

IOS-XR EVPN – Multipoint L2VPN Service with MPLS Data Plane

Table of Content:

1. Overview.....	2
2. Transport.....	3
3. Services EVPN – L2VPN Multipoint.....	4
3.1. EVPN Configuration	4
3.2. BGP Configuration	6
4. Verification.....	7
4.1. Transport	7
4.2. EVPN Service.....	7
5. Summary.....	10


1. Overview

This document will briefly describe IOS-XR configuration of EVPN – L2VPN multipoint service with MPLS data plane, which is applicable in Service Provider Data Center as well as in Service Provider Access and Edge network.

See also Cisco Metro Design <https://xrdocs.io/design/> to understand Transport and Services evolution in Service Provider Access, Aggregation and Edge network.

This document will focus on layer2 and we will add layer3 in a separate document.

Diagram: “lab topology” displays lab topology, which emulates small Service Provider Data Center.



lab topology

Devices in the lab:

- Two Spines: R34 and R35 – **Cisco NCS 5501-SE**
- Four Leaves: R36, R37, R38 and R39 – **Cisco NCS 5501-SE**
- Two Router Reflectors: RR103 and RR104 – **Cisco IOS XRv 9000**

- Two Hosts: H1 and H2, both hosts are dual-homed
 - H1 to R36 and R37
 - H2 to R38 and R39

All interfaces are 10Gbps Ethernet connected by direct fiber.

All devices run **Cisco IOS-XR 6.5.1** operating system

2. Transport

All devices except H1 and H2 are in the same IGP domain. H1 and H2 are connected to transparent EVPN – L2VPN multipoint service.

In our testbed, we use ISIS IGP protocol with Segment Routing extension and MPLS Data Plane. Please note that transport is independent of service and therefore ISIS can be replaced easily by OSPF without any changes in EVPN configuration.

In the configuration example you can also see Topology Independent Loop Free Alternate (TI-LFA) for fast convergence in case of any core link failure or Spine node failure.

ISIS Configuration of Leaf and Spine:

```
router isis 3
  is-type level-2-only
  net 49.0003.0000.0000.00[two digits router-id].00
  address-family ipv4 unicast
    metric-style wide
    segment-routing mpls
  !
  interface Loopback0
    address-family ipv4 unicast
    prefix-sid index [two digits router-id]
  !
  !
  interface TenGigE[interface-id] //core interface
    point-to-point
    address-family ipv4 unicast
    fast-reroute per-prefix
    fast-reroute per-prefix ti-lfa
  !
  [Add all core interface]
  !
```

Route reflectors configuration is without Segment Routing extension, because Route Reflectors don't participate in Data Plane forwarding.

ISIS Configuration of Route Reflector:

```

router isis 3
 is-type level-2-only
 net 49.0003.0000.0000.0[three digits router-id].00
 address-family ipv4 unicast
  metric-style wide
 !
 interface Loopback0
  address-family ipv4 unicast
  !
 !
 interface TenGigE[interface-id] //core interface
  point-to-point
  address-family ipv4 unicast
  !
 [Add all core interface]
 !

```

Please note that it is also highly recommended to configure high metric on the link from Spines to Route Reflectors to exclude Route Reflectors from the forwarding path, see configuration example below. Another option is to single-home each Route Reflector to a different Spine. For example RR103 to R34 and R104 to R35.

```

router isis 3
 interface TenGigE0/0/0/21
  point-to-point
  address-family ipv4 unicast
  metric [high-value in our example 20000]

```

3. Services EVPN – L2VPN Multipoint

EVPN uses BGP Control Plane, therefore we will split this section into two parts: EVPN configuration and BGP configuration.

3.1. EVPN Configuration

This section will focus on EVPN specific configuration together with Bridge-Domain configuration.

Host H1 is dual-homed by Leaves R36 and R37. Similarly, Host H2 is dual-homed by Leaves R38 and R39. Because both Hosts are transparently connected via LACP Port-Channel, each Leaf pair has to use the same LACP System MAC Address. See configuration example for R36 and R37 below.

Leaf LACP System MAC Configuration:

```

lacp system mac 3637.3637.3637

```

Interface Bundle-Ether Configuration:

```
interface TenGigE0/0/0/0
  bundle id 100 mode active
!

interface Bundle-Ether100
  l2transport
!
```

The below example shows the configuration of EVPN instance 100 (evi 100). Advertise-mac command causes BGP to generate EVPN Route-Type 2 (RT-2), which advertises host MAC address learned through access interface by Leaves' Data Plane. BGP EVPN Control Plane detail will be described in a separate document. To prevent core isolation, all MPLS core interfaces are in core-isolation-group 1. Interface Bundle-Ether100 is the access interface and the configuration below shows how to manually specify Ethernet Service Identifier (ESI) which is important for the multi-homed scenario covered in our testbed. Please note that in the case of a single-homed device, the ESI value is most of the time 0 (default value), therefore there is no ESI configuration.

