

Nexus L2 EVPN via segmentering routing MPLS implementeren

Inhoud

[Inleiding](#)

[Voorwaarden](#)

[Vereisten](#)

[Gebruikte componenten](#)

[Achtergrond](#)

[Configureren](#)

[Netwerkdigram](#)

[Verificatie](#)

[Referenties](#)

Inleiding

Dit document beschrijft hoe u Layer 2 EVPN kunt implementeren/configureren via Segment Routing MPLS op Cisco Nexus 9000 Series-switches.

Voorwaarden

Vereisten

Vereist kennis van BGP, OSPF, MPLS, LDP, RSVP, EVPN, Segment Routing (SR)

Gebruikte componenten

Cisco Nexus-switch 93360YC-FX2 actief met 9.3.(3)

Cisco Nexus-switch 93240YC-FX2 actief met 9.3.(3)

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.

Achtergrond

Definieert Layer 2 VPN, VPLS/L2-EVPN is een Multipoint-to-Multipoint Layer 2 VPN-service die meerdere takken van een klant verbindt in één logische switched architectuur via een IP/MPLS-netwerk.

Layer 2 EVPN-MPLS SR:

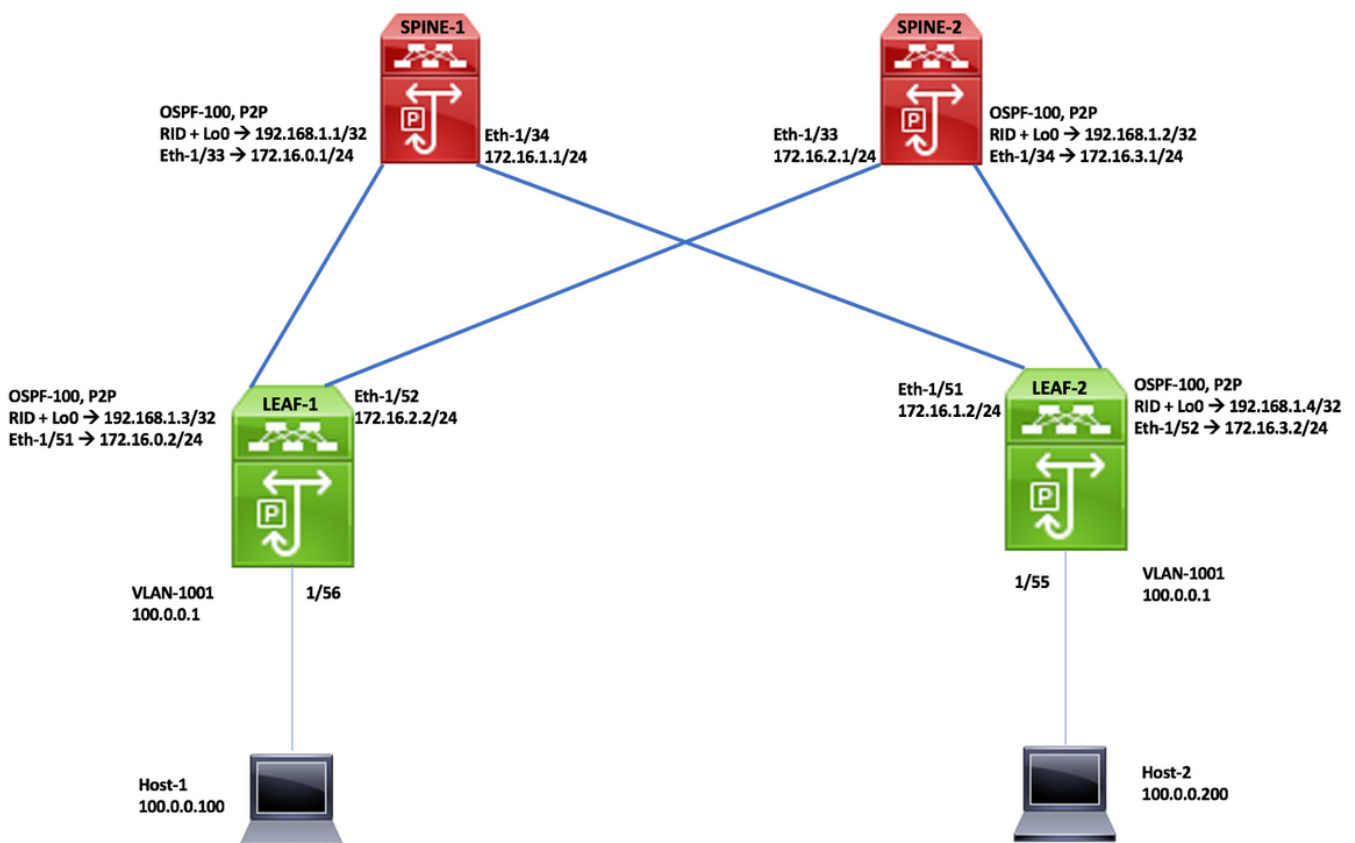
- EVPN (RFC 7432) is een op BGP MPLS gebaseerde oplossing die is gebruikt voor Ethernet-services van de volgende generatie in gevirtualiseerde datacenternetwerken
- EVPN maakt gebruik van verschillende bouwstenen zoals "RD", "RT" en "VRF" van bestaande MPLS-technologieën
- EVPN werkt in tegenstelling tot de bestaande VPLS door op het besturingsplane gebaseerde MAC-training in de kern mogelijk te maken
- In EVPN leren PE's die aan de EVPN instanties deelnemen de klant MAC-routes in besturingsplane met behulp van MP-BGP protocol
- Het leren van MAC van het besturingsplane biedt een aantal voordelen die EVPN in staat stellen om de tekortkomingen van VPLS aan te pakken, waaronder ondersteuning voor multihoming met taakverdeling per-flow
- SR L2 EVPN is nieuwe optie beschikbaar in NXOS 9.3(1) en wordt ondersteund op Nexus 9300 FX2 Series platform

