

# Deploy Layer3 EVPN over Segment Routing MPLS in Nexus 9300

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[MPLS L3VPN Recap](#)

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

[Configure](#)

[Topology](#)

[High-Level Configuration](#)

[Verify](#)

[Troubleshoot](#)

[Related Information](#)

## Introduction

This document describes how to deploy/configure Layer3 (L3) Ethernet VPN (EVPN) over Segment Routing (SR) Multiprotocol Label Switching (MPLS) [Open Shortest Path First (OSPF) / Internal Border Gateway Protocol (iBGP)] on Nexus 9300 products.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Border Gateway Protocol (BGP)
- L3VPN
- EVPN
- SR

### Components Used

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

- SPINE Hardware - 93360YC-FX2 that runs Release 9.3.(3)
- LEAF Hardware - 93240YC-FX2 that runs Release 9.3.(3)
- CLIENT - 93216TC-FX2

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 configuring 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 the addition of a new site requires a change to each edge device in the VPN.

MPLS-based VPNs are created in L3 and are based on the peer model. The peer model enables the service provider and the customer to exchange L3 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. P 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 SR.

SR is adopted for its benefits such as:

- Unified IGP and MPLS control planes
- Simpler traffic engineering methods
- Easier configuration

- Software-Defined Networking (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. It uses several building blocks such as Route Distinguisher (RD), Route Target (RT), and Virtual Routing and Forwarding (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. It 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

