

Deploy Layer3 EVPN over SR MPLS [Ospf / iBGP] [PE-CE is OSPF] in Nexus 9300

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[MPLS L3VPN Recap](#)

[Overview of EVPN with L3VPN \(MPLS SR\)](#)

[Network Diagram](#)

[Configuration](#)

[Verify](#)

[Related Information](#)

Introduction

This document describes how to deploy/configure Layer3 EVPN over Segment Routing (SR) Multiprotocol Label Switching (MPLS) on Nexus 9300 products with PE-CE protocol as Open Shortest Path First (OSPF).

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Border Gateway Protocol (BGP)
- Open Shortest Path First (OSPF)
- L3VPN
- EVPN
- Segment Routing (SR)

Components Used

The information in this document is based on these software and hardware versions:

- SPINE Hardware - 93360YC-FX2 running with 9.3.(3)
- LEAF Hardware - 93240YC-FX2 running with 9.3.(3)
- CLIENT - 93216TC-FX2 (Host-1), Catlyst-3750 (Host-2)

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 Information

MPLS L3VPN Recap

A VPN is:

- An IP-based network delivering private network services over a public infrastructure.
- A set of sites that are allowed to communicate with each other privately over the Internet or other public or private networks.

Conventional VPNs are created by the configuration of a full mesh of tunnels or permanent virtual circuits (PVCs) to all sites in a VPN. This type of VPN is not easy to maintain or expand, as adding a new site requires changing each edge device in the VPN.

MPLS-based VPNs are created in Layer 3 and are based on the peer model. The peer model enables the service provider and the customer to exchange Layer 3 routing information. The service provider relays the data between the customer sites without customer involvement.

MPLS VPNs are easier to manage and expand than conventional VPNs. When a new site is added to an MPLS VPN, only the edge router of the service provider that provides services to the customer site needs to be updated.

These are the components of the MPLS VPN:

- Provider (P) router- Router in the core of the provider network. PE routers run MPLS switching and do not attach VPN labels to routed packets. VPN labels are used to direct data packets to the correct private network or customer edge router.
- PE router- Router that attaches the VPN label to incoming packets based on the interface or subinterface on which they are received, and also attaches the MPLS core labels. A PE router attaches directly to a CE router.
- Customer (C) router- Router in the Internet service provider (ISP) or enterprise network.
- Customer Edge (CE) router- Edge router on the network of the ISP that connects to the PE router on the network. A CE router must interface with a PE router.

Overview of EVPN with L3VPN (MPLS SR)

Data Center (DC) deployments have adopted VXLAN EVPN (or) MPLS EVPN for its benefits such as EVPN control-plane learning, multitenancy, seamless mobility, redundancy, and easier POD additions. Similarly, the CORE is either a Label Distribution Protocol (LDP)-based MPLS L3VPN network or transitioning from the traditional MPLS L3VPN LDP-based underlay to a more sophisticated solution like Segment Routing (SR).

Segment Routing is adopted for its benefits such as:

- Unified IGP and MPLS control planes
- Simpler traffic engineering methods
- Easier configuration
- SDN adoption
- EVPN (RFC 7432) is BGP MPLS-based solution that has been used for next-generation Ethernet services in a virtualized data center network.
- EVPN uses several building blocks such as RD, RT, and VRF from MPLS technologies that exist.
- L3 EVPN over SR which was introduced in NXOS 7.0(3)I6(1) release uses the EVPN Type-5 route with MPLS encapsulation.
- L3 EVPN over SR offers Multi-tenant, Scalability, and High Performance for evolved data center services.

Note: In DC, the data plane can be VXLAN or MPLS.

