在Catalyst 9000交换机上配置服务VRF EVPN VxLAN

目录

简介

本文档介绍不同场景下EVPN （以太网VPN） VXLAN （虚拟可扩展LAN）的路由泄漏配置。

先决条件

建议您熟悉单播EVPN VxLAN功能BGP。

要求

本指南假设BGP、NVE对等体已经正确。如果出现基本EVPN VxLAN启动问题（单播ping故障、BGP、NVE对等体关闭等），请根据需要参考BGP、EVPN、路由/交换机故障排除指南。

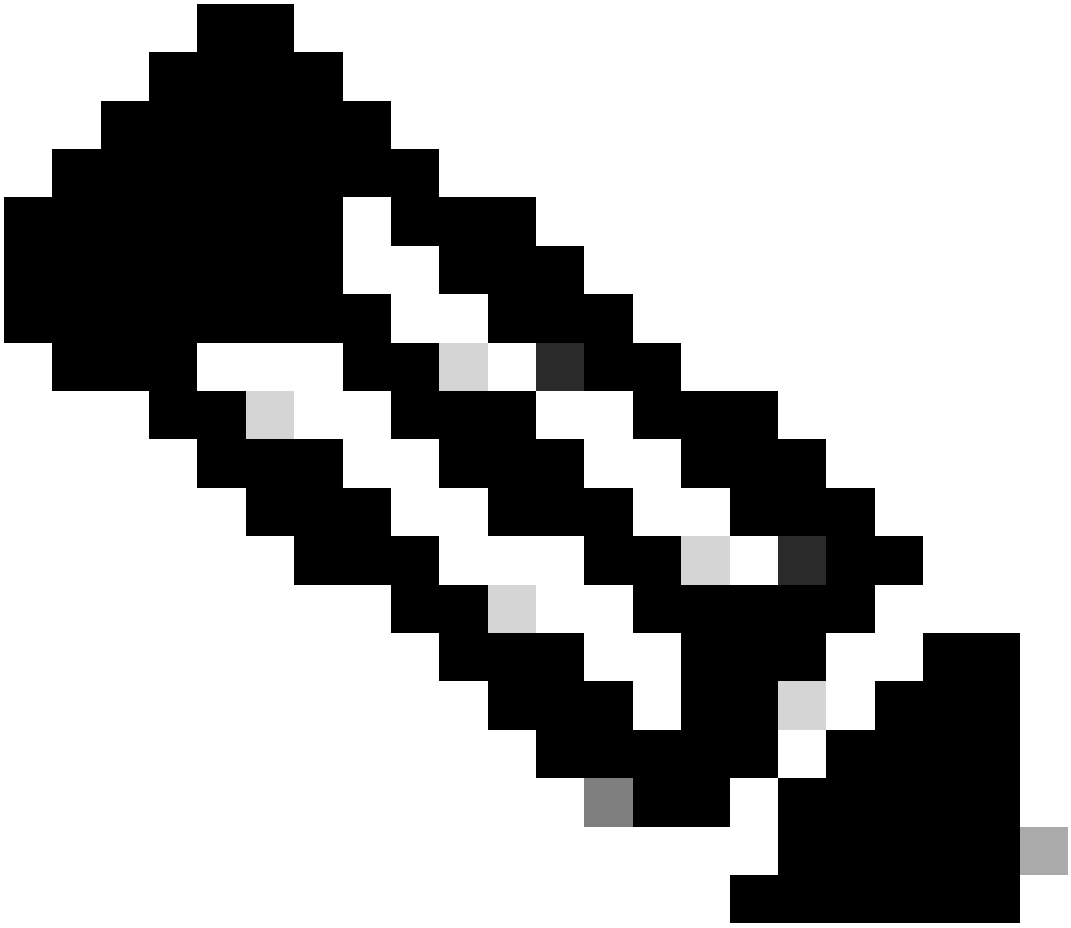

注意：仅IPv4支持服务VRF配置示例。

使用的组件

本文档中的信息基于以下软件和硬件版本：

- C9300
- C9400
- C9500
- C9600

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。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

