

Layer 3 EVPN implementeren via SR MPLS [SFP/iBGP] [PE-CE is OSPF] in Nexus 9300

Inhoud

[Inleiding](#)

[Voorwaarden](#)

[Vereisten](#)

[Gebruikte componenten](#)

[Achtergrondinformatie](#)

[MPLS L3VPN-Recap](#)

[Overzicht van EVPN met L3VPN \(MPLS SR\)](#)

[Netwerkdigram](#)

[Configuratie](#)

[Verifiëren](#)

[Gerelateerde informatie](#)

Inleiding

Dit document beschrijft hoe u Layer 3 EVPN kunt implementeren/configureren via Segment Routing (SR) Multiprotocol Label Switching (MPLS) op Nexus 9300-producten met PE-CE protocol als Open Shortest Path First (OSPF).

Voorwaarden

Vereisten

Cisco raadt kennis van de volgende onderwerpen aan:

- Border Gateway Protocol (BGP)
- Open kortste pad eerst (OSPF)
- L3VPN
- EVPN
- Segmentrouting (SR)

Gebruikte componenten

De informatie in dit document is gebaseerd op de volgende software- en hardware-versies:

- SPINE hardware-93360YC-FX2 met 9.3.(3)
- LEAF-hardware - 93240YC-FX2 actief met 9.3.(3)
- CLIENT-93216TC-FX2 (host-1), Catalyst-3750 (host-2)

De informatie in dit document is gebaseerd op de apparaten in een specifieke laboratoriumomgeving. Alle apparaten die in dit document worden beschreven, hadden een

opgeschoonde (standaard)configuratie. Als uw netwerk live is, moet u de potentiële impact van elke opdracht begrijpen.

Achtergrondinformatie

MPLS L3VPN-Recap

Een VPN is:

- Een op IP gebaseerd netwerk dat privé netwerkservices aanbiedt via een openbare infrastructuur.
- Een reeks sites die via het internet of andere openbare of particuliere netwerken met elkaar mogen communiceren.

Traditionele VPN's worden gecreëerd door de configuratie van een volledig netwerk van tunnels of permanente virtuele circuits (PVC's) naar alle locaties in een VPN. Dit type VPN is niet gemakkelijk te onderhouden of uit te breiden, omdat het toevoegen van een nieuwe site vereist dat u elk randapparaat in VPN wijzigt.

Op MPLS gebaseerde VPN's worden in Layer 3 gecreëerd en zijn gebaseerd op het peer-model. Het peer model stelt de serviceprovider en de klant in staat om Layer 3-routing informatie uit te wisselen. De dienstverlener verstrekt de gegevens tussen de locaties van de klant zonder betrokkenheid van de klant.

MPLS VPN's zijn gemakkelijker te beheren en uit te breiden dan conventionele VPN's. Wanneer een nieuwe site aan een MPLS VPN wordt toegevoegd, hoeft alleen de randrouter van de serviceprovider die services levert aan de klantensite te worden bijgewerkt.

Dit zijn de componenten van MPLS VPN:

- IP-router (Provider) in de kern van het providernetwerk. PE-routers voeren MPLS-switching uit en voegen VPN-labels niet toe aan Routed Packets. VPN-labels worden gebruikt om gegevenspakketten naar de juiste privé-netwerk- of clientrouter te sturen.
- PE router - router die het VPN-label aan inkomende pakketten vastlegt op basis van de interface of subinterface waarop ze worden ontvangen, en ook de MPLS kernetiketten aanbrengt. Een PE router hecht rechtstreeks aan een CE router.
- Router van de klant (C) - router in de Internet Service Provider (ISP) of het ondernemingsnetwerk.
- Router van Customer Edge (CE) - Edge-router op het netwerk van de ISP die verbinding maakt met de PE-router op het netwerk. Een CE router moet met een PE router een interface maken.

Overzicht van EVPN met L3VPN (MPLS SR)

Data Center (DC)-implementaties hebben VXLAN EVPN (of) MPLS EVPN goedgekeurd voor de voordelen ervan, zoals EVPN-besturingsplane learning, multi-tenancy, naadloze mobiliteit,

redundantie en makkelijke POD-toevoegingen. Op dezelfde manier is de CORE een Op Label Distribution Protocol (LDP) gebaseerd MPLS L3VPN-netwerk of een transitie van de traditionele op MPLS L3VPN LDP gebaseerde basis naar een geavanceerdere oplossing zoals Segment Routing (SR).