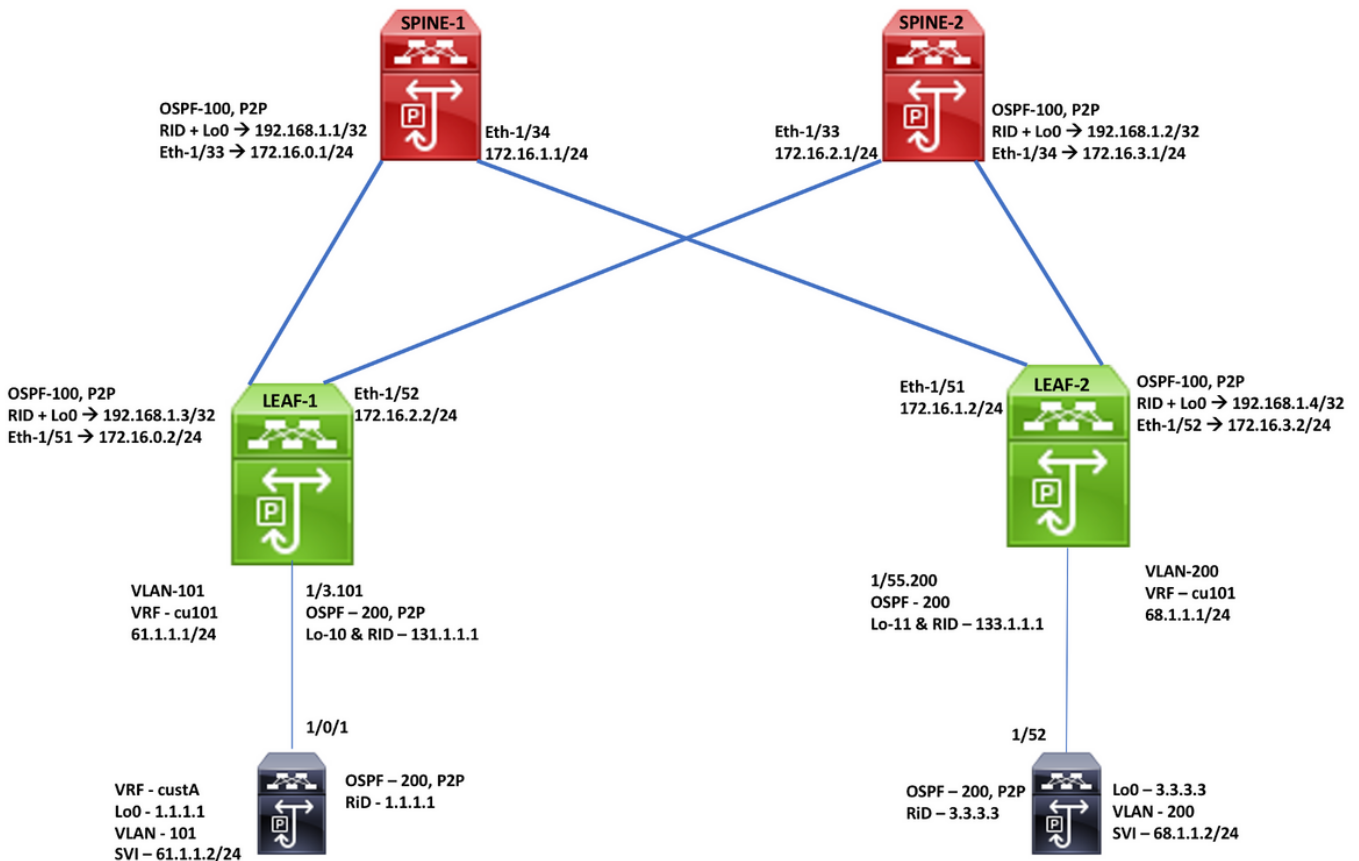
Traditional MPLS L3 VPN

Main build blocks: RD, RT, and VRF
 Underlay Layer for Transport: IGP, LDP, and RSVP-TE
 Overlay Layer for Service: VPNv4 and VPNv6

MPLS L3 VPN over SR

Main build blocks: RD, RT, and VRF
 Underlay Layer for Transport: IGP/BGP-LU and SR-TE
 Overlay Layer for Service: EVPN

Network Diagram



Configuration

SPINE-1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
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 template peer EVPN remote-as 65001 update-source loopback0 address-family l2vpn evpn send-community extended route-reflector-client encapsulation mpls 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 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
mpls label range 5000 450000 segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.1/32 index 211	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	
route-map label-index-spine1 permit 10 set label-index 211	interface loopback0 ip address 192.168.1.1/32 ip router ospf 100 area 0.0.0.0	
	router ospf 100 segment-routing mpls router-id 192.168.1.1	

SPINE-2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
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 template peer EVPN remote-as 65001 update-source loopback0 address-family l2vpn evpn send-community extended route-reflector-client encapsulation mpls 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 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
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	
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	
route-map label-index-spine2 permit 10 set label-index 221	router ospf 100 segment-routing mpls router-id 192.168.1.2	

