrf CustD ipv4  | begin Gate
Thu Apr 23 21:01:11.853 UTC
Gateway of last resort is not set

S    10.4.255.1/32 [1/0] via 172.19.1.1, 14:15:56, GigabitEthernet0/0/0/9
C    172.19.1.0/24 is directly connected, 14:15:56, GigabitEthernet0/0/0/9
L    172.19.1.12/32 is directly connected, 14:15:56, GigabitEthernet0/0/0/9



RP/0/0/CPU0:XRv12(config)#do sh bgp vpnv4 uni rd 20000:4 10.4.255.1
Thu Apr 23 20:59:14.081 UTC
```

```
BGP routing table entry for 10.4.255.1/32, Route Distinguisher: 20000:4
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                55          55
    Local Label: 24012
Last Modified: Apr 23 06:47:21.635 for 14:11:52
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.2
  Local
    172.19.1.1 from 0.0.0.0 (10.255.255.12)
      Origin incomplete, metric 110, localpref 100, weight 32768, valid, redistributed, best, group-best,
import-candidate
      Received Path ID 0, Local Path ID 0, version 55
      Extended community: RT:10.255.255.12:4
```

Let's configure the other PE:

XRv15:

```
vrf CustD
 address-family ipv4 unicast
  import route-target
   10.255.255.12:4
  !
  export route-target
   10.255.255.15:4
 !
!
interface GigabitEthernet0/0/0/8
 vrf CustD
 ipv4 address 172.19.3.15 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/9
```

```
 vrf CustD
 ipv4 address 172.19.2.15 255.255.255.0
 no shutdown
!
router static
 vrf CustD
  address-family ipv4 unicast
   10.4.255.2/32 GigabitEthernet0/0/0/9 172.19.2.2
   10.4.255.3/32 GigabitEthernet0/0/0/8 172.19.3.3
  !
 !
!
router bgp 20000
 vrf CustD
  rd 20000:4
  address-family ipv4 unicast
   redistribute connected metric 110
   redistribute static metric 110
  !
 commit
```

XRv15 is configured the same way as XRv12.

Verification:

```
RP/0/0/CPU0:XRv12(config)#do sh route vrf CustD ipv4  | begin Gate
Thu Apr 23 21:17:51.554 UTC
Gateway of last resort is not set


S    10.4.255.1/32 [1/0] via 172.19.1.1, 14:32:35, GigabitEthernet0/0/0/9
B    10.4.255.2/32 [200/110] via 10.255.255.15 (nexthop in vrf default), 00:00:35
B    10.4.255.3/32 [200/110] via 10.255.255.15 (nexthop in vrf default), 00:00:35
C    172.19.1.0/24 is directly connected, 14:32:35, GigabitEthernet0/0/0/9
L    172.19.1.12/32 is directly connected, 14:32:35, GigabitEthernet0/0/0/9
B    172.19.2.0/24 [200/110] via 10.255.255.15 (nexthop in vrf default), 13:22:03
B    172.19.3.0/24 [200/110] via 10.255.255.15 (nexthop in vrf default), 13:22:03
```

```
RP/0/0/CPU0:XRv12(config)#do sh cef vrf CustD
Thu Apr 23 21:18:10.543 UTC

Prefix              Next Hop            Interface
------------------- ------------------- ------------------
0.0.0.0/0           drop                default handler
0.0.0.0/32          broadcast
10.4.255.1/32       172.19.1.1/32       GigabitEthernet0/0/0/9
10.4.255.2/32       10.255.255.15/32    <recursive>
10.4.255.3/32       10.255.255.15/32    <recursive>
172.19.1.0/24       attached            GigabitEthernet0/0/0/9
172.19.1.0/32       broadcast           GigabitEthernet0/0/0/9
172.19.1.12/32      receive             GigabitEthernet0/0/0/9
172.19.1.255/32     broadcast           GigabitEthernet0/0/0/9
172.19.2.0/24       10.255.255.15/32    <recursive>
172.19.3.0/24       10.255.255.15/32    <recursive>
224.0.0.0/4         0.0.0.0/32
224.0.0.0/24        receive
255.255.255.255/32  broadcast



RP/0/0/CPU0:XRv12(config)#do sh cef vrf CustD 10.4.255.2/32
Thu Apr 23 21:18:38.951 UTC
10.4.255.2/32, version 131, internal 0x1000001 0x0 (ptr 0xa14051f4) [1], 0x0 (0x0), 0x208 (0xa1583500)
 Updated Apr 23 21:17:15.906
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 10.255.255.15/32, 3 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa15ec4f4 0x0]
    recursion-via-/32
    next hop VRF - 'default', table - 0xe0000000
    next hop 10.255.255.15/32 via 16015/0/21
     next hop 10.12.99.99/32 Gi0/0/0/2    labels imposed {16015 150006}


RP/0/0/CPU0:XRv12(config)#do sh bgp vpnv4 unicast rd 20000:4 10.4.255.2/32
Thu Apr 23 21:19:21.048 UTC
```

```
BGP routing table entry for 10.4.255.2/32, Route Distinguisher: 20000:4
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 597         597
Last Modified: Apr 23 21:17:15.635 for 00:02:05
Paths: (2 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local
    10.255.255.15 (metric 20) from 10.255.255.11 (10.255.255.15)
      Received Label 150006
      Origin incomplete, metric 110, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 0, version 597
      Extended community: RT:10.255.255.15:4 RT:10.255.255.15:20004
      Originator: 10.255.255.15, Cluster list: 10.255.255.11
      Source AFI: VPNv4 Unicast, Source VRF: CustD, Source Route Distinguisher: 20000:4
  Path #2: Received by speaker 0
  Not advertised to any peer
  Local
    10.255.255.15 (metric 20) from 10.255.255.13 (10.255.255.15)
      Received Label 150006
      Origin incomplete, metric 110, localpref 100, valid, internal, import-candidate, imported
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: RT:10.255.255.15:4 RT:10.255.255.15:20004
      Originator: 10.255.255.15, Cluster list: 10.255.255.13
      Source AFI: VPNv4 Unicast, Source VRF: CustD, Source Route Distinguisher: 20000:4
```

All Networks has been learned using BGP VPNv4 Unicast Address-Family.

Task 07:

- Configure Route-Leaking between two VRFs:

CustD-R2 and CustD-R3 should be able to reach Loopback 100 of CustB-R3 using their Loopback 0 interfaces.

Solution:

It is possible to leak routes between VRFs (Extranet), this is possible using Route-Target Values. We can even leak routes between a specific VRF and the Global routing table (GRT or Default VRF).

R2 and R3 routers of CustD want to access a prefix in CustB Site 3 (192.168.2.254/32).

```
CSR16:
ip prefix-list CustB_R3_Lo100: 1 entries
   seq 5 permit 192.168.2.254/32
ip prefix-list CustD_Lo0s: 1 entries
   seq 5 permit 10.4.255.0/24 ge 32
!
route-map CustB_Export_Map permit 50
 match ip address prefix-list CustB_R3_Lo100
 set extcommunity rt 10.255.255.16:24
route-map CustB_Export_Map permit 60
 set extcommunity rt 10.255.255.16:2
!
vrf definition CustB
address-family ipv4
  no route-target export 10.255.255.16:2
  export map CustB_Export_Map
  route-target import 10.255.255.15:20004
!
XRv15:
route-policy CustD_Export_Policy
  if destination in (10.4.255.0/24 eq 32) then
    set extcommunity rt (10.255.255.15:4) additive
    set extcommunity rt (10.255.255.15:20004) additive
```

```
    pass
  else
    set extcommunity rt (10.255.255.15:4) additive
    pass
  endif
end-policy
!
vrf CustD
 address-family ipv4 unicast
  import route-target
   10.255.255.12:4
   10.255.255.16:24
  !
  no export route-target
  export route-policy CustD_Export_Policy
 !
commit
```

Configuration seems complex but it is straight forward:

CSR16 uses a route-map to set different RT extended community values to different prefixes (Loopback 0 and Loopback 100). Then it is applied to an export map under VRF CustB.

XRv15 imports the VRF CustB exported routes with the RT value of 10.255.255.16:24, and also exports Loopback 0 interface addresses of Customer devices with the RT values of 10.255.255.15:4 and 10.255.255.15:20004! Yes routes can have multiple RT values, XRv12 uses 10.255.255.15:4 to import these prefixes and CSR16 uses 10.255.255.15:20004 to import the same prefixes to the CustB vrf.

Verification:

```
CSR16(config)#do sh bgp vpnv4 unicast rd 20000:2 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:2 (default for vrf CustB)
 *>i  10.2.255.1/32    10.255.255.12            2    100      0 ?
 * i                   10.255.255.12            2    100      0 ?
 *>i  10.2.255.2/32    10.255.255.12            2    100      0 ?
 * i                   10.255.255.12            2    100      0 ?
 *>   10.2.255.3/32    172.17.3.3               2         32768 ?
```

```
 *>i  10.4.255.2/32    10.255.255.15          110    100      0 ?
 *>i  10.4.255.3/32    10.255.255.15          110    100      0 ?
 *>   10.130.255.3/32  172.17.3.3              20         32768 ?
 *>i  172.17.1.0/24    10.255.255.12            0    100      0 ?
 * i                   10.255.255.12            0    100      0 ?
 *>i  172.17.2.0/24    10.255.255.12            0    100      0 ?
 * i                   10.255.255.12            0    100      0 ?
 *>   172.17.3.0/24    0.0.0.0                  0         32768 ?
 *>   192.168.2.254/32 172.17.3.3               2         32768 ?


CSR16(config)#do sh bgp vpnv4 unicast rd 20000:2 10.4.255.2/32
BGP routing table entry for 20000:2:10.4.255.2/32, version 54
Paths: (1 available, best #1, table CustB)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from 20000:4:10.4.255.2/32 (global)
    10.255.255.15 (metric 10) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 110, localpref 100, valid, internal, best
      Extended Community: RT:10.255.255.15:4 RT:10.255.255.15:20004
      Originator: 10.255.255.15, Cluster list: 10.255.255.11
      mpls labels in/out nolabel/150006
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 21:17:24 UTC

CSR16(config)#do sh bgp vpnv4 unicast rd 20000:2 10.4.255.3/32
BGP routing table entry for 20000:2:10.4.255.3/32, version 55
Paths: (1 available, best #1, table CustB)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from 20000:4:10.4.255.3/32 (global)
    10.255.255.15 (metric 10) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 110, localpref 100, valid, internal, best
      Extended Community: RT:10.255.255.15:4 RT:10.255.255.15:20004
```

```
      Originator: 10.255.255.15, Cluster list: 10.255.255.11
      mpls labels in/out nolabel/150007
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 21:17:24 UTC



CustB-R3(config)#do sh ip route ospf  | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 6 subnets
O IA     10.2.255.1 [110/12] via 172.17.3.16, 02:02:38, Ethernet0/0
O IA     10.2.255.2 [110/12] via 172.17.3.16, 02:02:38, Ethernet0/0
O E2     10.4.255.2 [110/1] via 172.17.3.16, 02:22:14, Ethernet0/0
O E2     10.4.255.3 [110/1] via 172.17.3.16, 02:22:14, Ethernet0/0
      172.17.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA     172.17.1.0/24 [110/11] via 172.17.3.16, 02:02:38, Ethernet0/0
O IA     172.17.2.0/24 [110/11] via 172.17.3.16, 02:02:38, Ethernet0/0



RP/0/0/CPU0:XRv15(config)#do sh bgp vpnv4 uni rd 20000:4 | begin Net
Thu Apr 23 21:48:41.738 UTC
   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:4 (default for vrf CustD)
*>i10.4.255.1/32      10.255.255.12          110    100      0 ?
* i                   10.255.255.12          110    100      0 ?
*> 10.4.255.2/32      172.19.2.2             110         32768 ?
*> 10.4.255.3/32      172.19.3.3             110         32768 ?
*>i172.19.1.0/24      10.255.255.12          110    100      0 ?
* i                   10.255.255.12          110    100      0 ?
*> 172.19.2.0/24      0.0.0.0                110         32768 ?
*> 172.19.3.0/24      0.0.0.0                110         32768 ?
*>i192.168.2.254/32   10.255.255.16            2    100      0 ?
```

CustD-R2 and CustD-R3 (CustD VPN) can ping Loopback 100 interface of CustB-R3 (CustB VPN).

This is called Inter-VRF Route-Leaking or Extranet.

We are done with L3VPNs. Let's begin L2VPN configuration.

L2VPNs are beyond of CCIE Enterprise Infrastructure blueprint, but Orhan Ergun LLC wants you to learn them as well. Because L2VPNs are used even in enterprises.

L2VPNs are an important topic of CCIE Service Provider.

In previous tasks and also MPLS-Lab-1 we have configured MPLS L3VPNs, the Customer Edge devices used a layer 3 connection to form a Dynamic PE-CE routing protocol (Or Static-Routes to learn networks statically), by doing this, Service Provider participated in the routing domain of customer, and the SP was aware of every routing information of the Customer.

There is another type of PE-CE connection, this time an L2 connection. In L2VPN services, the SP simulates a point-to-point Layer 2 connection or even a point-to-multipoint (multi-access) Layer 2 connection for the customer.

P2P is also called VPWS (Virtual Private Wire Service).

P2MP is also called VPLS (Virtual Private Lan Service).

If the SP provides VPWS for the customer, customer edge device thinks that it is connected to the other CE device using a back to back cable. Exactly like connecting them using an ethernet cable.

If the SP provides VPLS for the customer, customer edge device thinks that it is connected to an ethernet switch. SP simulates an Ethernet Switch for Customer Edge devices, the ports of this gigantic switch can be in different cities or even different coutries, also think about different continents!

The VPWS service is also called Cross Connect (xConnect/AToM) or in Juniper terms CCC.

VPWS is easy to configure in Cisco, Juniper and any other routers like Huawei, MikroTik , ZTE, etc… .

XRv 6.1.3 does not support any L2VPN forwarding, we will configure all L2VPN technologies on CSR1000v (IOS-XE).

Task 08:

- Provide VPWS (xConnect) service for CustE

Solution:

CSR13:

```
mpls ldp router-id lo0
interface GigabitEthernet4
 no ip address
 no shutdown
 service instance 5 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
!
l2vpn xconnect context CustE-VPWS
 member GigabitEthernet4 service-instance 5
 member 10.255.255.11 20005 encapsulation mpls
!
exit
```

CSR11:

```
mpls ldp router-id lo0
interface GigabitEthernet2
 no ip address
 no shutdown
 service instance 5 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
!
l2vpn xconnect context CustE-VPWS
 member GigabitEthernet2 service-instance 5
 member 10.255.255.13 20005 encapsulation mpls
```

This is the new method of configuring L2 VPNs, before the service instances, xconnect command must be entered under the interface, it is possible to configure this way today but using service instances are prefered.

The previous method:

```
CSR13:
interface GigabitEthernet4
 xconnect 10.255.255.11 20005 encapsulation mpls
!

