

# Configurar o Service VRF EVPN VxLAN nos Switches Catalyst 9000

## Contents

---

[Introdução](#)

[Pré-requisitos](#)

[Requisitos](#)

[Componentes Utilizados](#)

[Configurar](#)

[Diagrama de Rede](#)

[Vazamento de rota genérica](#)

[Vazamento de rota com filtragem](#)

[Somente anúncio de rota padrão e rastreamento de rota padrão](#)

[Anúncio de rota padrão somente com bordas redundantes](#)

[Informações Relacionadas](#)

---

## Introdução

Este documento descreve a configuração de vazamento de rota para EVPN (Ethernet VPN) VXLAN (Virtual Extensible LAN) em diferentes cenários.

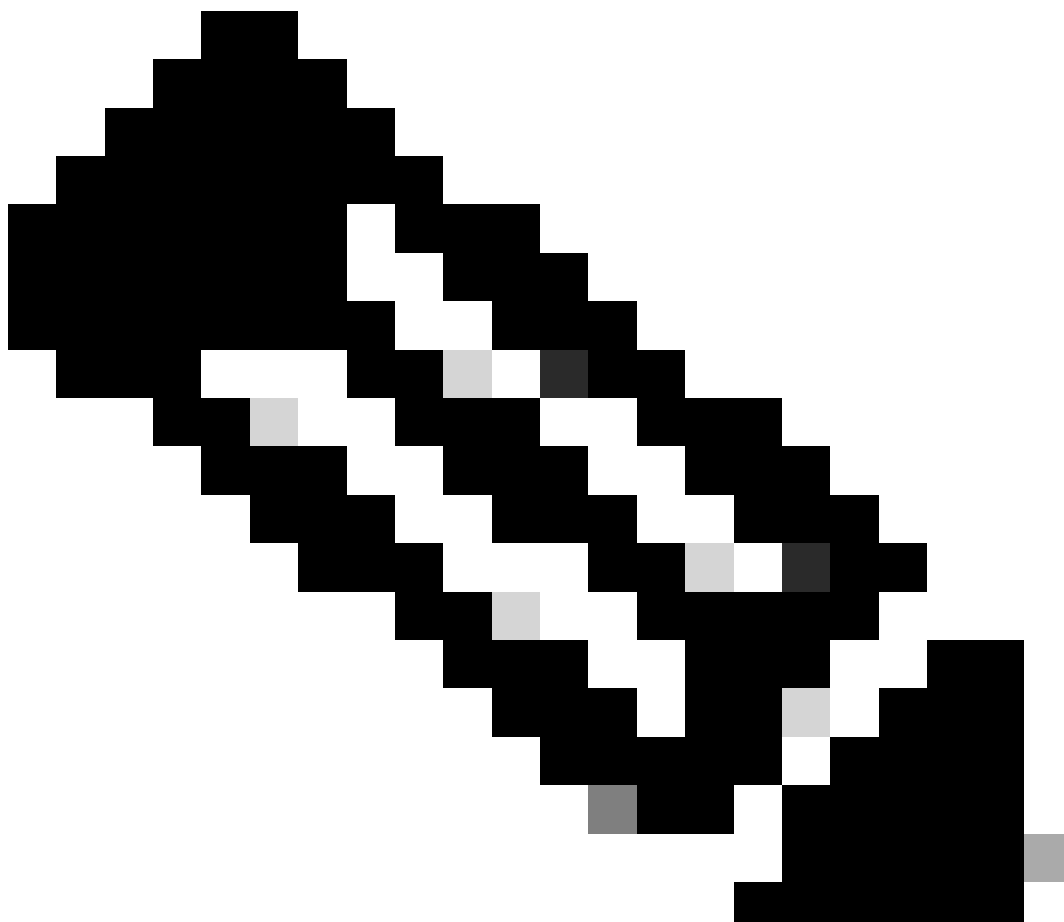
## Pré-requisitos

Recomenda-se que você esteja familiarizado com o recurso Unicast EVPN VxLAN, BGP.

## Requisitos

Este guia supõe que os peers BGP, NVE já estão corretos. Se houver problemas com a ativação básica da EVPN VxLAN (falha de ping unicast, BGP, peers NVE inativos e assim por diante) consulte os guias de solução de problemas de BGP, EVPN, rota/switch conforme necessário.