LEAF-1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index

```

install feature-set mpls
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

mpls label range 5000 450000
segment-routing
mpls
  global-block 16000 25000
  connected-prefix-sid-map
  address-family ipv4
    192.168.1.3/32 index 311

ip prefix-list test1 seq 5 permit 61.1.1.0/24
ip prefix-list test1 seq 10 permit 131.1.1.1/32

ip prefix-list test3 seq 5 permit 1.1.1.1/32

route-map bgp65001 permit 10
  match route-type internal
route-map direct1 permit 10
  match ip address prefix-list test1
  set community 65001:10
route-map label-index-leaf-1 permit 10
  set label-index 311
route-map ospf200 permit 10
  match ip address prefix-list test3

vrf context cu101
  rd auto
  address-family ipv4 unicast
  route-target import 1:101
  route-target import 1:101 evpn
  route-target export 1:101
  route-target export 1:101 evpn
  
```

OSPF Configuration

```

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 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

PE-CE
vrf cu101
  address-family ipv4 unicast

interface Ethernet1/3
  no shutdown
interface Ethernet1/3.101
  encapsulation dot1q 101
  vrf member cu101
  ip address 61.1.1.1/24
  ip ospf network point-to-point
  ip router ospf 200 area 0.0.0.0
  no shutdown

interface loopback10
  vrf member cu101
  ip address 131.1.1.1/32
  ip router ospf 200 area 0.0.0.0

router ospf 200
  vrf cu101
  router-id 131.1.1.1
  redistribute bgp 65001 route-map bgp65001
  
```

BGP/EVPN Configuration

```

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 cu101
  router-id 131.1.1.1
  address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map direct1
  redistribute ospf 200 route-map ospf200
  
```

LEAF-2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index

```

install feature-set mpls
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

mpls label range 5000 450000
segment-routing
mpls
  global-block 16000 25000
  connected-prefix-sid-map
  address-family ipv4
    192.168.1.4/32 index 321

ip prefix-list new seq 5 permit 68.1.1.0/24
ip prefix-list new seq 10 permit 133.1.1.1/32

ip prefix-list new1 seq 5 permit 3.3.3.3/32

ip prefix-list redtoospf seq 5 permit 61.1.1.0/24
ip prefix-list redtoospf seq 10 permit 1.1.1.1/32

route-map bgp65001 permit 10
  match route-type internal
route-map direct1 permit 10
  match ip address prefix-list new
route-map label-index-Leaf2 permit 10
  set label-index 321
route-map ospf200 permit 10
  match ip address prefix-list new1

vrf context cu101
rd auto
address-family ipv4 unicast
route-target import 1:101
route-target import 1:101 evpn
route-target export 1:101
route-target export 1:101 evpn
  
```

OSPF Configuration

```

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 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

PE-CE
vrf cu101
  address-family ipv4 unicast
interface Ethernet1/55
no shutdown
interface Ethernet1/55.200
encapsulation dot1q 200
vrf member cu101
ip address 68.1.1.1/24
ip ospf network point-to-point
ip router ospf 200 area 0.0.0.0
no shutdown

interface loopback11
vrf member cu101
ip address 133.1.1.1/32
ip router ospf 200 area 0.0.0.0

router ospf 200
vrf cu101
  router-id 133.1.1.1
  redistribute bgp 65001 route-map bgp65001
  
```

BGP/EVPN Configuration

```

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 cu101
router-id 133.1.1.1
address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map direct1
  redistribute ospf 200 route-map ospf200
  
```

End-Host Configuration

Host-1 / Cat-3750

```
vrf definition custA
rd 101:1
!
address-family ipv4
exit-address-family
!

interface Loopback0
vrf forwarding custA
ip address 1.1.1.1 255.255.255.255

interface GigabitEthernet1/0/1
switchport trunk allowed vlan 101
switchport trunk encapsulation dot1q
switchport mode trunk
!

interface Vlan101
vrf forwarding custA
ip address 61.1.1.2 255.255.255.0
ip ospf network point-to-point
ip ospf 200 area 0.0.0.0

router ospf 200 vrf custA
router-id 1.1.1.1
network 1.1.1.1 0.0.0.0 area 0.0.0.0
network 61.1.1.0 0.0.0.255 area 0.0.0.0
```