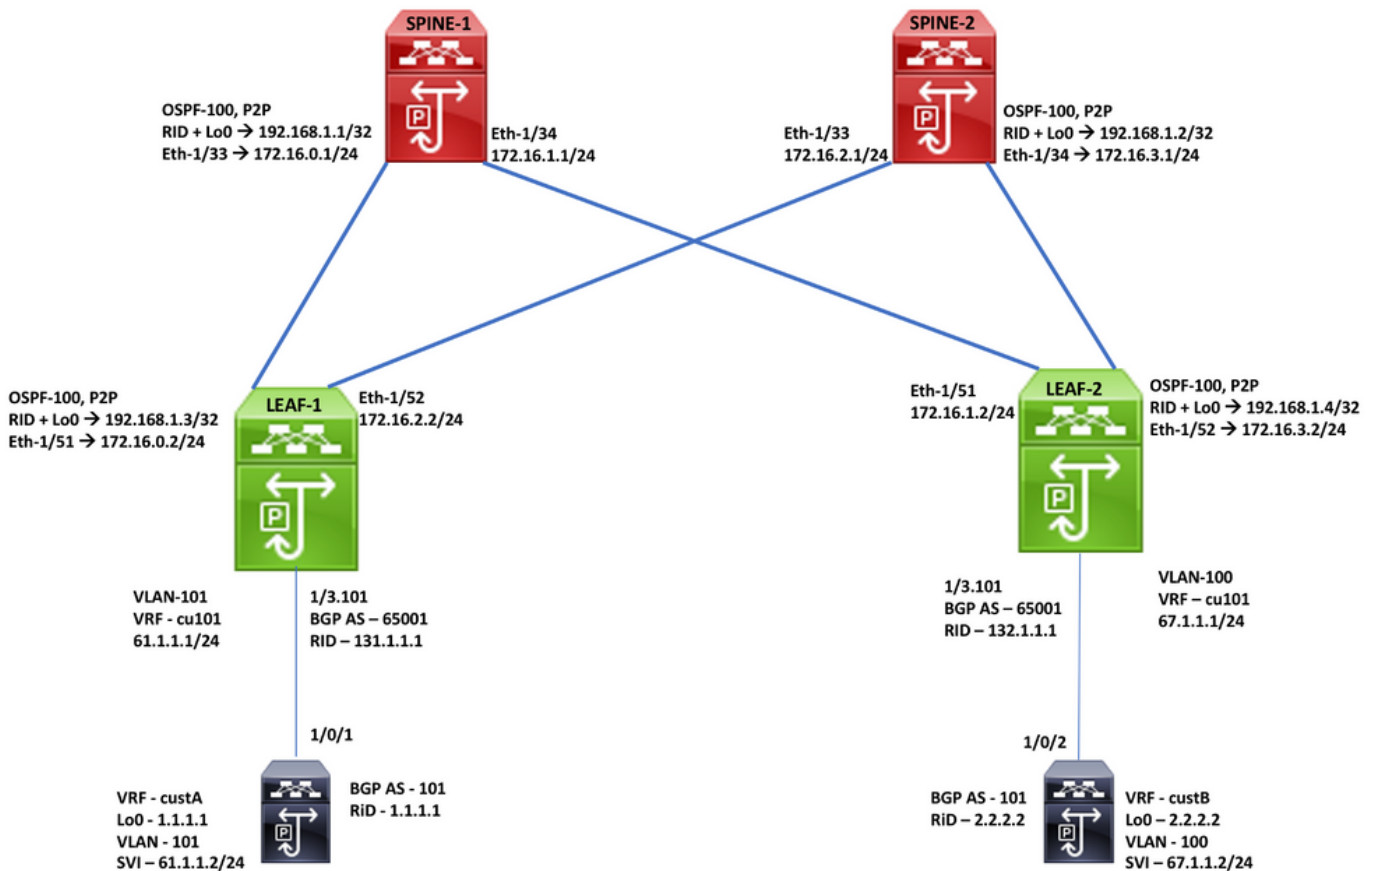
Main Build Blocks: RD, RT, and VRF

Underlay Layer for Transport: IGP/BGP-LU and S

Overlay Layer for Service: EVPN

## Configure

### Topology



### High-Level Configuration

1. Install Features
2. Configure IP address - Underlay

3. Configure IGP - OSPF
4. Configure MP - BGP
5. Configure VLAN and EVPN Overlay
6. Configure e-BGP between Hosts and LEAFs

SPINE-1 Configuration		
Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
feature-set mpls	interface Ethernet1/33	router bgp 65001
feature ospf	ip address 172.16.0.1/24	router-id 192.168.1.1
feature bgp	ip ospf network point-to-point	address-family ipv4 unicast
feature mpls segment-routing	ip router ospf 100 area 0.0.0.0	network 192.168.1.1/32 route-map label-index-spine1
feature mpls evpn	mpls ip forwarding	allocate-label all
feature interface-vlan	no shutdown	address-family ipv4 labeled-unicast
feature mpls oam		address-family l2vpn evpn
	interface Ethernet1/34	template peer EVPN
	ip address 172.16.1.1/24	remote-as 65001
	ip ospf network point-to-point	update-source loopback0
mpls label range 5000 45000	ip router ospf 100 area 0.0.0.0	address-family l2vpn evpn
segment-routing	mpls ip forwarding	send-community extended
mpls	no shutdown	route-reflector-client
global-block 16000 25000		encapsulation mpls
connected-prefix-sid-map		template peer Labeled-unicast
address-family ipv4	interface loopback0	remote-as 65001
192.168.1.1/32 index 211	ip address 192.168.1.1/32	address-family ipv4 labeled-unicast
	ip router ospf 100 area 0.0.0.0	send-community extended
		route-reflector-client
route-map label-index-spine1 permit 10		next-hop-self
set label-index 211	router ospf 100	soft-reconfiguration inbound always
	segment-routing mpls	neighbor 172.16.0.2
	router-id 192.168.1.1	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
feature-set mpls	interface Ethernet1/33	router bgp 65001
feature ospf	ip address 172.16.2.1/24	router-id 192.168.1.2
feature bgp	ip ospf network point-to-point	address-family ipv4 unicast
feature mpls segment-routing	ip router ospf 100 area 0.0.0.0	network 192.168.1.2/32 route-map label-index-spine2
feature mpls evpn	mpls ip forwarding	allocate-label all
feature interface-vlan	no shutdown	address-family ipv4 labeled-unicast
feature mpls oam		address-family l2vpn evpn
	interface Ethernet1/34	template peer EVPN
mpls label range 5000 45000	ip address 172.16.3.1/24	remote-as 65001
	ip ospf network point-to-point	update-source loopback0
segment-routing	ip router ospf 100 area 0.0.0.0	address-family l2vpn evpn
mpls	mpls ip forwarding	send-community extended
global-block 16000 25000	no shutdown	route-reflector-client
connected-prefix-sid-map		encapsulation mpls
address-family ipv4	interface loopback0	template peer Labeled-unicast
192.168.1.2/32 index 221	ip address 192.168.1.2/32	remote-as 65001
	ip router ospf 100 area 0.0.0.0	address-family ipv4 labeled-unicast
		send-community extended
route-map label-index-spine2 permit 10		route-reflector-client
set label-index 221	router ospf 100	next-hop-self
	segment-routing mpls	soft-reconfiguration inbound always
	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
feature-set mpls	interface Ethernet1/3.101	router bgp 65001
feature ospf	encapsulation dot1q 101	router-id 192.168.1.3
feature bgp	vrf member cu101	address-family ipv4 unicast
feature mpls segment-routing	ip address 61.1.1.1/24	network 192.168.1.3/32 route-map label-index-leaf-1
feature mpls evpn	no shutdown	allocate-label all
feature interface-vlan		address-family ipv4 labeled-unicast
feature mpls oam	interface Ethernet1/51	address-family l2vpn evpn
	ip address 172.16.0.2/24	template peer EVPN
	ip ospf network point-to-point	remote-as 65001
mpls label range 5000 450000	ip router ospf 100 area 0.0.0.0	update-source loopback0
	mpls ip forwarding	address-family l2vpn evpn
	no shutdown	send-community extended
segment-routing		encapsulation mpls
mpls	interface Ethernet1/52	template peer Labeled-unicast
global-block 16000 25000	ip address 172.16.2.2/24	remote-as 65001
connected-prefix-sid-map	ip ospf network point-to-point	address-family ipv4 labeled-unicast
address-family ipv4	ip router ospf 100 area 0.0.0.0	send-community extended
192.168.1.3/32 index 311	mpls ip forwarding	soft-reconfiguration inbound always
	no shutdown	template peer cu1
route-map label-index-leaf-1 permit 10		address-family ipv4 unicast
set label-index 311		as-override
	interface loopback0	send-community
vrf context cu101	ip address 192.168.1.3/32	soft-reconfiguration inbound always
rd auto	ip router ospf 100 area 0.0.0.0	neighbor 172.16.0.1
address-family ipv4 unicast		inherit peer Labeled-unicast
route-target import 1:101	router ospf 100	neighbor 172.16.2.1
route-target import 1:101 evpn	segment-routing mpls	inherit peer Labeled-unicast
route-target export 1:101	router-id 192.168.1.3	neighbor 192.168.1.1
route-target export 1:101 evpn		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
		neighbor 61.1.1.2
		inherit peer cu1
		remote-as 101

