

Configurazione del servizio VRF VPN VxLAN sugli switch Catalyst 9000

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Configurazione](#)

[Esempio di rete](#)

[Perdite route generica](#)

[Perdite di route con filtro](#)

[Solo annuncio route predefinita e rilevamento route predefinita](#)

[Pubblicità route predefinita solo con bordi ridondanti](#)

[Informazioni correlate](#)

Introduzione

Questo documento descrive la configurazione di route-leaking per la VxLAN (Virtual Extensible LAN) VPN (Ethernet VPN) in diversi scenari.

Prerequisiti

Si consiglia di avere familiarità con la funzione VxLAN di VPN unicast, BGP.

Requisiti

In questa guida si presume che i peer BGP e NVE siano già corretti. In caso di problemi con la VxLAN di base dell'EVPN (errore ping unicast, BGP, peer NVE inattivi e così via), fare riferimento alle guide alla risoluzione dei problemi di BGP, EVPN, route/switch, se necessario.



Nota: gli esempi di configurazione VRF del servizio sono supportati solo per IPv4.

Componenti usati

Le informazioni fornite in questo documento si basano sulle seguenti versioni software e hardware:

- C9300
- C9400
- C9500
- C9600

Le informazioni discusse in questo documento fanno riferimento a dispositivi usati in uno specifico ambiente di emulazione. Su tutti i dispositivi menzionati nel documento la configurazione è stata ripristinata ai valori predefiniti. Se la rete è operativa, valutare attentamente eventuali conseguenze derivanti dall'uso dei comandi.



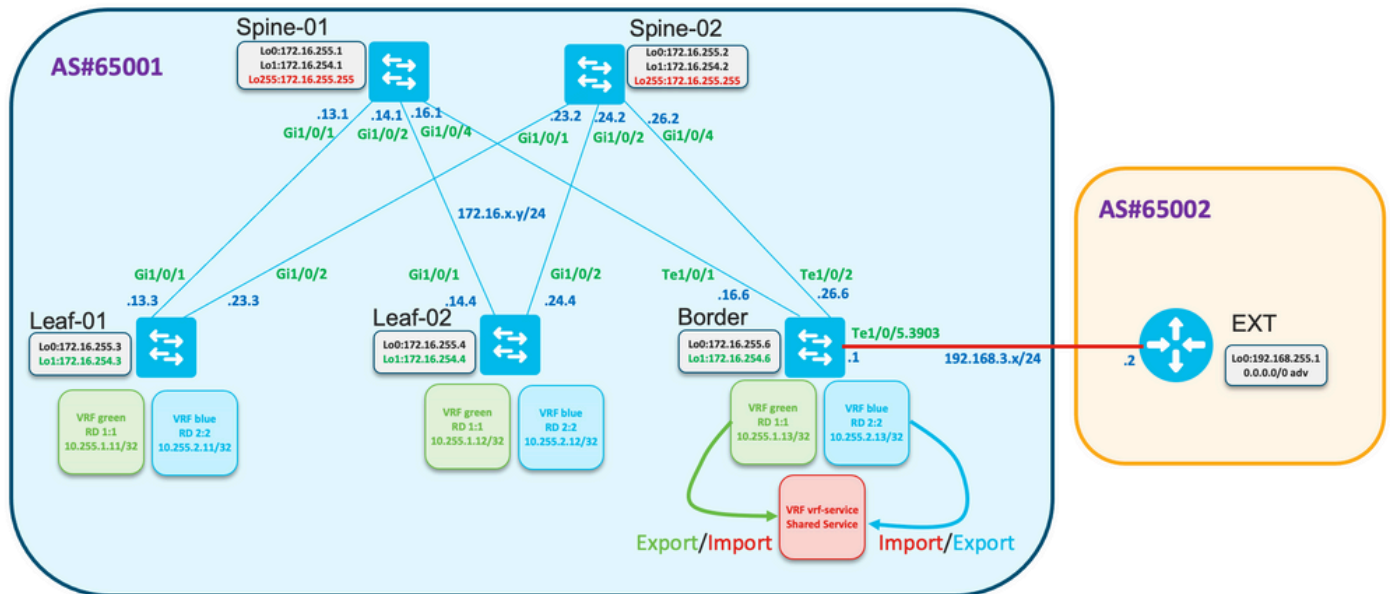
Nota: per i comandi che vengono usati per abilitare queste funzionalità su altre piattaforme Cisco, consultare la guida alla configurazione appropriata.

Configurazione

La funzione di perdita di percorso è ampiamente utilizzata nel caso della creazione di servizi "VRF condivisi" o della connessione di nodi di confine al firewall. In genere, i fogli di stile Bordo sono i nodi in cui è configurata la perdita di route.