CSR11:
interface GigabitEthernet2
 xconnect 10.255.255.13 20005 encapsulation mpls
!
```

A Pseudowire is going to be created between two PE routers (between their loopback 0 interfaces), the Virtual Circuit-ID will be 20005:

```
CSR11(config)#do sh l2vpn atom vc vcid 20005

                                    Service
Interface Peer ID         VC ID      Type   Name                      Status
--------- --------------- ---------- ------ ------------------------ ----------
pw100002  10.255.255.13   20005      p2p    CustE-VPWS                UP

CSR11(config)#do sh l2vpn atom vc vcid 20005 detail
pseudowire100002 is up, VC status is up PW type: Ethernet
  Create time: 14:43:56, last status change time: 14:43:36
    Last label FSM state change time: 14:43:36
  Destination address: 10.255.255.13 VC ID: 20005
    Output interface: Gi10, imposed label stack {16013 130001}
    Preferred path: not configured
    Default path: active
    Next hop: 10.11.12.12
  Member of xconnect service CustE-VPWS
    Associated member Gi2 is up, status is up
    Interworking type is Like2Like
```

```
  Service id: 0xdf000001
Signaling protocol: LDP, peer 10.255.255.13:0 up
  Targeted Hello: 10.255.255.11(LDP Id) -> 10.255.255.13, LDP is UP
  Graceful restart: not configured and not enabled
  Non stop routing: not configured and not enabled
  PWid FEC (128), VC ID: 20005
  Status TLV support (local/remote)         : enabled/supported
    LDP route watch                         : enabled
    Label/status state machine              : established, LruRru
    Local dataplane status received         : No fault
    BFD dataplane status received           : Not sent
    BFD peer monitor status received        : No fault
    Status received from access circuit     : No fault
    Status sent to access circuit           : No fault
    Status received from pseudowire i/f     : No fault
    Status sent to network peer             : No fault
    Status received from network peer       : No fault
    Adjacency status of remote peer         : No fault
Sequencing: receive disabled, send disabled
Bindings
  Parameter    Local                          Remote
  ------------ ------------------------------ ------------------------------
  Label        110012                         130001
  Group ID     7                              9
  Interface
  MTU          1500                           1500
  Control word on (configured: autosense)     on
  PW type      Ethernet                       Ethernet
  VCCV CV type 0x02                           0x02
                 LSPV [2]                       LSPV [2]
  VCCV CC type 0x07                           0x07
                 CW [1], RA [2], TTL [3]       CW [1], RA [2], TTL [3]
  Status TLV   enabled                        supported
SSO Descriptor: 10.255.255.13/20005, local label: 110012
```

When we configure xConnect, A targeted LDP connection has been formed between two PE devices, and they exchange labels for the Virtual Circuits using LDP:

CustE-R1:

```
interface Loopback0
 ip address 10.5.255.1 255.255.255.255
!
interface Ethernet0/0
 mac-address aaaa.5555.1111
 ip address 172.16.5.1 255.255.255.0
 duplex auto
!
router eigrp 5
 network 10.5.255.1 0.0.0.0
 network 172.16.5.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0
!
```

CustE-R2:

```
interface Loopback0
 ip address 10.5.255.2 255.255.255.255
!
interface Ethernet0/0
 mac-address aaaa.5555.2222
 ip address 172.16.5.2 255.255.255.0
 duplex auto
!
router eigrp 5
 network 10.5.255.2 0.0.0.0
 network 172.16.5.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0
!
```

Configuring EIGRP is not necessary, we just configured it to show that BUM (Broadcast, Unknown Unicast, Multicast) traffic passes through the SP network using MPLS network as an underlay network.

Verification:

```
CustE-R2(config)#do sh ip eigrp neigh
EIGRP-IPv4 Neighbors for AS(5)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   172.16.5.1              Et0/0                    13 14:57:43   14   100  0  9

CustE-R2(config)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 2 subnets
D        10.5.255.1 [90/409600] via 172.16.5.1, 14:57:47, Ethernet0/0



CustE-R2(config)#do ping 172.16.5.255
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.5.255, timeout is 2 seconds:

Reply to request 0 from 172.16.5.1, 10 ms
Reply to request 1 from 172.16.5.1, 24 ms
Reply to request 2 from 172.16.5.1, 27 ms
Reply to request 3 from 172.16.5.1, 19 ms
Reply to request 4 from 172.16.5.1, 7 ms

CustE-R2(config)#do sh cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
CustE-R1         Eth 0/0           130              R B   Linux Uni Eth 0/0
```

If we capture the ping between two CE devices, there will be 2 labels in the stack and also ethernet will be encapsulated inside another ethernet frame:

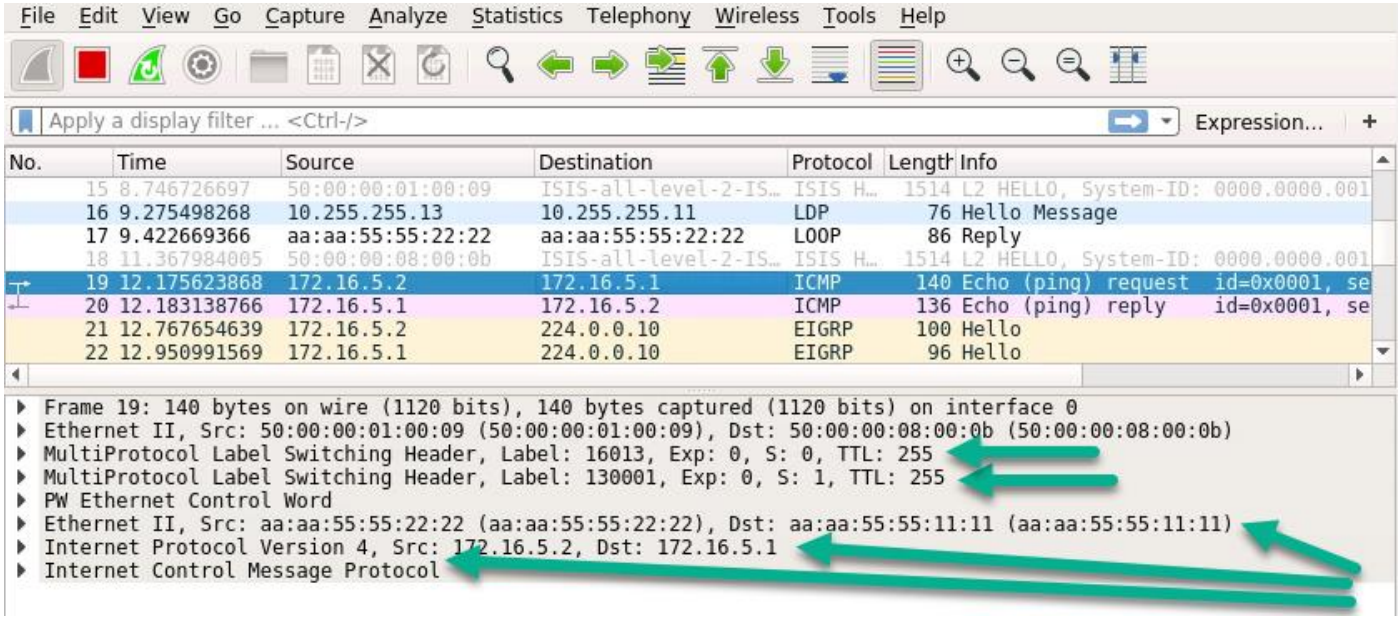

Task 09:

- Provide VPLS (Kompella) service for CustI.

Solution:

VPLS (Virtual Private Lan Service) is a Point-to-Multipoint Layer 2 VPN service.

It can be manually configured in the same way as VPWS (xConnect), it means that we can define members (other PE devices) manually. But it is not scalable, when there are a lot of Provider Edge devices that serves customers with VPLS, we can assist MP-BGP to automatically discover members (PE devices) and customer sites, then Virtual Circuits can be created automatically and BGP will take care of that, BGP also is going to be used for Label distribution/advertisement between PE devices.

This method is called Kompella.

Let's configure the lab, then you will figure out how it works:

We need to enable an address-family in BGP, that AF is called L2VPN VPLS, by enbling that AF for a neighbor, BGP is going to be used as the Auto Discovery (PE devices can find about each other when needed. also LDP (by default in IOS-XE) is going to be used for Label advertisement/Distribution between PEs (Signaling). We need to change that default behavior of IOS-XE, if we suppress LDP, BGP is going to be used for signaling purpose (label advertisement).

First, the CE device configurations, (we use EIGRP on CE devices to show you that the BUM traffic is going to be exchanged between Customer sites, it is not a must to configure a dunamic routing protocol on CE devices):

```
CustI-R1:
router eigrp 9
 network 172.16.9.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0
!

CustI-R2:
router eigrp 9
 network 10.9.255.2 0.0.0.0
 network 172.16.9.0 0.0.0.255
```

```
 passive-interface default
 no passive-interface Ethernet0/0
!