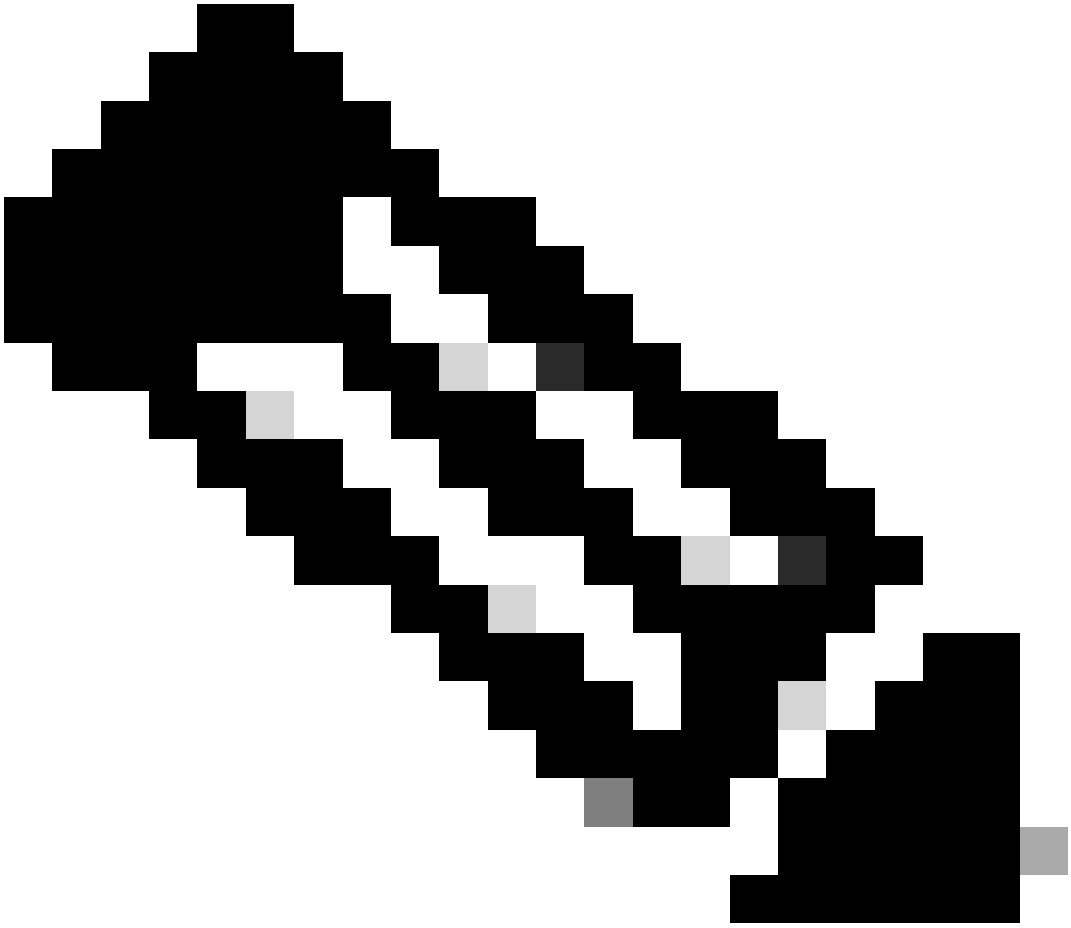

注意：有关在其他思科平台上启用这些功能的命令，请参阅相应的配置指南。

配置

在构建“共享VRF”服务或边界节点与防火墙的连接时，路由泄漏功能被广泛使用。通常，边界枝叶是配置路由泄漏的节点。

- Cisco IOS® XE上的EVPN/VXLAN的VRF之间的路由泄漏在BGP级别不照常执行。而应使用EVN（简易虚拟网络）功能。

网络图

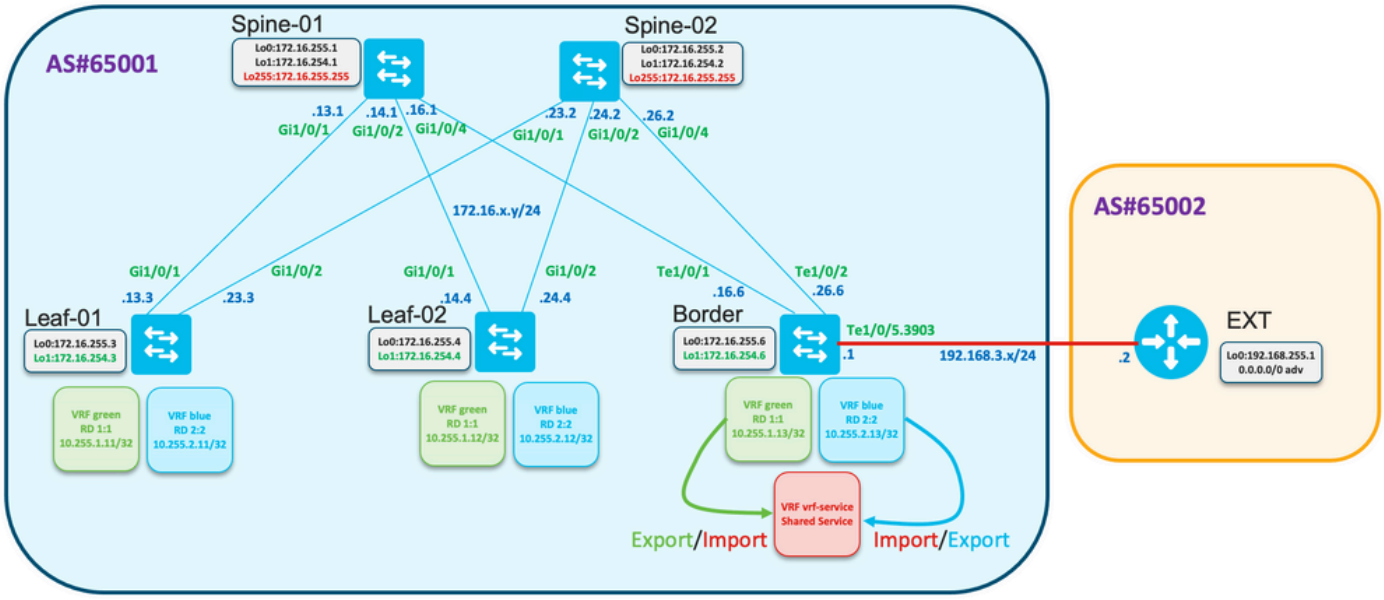

通用路由泄漏

在本示例中，计划在边界节点上配置从VRF“绿色”和“蓝色”到VRF“vrf服务”的路由渗透。

检查路由表中是否有边界上的VRF“绿色”和“蓝色”：

<#root>

Border#

show ip route vrf green

<…snip…>

```
      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B        10.1.1.0/24 [200/0] via 172.16.254.3, 01:19:43, Vlan901
B        10.1.2.0/24 [200/0] via 172.16.254.3, 01:19:43, Vlan901
B        10.255.1.11/32 [200/0] via 172.16.254.3, 01:19:43, Vlan901
B        10.255.1.12/32 [200/0] via 172.16.254.4, 01:19:43, Vlan901
C        10.255.1.13/32 is directly connected, Loopback11
```