## LEAF-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 Ethernet1/3.101 encapsulation dot1q 100 vrf member cu101 ip address 67.1.1.1/24 no shutdown	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
mpls label range 5000 450000	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	template peer Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended soft-reconfiguration inbound always
segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.4/32 index 321	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	template peer cu1 address-family ipv4 unicast as-override send-community soft-reconfiguration inbound always
route-map label-index-Leaf2 permit 10 set label-index 321	interface loopback0 ip address 192.168.1.4/32 ip router ospf 100 area 0.0.0.0	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 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	router ospf 100 segment-routing mpls router-id 192.168.1.4	vrf cu101 router-id 132.1.1.1 address-family ipv4 unicast advertise l2vpn evpn neighbor 67.1.1.2 inherit peer cu1 remote-as 101

## End-Host Configuration

VRF, Loopback Configuration	Interface, SVI Configuration	BGP Configuration
vrf definition custA rd 101:1 ! address-family ipv4 exit-address-family !	interface GigabitEthernet1/0/1 switchport trunk allowed vlan 101 switchport trunk encapsulation dot1q switchport mode trunk !	router bgp 101 bgp log-neighbor-changes no bgp default ipv4-unicast !
vrf definition custB rd 101:2 ! address-family ipv4 exit-address-family	interface GigabitEthernet1/0/2 switchport trunk allowed vlan 100 switchport trunk encapsulation dot1q switchport mode trunk	address-family ipv4 vrf custA bgp router-id 1.1.1.1 network 1.1.1.1 mask 255.255.255.255 redistribute connected neighbor 61.1.1.1 remote-as 65001 neighbor 61.1.1.1 activate neighbor 61.1.1.1 send-community neighbor 61.1.1.1 soft-reconfiguration inbound exit-address-family
interface Loopback0 vrf forwarding custA ip address 1.1.1.1 255.255.255.255 !	interface Vlan100 vrf forwarding custB ip address 67.1.1.2 255.255.255.0 !	address-family ipv4 vrf custB bgp router-id 2.2.2.2 network 2.2.2.2 mask 255.255.255.255 redistribute connected
interface Loopback1 vrf forwarding custB ip address 2.2.2.2 255.255.255.255	interface Vlan101 vrf forwarding custA ip address 61.1.1.2 255.255.255.0 !	neighbor 67.0.0.1 soft-reconfiguration inbound neighbor 67.1.1.1 remote-as 65001 neighbor 67.1.1.1 activate neighbor 67.1.1.1 send-community neighbor 67.1.1.1 soft-reconfiguration inbound exit-address-family

## Verify

Use this section in order to confirm that your configuration works properly.