- La route-leaking tra VRF per VPN/VXLAN su Cisco IOS® XE non viene eseguita al livello BGP come di consueto. Viene invece utilizzata la funzionalità EVN (Easy Virtual Network).

Esempio di rete



Perdite route generica

Nell'esempio, il routing della perdita da "verde" e "blu" VRF a "vrf-service" VRF dovrebbe essere configurato sul nodo Border.

Controllare se la tabella di routing contiene VRF "verde" e "blu" sul bordo:

```
<#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

```

Configurazione per importare tutti i percorsi dal VRF "verde" al 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
```

Verificare che la tabella di routing del VRF "vrf-service" sul bordo contenga percorsi da VRF "green"

```
<#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
```

Notare che i percorsi da VRF "green" vengono replicati nel VRF "vrf-service" e contrassegnati nella tabella di routing con "+".

Perdite di route con filtro

La replica route può essere eseguita con il filtro. A tale scopo vengono utilizzate le route map.

Replicare solo il prefisso 10.255.2.12 da "blu" VRF a "vrf-service" VRF.

```
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
```

Configura replica con 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
```

Osservare che la tabella di routing per il VRF "vrf-service" contiene il prefisso 10.255.2.12/32 originario del VRF "blu":

<#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

Solo annuncio route predefinita e rilevamento route predefinita

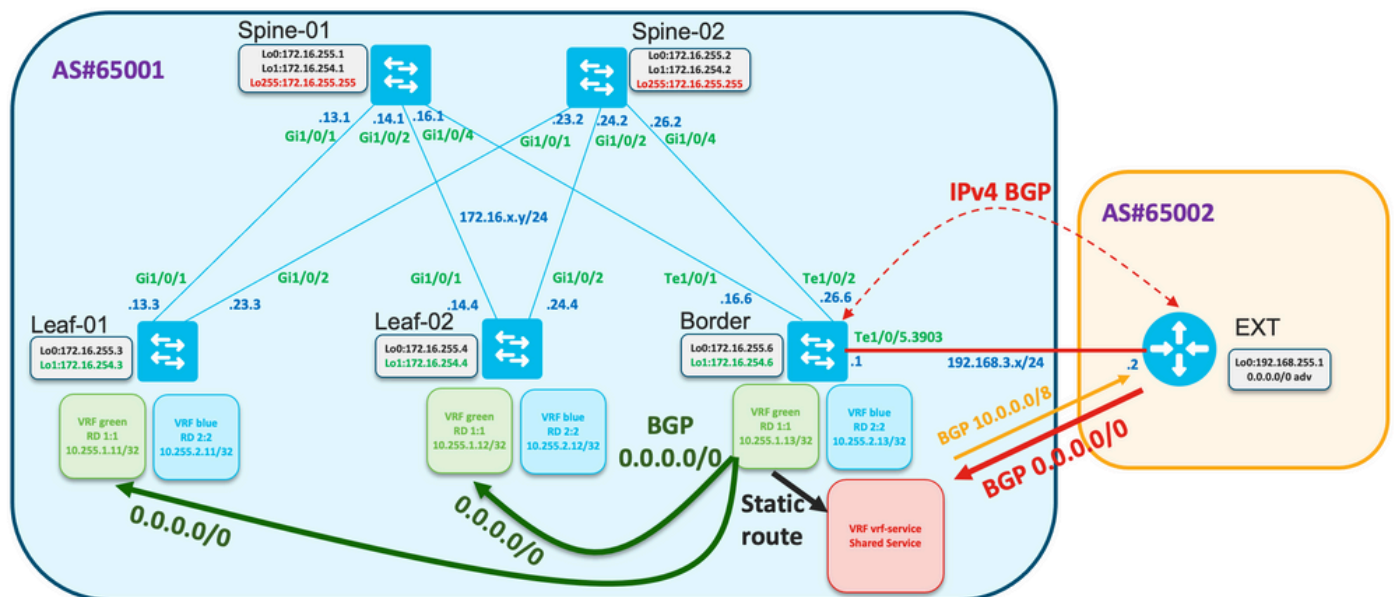
La connettività tra gli host collegati a Leafs con una rete esterna è costruita sopra il confine.

- In genere, il bordo riceve solo la route predefinita o la route predefinita più le route di riepilogo.
- Per ottimizzare la tabella di routing su foglie, è possibile annunciare solo il percorso predefinito dal bordo.

Il percorso predefinito viene ricevuto in VRF "vrf-service"/"Shared service".

