在Nexus 9300中通过分段路由MPLS部署第3层EVPN

目录

简介

本文档介绍如何在Nexus 9300产品上部署/配置第3层(L3)以太网VPN(EVPN)网段路由(SR)多协议标签交换(MPLS)[开放最短路径优先(OSPF)/内部边界网关协议(iBGP)]。

先决条件

要求

Cisco 建议您了解以下主题：

- 边界网关协议 (BGP)
- L3VPN
- EVPN
- SR

使用的组件

本文档中的信息基于以下软件和硬件版本：

- 主干硬件 — 运行版本9.3.(3)的93360YC-FX2
- 枝叶硬件 — 93240YC-FX2，运行版本9.3.(3)
- 客户端 — 93216TC-FX2

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

背景信息

MPLS L3VPN重述

VPN是：

- 基于IP的网络，通过公共基础设施提供专用网络服务。
- 允许通过Internet或其他公共或专用网络彼此私下通信的一组站点。

传统VPN是通过为VPN中的所有站点配置全网状隧道或永久虚电路(PVC)而创建的。这种类型的VPN不易维护或扩展，因为添加新站点需要更改VPN中的每个边缘设备。

基于MPLS的VPN在L3中创建，并基于对等体模型。对等模式使服务提供商和客户能够交换第3层路由信息。服务提供商在客户地点之间中继数据，无需客户参与。

MPLS VPN比传统VPN更易于管理和扩展。将新站点添加到MPLS VPN时，只需更新为客户站点提供服务的服务提供商的边缘路由器。

以下是MPLS VPN的组件：

- 提供商(P)路由器 — 提供商网络核心中的路由器。P路由器运行MPLS交换，不将VPN标签附加到路由的数据包。VPN标签用于将数据包转发到正确的专用网络或客户边缘路由器。
- PE路由器 — 根据接收VPN标签的接口或子接口将VPN标签附加到传入数据包，并附加MPLS核心标签的路由器。PE路由器直接连接到CE路由器。
- 客户(C)路由器 — Internet服务提供商(ISP)或企业网络中的路由器。
- 客户边缘(CE)路由器 — ISP网络上连接到网络上PE路由器的边缘路由器。CE路由器必须与PE路由器进行接口。

带L3VPN的EVPN概述(MPLS SR)

数据中心(DC)部署采用VXLAN EVPN或MPLS EVPN，其优势包括EVPN控制平面学习、多租户、无缝移动、冗余和更简单的POD添加。同样，CORE是基于标签分发协议(LDP)的MPLS L3VPN网络，或从传统的基于MPLS L3VPN LDP的底层过渡到更复杂的解决方案（如SR）。

SR的优势包括：

- 统一IGP和MPLS控制平面
- 更简单的流量工程方法
- 更轻松的配置
- 软件定义网络(SDN)采用

EVPN(RFC 7432)是基于BGP MPLS的解决方案，已用于虚拟化数据中心网络中的下一代以太网服务。它使用多个构建块，例如来自现有MPLS技术的路由识别器(RD)、路由目标(RT)和虚拟路由和转发(VRF)。

NXOS 7.0(3)I6(1)版本中引入的基于SR的L3 EVPN使用带MPLS封装的EVPN第5类路由。它为演进的数据中心服务提供多租户、可扩展性和高性能。

注意：在DC中，数据平面可以是VXLAN或MPLS。

传统MPLS L3 VPN **基于SR的MPLS L3 VPN**

主要构建块：RD、RT和VRF　　主要构建块：RD、RT和VRF
传输的底层：IGP、LDP和RSVP-TE　　传输的底层：IGP/BGP-LU和SR-TE
服务的重叠层：VPNv4和VPNv6　　服务的重叠层：EVPN