**Leaf 1 Captures : Control Plane and MPLS Data Plane:**

**Leaf1(config)# show ip bgp 1.1.1.1 vrf cu101**

```
BGP routing table information for VRF cu101, address family IPv4 Unicast
BGP routing table entry for 1.1.1.1/32, version 4
Paths: (2 available, best #1)
Flags: (0x880c0014) (high32 0x000020) on xmit-list, is in urib, is best urib route, is in HW, exported, has label
vpn: version 3, (0x00000000100002) on xmit-list
local label: 492288

Advertised path-id 1, VFN AF advertised path-id 1
Path type: external, path is valid, is best path, no labeled nexthop, in rib
AS-Path: 101 , path sourced external to AS
61.1.1.2 (metric 0) from 61.1.1.2 (1.1.1.1)
Origin IGP, MED 0, localpref 100, weight 0
Extcommunity: RT:1:101

Path type: external, path is valid, received only, no labeled nexthop
AS-Path: 101 , path sourced external to AS
61.1.1.2 (metric 0) from 61.1.1.2 (1.1.1.1)
Origin IGP, MED 0, localpref 100, weight 0

VRF advertise information:
Path-id 1 not advertised to any peer

VFN AF advertise information:
Path-id 1 not advertised to any peer
```

**Leaf1(config)# show bgp l2vpn evpn 1.1.1.1**

```
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 192.168.1.3:3
BGP routing table entry for [5]:[0]:[0]:[32]:[1.1.1.1]/224, version 6
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn, has label
local label: 492288

Advertised path-id 1
Path type: local, path is valid, is best path, no labeled nexthop
Gateway IP: 0.0.0.0
AS-Path: 101 , path sourced external to AS
0.0.0.0 (metric 0) from 0.0.0.0 (192.168.1.3)
Origin IGP, MED 0, localpref 100, weight 0
Received label 0
Extcommunity: RT:1:101

Path-id 1 advertised to peers:
192.168.1.1 192.168.1.2
```

**Leaf1(config)# show bgp ipv4 labeled-unicast 192.168.1.3**

```
BGP routing table information for VRF default, address family IPv4 Label Unicast
BGP routing table entry for 192.168.1.3/32, version 8
Paths: (1 available, best #1)
Flags: (0x20c0002) (high32 00000000) on xmit-list, is not in urib, has label
label af: version 11, (0x00000000100002) on xmit-list
local label: 3

Advertised path-id 1, Label AF advertised path-id 1
Path type: local, path is valid, is best path, no labeled nexthop
AS-Path: NONE, path locally originated
0.0.0.0 (metric 0) from 0.0.0.0 (192.168.1.3)
Origin IGP, MED not set, localpref 100, weight 32768
Prefix-SID Attribute: Length: 10
Label Index TLV: Length 7, Flags 0x0 Label Index 311

Path-id 1 not advertised to any peer

Label AF advertisement
Path-id 1 advertised to peers:
172.16.0.1 172.16.2.1
```

**Leaf1(config)# show forwarding mpls 192.168.1.4/32**

```
slot 1
-----
Local |Prefix |FEC |Next-Hop |Interface |Out
Label |Table Id |(Prefix/Tunnel id) | | |Label
-----|-----|-----|-----|-----|-----
16321 |0x1 |192.168.1.4/32 |172.16.0.1 |Eth1/51 |16321 SWAP
" |0x1 |192.168.1.4/32 |172.16.2.1 |Eth1/52 |16321 SWAP
```

**Leaf 2 Captures : Control Plane and MPLS Data Plane:**

**Leaf2# show forwarding 1.1.1.1/32 vrf cu101**

```
slot 1
-----
IPv4 routes for table cu101/base
-----
Prefix | Next-hop | Interface | Labels | Partial Install
-----|-----|-----|-----|-----
1.1.1.1/32 | 172.16.1.1 | Ethernet1/51 | POHS 14311 492288
172.16.1.1/32 | 172.16.3.1 | Ethernet1/52 | POHS 14311 492288

Leaf2#
Leaf2#
```

**Leaf2# show forwarding 172.16.1.1/24**

```
slot 1
-----
IPv4 routes for table default/base
-----
Prefix | Next-hop | Interface | Labels | Partial Install
-----|-----|-----|-----|-----
172.16.1.0/24 | Attached | Ethernet1/51 | |

Leaf2#
Leaf2#
```