Border#

show ip route vrf blue

<…snip…>

```
      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B        10.2.1.0/24 [200/0] via 172.16.254.3, 01:20:28, Vlan902
B        10.2.2.0/24 [200/0] via 172.16.254.3, 01:20:28, Vlan902
B        10.255.2.11/32 [200/0] via 172.16.254.3, 01:20:28, Vlan902
B        10.255.2.12/32 [200/0] via 172.16.254.4, 01:20:28, Vlan902
C        10.255.2.13/32 is directly connected, Loopback12
```

将从VRF“绿色”导入到VRF“vrf服务”的所有路由的配置

<#root>

```
vrf definition vrf-service
 rd 3:3
 !
 address-family ipv4

  route-replicate from vrf green unicast all

  route-target export 3:3
  route-target import 3:3
 exit-address-family
```

验证边界上VRF“vrf-service”的路由表是否包含来自VRF“green”的路由

<#root>

```
Border#

show ip route vrf vrf-service


Routing Table: vrf-service
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks

B   +    10.1.1.0/24 [200/0] via 172.16.254.3, 00:00:42, Vlan901
B   +    10.1.1.11/32 [200/0] via 172.16.254.3, 00:00:28, Vlan901
B   +    10.255.1.11/32 [200/0] via 172.16.254.3, 01:32:49, Vlan901
B   +    10.255.1.12/32 [200/0] via 172.16.254.4, 01:32:49, Vlan901

C   +    10.255.1.13/32 is directly connected, Loopback11
C        10.255.3.13/32 is directly connected, Loopback13
```

观察从VRF“绿色”的路由复制到VRF“vrf服务”，并在路由表中标记为“+”。

路由渗透（带过滤）

路由复制可以通过过滤来完成。路由映射用于此目的。

仅将前缀10.255.2.12从VRF“blue”复制到VRF“vrf-service”。

```
ip prefix-list PL-BLUE-2-VRF-SERVICE permit 10.255.2.12/32
!
route-map RM-BLUE-2-VRF-SERVICE permit 10
 match ip adderess prefix-list PL-BLUE-2-VRF-SERVICE
```

配置带过滤器的复制

<#root>

```
vrf definition vrf-service
 rd 3:3
 !
 address-family ipv4

  route-replicate from vrf green unicast all


  route-replicate from vrf blue unicast all route-map RM-BLUE-2-VRF-SERVICE


  route-target export 3:3
  route-target import 3:3
exit-address-family
```

观察VRF“vrf-service”的路由表包含源自VRF“blue”的前缀10.255.2.12/32：

<#root>

```
Border#

show ip route vrf VRF-SERVICE


<…snip…>

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
B   +    10.1.1.0/24 [200/0] via 172.16.254.3, 00:09:38, Vlan901
B   +    10.1.1.11/32 [200/0] via 172.16.254.3, 00:09:24, Vlan901
B   +    10.255.1.11/32 [200/0] via 172.16.254.3, 01:41:45, Vlan901
B   +    10.255.1.12/32 [200/0] via 172.16.254.4, 01:41:45, Vlan901
C   +    10.255.1.13/32 is directly connected, Loopback11

B   +    10.255.2.12/32 [200/0] via 172.16.254.4, 01:41:45, Vlan902 <--

C        10.255.3.13/32 is directly connected, Loopback13
```

仅默认路由通告和默认路由跟踪

连接到具有外部网络的枝叶的主机之间的连接是在边界上建立的。

- 通常，Border仅接收默认路由或默认路由加上汇总路由。
- 为了优化枝叶上的路由表，可以只通告来自边界的默认路由。

默认路由在VRF“vrf-service”/“Shared service”中接收。

- 此路由可以复制到VRF“绿色”中，但复制的路由不能重新通告。需要在BGP中为VRF“绿色”配置默认路由通告。
- 当通告VRF“绿色”中的默认路由，但VRF“vrf-service”中的默认路由不存在时，可以配置带有跟踪对象的静态路由来避免黑洞情况。