---



Observação: os exemplos de configuração de VRF de serviço são suportados apenas para IPv4.

---

## Componentes Utilizados

As informações neste documento são baseadas nestas versões de software e hardware:

- C9300
- C9400
- C9500
- C9600

As informações neste documento foram criadas a partir de dispositivos em um ambiente de laboratório específico. Todos os dispositivos utilizados neste documento foram iniciados com uma configuração (padrão) inicial. Se a rede estiver ativa, certifique-se de que você entenda o impacto potencial de qualquer comando.



Observação: consulte o guia de configuração apropriado para obter os comandos que são usados para habilitar esses recursos em outras plataformas Cisco.

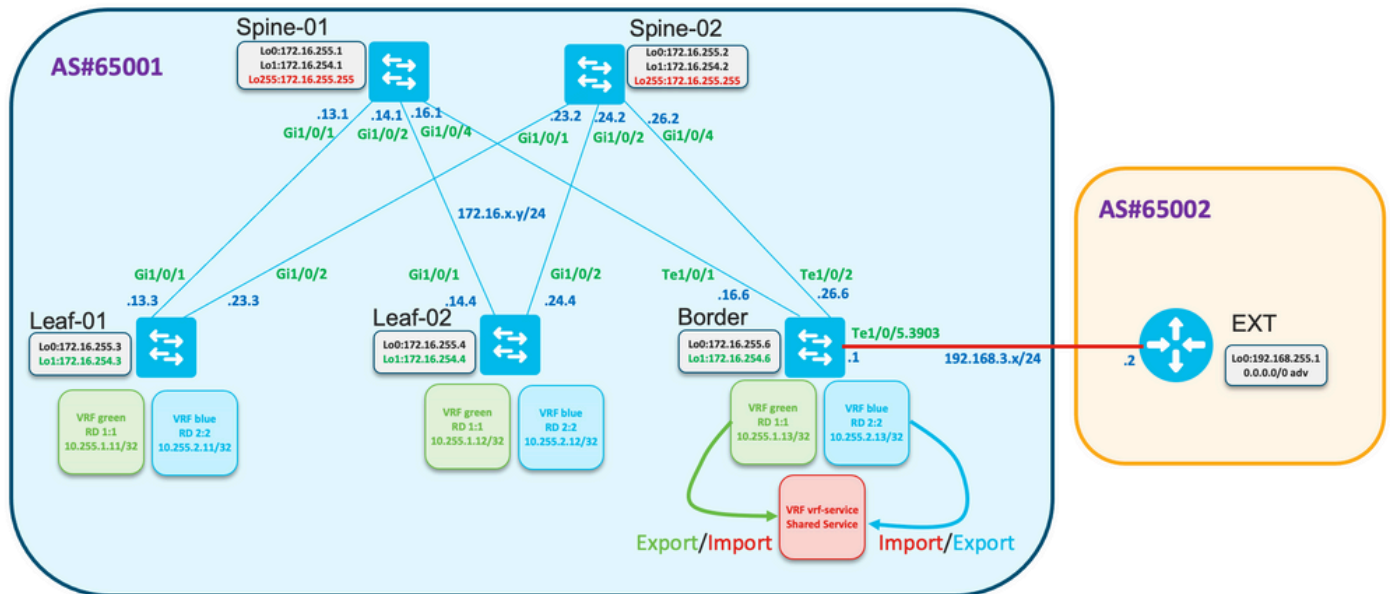
---

## Configurar

O recurso de vazamento de rota é amplamente usado no caso da criação de serviços de "VRF compartilhado" ou da conexão de nós de borda ao firewall. Geralmente, os leafs de borda são os nós onde o vazamento de rota é configurado.

- O vazamento de rota entre VRFs para EVPN/VXLAN no Cisco IOS® XE não é executado no nível de BGP como de costume. Em vez disso, é usado o recurso EVN (Easy Virtual Network).

## Diagrama de Rede



## Vazamento de rota genérica

Neste exemplo, o vazamento de rota de VRF "verde" e "azul" para VRF "vrf-service" planejado para ser configurado no nó de borda.