CustI-R3:

router eigrp 9
 network 10.9.255.3 0.0.0.0
 network 172.16.9.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0
!
```

Second, enable L2VPN VPLS (Suppress LDP as being the signaling protocol) on PE devices:

```
CSR11:
router bgp 20000
 address-family l2vpn vpls
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community extended
  neighbor 10.255.255.13 suppress-signaling-protocol ldp
  neighbor 10.255.255.14 activate
  neighbor 10.255.255.14 send-community extended
  neighbor 10.255.255.14 route-reflector-client
  neighbor 10.255.255.14 suppress-signaling-protocol ldp
 exit-address-family
!
CSR13:
router bgp 20000
 address-family l2vpn vpls
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community extended
  neighbor 10.255.255.11 suppress-signaling-protocol ldp
  neighbor 10.255.255.14 activate
  neighbor 10.255.255.14 send-community extended
  neighbor 10.255.255.14 route-reflector-client
  neighbor 10.255.255.14 suppress-signaling-protocol ldp
```

```
 exit-address-family
 !
CSR14:
router bgp 20000
 address-family l2vpn vpls
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community extended
  neighbor 10.255.255.11 suppress-signaling-protocol ldp
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community extended
  neighbor 10.255.255.13 suppress-signaling-protocol ldp
 exit-address-family
 !
```

Verification:

```
CSR14(config)#do sh bgp l2vpn vpls all summ | begin Neig
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.11   4        20000    3883    2734        3    0    0 1d05h           0
10.255.255.13   4        20000    3858    2728        3    0    0 1d05h           0

CSR11(config)#do sh bgp l2vpn vpls all summ | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.13   4        20000    3875    3891        4    0    0 1d05h           0
10.255.255.14   4        20000    2735    3884        4    0    0 1d05h           0

CSR13(config)#do sh bgp l2vpn vpls all summ | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.11   4        20000    3891    3875        4    0    0 1d05h           0
10.255.255.14   4        20000    2729    3859        4    0    0 1d05h           0
```

L2VPN VPLS iBGP peering has been formed.

Time to configure PE devices:

CSR11:

```
interface GigabitEthernet4
 no ip address
 negotiation auto
 no shutdown
 service instance 9 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
l2vpn vfi context CustI-VPLS
 vpn id 20009
 autodiscovery bgp signaling bgp
  ve id 11
  rd 20000:9
!
bridge-domain 9
 member GigabitEthernet4 service-instance 9
 member vfi CustI-VPLS
!
```

The last step is to create Bridge-Domain, think of a bridge domain being a virtual Switch on a PE device. We can add members to that switch, the members can be an interface with the service-instance for a specific customer, or it can be a VFI. Think of adding a VFI to the bridge-domain like a Virtual Switch between PE devices.

Now VFI will monitor the MP-BGP for the other PE devices with the same VPN-ID and establish pseudowires with them, MPLS labels will be advertesed/received to/from other PEs.

Configuration of the other PE:

CSR14:

```
interface GigabitEthernet7
 no ip address
 negotiation auto
 no shutdown
 service instance 9 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
```

```
interface GigabitEthernet8
 no ip address
 negotiation auto
 no shutdown
 service instance 9 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
l2vpn vfi context CustI-VPLS
 vpn id 20009
 autodiscovery bgp signaling bgp
  ve id 14
  rd 20000:9
!
bridge-domain 9
 member GigabitEthernet7 service-instance 9
 member GigabitEthernet8 service-instance 9
 member vfi CustI-VPLS
!
```

Verification:

```
CSR11(config)#do sh bgp l2vpn vpls all summ | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.13   4        20000    3916    3928        4    0    0 1d05h           1
10.255.255.14   4        20000    2764    3925        4    0    0 1d05h           1



CSR13(config)#do sh bgp l2vpn vpls all summ | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.11   4        20000    3929    3916        4    0    0 1d05h           2
10.255.255.14   4        20000    2759    3902        4    0    0 1d05h           1

CSR14(config)#do sh bgp l2vpn vpls all summ | begin Neig
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.255.255.11   4        20000    3926    2765        3    0    0 1d05h           1
```

```
10.255.255.13   4        20000    3902    2759        3    0    0 1d05h           1



CSR11(config)#do sh bgp l2vpn vpls rd 20000:9 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:9
 *>   20000:9:VEID-11:Blk-10/136
                      0.0.0.0                            32768 ?
 * i  20000:9:VEID-14:Blk-10/136
                      10.255.255.14            0    100      0 ?
 *>i                   10.255.255.14            0    100      0 ?



CSR13(config)#do sh bgp l2vpn vpls rd 20000:9 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:9
 *>i  20000:9:VEID-11:Blk-10/136
                      10.255.255.11            0    100      0 ?
 * i  20000:9:VEID-14:Blk-10/136
                      10.255.255.14            0    100      0 ?
 *>i                   10.255.255.14            0    100      0 ?



CSR14(config)#do sh bgp l2vpn vpls rd 20000:9 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:9
 * i  20000:9:VEID-11:Blk-10/136
                      10.255.255.11            0    100      0 ?
 *>i                   10.255.255.11            0    100      0 ?
 *>   20000:9:VEID-14:Blk-10/136
                      0.0.0.0                            32768 ?
```

```
CSR11(config)#do sh bgp l2vpn vpls rd 20000:9 detail


Route Distinguisher: 20000:9
BGP routing table entry for 20000:9:VEID-11:Blk-10/136, version 2
  Paths: (1 available, best #1, table L2VPN-VPLS-BGP-Table)
  Advertised to update-groups:
     1          2
  Refresh Epoch 1
  Local
    0.0.0.0 from 0.0.0.0 (10.255.255.11)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      AGI version(0), VE Block Size(10) Label Base(110000)
      Extended Community: RT:20000:20009 L2VPN L2:0x0:MTU-1500
      mpls labels in/out exp-null/110000
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 07:54:25 UTC
BGP routing table entry for 20000:9:VEID-14:Blk-10/136, version 4
  Paths: (2 available, best #2, table L2VPN-VPLS-BGP-Table)
  Flag: 0x100
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    10.255.255.14 (metric 20) from 10.255.255.13 (10.255.255.13)
      Origin incomplete, metric 0, localpref 100, valid, internal
      AGI version(0), VE Block Size(10) Label Base(140000)
      Extended Community: RT:20000:20009 L2VPN L2:0x0:MTU-1500
      Originator: 10.255.255.14, Cluster list: 10.255.255.13
      mpls labels in/out exp-null/140000
      rx pathid: 0, tx pathid: 0
      Updated on Apr 23 2020 07:55:44 UTC
  Refresh Epoch 1
  Local, (Received from a RR-client)
    10.255.255.14 (metric 20) from 10.255.255.14 (10.255.255.14)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
```

```
      AGI version(0), VE Block Size(10) Label Base(140000)
      Extended Community: RT:20000:20009 L2VPN L2:0x0:MTU-1500
      mpls labels in/out exp-null/140000
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 07:55:44 UTC



CSR11(config)#do sh l2vpn vfi
Legend: RT=Route-target, S=Split-horizon, Y=Yes, N=No

VFI name: CustI-VPLS, state: up, type: multipoint, signaling: BGP
  VPN ID: 20009, VE-ID: 11, VE-SIZE: 10
  RD: 20000:9, RT: 20000:20009,
  Bridge-Domain 9 attachment circuits:
  Pseudo-port interface: pseudowire100001
  Interface          Peer Address    VE-ID  Local Label  Remote Label    S
  pseudowire100003   10.255.255.14   14     110004       140001          Y



CSR11(config)#do sh bridge-domain | begin Bridge-domain 9
Bridge-domain 9 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
Maximum address limit: 65536
    GigabitEthernet4 service instance 9
    vfi CustI-VPLS neighbor 10.255.255.14 20009
   AED MAC address    Policy  Tag       Age  Pseudoport
   0   AAAA.9999.2222 forward dynamic   299  CustI-VPLS.404014
   0   AAAA.9999.3333 forward dynamic   300  GigabitEthernet4.EFP9
   0   AAAA.9999.1111 forward dynamic   298  CustI-VPLS.404014
```

Pseudowire between two PE devices has been created, and inside the bridge-domain we can see the MAC address of the CE devices.