检查拓扑

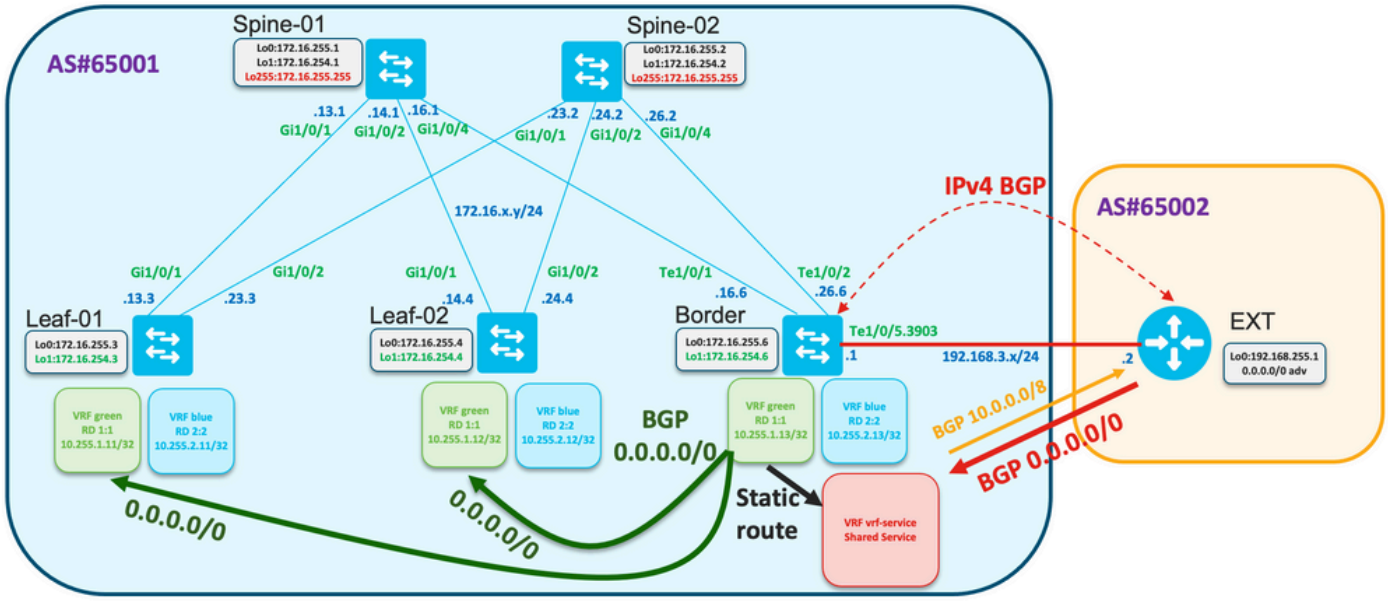

验证边界节点上是否收到默认路由：

<#root>

```
Border#
```

show ip route vrf vrf-service 0.0.0.0

```
Routing Table: red
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 65001", distance 20, metric 0, candidate default path
  Tag 65002, type external
  Last update from 192.168.3.2 00:13:32 ago
  Routing Descriptor Blocks:
  * 192.168.3.2, from 192.168.3.2, 00:13:32 ago
      opaque_ptr 0x7FA2A139FE50
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 65002
      MPLS label: none
```

<#root>

```
Border#
show ip cef vrf vrf-service 0.0.0.0/0

0.0.0.0/0
  nexthop 192.168.3.2 TenGigabitEthernet1/0/5.3903
```

跟踪1 检查VRF“vrf-service”中默认路由的可达性。

```
track 1 ip route 0.0.0.0 0.0.0.0 reachability
 ip vrf vrf-service
```

验证默认路由存在于VRF“vrf-service”中，且跟踪对象为“Up”。

```
<#root>
Border#
show track 1

Track 1
  IP route 0.0.0.0 0.0.0.0 reachability

  Reachability is Up (BGP)

    2 changes, last change 00:23:12
  VPN Routing/Forwarding table "vrf-service"
  First-hop interface is TenGigabitEthernet1/0/5.3903
  Tracked by:
    Static IP Routing 0
```

在VRF中配置默认路由“绿色”和跟踪选项

```
!
ip route vrf green 0.0.0.0 0.0.0.0 TenGigabitEthernet1/0/5.3903 192.168.3.2 track 1
!
```

```
<#root>
Border#
show ip route vrf green 0.0.0.0

Routing Table: green
Routing entry for 0.0.0.0/0, supernet
  Known via "static", distance 1, metric 0, candidate default path
  Redistributing via bgp 65001
```

```
  Advertised by bgp 65001
  Routing Descriptor Blocks:
  * 192.168.3.2, via TenGigabitEthernet1/0/5.3903
      Route metric is 0, traffic share count is 1
```

在VRF“绿色”的BGP进程下配置默认路由通告

<#root>

```
router bgp 65001
!
<...snip...>
 !
 address-family ipv4 vrf green
  advertise l2vpn evpn
  redistribute static
  redistribute connected

  default-information originate

 exit-address-family
!
<...snip...>
```

验证默认路由是否像路由类型5一样通告到L2VPN EVPN AF并在交换矩阵上传播

<#root>

```
Border#

show bgp l2vpn evpn rd 1:1 route-type 5 0 0.0.0.0 0

BGP routing table entry for [5][1:1][0][0][0.0.0.0]/17, version 622
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local, imported path from base
    192.168.3.2 (via vrf red) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, external, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, local vtep: 172.16.254.6, VNI Label 5090
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
      rx pathid: 0, tx pathid: 0x0
      Updated on Jul 8 2022 10:41:40 UTC
```

检查枝叶-01上的EVPN、路由、cef信息

<#root>

```
Leaf-01#

show bgp l2vpn evpn rd 1:1 route-type 5 0 0.0.0.0 0
```

```
BGP routing table entry for [5][1:1][0][0][0.0.0.0]/17, version 595
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 7
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Jul 8 2022 10:41:40 UTC
  Refresh Epoch 7
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jul 8 2022 10:41:40 UTC

Leaf-01#
```

show ip route vrf green 0.0.0.0

```
Routing Table: green
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 65001", distance 200, metric 0, candidate default path, type internal
  Last update from 172.16.254.6 on Vlan901, 02:07:17 ago
  Routing Descriptor Blocks:
  * 172.16.254.6 (default), from 172.16.255.1, 02:07:17 ago, via Vlan901
      opaque_ptr 0x7FC3606F4D80
      Route metri c is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none

Leaf-01#
```

show ip cef vrf green 0.0.0.0/0

```
0.0.0.0/0
  nexthop 172.16.254.6 Vlan901
```

从交换矩阵到外部网络的反向路由来源于BGP，类似于总结路由

<#root>

!

ip route vrf vrf-service 10.0.0.0 255.0.0.0 Null0

!
router bgp 65001
<…snip…>
!
address-family ipv4 vrf vrf-service
** advertise l2vpn evpn**

```
  aggregate-address 10.0.0.0 255.0.0.0 summary-only

  redistribute static
  redistribute connected
  neighbor 192.168.3.2 remote-as 65002
  neighbor 192.168.3.2 activate
 exit-address-family
!
<…snip…>
```