Verifique na tabela de roteamento se há VRF "verde" e "azul" na borda:

```
<#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

```

## Configuração para importar todas as rotas do VRF "verde" para o VRF "vrf-service"

```
<#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
```

Verificar se a tabela de roteamento do "vrf-service" do VRF na borda contém rotas do "verde" do VRF

```
<#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
```

Observe que as rotas do VRF "verde" são replicadas para o VRF "vrf-service" e marcadas na tabela de roteamento com "+".

## Vazamento de rota com filtragem

A replicação de rotas pode ser feita com filtragem. Os mapas de rotas são usados para essa finalidade.

Replicar somente o prefixo 10.255.2.12 de VRF "blue" para 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 address prefix-list PL-BLUE-2-VRF-SERVICE
```

Configurar a replicação com filtro

<#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
```

Observe que a tabela de roteamento para VRF "vrf-service" contém o prefixo 10.255.2.12/32 originado no VRF "blue":

<#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

## Somente anúncio de rota padrão e rastreamento de rota padrão

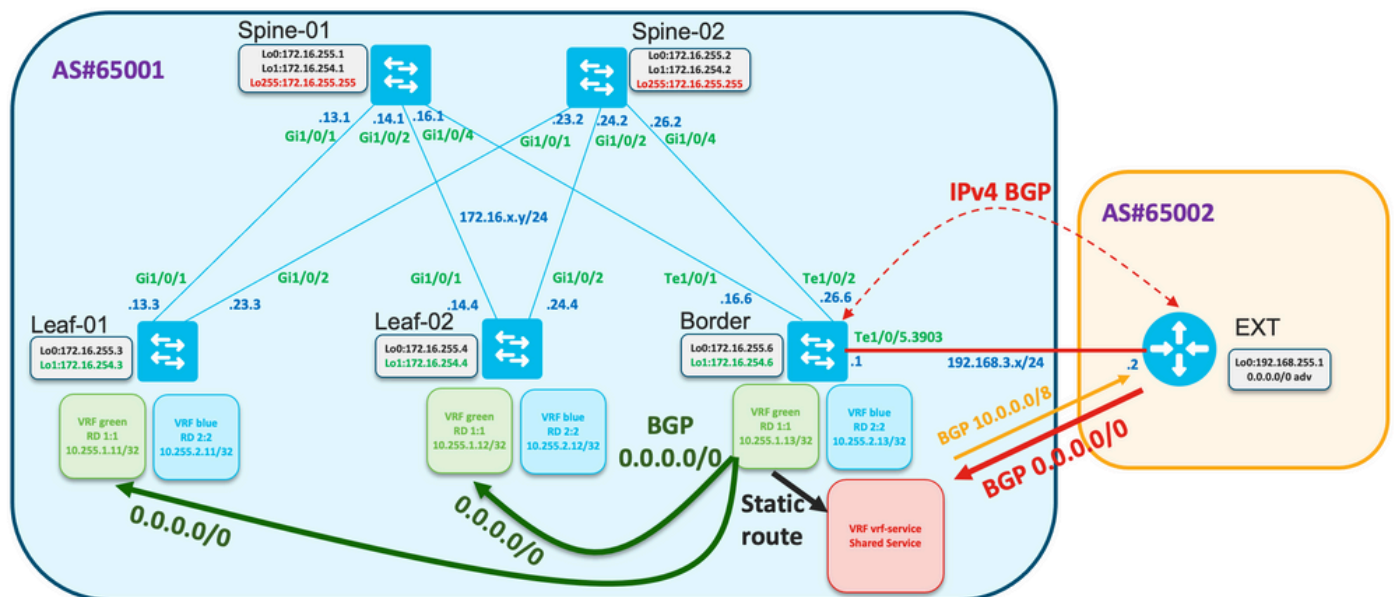
A conectividade entre hosts conectados a Leafs com uma rede externa é construída sobre a fronteira.

- Geralmente, Borda recebe apenas rota padrão ou rota padrão mais rotas de sumarização.
- Para otimizar a tabela de roteamento em Leafs, é possível anunciar apenas a rota padrão da borda.

A rota padrão é recebida em VRF "vrf-service"/"Shared service".

- Essa rota poderia ser replicada no VRF "verde", mas a rota replicada não pode ser anunciada novamente. É necessário configurar o anúncio de rota padrão no BGP para VRF "verde".
- A rota estática com objeto de trilha pode ser configurada para evitar uma situação de blackhole quando a rota padrão no VRF "verde" é anunciada, mas a rota padrão no VRF "vrf-service" não está presente.

