

# Implementing Nexus L2 EVPN over Segment Routing MPLS

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background](#)

[Configure](#)

[Network Diagram](#)

[Verification](#)

[References](#)

## Introduction

This document describes how to deploy/configure Layer2 EVPN over Segment Routing MPLS on Cisco Nexus 9000 series switches.

## Prerequisites

### Requirements

Required to have knowledge on BGP, OSPF, MPLS, LDP, RSVP, EVPN, Segment Routing(SR)

### Components Used

Cisco Nexus switch 93360YC-FX2 running with 9.3.(3)

Cisco Nexus switch 93240YC-FX2 running with 9.3.(3)

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

## Background

Define Layer 2 VPN, VPLS/L2-EVPN is a Multipoint-to-Multipoint Layer 2 VPN Service that connects multiple branches of a Customer, in a single logical switched architecture over an IP/MPLS network.

### Layer2 EVPN-MPLS SR:

- EVPN (RFC 7432) is BGP MPLS-based solution which has been used for next generation

Ethernet services in virtualized data center network

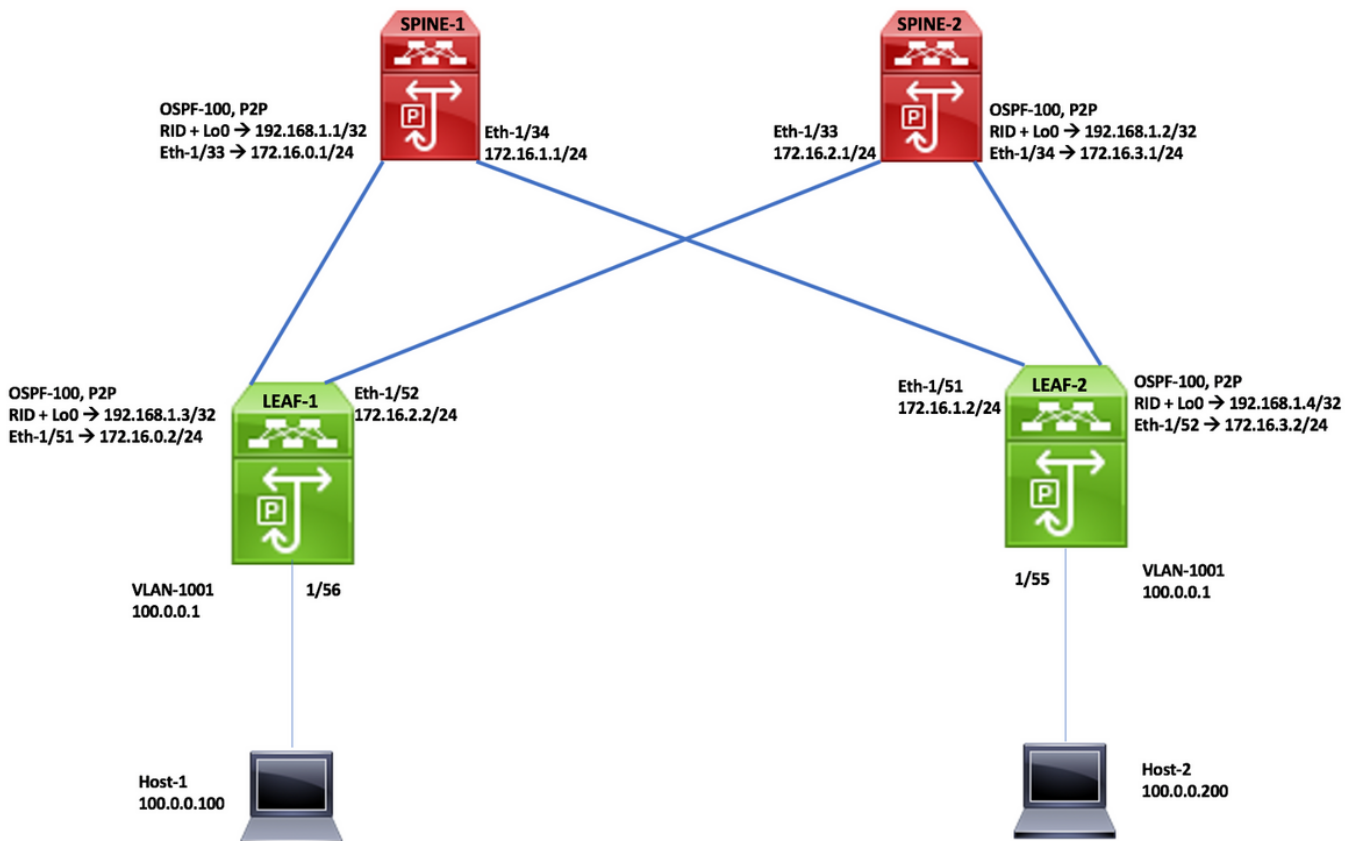
- EVPN uses several building blocks such as 'RD', 'RT' and 'VRF' from existing MPLS technologies
- EVPN operates in contrast to the existing VPLS by enabling control-plane based MAC learning in the core
- In EVPN, PEs participating in the EVPN instances learn customer MAC routes in control-plane using MP-BGP protocol
- Control-plane MAC learning provides a number of benefits that allow EVPN to address the VPLS shortcomings, including support for multihoming with per-flow load balancing
- SR L2 EVPN is new feature available in NXOS 9.3(1) and it's supported on Nexus 9300 FX2 series platform