**Leaf2# show forwarding mpls 192.168.1.3/32**

```
slot 1
-----
Local |Prefix |FEC |Next-Hop |Interface |Out
Label |Table Id |(Prefix/Tunnel id) | | |Label
-----|-----|-----|-----|-----|-----
16311 |0x1 |192.168.1.3/32 |172.16.1.1 |Eth1/51 |16311 SWAP
" |0x1 |192.168.1.3/32 |172.16.3.1 |Eth1/52 |16311 SWAP
```

**Leaf2# show forwarding 192.168.1.3/32**

```
slot 1
-----
IPv4 routes for table default/base
-----
Prefix | Next-hop | Interface | Labels | Partial Install
-----|-----|-----|-----|-----
192.168.1.3/32 | 172.16.1.1 | Ethernet1/51 | POHS 14311
172.16.3.1 | Ethernet1/52 | POHS 14311
```

**Spine 1 Captures**

**spine1# show bgp ipv4 labeled-unicast 1.1.1.1**

```
spine1# show bgp l2vpn evpn 1.1.1.1
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 192.168.1.3:3
BGP routing table entry for [5]:[0]:[0]:[32]:[1.1.1.1]/224, version 5
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn, is not in HW

Advertised path-id 1
Path type: internal, path is valid, is best path
Gateway IP: 0.0.0.0
AS-Path: 101 , path sourced external to AS
192.168.1.3 (metric 0) from 192.168.1.3 (192.168.1.3)
Origin IGP, MED 0, localpref 100, weight 0
Received label 492288
Extcommunity: RT:1:101

Path-id 1 advertised to peers:
192.168.1.4
```

**spine1# show bgp ipv4 labeled-unicast 192.168.1.3**

```
BGP routing table information for VRF default, address family IPv4 Label Unicast
BGP routing table entry for 192.168.1.3/32, version 5
Paths: (1 available, best #1)
Flags: (0x820c0012) (high32 00000000) on xmit-list, is in urib, is backup urib route, is in HW, has label
label af: version 7, (0x00000000100002) on xmit-list
local label: 16311

Advertised path-id 1, Label AF advertised path-id 1
Path type: internal, path is valid, received and used, is best path, no labeled nexthop, in rib
AS-Path: NONE, path sourced internal to AS
172.16.0.2 (metric 0) from 172.16.0.2 (192.168.1.3)
Origin IGP, MED not set, localpref 100, weight 0
Received label 3
Prefix-SID Attribute: Length: 10
Label Index TLV: Length 7, Flags 0x0 Label Index 311

Path-id 1 not advertised to any peer

Label AF advertisement
Path-id 1 advertised to peers:
172.16.1.2
```

**spine1# show forwarding mpls 192.168.1.4/32**

```
slot 1
-----
Local |Prefix |FEC |Next-Hop |Interface |Out
Label |Table Id |(Prefix/Tunnel id) | | |Label
-----|-----|-----|-----|-----|-----
16321 |0x1 |192.168.1.4/32 |172.16.1.2 |Eth1/34 |0 SWAP
```

## End-Host Captures

```
endhost#show ip int brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
Vlan1	unassigned	YES	NVRAM	up	up
Vlan100	67.1.1.2	YES	manual	up	up
Vlan101	61.1.1.2	YES	manual	up	up
Loopback0	1.1.1.1	YES	manual	up	up
Loopback1	2.2.2.2	YES	manual	up	up

```
endhost#ping vrf custB 1.1.1.1
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:  
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/7/17 ms
```

```
endhost#ping vrf custA 2.2.2.2
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:  
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/8/17 ms
```

```
endhost#traceroute vrf custB 1.1.1.1
```

```
Type escape sequence to abort.  
Tracing the route to 1.1.1.1  
VRF info: (vrf in name/id, vrf out name/id)  
 1 67.1.1.1 0 msec 8 msec 0 msec  
 2 172.16.3.1 0 msec 0 msec 0 msec  
 3 172.16.0.2 0 msec  
   172.16.2.2 0 msec  
   172.16.0.2 8 msec  
 4 61.1.1.2 0 msec * 0 msec
```

```
endhost#traceroute vrf custA 2.2.2.2
```

```
Type escape sequence to abort.  
Tracing the route to 2.2.2.2  
VRF info: (vrf in name/id, vrf out name/id)  
 1 61.1.1.1 0 msec 17 msec 0 msec  
 2 172.16.2.1 17 msec  
   172.16.0.1 0 msec  
   172.16.2.1 9 msec  
 3 172.16.3.2 0 msec  
   172.16.1.2 0 msec  
   172.16.3.2 17 msec  
 4 67.1.1.2 8 msec * 0 msec  
endhost#
```

## Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

## 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](#)
- [Technical Support & Documentation - Cisco Systems](#)