```
CustI-R1(config)#do sh ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(9)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.16.9.3              Et0/0                    13 1d05h      18   108  0  22
0   172.16.9.2              Et0/0                    13 1d05h       4   100  0  21

CustI-R2(config)#do sh ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(9)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.16.9.3              Et0/0                    10 1d05h      16   100  0  21
0   172.16.9.1              Et0/0                    13 1d05h       5   100  0  23
CustI-R3(config)#do sh ip eigrp neigh
EIGRP-IPv4 Neighbors for AS(9)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.16.9.1              Et0/0                    14 1d05h      23   138  0  23
0   172.16.9.2              Et0/0                    13 1d05h      18   108  0  21


CustI-R3(config)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 2 subnets
D        10.9.255.2 [90/409600] via 172.16.9.2, 1d05h, Ethernet0/0


CustI-R3(config)#do ping 172.16.9.255 re 2
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 172.16.9.255, timeout is 2 seconds:

Reply to request 0 from 172.16.9.2, 12 ms
Reply to request 0 from 172.16.9.1, 12 ms
Reply to request 1 from 172.16.9.2, 10 ms
```

```
Reply to request 1 from 172.16.9.1, 11 ms



CustI-R3(config)#do sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.9.1            122   aaaa.9999.1111  ARPA   Ethernet0/0
Internet  172.16.9.2            103   aaaa.9999.2222  ARPA   Ethernet0/0
Internet  172.16.9.3              -   aaaa.9999.3333  ARPA   Ethernet0/0


CustI-R1(config)#do sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.9.1              -   aaaa.9999.1111  ARPA   Ethernet0/0
Internet  172.16.9.2            109   aaaa.9999.2222  ARPA   Ethernet0/0
Internet  172.16.9.3            123   aaaa.9999.3333  ARPA   Ethernet0/0


CustI-R2(config)#do sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.9.1            110   aaaa.9999.1111  ARPA   Ethernet0/0
Internet  172.16.9.2              -   aaaa.9999.2222  ARPA   Ethernet0/0
Internet  172.16.9.3            103   aaaa.9999.3333  ARPA   Ethernet0/0


CustI-R3(config)#do sh cdp ne
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID

Total cdp entries displayed : 0
```

Everything works fine but CDP, in order to forward CDP or any other L2 protocol to other PEs (the default behavior of VFI is to not forward any L2Protocol (Like STP, CDP) from CE device

to remote AFIs:

```
CSR11, CSR14:
!
l2vpn vfi context CustI-VPLS
 l2protocol forward cdp
!
```

Verification:

```
CustI-R3(config)#do sh cdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
CustI-R2         Eth 0/0           168              R B   Linux Uni Eth 0/0
CustI-R1         Eth 0/0           179              R B   Linux Uni Eth 0/0

Total cdp entries displayed : 2

CustI-R1(config)#do sh cdp neigh
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
CustI-R3         Eth 0/0           139              R B   Linux Uni Eth 0/0
CustI-R2         Eth 0/0           147              R B   Linux Uni Eth 0/0

Total cdp entries displayed : 2
```

Task 10:

- Provide EVPN VPWS service for CustG.

Solution:

EVPN is a newer and better technology compared to xConnect and VPLS that we have seen in the tasks so far.

EVPN stands for Ethernet Virtual Private Network. In classic, VPWS and VPLS every BUM (Broadcast, Unknown unicast and Multicast) traffic generated from Customer Edge devices will be forwarded in Control Plane to another PE devices. Imagine the situation that we want to give a Customer redundant L2 WAN connections:

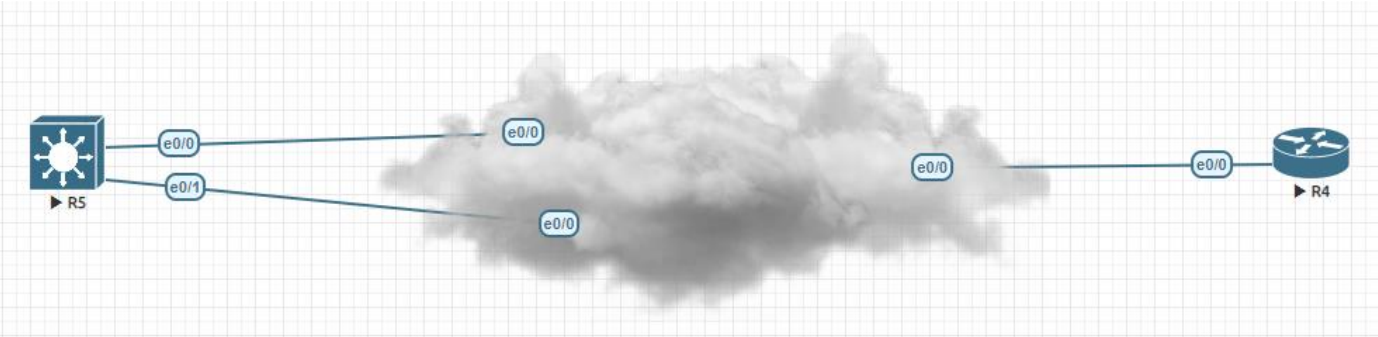

R5 (CE device) is connected to two PE routers using two redundant ethernet connections. If R5 sends BUM traffic using both links, It will cause a loop because two PE routers would send the BUM traffic to each other and the traffic send from e0/0 comes back to e0/1. We could make one link standby and the other one active, this would solve the problem, or we could tunnel the STP L2Protocol in VFI then R5 would block one link because of STP rules. But EVPN solves this problem, this is one of the advantages of EVPN. By using EVPN we can boundle e0/0 and e0/1 of R5 together (Port-Channel) and both links can be used at the same time (Active-Active), it is possible by a feature in EVPN called Active BUM forwarder. The election happens between two PE devices, and on of them will act as a BUM forwarder.

The other advantage of EVPN is learning MAC addresses and even IP addresses in Control Plane and advertise them using MP-BGP to other PEs. In classic VPLS and VPWS all the traffic would be send to every site even the site does not care about that traffic (the destination is not there). But in EVPN the traffic can be send to the exact site that the receiver resides on, because the MAC addresses and even IP addresses have been learned using Control Plane and MP-BGP.

EVPN can be used to provide VPWS service (P2P service) to the customers. Customer edge device is able to form an Ether-Channel between 2 PEs and both links can be used to forward traffic. Here, in this Example we only configure a simple EVPN VPWS service.

Customer Edge devices Configuration:

```
CustG-R1:


router eigrp 7
 network 10.7.255.1 0.0.0.0
 network 172.16.7.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0
!
CustG-R2:


router eigrp 7
 network 10.7.255.2 0.0.0.0
 network 172.16.7.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0
!
```

We need to enable L2VPN EVPN Address-Family on PE devices (For Auto-Discovery purpose):

```
CSR11:
router bgp 20000
 address-family l2vpn evpn
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community both
  neighbor 10.255.255.14 activate
  neighbor 10.255.255.14 send-community both
  neighbor 10.255.255.14 route-reflector-client
  neighbor 10.255.255.16 activate
  neighbor 10.255.255.16 send-community both
  neighbor 10.255.255.16 route-reflector-client
 exit-address-family
```

CSR13:

```
router bgp 20000
 address-family l2vpn evpn
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community both
  neighbor 10.255.255.14 activate
  neighbor 10.255.255.14 send-community both
  neighbor 10.255.255.14 route-reflector-client
  neighbor 10.255.255.16 activate
  neighbor 10.255.255.16 send-community both
  neighbor 10.255.255.16 route-reflector-client
 exit-address-family
```

CSR14:

```
router bgp 20000
 address-family l2vpn evpn
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community both
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community both
 exit-address-family
 !
```

CSR16:

```
router bgp 20000
 address-family l2vpn evpn
  neighbor 10.255.255.11 activate
  neighbor 10.255.255.11 send-community both
  neighbor 10.255.255.13 activate
  neighbor 10.255.255.13 send-community both
 exit-address-family
 !
```

EVPN VPWS PE configuration:

CSR14:

```
interface GigabitEthernet6
 no ip address
 negotiation auto
 no shutdown
 service instance 7 ethernet
  encapsulation untagged
 !
l2vpn evpn
 replication-type ingress
 mpls label mode per-ce
 router-id Loopback0
!
l2vpn evpn instance 7 point-to-point
 rd 20000:7
 vpws context CustG-VPWS
  service target 20016 source 20014
  member GigabitEthernet6 service-instance 7
 !
!
end
```

CSR16:

```
interface GigabitEthernet7
 no ip address
 negotiation auto
 no shutdown
 service instance 7 ethernet
  encapsulation untagged
 !
l2vpn evpn
 replication-type ingress
 mpls label mode per-ce
 router-id Loopback0
```

```
!
l2vpn evpn instance 7 point-to-point
 rd 20000:7
 vpws context CustG-VPWS
  service target 20014 source 20016
  member GigabitEthernet7 service-instance 7
 !
!
end
```

20014 and 20016 are the VPWS service identifiers, unlike VC Identifier in xConnect that both PEs use the same VC ID, EVPN VPWS Service Identifier must be different for source (local Service ID) and target (Remote Service ID).

Verification:

```
CSR14(config)#do sh l2vpn evpn vpws vc all
EVPN ID Source   Target   Type   Name/Interface                   Status
------- -------- -------- ------ -------------------------------- ----------
7       20014    20016    p2p    CustG-VPWS                       up
                                 Gi6:7                            up



CSR14(config)#do sh bgp l2vpn evpn rd 20000:7 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:7
 *>   [1][20000:7][00000000000000000000][20014]/23
                      ::                                    32768 ?
 *>i  [1][20000:7][00000000000000000000][20016]/23
                     10.255.255.16            0    100      0 ?
 * i                  10.255.255.16            0    100      0 ?
```

Task 11:

- Provide EVPN Single-Homed service for CustH

Solution:

Using EVPN we can have Single-Homed or Multi-Homed P2MP Layer 2 connectivity for the Customer sites.

Single homed means, the CE device is connected to a single PE device using only a single link. In Multi-Homed the CE device can have an Ether-Channel connection to two or more PE devices.

CSR16:

```
interface GigabitEthernet10
 no ip address
 no shutdown
 service instance 8 ethernet
  encapsulation dot1q 8
 !
interface GigabitEthernet11
 no ip address
 no shutdown
 service instance 8 ethernet
  encapsulation dot1q 8
 !
l2vpn evpn instance 8 vlan-based
 rd 20000:8
!
bridge-domain 8
 member GigabitEthernet10 service-instance 8
 member GigabitEthernet11 service-instance 8
 member evpn-instance 8
!
exit
```

```
CSR14:
interface GigabitEthernet10
 no ip address
 no shutdown
 service instance 8 ethernet
  encapsulation dot1q 8
 !
l2vpn evpn instance 8 vlan-based
 rd 20000:8
!
bridge-domain 8
 member GigabitEthernet10 service-instance 8
 member evpn-instance 8
!
exit
```

The EVPN Single-Homed configuration steps are very straight forward and it is just like VPLS, even easier that VPLS.

In the previous Tasks we have only used untagged traffic, but this time we are going to receive dot1q tagged traffic from CE devices.

The CE routers must use sub-interfaces in order to tag the traffic:

```
CustH-R1:
interface Ethernet0/0
 mac-address aaaa.8888.1111
 no ip address
 duplex auto
 no shutdown
!
interface Ethernet0/0.8
 encapsulation dot1Q 8
 ip address 172.16.8.1 255.255.255.0
!
router eigrp 8
 network 10.8.255.1 0.0.0.0
 network 172.16.8.0 0.0.0.255
```

```
 passive-interface default
 no passive-interface Ethernet0/0.8
!
CustH-R2:
interface Ethernet0/0
 mac-address aaaa.8888.2222
 no ip address
 duplex auto
 no shutdown
!
interface Ethernet0/0.8
 encapsulation dot1Q 8
 ip address 172.16.8.2 255.255.255.0
!
router eigrp 8
 network 10.8.255.2 0.0.0.0
 network 172.16.8.0 0.0.0.255
 passive-interface default
 no passive-interface Ethernet0/0.8
!
CustH-R3:
interface Ethernet0/0
 mac-address aaaa.8888.3333
 no ip address
 duplex auto
 no shutdown
!
interface Ethernet0/0.8
 encapsulation dot1Q 8
 ip address 172.16.8.3 255.255.255.0
!
router eigrp 8
 network 10.8.255.3 0.0.0.0
 network 172.16.8.0 0.0.0.255
```

```
 passive-interface default
 no passive-interface Ethernet0/0.8
!
```

This is the typical way that a service provider, provides service for their customers. For simplicity we have not used an access layer L2 device in the Lab. CE device does not connect to PE device directly, in fact PE device connects to an access layer switch (typically l2 switch) using a trunk link, and CE equipement connects to that L2 switch, and each customer is served by using a vlan.

Let's Verify the connectivity and also EVPN different route types:

```
CustH-R1(config)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
D        10.8.255.2 [90/409600] via 172.16.8.2, 00:11:00, Ethernet0/0.8
D        10.8.255.3 [90/409600] via 172.16.8.3, 00:10:02, Ethernet0/0.8

CustH-R1(config)#do sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.8.1               -   aaaa.8888.1111  ARPA   Ethernet0/0.8
Internet  172.16.8.2              10   aaaa.8888.2222  ARPA   Ethernet0/0.8
Internet  172.16.8.3              10   aaaa.8888.3333  ARPA   Ethernet0/0.8

CustH-R1(config)#do sh ip eigrp ne
EIGRP-IPv4 Neighbors for AS(8)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.16.8.3              Et0/0.8                  12 00:10:14   11   100  0  34
0   172.16.8.2              Et0/0.8                  13 00:11:10   13   100  0  37

CustH-R1(config)#do trace 10.8.255.2 source lo 0
Type escape sequence to abort.
Tracing the route to 10.8.255.2
VRF info: (vrf in name/id, vrf out name/id)
```

```
  1 172.16.8.2 10 msec 8 msec 7 msec
CustH-R1(config)#do trace 10.8.255.3 source lo 0
Type escape sequence to abort.
Tracing the route to 10.8.255.3
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.8.3 12 msec 10 msec 7 msec

CustH-R2(config)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
D        10.8.255.1 [90/409600] via 172.16.8.1, 00:11:29, Ethernet0/0.8
D        10.8.255.3 [90/409600] via 172.16.8.3, 00:10:29, Ethernet0/0.8

CustH-R2(config)#do sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.8.1             10   aaaa.8888.1111  ARPA   Ethernet0/0.8
Internet  172.16.8.2              -   aaaa.8888.2222  ARPA   Ethernet0/0.8
Internet  172.16.8.3             10   aaaa.8888.3333  ARPA   Ethernet0/0.8

CustH-R2(config)#do sh ip eigrp ne
EIGRP-IPv4 Neighbors for AS(8)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.16.8.3              Et0/0.8                  13 00:10:55    5   100  0  33
0   172.16.8.1              Et0/0.8                  12 00:11:50   11   100  0  38

CustH-R2(config)#do ping 172.16.8.255 re 1
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 172.16.8.255, timeout is 2 seconds:

Reply to request 0 from 172.16.8.3, 1 ms
Reply to request 0 from 172.16.8.1, 10 ms
```

```
CSR14(config)#do sh l2vpn evpn mac evi 8
MAC Address    EVI   BD    ESI                      Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- --------------
aaaa.8888.1111 8     8     0000.0000.0000.0000.0000 0          Gi10:8
aaaa.8888.2222 8     8     0000.0000.0000.0000.0000 0          10.255.255.16
aaaa.8888.3333 8     8     0000.0000.0000.0000.0000 0          10.255.255.16



CSR16(config-if)#do sh l2vpn evpn mac evi 8
MAC Address    EVI   BD    ESI                      Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- --------------
aaaa.8888.1111 8     8     0000.0000.0000.0000.0000 0          10.255.255.14
aaaa.8888.2222 8     8     0000.0000.0000.0000.0000 0          Gi10:8
aaaa.8888.3333 8     8     0000.0000.0000.0000.0000 0          Gi11:8
```

MAC addresses have been learned from CE device frames and advertised to the other PE using MP-BGP, Look at the ESI values, they are all zeroes, it is because in Single-Homed EVPN there is no need to Active BUM forwarder selection, ESI is the Ethernet Segment Identifier and it is being used to identify the PE-CE segment.