Leaf EVPN Configuration:

```
evpn
  evi 100
    advertise-mac //generate BGP EVPN RT-2
    !
  !
  group 1
    core interface TenGigE[interface-id]
    core interface TenGigE[interface-id]
    !
  interface Bundle-Ether100
    ethernet-segment
      identifier type 0 36.37.00.00.00.00.00.11.00
      //R36 and R37 ESI. R38 and R39 use different, but also common ESI
    !
    core-isolation-group 1
  !
```

The Bridge-Domain configuration example below shows how to specify new Bridge-Domain 100 and add access interface Bundle-Ether100 and evi 100.

Leaf Bridge-Domain Configuration:

```
l2vpn
  bridge group 100
  bridge-domain 100
  interface Bundle-Ether100
    !
  evi 100
```

```

!
!
!
!

```

3.2. BGP Configuration

Each Leaf has a BGP (Address Family EVPN) session with both dedicated Route Reflectors. Please note that it is highly recommended to also use dedicated Route Reflectors in your production network. Spine node should remain BGP-free, as should a core Provider router.

Leaf BGP Configuration:

```

router bgp 1
  bgp router-id 3.3.3.[two digits router-id]
  address-family l2vpn evpn
  !
  neighbor-group rr
    remote-as 1
    update-source Loopback0
    address-family l2vpn evpn
  !
  !
  neighbor 3.3.3.[three digits router-id] //Route Reflector
    use neighbor-group rr
  !
[Add all Route Reflectors]
!

```

Router Reflector BGP Configuration:

```

router bgp 1
  bgp router-id 3.3.3.[router-id]
  address-family l2vpn evpn
  !
  neighbor-group leaf
    remote-as 1
    update-source Loopback0
    address-family l2vpn evpn
    route-reflector-client
  !
  !
  neighbor 3.3.3.[leaf router-id]
    use neighbor-group leaf
  !
[Add all leaves]
!

```

This step was the last one to successfully configure Multi-Homed/Multipoint EVPN service with MPLS Data Plane.

4. Verification

This section shows how to quickly verify EVPN service end-to-end.

4.1. Transport

Each Leaf should see entry in FIB for all other Leaves. The example below shows entry for 3.3.3.37 which is loopback 0 of R37. R37 is reachable via the ECMP path through Spines R34 and R35. Links and Spine nodes are also protected by TI-LFA. Label imposition is very important because it shows particular transport prefix-SID (MPLS Data Plane).

```
RP/0/RP0/CPU0:R36#show cef 3.3.3.37
```

```
3.3.3.37/32, version 578, labeled SR, internal 0x1000001 0x81 (ptr 0x97c167a8) [1], 0x0
(0x97dd9ce8), 0xa28 (0x98c00378)
  Updaed Aug 13 22:48:54.216
  local adjacency 35.36.1.35
  Prefix Len 32, traffic index 0, precedence n/a, priority 1
    via 35.36.1.35/32, TenGigE0/0/0/39, 10 dependencies, weight 0, class 0, protected,
ECMP-backup (Local-LFA) [flags 0x600]
      path-idx 0 bkup-idx 1 NHID 0x0 [0x98b7f6d0 0x0]
      next hop 35.36.1.35/32
        local label 16037      labels imposed {16037}
          via 34.36.1.34/32, TenGigE0/0/0/38, 10 dependencies, weight 0, class 0, protected,
ECMP-backup (Local-LFA) [flags 0x600]
            path-idx 1 bkup-idx 0 NHID 0x0 [0x98b7f350 0x0]
            next hop 34.36.1.34/32
              local label 16037      labels imposed {16037}
```

Spines are responsible for label swapping, therefor LFIB entry is most important from transport point of view. The example below shows LFIB entry of 16036 (prefix-SID of loopback 0 on R36) executed on R34, together with backup bath created by TI-LFA.

```
RP/0/RP0/CPU0:R34#show mpls forwarding
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched	
16036	Pop	SR Pfx (idx 36)	Te0/0/0/16	34.36.1.36	0	
	16035	SR Pfx (idx 36)	Te0/0/0/17	34.37.1.37	0	(!)

4.2. EVPN Service

Each Leaf should have two active BGP EVPN sessions with each of both Route Reflectors.