- Questa route può essere replicata in "verde" VRF, ma la route replicata non può essere reannunciata. È necessario configurare l'annuncio della route predefinita in BGP per VRF "green".
- Il percorso statico con oggetto traccia può essere configurato in modo da evitare una situazione di blackhole quando viene annunciato il percorso predefinito in verde VRF, ma il percorso predefinito in servizio vrf non è presente.

Rivedere la topologia



Verificare che venga ricevuta la route predefinita sul nodo di bordo:

```
<#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
```

Il Track 1 verifica la raggiungibilità del percorso predefinito nel "vrf-service" VRF.

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

Verificare che il percorso predefinito sia presente nel VRF "vrf-service" e che l'oggetto del brano sia "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
```

Configurare il percorso predefinito nel VRF "green" con l'opzione track


```
!  
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
```

L'annuncio della route predefinita è configurato nel processo BGP per il 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...>
```

Verificare che la route predefinita sia annunciata a L2VPN EVPN AF come route-type 5 e propagata nell'infrastruttura

<#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
```

Controllare EVPN, routing, cef informazioni su 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

Il percorso inverso tra l'infrastruttura e la rete esterna ha origine da BGP come un percorso di riepilogo

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

Controllare la tabella di routing su Leaf-01 nel VRF "verde" ed eseguire il ping dell'indirizzo 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 si perde il percorso predefinito sul bordo nel VRF "vrf-service", l'oggetto di traccia scende, il percorso statico nel VRF "verde" viene rimosso dal RIB e il percorso predefinito annunciato nel BGP viene rimosso

```
<#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
```

Il percorso predefinito dal filtro "verde" VRF al filtro "vrf-service" VRF deve essere filtrato

```
<#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
```



Attenzione: a causa del ritardo tra lo smarrimento del percorso predefinito e il disinserimento dell'oggetto di traccia, il percorso predefinito statico viene replicato da "verde" VRF a "vrf-service" VRF e mantiene attivo l'oggetto di traccia. Di conseguenza, il percorso predefinito viene annunciato all'infrastruttura e il traffico viene interrotto.


Publicità route predefinita solo con bordi ridondanti

Questa sezione fornisce un esempio di caso in cui vengono utilizzati bordi ridondanti.



Nota: nell'esempio riportato viene utilizzata la funzionalità BGP additional-path. Un'altra opzione consiste nell'utilizzare un RD diverso su Border-01 e Border-02 per pubblicizzare ENTRAMBE le rotte predefinite da entrambi i Border alle foglie.

[Rivedere la topologia](#)



Sia Border-01 che Border-02 ricevono la route predefinita rispettivamente da EXT-01 e EXT-02.

Da bordo 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 bordo 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
```

Lo stesso approccio viene utilizzato nella configurazione a doppio bordo, come nell'esempio precedente - static default route with tracking.

Configurare Border-01/02 track, static route per il predefinito in vrf "green", bgp config per l'annuncio.

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

Verifica su Spine che vengano ricevute route predefinite da entrambi i bordi

<#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...>

Configura su spine per propagare ENTRAMBE le route predefinite BGP additional-path

<#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...>

```

Osservare che questa configurazione modifica la propagazione predefinita di sola lettura e annuncia invece BOTH route

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

Osservare su Leaf per individuare 4 route BGP predefinite

<#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
```

La configurazione sulla foglia è mostrata qui

<#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...>
```

Verificare che nella tabella di routing foglia siano presenti due percorsi per entrambi i bordi

<#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
```

Osservate cosa succede in caso di perdita della rotta predefinita da Border-01.

```
<#root>
```

```
Border-01#
```

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

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

Il brano si interrompe

```
<#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
```

Sugli aculei vediamo il percorso solo da 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...>

Sulla Foglia vediamo percorso solo da 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

Nella tabella di routing e nel CEF nel Leaf-01 è presente un solo percorso

<#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
```

Informazioni correlate

- [Documentazione e supporto tecnico – Cisco Systems](#)
- [Guida alla configurazione della VPN BGP, Cisco IOS XE Amsterdam 17.3.x \(switch Catalyst 9500\)](#)
- [Cronologia funzionalità per BGP VPN VXLAN](#)

Informazioni su questa traduzione

Cisco ha tradotto questo documento utilizzando una combinazione di tecnologie automatiche e umane per offrire ai nostri utenti in tutto il mondo contenuti di supporto nella propria lingua. Si noti che anche la migliore traduzione automatica non sarà mai accurata come quella fornita da un traduttore professionista. Cisco Systems, Inc. non si assume alcuna responsabilità per l'accuratezza di queste traduzioni e consiglia di consultare sempre il documento originale in inglese (disponibile al link fornito).