Segment-routing wordt gebruikt voor de voordelen ervan, zoals:

- Unified IGP- en MPLS-besturingsplanes
- Eenvoudige verkeerstechnische methoden
- Eenvoudige configuratie
- SDN-adoptie
- EVPN (RFC 7432) is een op BGP MPLS gebaseerde oplossing die is gebruikt voor Ethernet-services van de volgende generatie in een gevirtualiseerd datacenternetwerk.
- EVPN gebruikt verschillende bouwstenen zoals RD, RT, en VRF van MPLS technologieën die bestaan.
- L3 EVPN over SR die werd geïntroduceerd in NXOS 7.0(3)I6(1) release gebruikt de EVPN Type-5 route met MPLS-insluiting.
- L3 EVPN over SR biedt meerdere leveranciers, schaalbaarheid en hoge prestaties voor geavanceerde datacenterservices.

Opmerking: In DC kan het gegevensvliegtuig VXLAN of MPLS zijn.

Traditionele MPLS L3 VPN

Hoofdbouwstenen: RD, RT en VRF

Onderlaag voor transport: IGP, LDP en RSVP-TE

Overlay Layer voor service: VPNv4 en VPNv6

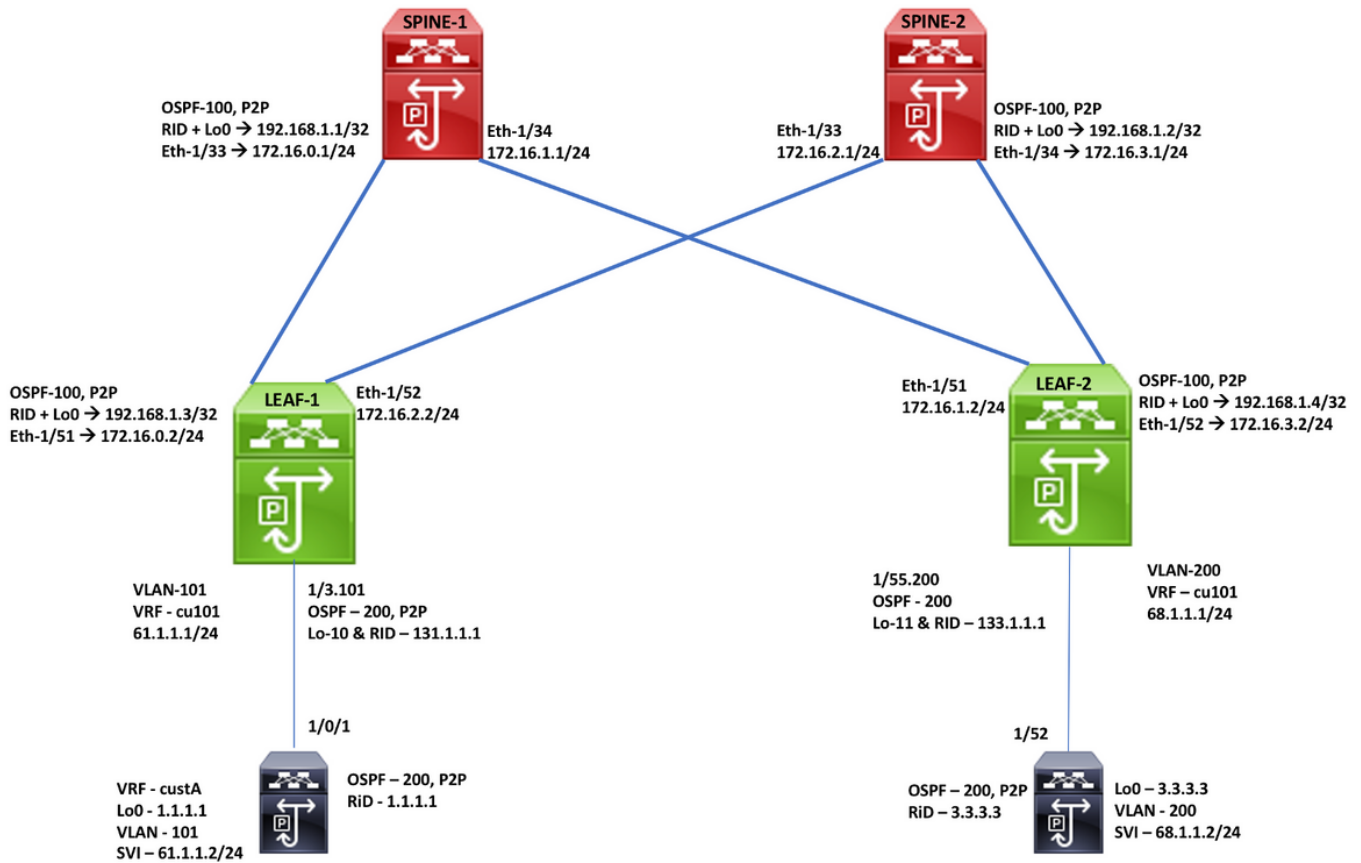
MPLS L3 VPN via SR

Hoofdbouwstenen: RD, RT en VRF

Onderlaag voor transport: IGP/BGP-LU en SR-TE

Overlay Layer voor service: EVPN

Netwerkdigram



Configuratie

SPINE-1 Configuration		
Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
feature-set mpls	interface Ethernet1/33	router bgp 65001
feature ospf	ip address 172.16.0.1/24	router-id 192.168.1.1
feature bgp	ip ospf network point-to-point	address-family ipv4 unicast
feature mpls segment-routing	ip router ospf 100 area 0.0.0.0	network 192.168.1.1/32 route-map label-index-spine1
feature mpls evpn	mpls ip forwarding	allocate-label all
feature interface-vlan	no shutdown	address-family ipv4 labeled-unicast
feature mpls oam		address-family l2vpn evpn
	interface Ethernet1/34	template peer EVPN
	ip address 172.16.1.1/24	remote-as 65001
	ip ospf network point-to-point	update-source loopback0
mpls label range 5000 450000	ip router ospf 100 area 0.0.0.0	address-family l2vpn evpn
segment-routing	mpls ip forwarding	send-community extended
mpls	no shutdown	route-reflector-client
global-block 16000 25000		encapsulation mpls
connected-prefix-sid-map	interface loopback0	template peer Labeled-unicast
address-family ipv4	ip address 192.168.1.1/32	remote-as 65001
192.168.1.1/32 index 211	ip router ospf 100 area 0.0.0.0	address-family ipv4 labeled-unicast
		send-community extended
		route-reflector-client
		next-hop-self
		soft-reconfiguration inbound always
route-map label-index-spine1 permit 10	router ospf 100	neighbor 172.16.0.2
set label-index 211	segment-routing mpls	inherit peer Labeled-unicast
	router-id 192.168.1.1	neighbor 172.16.1.2
		inherit peer Labeled-unicast
		neighbor 192.168.1.3
		inherit peer EVPN
		neighbor 192.168.1.4
		inherit peer EVPN

SPINE-2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
feature-set mpls feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam 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	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	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	OSPF Configuration	BGP/EVPN Configuration