检查VRF“绿色”中的枝叶-01上的路由表并ping远程IP地址192.168.255.1

<#root>

```
Leaf-01#

show ip route vrf green 192.168.255.1



Routing Table: green
% Network not in table

Leaf-01#

show ip route vrf green 0.0.0.0



Routing Table: green
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 65001", distance 200, metric 0, candidate default path, type internal
  Last update from 172.16.254.6 on Vlan901, 05:15:19 ago
  Routing Descriptor Blocks:
  * 172.16.254.6 (default), from 172.16.255.1, 05:15:19 ago, via Vlan901
      opaque_ptr 0x7FC3606F4D80
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none

Leaf-01#

show ip cef vrf green 0.0.0.0/0

0.0.0.0/0
  nexthop 172.16.254.6 Vlan901

Leaf-01#

ping vrf green 192.168.3.2 source 10.255.1.11


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.3.2, timeout is 2 seconds:
Packet sent with a source address of 10.255.1.11
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

如果VRF“vrf-service”中的边界上的默认路由丢失，则跟踪对象将关闭，VRF“绿色”中的静态路由将从RIB中删除，并且在BGP中通告的默认路由也将删除

<#root>

```
### Border ###


Border#

show ip route vrf vrf-service 0.0.0.0



Routing Table: vrf-service
% Network not in table

Border#

show track 1

Track 1
  IP route 0.0.0.0 0.0.0.0 reachability

  Reachability is Down (no ip route)     <-- Track object is down

    3 changes, last change 00:03:15
  VPN Routing/Forwarding table "vrf-service"
  First-hop interface is unknown
  Tracked by:
    Static IP Routing 0

Border#

show ip route vrf green 0.0.0.0


Routing Table: green
% Network not in table

Border#

show bgp l2vpn evpn rd 1:1 route-type 5 0 0.0.0.0 0

% Network not in table


### Leaf ###

Leaf-01#

show ip route vrf green 0.0.0.0


Routing Table: green
% Network not in table
```

必须过滤从VRF“绿色”到VRF“vrf服务”的默认路由

<#root>

