

在Nexus 9300中通過分段路由MPLS部署第3層EVPN

目錄

[簡介](#)

[必要條件](#)

[需求](#)

[採用元件](#)

[背景資訊](#)

[MPLS L3VPN重述](#)

[含L3VPN\(MPLS SR\)的EVPN概觀](#)

[設定](#)

[拓撲](#)

[高級配置](#)

[驗證](#)

[疑難排解](#)

[相關資訊](#)

簡介

本檔案介紹如何在Nexus 9300產品上透過區段路由(SR)多重協定標籤交換(MPLS)[開放最短路徑優先(OSPF)/內部邊界閘道通訊協定(iBGP)]部署/設定第3層(L3)乙太網路VPN(EVPN)。

必要條件

需求

思科建議您瞭解以下主題：

- 邊界閘道通訊協定(BGP)
- L3VPN
- EVPN
- SR

採用元件

本文中的資訊係根據以下軟體和硬體版本：

- 主乾硬體 — 93360行9.3.(3)版的YC-FX2
- 枝葉硬體 — 93240行9.3.(3)版的YC-FX2
- 客戶端 — 93216TC-FX2

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除（預設）的組態來啟動。如果您的網路正在作用，請確保您已瞭解任何指令可能造成的影響。

背景資訊

MPLS L3VPN重述

VPN是：

- 基於IP的網路，通過公共基礎設施提供專用網路服務。
- 允許通過Internet或其它公共或專用網路相互私下通訊的一組站點。

傳統VPN是通過配置到VPN中所有站點的全網狀隧道或永久虛擬電路(PVC)建立的。這種型別的VPN不易維護或擴展，因為新增新站點需要更改VPN中的每個邊緣裝置。

基於MPLS的VPN在第3層建立，並且基於對等體模型。對等體模型使服務提供商和客戶能夠交換L3路由資訊。服務提供商在客戶站點之間轉發資料，而無需客戶參與。

MPLS VPN比傳統VPN更易於管理和擴展。向MPLS VPN新增新站點時，僅需要更新向客戶站點提供服務的服務提供商的邊緣路由器。

以下是MPLS VPN的元件：

- 提供商(P)路由器 — 提供商網路核心的路由器。P路由器運行MPLS交換，並且不會將VPN標籤附加到路由資料包。VPN標籤用於將資料包定向到正確的專用網路或客戶邊緣路由器。
- PE路由器 — 根據接收傳入資料包的介面或子介面將VPN標籤附加到這些資料包的路由器，同時附加MPLS核心標籤。PE路由器直接連線到CE路由器。
- 客戶(C)路由器 — 網際網路服務提供商(ISP)或企業網路中的路由器。
- 客戶邊緣(CE)路由器 — ISP網路中的邊緣路由器，連線到網路上的PE路由器。CE路由器必須與PE路由器介面。

含L3VPN(MPLS SR)的EVPN概觀

資料中心(DC)部署已採用VXLAN EVPN或MPLS EVPN，因為它具有多種優勢，例如EVPN控制平面學習、多租戶、無縫移動性、冗餘和更輕鬆的POD新增。同樣，CORE是基於標籤分發協定(LDP)的MPLS L3VPN網路，或從傳統的基於MPLS L3VPN LDP的底層過渡到更複雜的解決方案，如SR。

SR的優點如下：

- 整合IGP和MPLS控制平面
- 更簡單的流量工程方法
- 更輕鬆的配置
- 軟體定義網路(SDN)採用

EVPN(RFC 7432)是基於BGP MPLS的解決方案，已用於虛擬化資料中心網路中的下一代乙太網服務。它使用多個構建塊，例如路由區分器(RD)、路由目標(RT)以及虛擬路由和轉發(VRF)，而不是現有的MPLS技術。

NXOS 7.0(3)I6(1)版本中引入的L3 EVPN over SR使用帶MPLS封裝的EVPN第5類路由。它為演化的資料中心服務提供多租戶、可擴充性和高效能。