```
RP/0/RP0/CPU0:R36#show bgp l2vpn evpn summary
```

```
BGP router identifier 3.3.3.36, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table stat: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 14814
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process Speaker	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
	14814	14814	14814	14814	14814	0

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
3.3.3.103	0	1	25826	21481	14814	0	0	00:04:35	12
3.3.3.104	0	1	25850	21482	14814	0	0	00:04:42	12

The example below shows how to verify Ethernet Segment Identifier (ESI) configuration. ESI 0036.3700.0000.0000.1100 is Multi-Homed via 3.3.3.36 (R36) and 3.3.3.37 (R37).

```
RP/0/RP0/CPU0:R36#show evpn ethernet-segment
```

Ethernet Segment Id	Interface	Nexthops
0036.3700.0000.0000.1100	BE100	3.3.3.36 3.3.3.37

The next example shows detail of ESI 0036.3700.0000.0000.1100. This ESI is in Multi-Homed (MH) All-Active mode.

```
RP/0/RP0/CPU0:R36#show evpn ethernet-segment esi 0036.3700.0000.0000.1100 detail
```

Legend:

- B - No Forwarders EVPN-enabled,
- C - Backbone Source MAC missing (PBB-EVPN),
- RT - ES-Import Route Target missing,
- E - ESI missing,
- H - Interface handle missing,
- I - Name (Interface or Virtual Access) missing,
- M - Interface in Down state,
- O - BGP End of Download missing,
- P - Interface already Access Protected,
- Pf - Interface forced single-homed,
- R - BGP RID not received,
- S - Interface in redundancy standby state,
- X - ESI-extracted MAC Conflict
- SHG - No local split-horizon-group label allocated

Ethernet Segment Id	Interface	Nexthops
---------------------	-----------	----------


```

-----
0036.3700.0000.0000.1100 BE100                                3.3.3.36
                                                                3.3.3.37

ES to BGP Gates      : Ready
ES to L2FIB Gates   : Ready
Main port            :
  Interface name     : Bundle-Ether100
  Interface MAC      : 008a.9644.d8dd
  IfHandle           : 0x0800001c
  State              : Up
  Redundancy         : Not Defined
ESI type             : 0
  Value              : 36.3700.0000.0000.1100
ES Import RT        : 3637.0000.0000 (from ESI)
Source MAC          : 0000.0000.0000 (N/A)
Topology            :
  Operational        : MH, All-active
  Configured         : All-active (AApF) (default)
Service Carving     : Auto-selection
Peering Details     : 3.3.3.36[MOD:P:00] 3.3.3.37[MOD:P:00]
Service Carving Results:
  Forwarders        : 1
  Permanent         : 0
  Elected          : 1
  Not Elected      : 0
MAC Flushing mode   : STP-TCN
Peering timer       : 3 sec [not running]
Recovery timer      : 30 sec [not running]
Carving timer       : 0 sec [not running]
Local SHG label     : 64005
Remote SHG labels : 1
                    64005 : nexthop 3.3.3.37
    
```

The next example shows MAC addresses learned locally or via BGP Control Plane. For Example MAC address of H2 0062.ec71.1000 is available via 3.3.3.38 (Leaf R38) and 3.3.3.39 (Leaf 3.3.3.39).

```

RP/0/RP0/CPU0:R36#show evpn evi mac
-----
VPN-ID      Encap  MAC address      IP address      Nexthop          Label
-----
100         MPLS  0062.ec71.1000  ::             3.3.3.38        64004
100         MPLS  0062.ec71.1000  ::             3.3.3.39        64004
100         MPLS  0062.ec71.fbd7  ::             Bundle-Ether100 64004
100         MPLS  0062.ec71.fbd8  ::             Bundle-Ether100 64004
100         MPLS  0062.ec71.fbd9  ::             3.3.3.37        64004
100         MPLS  0062.ec71.fbe0  ::             3.3.3.38        64004
100         MPLS  0062.ec71.fbe0  ::             3.3.3.39        64004
100         MPLS  0062.ec71.fbe1  ::             3.3.3.38        64004
100         MPLS  0062.ec71.fbe1  ::             3.3.3.39        64004
65535      N/A   008a.9644.d8d8  ::             Local           0
    
```

The next example shows EVPN label allocation. In this example local label for EVI 100 and ESI 0038.3900.0000.0000.1100 is 64007, outgoing labels 64004 via 3.3.3.38 and 64004 via 3.3.3.39.

```
RP/0/RP0/CPU0:R36#show evpn internal-label
Tue Aug 21 20:33:39.453 UTC
```

VPN-ID	Encap	Ethernet Segment Id	EtherTag	Label
100	MPLS	0036.3700.0000.0000.1100	0	64006
		Summary pathlist:		
		0x02000001 3.3.3.37		64004
100	MPLS	0038.3900.0000.0000.1100	0	64007
		Summary pathlist:		
		0x02000002 3.3.3.38		64004
		0x02000003 3.3.3.39		64004

To see all the BGP EVPN routes you can use the example below.

```
RP/0/RP0/CPU0:R36#show bgp l2vpn evpn
```

BGP EVPN Control Plane detail will be covered in a separate document.

5. Summary

This document shows how EVPN – L2VPN Multipoint service can be easily configured with MPLS Data Plane on Cisco IOS-XR operating system. The next document will focus on L3 integration for optimal forwarding and later we will also cover integration with L3VPN services or with legacy L2VPN services such as VPLS, EoMPLS, etc.