```
vrf definition vrf-service
 rd 3:3
 !
 address-family ipv4

  route-replicate from vrf green unicast all route-map RM-GREEN-2-VRF-SERVICE

  route-target export 3:3
  route-target import 3:3
 exit-address-family

ip prefix-list PL-DEFAULT seq 5 permit 0.0.0.0/0
!
route-map RM-GREEN-2-VRF-SERVICE deny 10
 match ip address prefix-list PL-DEFAULT
!
route-map RM-GREEN-2-VRF-SERVICE permit 20
```

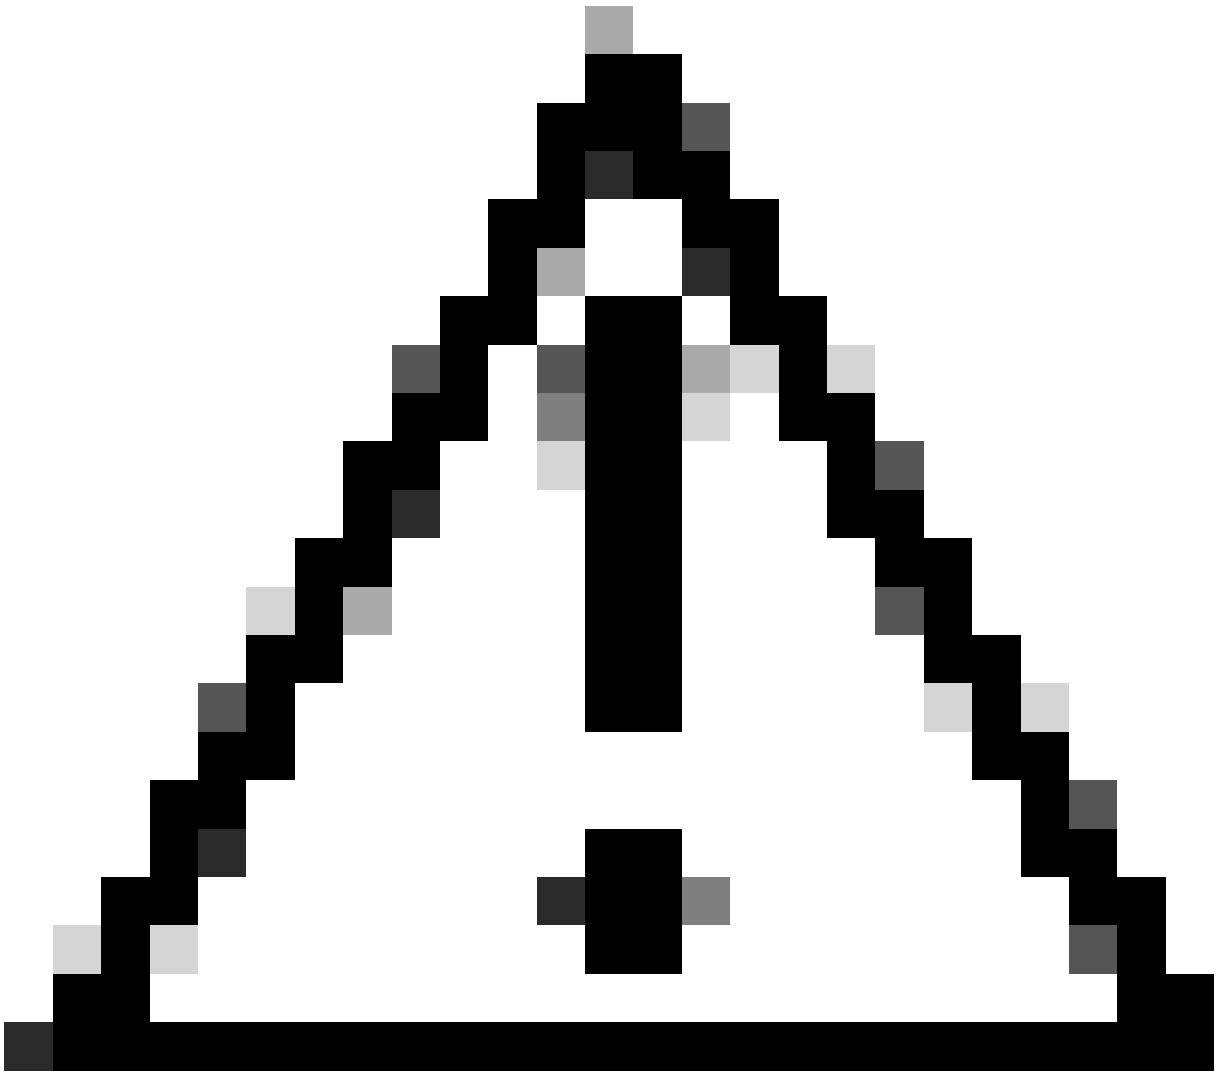

注意：由于丢失默认路由与跟踪对象关闭之间的延迟，因此静态默认路由将从VRF“绿色”复制到VRF“vrf服务”，并保持跟踪对象正常运行。因此，默认路由会通告给交换矩阵并将流量黑洞化。

仅包含冗余边界的默认路由通告

本节提供了一个使用冗余边框的示例。

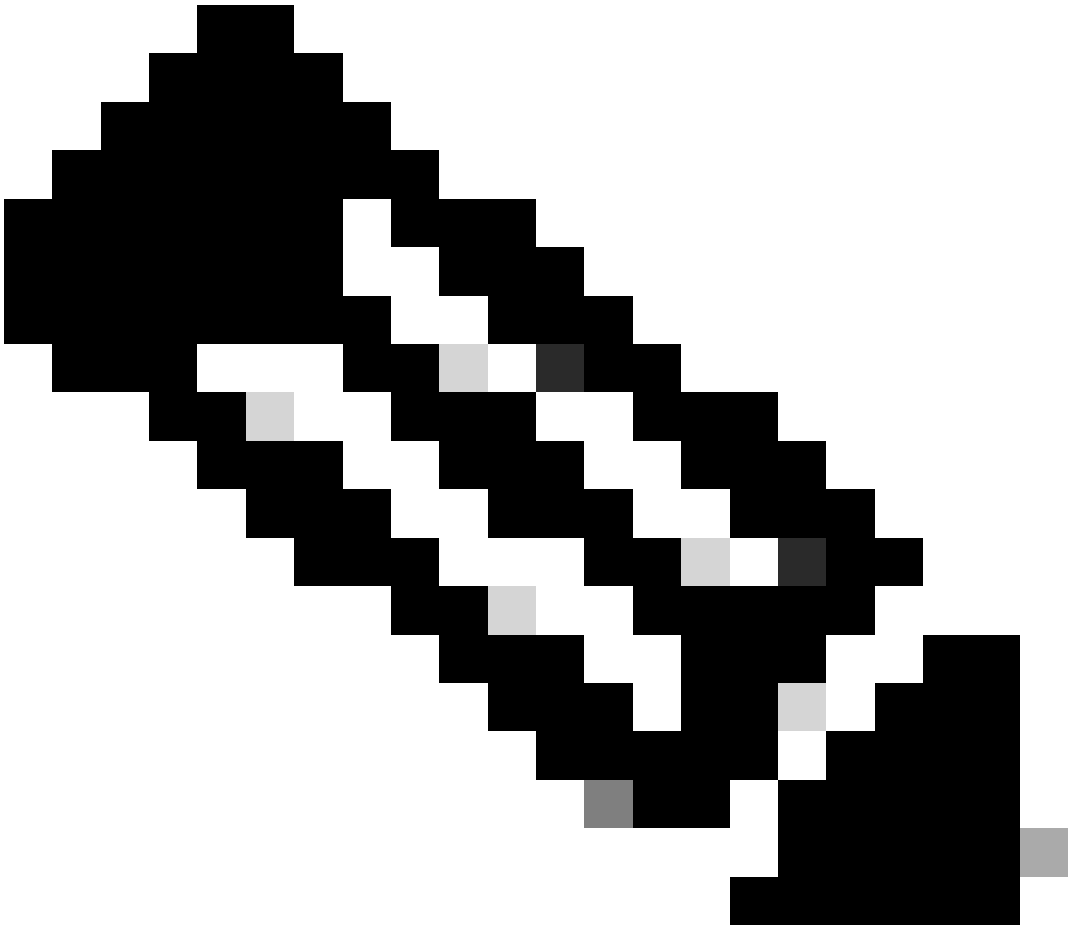

> 注意：在本示例中，我们使用了BGP附加路径功能。另一种方案是在Border-01上使用不同的RD，Border-02可用于将两个默认路由从两个边界通告到枝叶。

检查拓扑

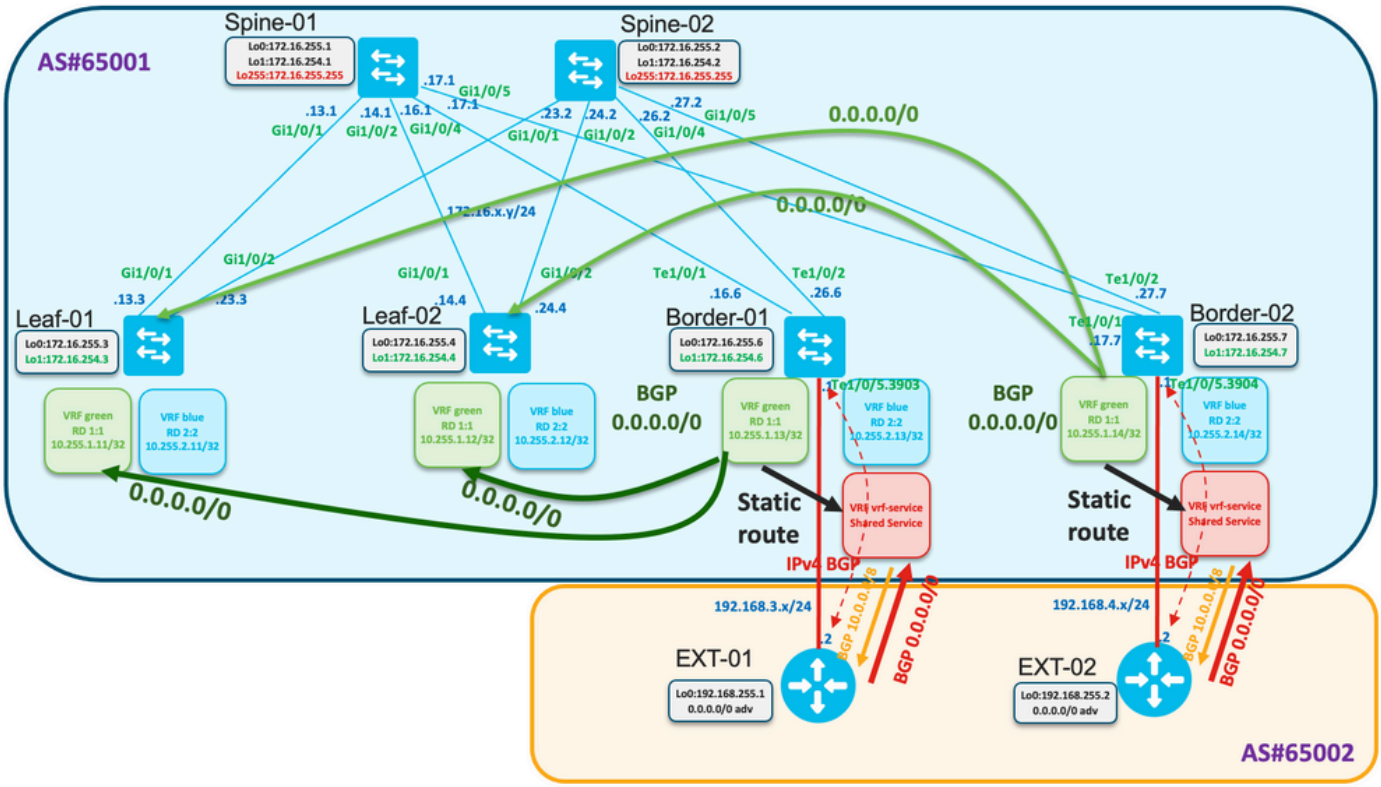

Border-01和Border-02分别接收来自EXT-01和EXT-02的默认路由。

从Border-01

<#root>