Beperkingen voor L2 EVPN via SR MPLS:

- Segment-routing Layer 2 EVPN-overstroming is gebaseerd op het ingangsreplatiemechanisme
- Hiermee wordt EVPN Type 3-route voor BUM-verkeer gebruikt
- MPLS-kern ondersteunt multicast niet
- ARP-suppressie wordt niet ondersteund
- Consistentie-controle op VPC wordt niet ondersteund
- Dezelfde L2 EVI en L3 EVI kunnen niet samen worden geconfigureerd

Configureren

Netwerkdigram



Configuratiestappen op hoog niveau:

- Installatiefuncties
- IP-adres configureren
- IGP-OSPF configureren
- MP-BGP configureren
- VLAN- en EVPN-overlay configureren
- End-of-life host voor Layer 2 configureren

SPINE -1 Configuration		
Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
install feature-set mpls	interface Ethernet1/33	router bgp 65001
feature-set mpls	ip address 172.16.0.1/24	router-id 192.168.1.1
feature ospf	ip ospf network point-to-point	address-family ipv4 unicast
feature bgp	ip router ospf 100 area 0.0.0.0	network 192.168.1.1/32 route-map label-index-spine1
feature mpls segment-routing	mpls ip forwarding	allocate-label all
feature mpls evpn	no shutdown	address-family ipv4 labeled-unicast
feature interface-vlan	interface Ethernet1/34	address-family I2vpn evpn
feature mpls oam	ip address 172.16.1.1/24	template peer EVPN
mpls label range 5000 450000	ip ospf network point-to-point	remote-as 65001
segment-routing	ip router ospf 100 area 0.0.0.0	update-source loopback0
mpls	mpls ip forwarding	address-family I2vpn evpn
global-block 16000 25000	no shutdown	send-community extended
connected-prefix-sid-map	interface loopback0	route-reflector-client
address-family ipv4	ip address 192.168.1.1/32	encapsulation mpls
192.168.1.1/32 index 211	ip router ospf 100 area 0.0.0.0	template peer Labeled-unicast
route-map label-index-spine1 permit 10	router ospf 100	remote-as 65001
set label-index 211	segment-routing mpls	address-family ipv4 labeled-unicast
	router-id 192.168.1.1	send-community extended
		route-reflector-client
		next-hop-self
		soft-reconfiguration inbound always
		neighbor 172.16.0.2
		inherit peer Labeled-unicast
		neighbor 172.16.1.2
		inherit peer Labeled-unicast
		neighbor 192.168.1.3
		inherit peer EVPN
		neighbor 192.168.1.4
		inherit peer EVPN

SPINE -2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
<pre>install feature-set mpls 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</pre>	<pre>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</pre>	<pre>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</pre>