install feature-set mpls feature-set mpls nv overlay evpn feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam feature nv overlay 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 ip prefix-list test1 seq 5 permit 61.1.1.0/24 ip prefix-list test1 seq 10 permit 131.1.1.1/32 ip prefix-list test3 seq 5 permit 1.1.1.1/32 route-map bgp65001 permit 10 match route-type internal route-map direct1 permit 10 match ip address prefix-list test1 set community 65001:10 route-map label-index-leaf-1 permit 10 set label-index 311 route-map ospf200 permit 10 match ip address prefix-list test3 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	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 PE-CE vrf cu101 address-family ipv4 unicast interface Ethernet1/3 no shutdown interface Ethernet1/3.101 encapsulation dot1q 101 vrf member cu101 ip address 61.1.1.1/24 ip ospf network point-to-point ip router ospf 200 area 0.0.0.0 no shutdown interface loopback10 vrf member cu101 ip address 131.1.1.1/32 ip router ospf 200 area 0.0.0.0 router ospf 200 vrf cu101 router-id 131.1.1.1 redistribute bgp 65001 route-map bgp65001	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 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 redistribute direct route-map direct1 redistribute ospf 200 route-map ospf200

LEAF-2 Configuration

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

```

install feature-set mpls
feature-set mpls
nv overlay evpn
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam
feature nv overlay

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

ip prefix-list new seq 5 permit 68.1.1.0/24
ip prefix-list new seq 10 permit 133.1.1.1/32

ip prefix-list new1 seq 5 permit 3.3.3.3/32

ip prefix-list redtoospf seq 5 permit 61.1.1.0/24
ip prefix-list redtoospf seq 10 permit 1.1.1.1/32

route-map bgp65001 permit 10
  match route-type internal
route-map direct1 permit 10
  match ip address prefix-list new
route-map label-index-Leaf2 permit 10
  set label-index 321
route-map ospf200 permit 10
  match ip address prefix-list new1

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 Ethernet1/51
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

PE-CE
vrf cu101
  address-family ipv4 unicast
interface Ethernet1/55
no shutdown
interface Ethernet1/55.200
encapsulation dot1q 200
vrf member cu101
ip address 68.1.1.1/24
ip ospf network point-to-point
ip router ospf 200 area 0.0.0.0
no shutdown

interface loopback11
vrf member cu101
ip address 133.1.1.1/32
ip router ospf 200 area 0.0.0.0

router ospf 200
vrf cu101
  router-id 133.1.1.1
  redistribute bgp 65001 route-map bgp65001
  
```

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

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 133.1.1.1
address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map direct1
  redistribute ospf 200 route-map ospf200
  
```

End-Host Configuration

Host-1 / Cat-3750

```
vrf definition custA
rd 101:1
!
address-family ipv4
exit-address-family
!