```
Border-01#
```

```
show ip route vrf vrf-service 0.0.0.0
```

```
Routing Table: vrf-service
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 65001", distance 20, metric 0, candidate default path
  Tag 65002, type external
  Last update from 192.168.3.2 00:00:06 ago
  Routing Descriptor Blocks:
  * 192.168.3.2, from 192.168.3.2, 00:00:06 ago
      opaque_ptr 0x7F68E5AC02A0
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 65002
      MPLS label: none

Border-01#
```

```
show ip cef vrf vrf-service 0.0.0.0/0
```

```
0.0.0.0/0
  nexthop 192.168.3.2 TenGigabitEthernet1/0/5.3903
```

从Border-02

<#root>

```
Border-02#
```

```
show ip route vrf vrf-service 0.0.0.0
```

```
Routing Table: vrf-service
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 65001", distance 20, metric 0, candidate default path
  Tag 65002, type external
  Last update from 192.168.4.2 01:22:08 ago
  Routing Descriptor Blocks:
  * 192.168.4.2, from 192.168.4.2, 01:22:08 ago
      opaque_ptr 0x7FE529FF3D48
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 65002
      MPLS label: none
```

```
Border-02#
```

```
show ip cef vrf vrf-service 0.0.0.0/0
```

```
0.0.0.0/0
  nexthop 192.168.4.2 TenGigabitEthernet1/0/5.3904
```

在双边界配置中使用相同的方法，如上一个示例-带跟踪的静态默认路由。

配置Border-01/02跟踪、vrf“绿色”中的默认静态路由以及用于通告的bgp配置。

<#root>

```
track 1 ip route 0.0.0.0 0.0.0.0 reachability
 ip vrf vrf-service
!
```

```
ip route vrf green 0.0.0.0 0.0.0.0 TenGigabitEthernet1/0/5.3903 192.168.3.2 track 1
```

```
!
router bgp 65001
!
<...snip...>
 !
 address-family ipv4 vrf green
  advertise l2vpn evpn
  redistribute static
  redistribute connected
```

```
  default-information originate
```

```
 exit-address-family
!
<...snip...>
```

验证主干上是否收到来自两个边界的默认路由

<#root>

```
Spine-01#

show bgp l2vpn evpn


BGP table version is 25, local router ID is 172.16.255.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1

 * ia [5][1:1][0][0][0.0.0.0]/17

                      172.16.254.7             0    100      0 ?
 *>i                   172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
<...snip...>

Spine-02#

show bgp l2vpn evpn


BGP table version is 75, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1

 * i  [5][1:1][0][0][0.0.0.0]/17

                      172.16.254.6             0    100      0 ?
 * ia                  172.16.254.7             0    100      0 ?
 *>i                   172.16.254.6             0    100      0 ?
<...snip...>
```

在主干上配置以传播两个默认路由BGP附加路径

<#root>

```
router bgp 65001
!
<...snip...>
!
 address-family l2vpn evpn
```

```
 bgp additional-paths select all best 2
  bgp additional-paths send receive

<...snip...>

  neighbor 172.16.255.3 advertise additional-paths best 2

<...snip...>

  neighbor 172.16.255.4 advertise additional-paths best 2

!
<...snip...>
```

观察此配置更改了默认的“仅最佳传播”，改为通告两个路由

<#root>

```
Spine-01#

show bgp l2vpn evpn neighbors 172.16.255.3 advertised-routes


BGP table version is 25, local router ID is 172.16.255.1
Status codes: s suppressed, d damped, h history, * valid,

> best

, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external,

a additional-path

, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1


*>i  [5][1:1][0][0][0.0.0.0]/17

                      172.16.254.6             0    100      0 ?