### Limitations for L2 EVPN over SR MPLS:

- Segment routing Layer 2 EVPN flooding is based on the ingress replication mechanism
- It uses EVPN Type 3 Route for BUM Traffic
- MPLS core does not support multicast
- ARP suppression is not supported
- Consistency checking on VPC is not supported
- The same L2 EVI and L3 EVI cannot be configured together

## Configure

### Network Diagram



### High Level Configuration Steps:

- Install Features
- Configure Ip address -Underlay
- Configure IGP -OSPF
- Configure MP-BGP
- Configure VLAN and EVPN Overlay
- Configure End Host for Layer 2

### SPINE -1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
install feature-set mpls feature-set mpls feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam	interface Ethernet1/33 ip address 172.16.0.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown	router bgp 65001 router-id 192.168.1.1 address-family ipv4 unicast network 192.168.1.1/32 route-map label-index-spine1 allocate-label all address-family ipv4 labeled-unicast address-family l2vpn evpn
mpls label range 5000 450000	interface Ethernet1/34 ip address 172.16.1.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown	template peer EVPN remote-as 65001 update-source loopback0 address-family l2vpn evpn send-community extended route-reflector-client encapsulation mpls
segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.1/32 index 211	interface loopback0 ip address 192.168.1.1/32 ip router ospf 100 area 0.0.0.0	template peer Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended route-reflector-client next-hop-self soft-reconfiguration inbound always
route-map label-index-spine1 permit 10 set label-index 211	router ospf 100 segment-routing mpls router-id 192.168.1.1	neighbor 172.16.0.2 inherit peer Labeled-unicast neighbor 172.16.1.2 inherit peer Labeled-unicast neighbor 192.168.1.3 inherit peer EVPN neighbor 192.168.1.4 inherit peer EVPN

### SPINE -2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
install feature-set mpls feature-set mpls feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam	interface Ethernet1/33 ip address 172.16.2.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown	router bgp 65001 router-id 192.168.1.2 address-family ipv4 unicast network 192.168.1.2/32 route-map label-index-spine2 allocate-label all address-family ipv4 labeled-unicast address-family l2vpn evpn
mpls label range 5000 450000	interface Ethernet1/34 ip address 172.16.3.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown	template peer EVPN remote-as 65001 update-source loopback0 address-family l2vpn evpn send-community extended route-reflector-client encapsulation mpls
segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.2/32 index 221	interface loopback0 ip address 192.168.1.2/32 ip router ospf 100 area 0.0.0.0	template peer Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended route-reflector-client next-hop-self soft-reconfiguration inbound always
route-map label-index-spine2 permit 10 set label-index 221	router ospf 100 segment-routing mpls router-id 192.168.1.2	neighbor 172.16.2.2 inherit peer Labeled-unicast neighbor 172.16.3.2 inherit peer Labeled-unicast neighbor 192.168.1.3 inherit peer EVPN neighbor 192.168.1.4 inherit peer EVPN