Host-2 / N9K

```
feature ospf
feature interface-vlan

interface Ethernet1/52
switchport
switchport mode trunk
switchport trunk allowed vlan 200
no shutdown

interface Vlan200
no shutdown
ip address 68.1.1.2/24
ip ospf network point-to-point
ip router ospf 200 area 0.0.0.0

interface loopback0
ip address 3.3.3.3/32
ip router ospf 200 area 0.0.0.0

router ospf 200
router-id 3.3.3.3
```

Verify

Host2# show ip int brief

```
IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Vlan200        68.1.1.2        protocol-up/link-up/admin-up
Vlan1001       100.0.0.100    protocol-down/link-down/admin-up
Lo0            3.3.3.3         protocol-up/link-up/admin-up
```

Host2# show ip route

```
IP Route Table for VRF "default"
*** denotes best ucast next-hop
**** denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

1.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
3.3.3.3/32, ubest/mbest: 2/0, attached
  *via 3.3.3.3, Lo0, [0/0], 20:16:34, local
  *via 3.3.3.3, Lo0, [0/0], 20:16:34, direct
61.1.1.0/24, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
68.1.1.0/24, ubest/mbest: 1/0, attached
  *via 68.1.1.2, Vlan200, [0/0], 20:20:55, direct
68.1.1.2/32, ubest/mbest: 1/0, attached
  *via 68.1.1.2, Vlan200, [0/0], 20:20:55, local
131.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
133.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/41], 20:15:32, ospf-200, intra
```

Host2# traceroute 1.1.1.1

```
traceroute to 1.1.1.1 (1.1.1.1), 30 hops max, 40 byte packets
 1 68.1.1.1 (68.1.1.1)  0.989 ms  0.585 ms  0.407 ms
 2 172.16.3.1 (172.16.3.1)  0.886 ms  172.16.1.1 (172.16.1.1)  0.765 ms  0.731 ms
   [Label=16311 E=0 TTL=1 S=0, Label=492289 E=0 TTL=1 S=1]
   [Label=16311 E=0 TTL=1 S=0, Label=492289 E=0 TTL=1 S=1]
 3 172.16.0.2 (172.16.0.2)  0.717 ms  172.16.2.2 (172.16.2.2)  0.509 ms  172.16.0.2 (172.16.0.2)  0.678 ms
   [Label=492289 E=0 TTL=1 S=1]
   [Label=492289 E=0 TTL=1 S=1]
 4 61.1.1.2 (61.1.1.2)  2.061 ms * 1.315 ms
```

Host2# ping 1.1.1.1 source 3.3.3.3

```
PING 1.1.1.1 (1.1.1.1) from 3.3.3.3: 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=251 time=5.538 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=251 time=1.338 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=251 time=2.201 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=251 time=2.217 ms
64 bytes from 1.1.1.1: icmp_seq=4 ttl=251 time=4.021 ms

--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.338/3.063/5.538 ms
```

```
Leaf1# show ip route 3.3.3.3/32 vrf cu101
```

```
IP Route Table for VRF "cu101"  
** denotes best ucast next-hop  
*** denotes best mcast next-hop  
'[x/y]' denotes [preference/metric]  
'<string>' in via output denotes VRF <string>  
  
3.3.3.3/32, ubest/mbest: 1/0  
  *via 192.168.1.4%default, [200/2], 00:44:27, bgp-65001, internal, tag 65001 (mpls-vpn)
```

```
Leaf1# show forwarding mpls 192.168.1.4/32
```

```
slot 1  
=====
```

Local Label	Prefix Table Id	FEC (Prefix/Tunnel id)	Next-Hop	Interface	Out Label
16321	0x1	192.168.1.4/32	172.16.0.1	Eth1/51	16321
"	0x1	192.168.1.4/32	172.16.2.1	Eth1/52	16321

```
SWAP  
SWAP
```

```
Leaf1# show forwarding 3.3.3.3/32 vrf cu101
```

```
slot 1  
=====
```

```
IPv4 routes for table cu101/base
```

Prefix	Next-hop	Interface	Labels	Partial Install
*3.3.3.3/32	172.16.0.1	Ethernet1/51	PUSH 16321 492288	
	172.16.2.1	Ethernet1/52	PUSH 16321 492288	

Related Information

- [Multiprotocol BGP MPLS VPN](#)
- [Segment Routing on Cisco Nexus 9500, 9300, 9200, 3200, and 3100 Platform Switches White paper](#)
- [Configuring Layer 3 EVPN and Layer 3 VPN over Segment Routing MPLS](#)