<-- best path

 * ia [5][1:1][0][0][0.0.0.0]/17

                      172.16.254.7             0    100      0 ?

<-- addtional path (note the a flag indicating this)

<...snip...>
```

观察枝叶上，我们看到4条BGP默认路由

<#root>

Leaf-01#

sh bgp l2vpn evpn

BGP table version is 63, local router ID is 172.16.255.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network   Next Hop   Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)

*** i [5][1:1][0][0][0.0.0.0]/17**
172.16.254.7   0   100   0 ?
*** ia   172.16.254.7   0   100   0 ?**
***>i   172.16.254.6   0   100   0 ?**
*** i   172.16.254.6   0   100   0 ?**

<...snip...>

Leaf-01#

sh bgp l2vpn evpn route-type 5 0 0.0.0.0 0

BGP routing table entry for [5][1:1][0][0][0.0.0.0]/17, version 64
Paths: (4 available, best #3, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 4
  Local
    172.16.254.7 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD68.E548
      Originator: 172.16.255.7, Cluster list: 172.16.255.2
      rx pathid: 0x1, tx pathid: 0
      Updated on Aug 24 2022 16:52:56 UTC
  Refresh Epoch 1
  Local
    172.16.254.7 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD68.E548
      Originator: 172.16.255.7, Cluster list: 172.16.255.1
      rx pathid: 0x1, tx pathid: 0
      Updated on Aug 24 2022 16:49:48 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0x0, tx pathid: 0x0
      Updated on Aug 24 2022 16:49:48 UTC
  Refresh Epoch 4
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)

```
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0x0, tx pathid: 0
      Updated on Aug 24 2022 16:52:56 UTC
```

枝叶上的配置如下所示

<#root>

```
router bgp 65001
!
<...snip...>
!
 address-family l2vpn evpn

  bgp additional-paths receive

<...snip...>
!
 address-family ipv4 vrf green

  import path selection all
  maximum-paths ibgp 2

<...snip...>
```

验证在枝叶路由表上，我们看到两条通向两个边界的路由

<#root>

```
Leaf-01#

show ip route vrf green


Routing Table: green
<...snip...>

Gateway of last resort is 172.16.254.7 to network 0.0.0.0

B*    0.0.0.0/0 [200/0] via 172.16.254.7, 00:02:15, Vlan901
                [200/0] via 172.16.254.6, 00:02:15, Vlan901
<...snip...>

Leaf-01#

show ip cef vrf green 0.0.0.0/0

0.0.0.0/0
  nexthop 172.16.254.6 Vlan901
  nexthop 172.16.254.7 Vlan901
```

观察如果从Border-01丢失默认路由会发生什么情况。

<#root>

Border-01#

show ip route vrf vrf-service 0.0.0.0

```
Routing Table: vrf-service
% Network not in table
```

跟踪关闭

<#root>

Border-01#

show track 1

```
Track 1
  IP route 0.0.0.0 0.0.0.0 reachability
  Reachability is Down (no ip route)
    5 changes, last change 00:00:56
  VPN Routing/Forwarding table "vrf-service"
  First-hop interface is unknown
  Tracked by:
    Static IP Routing 0
```

在脊柱上，我们只能看到来自Border-02的路由

<#root>

Spine-01#

show bgp l2vpn evpn

```
BGP table version is 27, local router ID is 172.16.255.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][0][0][0.0.0.0]/17
                     172.16.254.7             0    100      0 ?
 *>i                  172.16.254.7             0    100      0 ?
<...snip...>
```

在枝叶上，我们只能看到来自Border-02的路由

<#root>

```
Leaf-01#
```

show bgp l2vpn evpn

```
BGP table version is 68, local router ID is 172.16.255.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][0][0][0.0.0.0]/17
                      172.16.254.7             0    100      0 ?
 * i                   172.16.254.7             0    100      0 ?
<...snip...>

Leaf-01#
```

sh bgp l2vpn evpn route-type 5 0 0.0.0.0 0

```
BGP routing table entry for [5][1:1][0][0][0.0.0.0]/17, version 68
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.7 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD68.E548
      Originator: 172.16.255.7, Cluster list: 172.16.255.1
      rx pathid: 0x0, tx pathid: 0x0
      Updated on Aug 24 2022 17:17:31 UTC
  Refresh Epoch 4
  Local
    172.16.254.7 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN Label 0
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD68.E548
      Originator: 172.16.255.7, Cluster list: 172.16.255.2
      rx pathid: 0x0, tx pathid: 0
      Updated on Aug 24 2022 17:17:31 UTC
```

路由表中和Leaf-01上的CEF中仅存在一个路由

<#root>

```
Leaf-01#
```

show ip route vrf green

```
Routing Table: green
<...snip...>

Gateway of last resort is 172.16.254.7 to network 0.0.0.0

B*    0.0.0.0/0 [200/0] via 172.16.254.7, 00:04:02, Vlan901
<...snip...>

Leaf-01#
```

`show ip cef vrf green 0.0.0.0/0`

```
0.0.0.0/0
  nexthop 172.16.254.7 Vlan901
```

相关信息

- 技术支持和文档 - Cisco Systems
- BGP EVPN VXLAN配置指南，Cisco IOS XE Amsterdam 17.3.x（Catalyst 9500交换机）
- BGP EVPN VXLAN的功能历史记录

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。