## Leaf-1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
<pre>install feature-set mpls nv overlay evpn feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam feature nv overlay  fabric forwarding anycast-gateway-mac 0000.0000.1111 mpls label range 5000 450000  vlan 1,1001 segment-routing   mpls     global-block 16000 25000     connected-prefix-sid-map     address-family ipv4       192.168.1.3/32 index 311 vlan 1001 evi auto  route-map label-index-leaf-1 permit 10 set label-index 311 vrf context Tenant-A evi 30001  interface Vlan1001 no shutdown vrf member Tenant-A ip address 100.0.0.1/24 fabric forwarding mode anycast-gateway</pre>	<pre>interface Ethernet1/51 ip address 172.16.0.2/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown  interface Ethernet1/52 ip address 172.16.2.2/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown  interface Ethernet1/56 switchport switchport mode trunk switchport trunk allowed vlan 1001 no shutdown  interface loopback0 ip address 192.168.1.3/32 ip router ospf 100 area 0.0.0.0  router ospf 100 segment-routing mpls router-id 192.168.1.3</pre>	<pre>router bgp 65001 router-id 192.168.1.3 address-family ipv4 unicast   network 192.168.1.3/32 route-map label-index-leaf-1   allocate-label all address-family ipv4 labeled-unicast address-family l2vpn evpn template peer EVPN   remote-as 65001   update-source loopback0 address-family l2vpn evpn   send-community extended   encapsulation mpls template peer Labeled-unicast   remote-as 65001   address-family ipv4 labeled-unicast   send-community extended   soft-reconfiguration inbound always  neighbor 172.16.0.1 inherit peer Labeled-unicast neighbor 172.16.2.1 inherit peer Labeled-unicast neighbor 192.168.1.1 inherit peer EVPN neighbor 192.168.1.2 inherit peer EVPN vrf Tenant-A  evpn encapsulation mpls source-interface loopback0</pre>

## Leaf-2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
<pre>install feature-set mpls nv overlay evpn feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam feature nv overlay  fabric forwarding anycast-gateway-mac 0000.0000.1111 mpls label range 5000 450000  vlan 1,1001 segment-routing   mpls     global-block 16000 25000     connected-prefix-sid-map     address-family ipv4       192.168.1.4/32 index 321 vlan 1001 evi auto  route-map label-index-Leaf2 permit 10 set label-index 321 vrf context Tenant-A evi 30001  interface Vlan1001 no shutdown vrf member Tenant-A ip address 100.0.0.1/24 fabric forwarding mode anycast-gateway</pre>	<pre>interface Ethernet1/51 ip address 172.16.1.2/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown  interface Ethernet1/52 ip address 172.16.3.2/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown  interface Ethernet1/55 switchport switchport mode trunk switchport trunk allowed vlan 1001 no shutdown  interface loopback0 ip address 192.168.1.4/32 ip router ospf 100 area 0.0.0.0  router ospf 100 segment-routing mpls router-id 192.168.1.4</pre>	<pre>router bgp 65001 router-id 192.168.1.4 address-family ipv4 unicast   network 192.168.1.4/32 route-map label-index-Leaf2   allocate-label all address-family ipv4 labeled-unicast address-family l2vpn evpn template peer EVPN   remote-as 65001   update-source loopback0 address-family l2vpn evpn   send-community extended   encapsulation mpls template peer Labeled-unicast   remote-as 65001   address-family ipv4 labeled-unicast   send-community extended   soft-reconfiguration inbound always  neighbor 172.16.1.1 inherit peer Labeled-unicast neighbor 172.16.3.1 inherit peer Labeled-unicast neighbor 192.168.1.1 inherit peer EVPN neighbor 192.168.1.2 inherit peer EVPN vrf Tenant-A  evpn encapsulation mpls source-interface loopback0</pre>