配置

拓扑

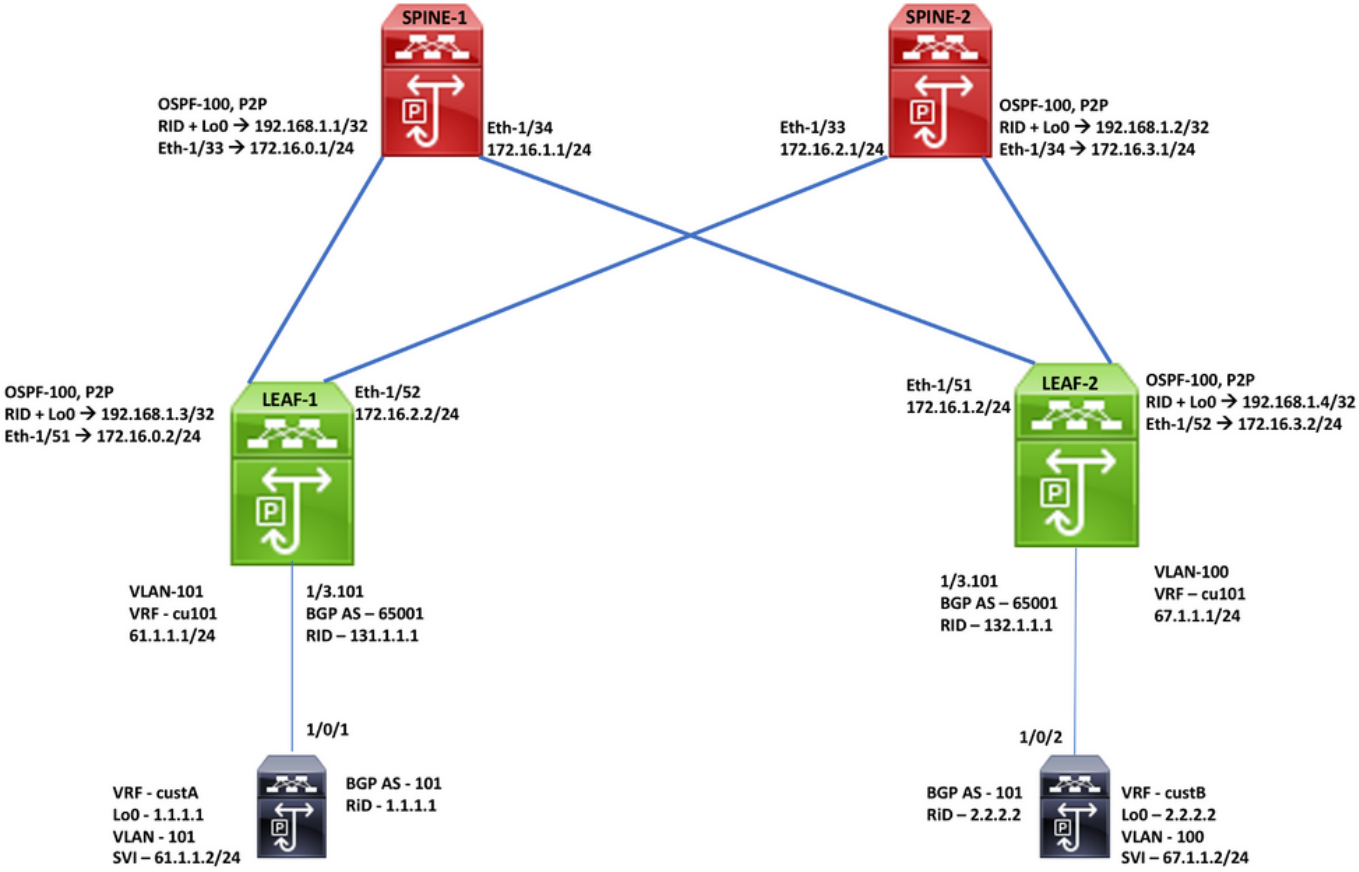

高级配置

1. 安装功能
2. 配置IP地址 — 底层
3. 配置IGP - OSPF
4. 配置MP - BGP
5. 配置VLAN和EVPN重叠
6. 在主机和枝叶之间配置e-BGP

SPINE-1 Configuration

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

```
feature-set mpls
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam


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


route-map label-index-spine1 permit 10
 set label-index 211
```

OSPF Configuration

```
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

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


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
```

BGP/EVPN Configuration

```
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
```

SPINE-2 Configuration

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

```
feature-set mpls
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam

mpls label range 5000 450000

segment-routing
 mpls
  global-block 16000 25000
  connected-prefix-sid-map
   address-family ipv4
    192.168.1.2/32 index 221


route-map label-index-spine2 permit 10
 set label-index 221
```

OSPF Configuration

```
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

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


interface loopback0
 ip address 192.168.1.2/32
 ip router ospf 100 area 0.0.0.0


router ospf 100
 segment-routing mpls
 router-id 192.168.1.2
```

BGP/EVPN Configuration

```
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
```

LEAF-1 Configuration

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

```
feature-set mpls
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam


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

route-map label-index-leaf-1 permit 10
 set label-index 311


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 Ethernet1/3.101
 encapsulation dot1q 101
 vrf member cu101
 ip address 61.1.1.1/24
 no shutdown

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
```

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
 template peer cu1
  address-family ipv4 unicast
   as-override
   send-community
   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
  neighbor 61.1.1.2
   inherit peer cu1
   remote-as 101
```

LEAF-2 Configuration

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

```
feature-set mpls
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam


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

route-map label-index-Leaf2 permit 10
 set label-index 321


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 Ethernet1/3.101
 encapsulation dot1q 100
 vrf member cu101
 ip address 67.1.1.1/24
 no shutdown

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
```

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
 template peer cu1
  address-family ipv4 unicast
   as-override
   send-community
   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 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

```
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 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, SVI Configuration

```
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


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
 !
```

BGP Configuration

```
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
```

验证

使用本部分可确认配置能否正常运行。

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 6
Paths: (2 available, best #1)
Flags: (0x880c001a) (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, VPN 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

  VPN 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      PUSH 16311 492288
                  172.16.3.1                                Ethernet1/52      PUSH 16311 492288
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 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 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      PUSH 16311
                  172.16.3.1                                Ethernet1/52      PUSH 16311
```

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#
```

故障排除

目前没有针对此配置的故障排除信息。

相关信息

- 多协议BGP MPLS VPN
- Cisco Nexus 9500、9300、9200、3200和3100平台交换机上的分段路由白皮书
- 配置第3层EVPN和第3层VPN over Segment Routing MPLS
- 技术支持和文档 - Cisco Systems