interface Loopback0
vrf forwarding custA
ip address 1.1.1.1 255.255.255.255

interface GigabitEthernet1/0/1
switchport trunk allowed vlan 101
switchport trunk encapsulation dot1q
switchport mode trunk
!

interface Vlan101
vrf forwarding custA
ip address 61.1.1.2 255.255.255.0
ip ospf network point-to-point
ip ospf 200 area 0.0.0.0

router ospf 200 vrf custA
router-id 1.1.1.1
network 1.1.1.1 0.0.0.0 area 0.0.0.0
network 61.1.1.0 0.0.0.255 area 0.0.0.0
```

Host-2 / N9K

```
feature ospf
feature interface-vlan

interface Ethernet1/52
switchport
switchport mode trunk
switchport trunk allowed vlan 200
no shutdown

interface Vlan200
no shutdown
ip address 68.1.1.2/24
ip ospf network point-to-point
ip router ospf 200 area 0.0.0.0

interface loopback0
ip address 3.3.3.3/32
ip router ospf 200 area 0.0.0.0

router ospf 200
router-id 3.3.3.3
```

Verifiären

Host2# show ip int brief

```
IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Vlan200        68.1.1.2        protocol-up/link-up/admin-up
Vlan1001       100.0.0.100    protocol-down/link-down/admin-up
Lo0            3.3.3.3         protocol-up/link-up/admin-up
```

Host2# show ip route

```
IP Route Table for VRF "default"
*** denotes best ucast next-hop
**** denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

1.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
3.3.3.3/32, ubest/mbest: 2/0, attached
  *via 3.3.3.3, Lo0, [0/0], 20:16:34, local
  *via 3.3.3.3, Lo0, [0/0], 20:16:34, direct
61.1.1.0/24, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
68.1.1.0/24, ubest/mbest: 1/0, attached
  *via 68.1.1.2, Vlan200, [0/0], 20:20:55, direct
68.1.1.2/32, ubest/mbest: 1/0, attached
  *via 68.1.1.2, Vlan200, [0/0], 20:20:55, local
131.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
133.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/41], 20:15:32, ospf-200, intra
```

Host2# traceroute 1.1.1.1

```
traceroute to 1.1.1.1 (1.1.1.1), 30 hops max, 40 byte packets
 1 68.1.1.1 (68.1.1.1) 0.989 ms 0.585 ms 0.407 ms
 2 172.16.3.1 (172.16.3.1) 0.886 ms 172.16.1.1 (172.16.1.1) 0.765 ms 0.731 ms
   [Label=16311 E=0 TTL=1 S=0, Label=492289 E=0 TTL=1 S=1]
   [Label=16311 E=0 TTL=1 S=0, Label=492289 E=0 TTL=1 S=1]
 3 172.16.0.2 (172.16.0.2) 0.717 ms 172.16.2.2 (172.16.2.2) 0.509 ms 172.16.0.2 (172.16.0.2) 0.678 ms
   [Label=492289 E=0 TTL=1 S=1]
   [Label=492289 E=0 TTL=1 S=1]
 4 61.1.1.2 (61.1.1.2) 2.061 ms * 1.315 ms
```

Host2# ping 1.1.1.1 source 3.3.3.3

```
PING 1.1.1.1 (1.1.1.1) from 3.3.3.3: 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=251 time=5.538 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=251 time=1.338 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=251 time=2.201 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=251 time=2.217 ms
64 bytes from 1.1.1.1: icmp_seq=4 ttl=251 time=4.021 ms

--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.338/3.063/5.538 ms
```

```
Leaf1# show ip route 3.3.3.3/32 vrf cu101
```

```
IP Route Table for VRF "cu101"
** denotes best ucast next-hop
*** denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

3.3.3.3/32, ubest/mbest: 1/0
  *via 192.168.1.4%default, [200/2], 00:44:27, bgp-65001, internal, tag 65001 (mpls-vpn)
```

```
Leaf1# 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
      |"      |192.168.1.4/32    |172.16.2.1 |Eth1/52 |16321 | SWAP
```

```
Leaf1# show forwarding 3.3.3.3/32 vrf cu101
```

```
slot 1
=====
IPv4 routes for table cu101/base
```

```
-----+-----+-----+-----+-----+
Prefix | Next-hop | Interface | Labels | Partial Install |
-----+-----+-----+-----+-----+
*3.3.3.3/32 | 172.16.0.1 | Ethernet1/51 | PUSH 16321 492288 |
           | 172.16.2.1 | Ethernet1/52 | PUSH 16321 492288 |
```

Gerelateerde informatie

- [Multiprotocol BGP MPLS VPN](#)
- [Segment-routing op Cisco Nexus 9500, 9300, 9200, 3200 en 3100 platform switchingmodule in witboek](#)
- [Layer 3 EVPN en Layer 3 VPN configureren via segment routing MPLS](#)