附註：在DC中，資料平面可以是VXLAN或MPLS。

傳統MPLS第3層VPN

主構建塊：RD、RT和VRF

用於傳輸的底層層：IGP、LDP和RSVP-TE

服務的覆蓋層：VPNv4和VPNv6

使用SR的MPLS第3層VPN

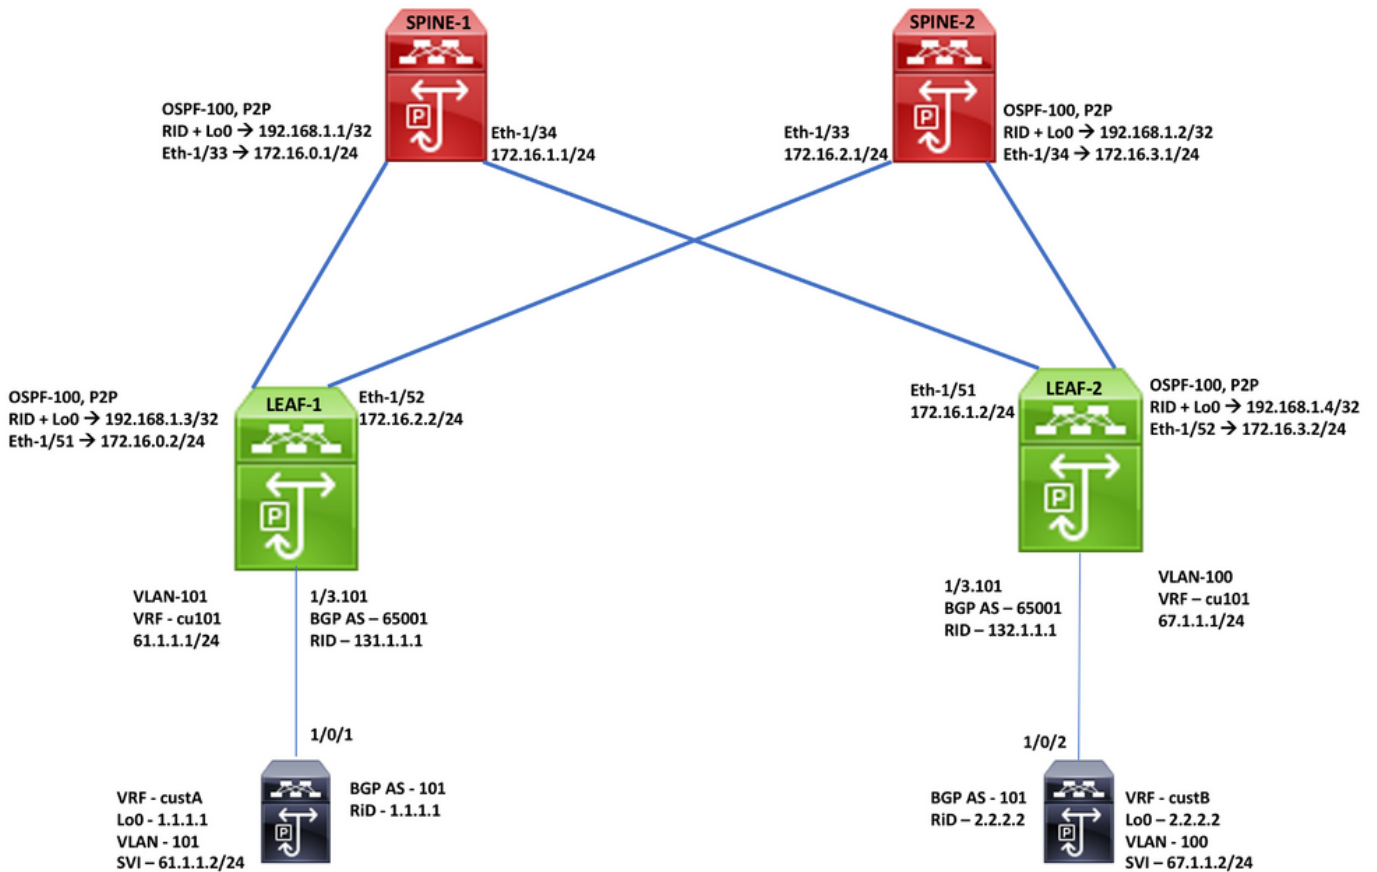
主構建塊：RD、RT和VRF

用於傳輸的底層層：IGP/BGP-LU和SR-TE

服務的覆蓋層：EVPN

設定

拓撲



高級配置

1. 安裝功能
2. 配置IP地址 — 襯底
3. 配置IGP - OSPF
4. 設定MP - BGP
5. 配置VLAN和EVPN重疊
6. 在主機和枝葉之間配置e-BGP

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
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	template peer Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended route-reflector-client
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	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
	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
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
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 Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended route-reflector-client
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	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
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	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 template peer cu1 address-family ipv4 unicast as-override send-community 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	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
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	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
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	

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

驗證

使用本節內容，確認您的組態是否正常運作。

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, (0x0000000100002) 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 | 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
```

```
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)
 0 67.1.1.1 0 msec 8 msec 0 msec
 1 172.16.3.1 0 msec 0 msec 0 msec
 2 172.16.0.2 0 msec
 3 172.16.2.2 0 msec
 4 172.16.0.2 8 msec
 5 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)
 0 61.1.1.1 0 msec 17 msec 0 msec
 1 172.16.2.1 17 msec
 2 172.16.0.1 0 msec
 3 172.16.2.1 9 msec
 4 172.16.3.2 0 msec
 5 172.16.1.2 0 msec
 6 172.16.3.2 17 msec
 7 67.1.1.2 8 msec * 0 msec
endhost#
```

疑難排解

目前尚無適用於此組態的具體疑難排解資訊。

相關資訊

- [多重通訊協定BGP MPLS VPN](#)
- [Cisco Nexus 9500、9300、9200、3200和3100平台交換機上的分段路由白皮書](#)
- [在分段路由MPLS上配置第3層EVPN和第3層VPN](#)
- [技術支援與文件 - Cisco Systems](#)