```
CSR14(config)#do sh bridge-domain 8
Bridge-domain 8 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
Maximum address limit: 65536
    GigabitEthernet10 service instance 8
    EVPN Instance 8
   AED MAC address    Policy  Tag       Age  Pseudoport
   -   AAAA.8888.2222 forward static_r  0    OCE_PTR:0xe93bbd20
   -   AAAA.8888.3333 forward static_r  0    OCE_PTR:0xe93bbcc0
   -   AAAA.8888.1111 forward dynamic_c 298  GigabitEthernet10.EFP8

CSR16(config-if)#do sh bridge-domain 8
Bridge-domain 8 (3 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
```

```
Maximum address limit: 65536
    GigabitEthernet10 service instance 8
    GigabitEthernet11 service instance 8
    EVPN Instance 8
   AED MAC address    Policy  Tag       Age  Pseudoport
   -   AAAA.8888.2222 forward dynamic_c 297  GigabitEthernet10.EFP8
   -   AAAA.8888.3333 forward dynamic_c 298  GigabitEthernet11.EFP8
   -   AAAA.8888.1111 forward static_r  0    OCE_PTR:0xe93ac4e0


CSR14(config)#do sh bgp l2vpn evpn rd 20000:8 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:8
 *>   [2][20000:8][0][48][AAAA88881111][0][*]/20
                      ::                                    32768 ?
 *>i  [2][20000:8][0][48][AAAA88882222][0][*]/20
                      10.255.255.16            0    100      0 ?
 * i                   10.255.255.16            0    100      0 ?
 *>i  [2][20000:8][0][48][AAAA88883333][0][*]/20
                      10.255.255.16            0    100      0 ?
 * i                   10.255.255.16            0    100      0 ?
 *>   [3][20000:8][0][32][10.255.255.14]/17
                      ::                                    32768 ?
 *>i  [3][20000:8][0][32][10.255.255.16]/17
                      10.255.255.16            0    100      0 ?
 * i                   10.255.255.16            0    100      0 ?


CSR16(config-if)#do sh bgp l2vpn evpn rd 20000:8 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:8
 *>i  [2][20000:8][0][48][AAAA88881111][0][*]/20
                      10.255.255.14            0    100      0 ?
 * i                   10.255.255.14            0    100      0 ?
```

```
 *>   [2][20000:8][0][48][AAAA88882222][0][*]/20
                      ::                                32768 ?
 *>   [2][20000:8][0][48][AAAA88883333][0][*]/20
                      ::                                32768 ?
 *>i  [3][20000:8][0][32][10.255.255.14]/17
                      10.255.255.14            0    100      0 ?
 * i                   10.255.255.14            0    100      0 ?
 *>   [3][20000:8][0][32][10.255.255.16]/17
                      ::                                32768 ?
```

Let's take a look at detailed BGP learned routes:

```
CSR14(config)#do sh bgp l2vpn evpn rd 20000:8 detail


Route Distinguisher: 20000:8
BGP routing table entry for [2][20000:8][0][48][AAAA88881111][0][*]/20, version 6433
  Paths: (1 available, best #1, table evi_8)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.255.255.14)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 140011
      Extended Community: RT:20000:8
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 26 2020 12:25:32 UTC
BGP routing table entry for [2][20000:8][0][48][AAAA88882222][0][*]/20, version 6436
  Paths: (2 available, best #1, table evi_8)
  Not advertised to any peer
  Refresh Epoch 5
  Local
    10.255.255.16 (metric 20) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 160001
```

```
      Extended Community: RT:20000:8
      Originator: 10.255.255.16, Cluster list: 10.255.255.11
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 26 2020 12:27:07 UTC
  Refresh Epoch 5
  Local
    10.255.255.16 (metric 20) (via default) from 10.255.255.13 (10.255.255.13)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 160001
      Extended Community: RT:20000:8
      Originator: 10.255.255.16, Cluster list: 10.255.255.13
      rx pathid: 0, tx pathid: 0
      Updated on Apr 26 2020 12:27:07 UTC
BGP routing table entry for [2][20000:8][0][48][AAAA88883333][0][*]/20, version 6444
  Paths: (2 available, best #1, table evi_8)
  Not advertised to any peer
  Refresh Epoch 5
  Local
    10.255.255.16 (metric 20) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 160002
      Extended Community: RT:20000:8
      Originator: 10.255.255.16, Cluster list: 10.255.255.11
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 26 2020 12:28:02 UTC
  Refresh Epoch 5
  Local
    10.255.255.16 (metric 20) (via default) from 10.255.255.13 (10.255.255.13)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 160002
      Extended Community: RT:20000:8
      Originator: 10.255.255.16, Cluster list: 10.255.255.13
      rx pathid: 0, tx pathid: 0
      Updated on Apr 26 2020 12:28:02 UTC
```

```
BGP routing table entry for [3][20000:8][0][32][10.255.255.14]/17, version 5
  Paths: (1 available, best #1, table evi_8)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.255.255.14)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      Extended Community: RT:20000:8
      PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, label:140012 tunnel identifier: 0000 0000
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 07:54:26 UTC
BGP routing table entry for [3][20000:8][0][32][10.255.255.16]/17, version 74
  Paths: (2 available, best #1, table evi_8)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 5
  Local
    10.255.255.16 (metric 20) (via default) from 10.255.255.11 (10.255.255.11)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:20000:8
      Originator: 10.255.255.16, Cluster list: 10.255.255.11
      PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, label:160003 tunnel identifier: < Tunnel Endpoint:
10.255.255.16 >
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 23 2020 08:16:58 UTC
  Refresh Epoch 5
  Local
    10.255.255.16 (metric 20) (via default) from 10.255.255.13 (10.255.255.13)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: RT:20000:8
      Originator: 10.255.255.16, Cluster list: 10.255.255.13
      PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, label:160003 tunnel identifier: < Tunnel Endpoint:
10.255.255.16 >
      rx pathid: 0, tx pathid: 0
```

There are different routes types, in above output [2] for type 2 and [3] for type 3.

Type 2 routes are used for MAC/IP advertisement, Type 3 routes are used for forwarding BUM traffic (Ingress Replication in this example).

We will discuss about route types in detail in the VXLAN, MPLS EVPN lab.

Task 12:

- Provide EVPN Multi-Homed service for CustF

Solution:

This time we are going to configure Multi-Homed EVPN service for the customer, CE device will be connected to two PE devices using a single Ethernet Boundle link (A Port-Channel). First, we need to configure a remote port-channel and also an ethernet-segment between two PE devices that are connected to the same CE device, then a service instance must be created to serve the customer traffic:

```
CSR11:
interface Port-channel6
 !
 no shutdown
 evpn ethernet-segment 1
  identifier type 3 system-mac 0000.1111.1111
  redundancy all-active
 service instance 6 ethernet
  encapsulation untagged
 !
interface GigabitEthernet3
 no ip address
 no shutdown
 channel-group 6
!
l2vpn evpn
 replication-type ingress
 mpls label mode per-ce
 router-id Loopback0
!
l2vpn evpn instance 6 vlan-based
 rd 20000:6
!
bridge-domain 6
 member Port-channel6 service-instance 6
```