# Verification

## Host1# show ip int brief

```
IP Interface Status for VRF "default"(1)
Interface IP Address Interface Status
Vlan1001 100.0.0.200 protocol-up/link-up/admin-up
```

### Mhost1# ping 100.0.0.100

```
PING 100.0.0.100 (100.0.0.100): 56 data bytes
64 bytes from 100.0.0.100: icmp_seq=0 ttl=253 time=0.84 ms
64 bytes from 100.0.0.100: icmp_seq=1 ttl=253 time=0.45 ms
64 bytes from 100.0.0.100: icmp_seq=2 ttl=253 time=0.443 ms
64 bytes from 100.0.0.100: icmp_seq=3 ttl=253 time=0.438 ms
64 bytes from 100.0.0.100: icmp_seq=4 ttl=253 time=0.431 ms
```

--- 100.0.0.100 ping statistics ---

```
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.431/0.52/0.84 ms
```

## Host2# show ip int brief

```
IP Interface Status for VRF "default"(1)
Interface IP Address Interface Status
Vlan1001 100.0.0.100 protocol-up/link-up/admin-up
```

### Mhost2# ping 100.0.0.200

```
PING 100.0.0.200 (100.0.0.200): 56 data bytes
64 bytes from 100.0.0.200: icmp_seq=0 ttl=253 time=0.854 ms
64 bytes from 100.0.0.200: icmp_seq=1 ttl=253 time=0.46 ms
64 bytes from 100.0.0.200: icmp_seq=2 ttl=253 time=0.451 ms
64 bytes from 100.0.0.200: icmp_seq=3 ttl=253 time=0.427 ms
64 bytes from 100.0.0.200: icmp_seq=4 ttl=253 time=0.418 ms
```

--- 100.0.0.200 ping statistics ---

```
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.418/0.522/0.854 ms
```

Mhost2#

## Leaf1# show bgp l2vpn evpn

BGP routing table information for VRF default, address family L2VPN EVPN

BGP table version is 57, Local Router ID is 192.168.1.3

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, \*-valid, >-best

Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, l-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
<b>Route Distinguisher: 192.168.1.3:37864 (L2VNI 1001)</b>					
*>[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[0]:[0.0.0.0]/216	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	32768	i	
<b>*&gt;[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[32]:[100.0.0.100]/248</b>					
192.168.1.4	100	0	i		
<b>*&gt;[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272</b>					
192.168.1.3	100	32768	i		
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	32768	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	0	i	
Route Distinguisher: 192.168.1.4:37864					
*[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[0]:[0.0.0.0]/216	192.168.1.4	100	0	i	
*i	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[32]:[100.0.0.100]/248	192.168.1.4	100	0	i	
*i	192.168.1.4	100	0	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	0	i	
*i	192.168.1.4	100	0	i	

## Leaf2# show bgp l2vpn evpn

BGP routing table information for VRF default, address family L2VPN EVPN

BGP table version is 40, Local Router ID is 192.168.1.4

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, \*-valid, >-best

Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, l-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
<b>Route Distinguisher: 192.168.1.3:37864</b>					
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	0	i	
*i	192.168.1.3	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272	192.168.1.3	100	0	i	
*i	192.168.1.3	100	0	i	
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	0	i	
*i	192.168.1.3	100	0	i	
<b>Route Distinguisher: 192.168.1.4:37864 (L2VNI 1001)</b>					
*>[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[0]:[0.0.0.0]/216	192.168.1.4	100	32768	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[32]:[100.0.0.100]/248	192.168.1.4	100	32768	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272	192.168.1.3	100	0	i	
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	0	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	32768	i	

# References

[Segment Routing on Cisco Nexus 9500, 9300, 9200, 3200, and 3100 Platform Switches White paper](#)

[Configuring Layer2 EVPN over Segment Routing MPLS](#)