Revisar a topologia



Verifique se a rota padrão é recebida no nó de borda:

```
<#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
```

O Track 1 verifica a acessibilidade da rota padrão no "vrf-service" do VRF.

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

Verifique se a rota padrão está presente no "vrf-service" do VRF e se o objeto de controle está "Ativo".

<#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
```

Configure a rota padrão no VRF "verde" com opção de faixa



```
!  
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
```

O anúncio de rota padrão é configurado no processo BGP para o VRF "verde"

<#root>

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

Verifique se a rota padrão é anunciada para L2VPN EVPN AF como o tipo de rota 5 e propagada pela estrutura

<#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 50901
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
```

## Verifique as informações de EVPN, roteamento e cef no Leaf-01

<#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 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

A rota inversa da estrutura para a rede externa é originada do BGP como uma rota sumarizada

```
<#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...>
```

Verifique a tabela de roteamento na folha 01 no VRF "verde" e faça ping no endereço IP remoto 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
```

Se a rota padrão for perdida na borda no "vrf-service" do VRF, o objeto de trilha será desativado, a rota estática no "verde" do VRF será removida do RIB e a rota padrão anunciada no BGP será removida

```
<#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
```

A rota padrão do VRF "verde" para o VRF "vrf-service" deve ser filtrada

```
<#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
```




Cuidado: devido ao atraso entre a perda da rota padrão e o desligamento do objeto de controle, a rota padrão estática é replicada do VRF "verde" para o VRF "vrf-service" e mantém o objeto de controle ativo. Como resultado, a rota padrão é anunciada à estrutura e o tráfego é bloqueado.

---

## Anúncio de rota padrão somente com bordas redundantes


Esta seção forneceu um exemplo de caso em que são usadas bordas redundantes.



Observação: neste exemplo, usamos o recurso de caminho adicional BGP. Outra opção é usar RD diferente em Border-01 e Border-02 poderia ser usado para anunciar AMBAS as rotas padrão de ambas as Bordas para os Leafs.

---

Revisar a topologia



Tanto Border-01 como Border-02 recebem rota padrão de EXT-01 e EXT-02, respectivamente.

Da Borda-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
```



## Da Borda-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
```

A mesma abordagem é usada na configuração de borda dupla, como no exemplo anterior - rota padrão estática com rastreamento.

Configure a faixa Border-01/02, rota estática para o padrão em vrf "verde", configuração de bgp para anúncio.

```
<#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...>
```

Verifique em Spines se as rotas padrão de ambas as Bordas foram recebidas

<#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...>

Configure em Spines para propagar AMBAS as rotas padrão, caminho adicional BGP

<#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...>

```

Observe que essa configuração altera a propagação padrão apenas melhor e, em vez disso, anuncia AMBAS as rotas

```
<#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 ?
```

```
<-- additional path (note the a flag indicating this)
```

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

Observe na folha que vemos 4 rotas padrão de BGP

<#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
```

A configuração na folha é mostrada aqui

<#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...>
```

Verifique se na tabela de roteamento Leaf vemos duas rotas para ambas as Bordas

<#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
```

Observe o que acontece em caso de perda da rota padrão de Border-01.

```
<#root>
```

```
Border-01#
```

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

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

A trilha cai

```
<#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
```

Nas lombadas, vemos a rota apenas de Border-02

```
<#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...>

Na folha vemos rota somente de Border-02

<#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

Somente uma rota está presente na tabela de roteamento e no CEF na folha 01

<#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
```

## Informações Relacionadas

- [Suporte Técnico e Documentação - Cisco Systems](#)
- [Guia de Configuração de BGP EVPN VXLAN, Cisco IOS XE Amsterdam 17.3.x \(Switches Catalyst 9500\)](#)
- [Histórico de recursos para BGP EVPN VXLAN](#)



## Sobre esta tradução

A Cisco traduziu este documento com a ajuda de tecnologias de tradução automática e humana para oferecer conteúdo de suporte aos seus usuários no seu próprio idioma, independentemente da localização.

Observe que mesmo a melhor tradução automática não será tão precisa quanto as realizadas por um tradutor profissional.

A Cisco Systems, Inc. não se responsabiliza pela precisão destas traduções e recomenda que o documento original em inglês ([link fornecido](#)) seja sempre consultado.