```
 member evpn-instance 6
!
exit
CSR16:
interface Port-channel6
 !
 no shutdown
 evpn ethernet-segment 1
  identifier type 3 system-mac 0000.1111.1111
  redundancy all-active
 service instance 6 ethernet
  encapsulation untagged
 !
interface GigabitEthernet9
 no ip address
 negotiation auto
 no mop enabled
 no shutdown
 channel-group 6
!
l2vpn evpn
 replication-type ingress
 mpls label mode per-ce
 router-id Loopback0
!
l2vpn evpn instance 6 vlan-based
 rd 20000:6
!
bridge-domain 6
 member Port-channel6 service-instance 6
 member evpn-instance 6
!
exit
```

CSR13:

```
interface Port-channel6
 no shutdown
 evpn ethernet-segment 2
  identifier type 3 system-mac 5000.1313.1313
  redundancy all-active
 service instance 6 ethernet
  encapsulation untagged
 !
interface GigabitEthernet5
 no ip address
 no shutdown
 channel-group 6
!
l2vpn evpn
 replication-type ingress
 mpls label mode per-ce
 router-id Loopback0
!
l2vpn evpn instance 6 vlan-based
 rd 20000:6
!
bridge-domain 6
 member Port-channel6 service-instance 6
 member evpn-instance 6
!
Exit
```

CSR14:

```
interface Port-channel6
 no ip address
 no shutdown
 evpn ethernet-segment 2
  identifier type 3 system-mac 5000.1313.1313
```

```
  redundancy all-active
 service instance 6 ethernet
  encapsulation untagged
 !
interface GigabitEthernet5
 no ip address
 negotiation auto
 no shutdown
 isis network point-to-point
 isis three-way-handshake cisco
 channel-group 6
!
l2vpn evpn instance 6 vlan-based
 rd 20000:6
!
bridge-domain 6
 member Port-channel6 service-instance 6
 member evpn-instance 6
!
exit
```

Now, the provider side configuration is done, let's configure the CE devices:

CustF-SW1:

```
interface Loopback0
 ip address 10.6.255.1 255.255.255.255
!
interface Ethernet0/0
 switchport access vlan 10
 switchport mode access
 switchport nonegotiate
!
interface Ethernet0/1
 switchport access vlan 10
```

```
 switchport mode access
 switchport nonegotiate
!
interface Vlan10
 ip address 172.16.6.1 255.255.255.0
!
!
router eigrp 6
 network 10.6.255.1 0.0.0.0
 network 172.16.6.0 0.0.0.255
!
```

CustF-SW2:

```
interface Loopback0
 ip address 10.6.255.2 255.255.255.255
!
interface Ethernet0/0
 switchport access vlan 10
 switchport mode access
 switchport nonegotiate
!
interface Ethernet0/1
 switchport access vlan 10
 switchport mode access
 switchport nonegotiate
!
interface Ethernet0/2
!
interface Ethernet0/3
!
interface Vlan10
 ip address 172.16.6.2 255.255.255.0
!
!
```

```
router eigrp 6
 network 10.6.255.2 0.0.0.0
 network 172.16.6.0 0.0.0.255
!
```

Verification:

```
CUstF-SW2(config-if)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 2 subnets
D        10.6.255.1 [90/130816] via 172.16.6.1, 00:05:41, Vlan10

CUstF-SW2(config-if)#do sh ip eigrp neighbor
EIGRP-IPv4 Neighbors for AS(6)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   172.16.6.1              Vl10                     10 00:05:59   24   144  0  7

CUstF-SW2(config-if)#do sh mac add
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  10    aabb.cc80.1400    DYNAMIC     Et0/0
Total Mac Addresses for this criterion: 1
CUstF-SW2(config-if)#do sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.6.1               6   aabb.cc80.1400  ARPA   Vlan10
Internet  172.16.6.2               -   aabb.cc80.1500  ARPA   Vlan10


CustF-SW1(config-if)#do sh ip route eigrp | begin Gate
Gateway of last resort is not set
```

```
      10.0.0.0/32 is subnetted, 2 subnets
D        10.6.255.2 [90/130816] via 172.16.6.2, 00:07:43, Vlan10



CustF-SW1(config-if)#do ping 10.6.255.2 source lo 0 re 100
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 10.6.255.2, timeout is 2 seconds:
Packet sent with a source address of 10.6.255.1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 6/8/21 ms



CustF-SW1(config-if)#do trace 10.6.255.2
Type escape sequence to abort.
Tracing the route to 10.6.255.2
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.6.2 9 msec 10 msec 8 msec


CSR16(config)#do sh l2vpn evpn evi 6 detail
EVPN instance:       6 (VLAN Based)
  RD:                20000:6 (cfg)
  Import-RTs:        20000:6
  Export-RTs:        20000:6
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     mpls
  Bridge Domain:     6
    Ethernet-Tag:    0
    BUM Label:       160001
    Per-BD Label:    none
    State:           Established
```

```
    Pseudoports (Labels):
      Port-channel6 service instance 6 (160000) (DF state: PE-to-CE BUM blocked)
        Routes: 2 MAC, 0 MAC/IP
    Peers:
      10.255.255.11
        Routes: 1 MAC, 0 MAC/IP, 1 IMET, 1 EAD
      10.255.255.13
        Routes: 1 MAC, 0 MAC/IP, 1 IMET, 1 EAD
      10.255.255.14
        Routes: 0 MAC, 0 MAC/IP, 1 IMET, 1 EAD



CSR11(config)#do show l2vpn evpn ethernet-segment forwarder
EVPN Ethernet Segment ID: 0300.0011.1111.1100.0001
  Forwarder List:         10.255.255.11 10.255.255.16



CSR13(config)#do show l2vpn evpn ethernet-segment forwarder
EVPN Ethernet Segment ID: 0350.0013.1313.1300.0002
  Forwarder List:         10.255.255.13 10.255.255.14



CSR16(config-if)#do sh bgp l2vpn evpn rd 20000:6 | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 20000:6
 * i  [1][20000:6][03000011111111000001][0]/23
                      10.255.255.11            0    100      0 ?
 *mi                   10.255.255.11            0    100      0 ?
 *>                    ::                                   32768 ?
 *>i  [1][20000:6][03500013131313000002][0]/23
                      10.255.255.14            0    100      0 ?
 *mi                   10.255.255.13            0    100      0 ?
 * i  [2][20000:6][0][48][AABBCC801400][0][*]/20
                      10.255.255.13            0    100      0 ?
```

```
 *>i                    10.255.255.13            0    100      0 ?
 * i  [2][20000:6][0][48][AABBCC801500][0][*]/20
                       10.255.255.11            0    100      0 ?
 *mi                    10.255.255.11            0    100      0 ?
 *>                     ::                                 32768 ?
 * i  [3][20000:6][0][32][10.255.255.11]/17
                       10.255.255.11            0    100      0 ?
 *>i                    10.255.255.11            0    100      0 ?
 * i  [3][20000:6][0][32][10.255.255.13]/17
                       10.255.255.13            0    100      0 ?
 *>i                    10.255.255.13            0    100      0 ?
 *>i  [3][20000:6][0][32][10.255.255.14]/17
                       10.255.255.14            0    100      0 ?
 * i                    10.255.255.14            0    100      0 ?
 *>   [3][20000:6][0][32][10.255.255.16]/17
                       ::                                 32768 ?


CSR11(config)#do sh l2vpn evpn summ
L2VPN EVPN
  EVPN Instances (excluding point-to-point): 1
    VLAN Aware:   0
    VLAN Based:   1
    VLAN Bundle:  0
  Bridge Domains: 1
  BGP: ASN 20000, address-family l2vpn evpn configured
  Router ID: 10.255.255.11
  Label Allocation Mode: Per-CE
  Global Replication Type: Ingress
  Connectivity to Core: UP
  MAC Duplication: seconds 180 limit 5
  MAC Addresses: 2
    Local:     1
    Remote:    1
    Duplicate: 0
```

```
Maximum number of Route Targets per EAD-ES route: 200
Multi-home aliasing: Enabled
```

We will discuss about the route types in detail later on (VXLAN, MPLS EVPN Lab).