Leaf-1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
<pre>install 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 fabric forwarding anycast-gateway-mac 0000.0000.1111 mpls label range 5000 450000 vlan 1,1001 segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.3/32 index 311 vlan 1001 evi auto route-map label-index-leaf-1 permit 10 set label-index 311 vrf context Tenant-A evi 30001 interface Vlan1001 no shutdown vrf member Tenant-A ip address 100.0.0.1/24 fabric forwarding mode anycast-gateway</pre>	<pre>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 Ethernet1/56 switchport switchport mode trunk switchport trunk allowed vlan 1001 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</pre>	<pre>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 Tenant-A evpn encapsulation mpls source-interface loopback0</pre>

Leaf-2 Configuration

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

```
install 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

fabric forwarding anycast-gateway-mac 0000.0000.1111
mpls label range 5000 450000

vlan 1,1001
segment-routing
mpls
  global-block 16000 25000
  connected-prefix-sid-map
  address-family ipv4
    192.168.1.4/32 index 321
vlan 1001
  evi auto

route-map label-index-Leaf2 permit 10
  set label-index 321
vrf context Tenant-A
  evi 30001

interface Vlan1001
  no shutdown
  vrf member Tenant-A
  ip address 100.0.0.1/24
  fabric forwarding mode anycast-gateway
```

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 Ethernet1/55
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 1001
  no shutdown

interface loopback0
  ip address 192.168.1.4/32
  ip router ospf 100 area 0.0.0.0

router ospf 100
  segment-routing mpls
  router-id 192.168.1.4
```

BGP/EVPN Configuration

```
router bgp 65001
  router-id 192.168.1.4
  address-family ipv4 unicast
    network 192.168.1.4/32 route-map label-index-Leaf2
    allocate-label all
  address-family ipv4 labeled-unicast
  address-family l2vpn evpn
  template peer EVPN
    remote-as 65001
  update-source loopback0
  address-family l2vpn evpn
    send-community extended
    encapsulation mpls
  template peer Labeled-unicast
    remote-as 65001
  address-family ipv4 labeled-unicast
    send-community extended
  soft-reconfiguration inbound always

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 Tenant-A

evpn
  encapsulation mpls
  source-interface loopback0
```

Verificatie

Host1# show ip int brief

```
IP Interface Status for VRF "default"(1)
Interface IP Address Interface Status
Vlan1001 100.0.0.200 protocol-up/link-up/admin-up
```

```
Mhost1# ping 100.0.0.100
PING 100.0.0.100 (100.0.0.100): 56 data bytes
64 bytes from 100.0.0.100: icmp_seq=0 ttl=253 time=0.84 ms
64 bytes from 100.0.0.100: icmp_seq=1 ttl=253 time=0.45 ms
64 bytes from 100.0.0.100: icmp_seq=2 ttl=253 time=0.443 ms
64 bytes from 100.0.0.100: icmp_seq=3 ttl=253 time=0.438 ms
64 bytes from 100.0.0.100: icmp_seq=4 ttl=253 time=0.431 ms

--- 100.0.0.100 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.431/0.52/0.84 ms
```

Host2# show ip int brief

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

```
Mhost2# ping 100.0.0.200
PING 100.0.0.200 (100.0.0.200): 56 data bytes
64 bytes from 100.0.0.200: icmp_seq=0 ttl=253 time=0.854 ms
64 bytes from 100.0.0.200: icmp_seq=1 ttl=253 time=0.46 ms
64 bytes from 100.0.0.200: icmp_seq=2 ttl=253 time=0.451 ms
64 bytes from 100.0.0.200: icmp_seq=3 ttl=253 time=0.427 ms
64 bytes from 100.0.0.200: icmp_seq=4 ttl=253 time=0.418 ms

--- 100.0.0.200 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.418/0.522/0.854 ms
Mhost2#
```

Leaf1# show bgp l2vpn evpn

```
BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 57, Local Router ID is 192.168.1.3
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 192.168.1.3:37864 (L2VNI 1001)					
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[0]:[0.0.0.0]/216	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	32768	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272	192.168.1.3	100	32768	i	
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	32768	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	0	i	
Route Distinguisher: 192.168.1.4:37864					
* [2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[0]:[0.0.0.0]/216	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248	192.168.1.4	100	0	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	0	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	0	i	

Leaf2# show bgp l2vpn evpn

```
BGP routing table information for VRF default, address family L2VPN EVPN
BGP table version is 40, Local Router ID is 192.168.1.4
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 192.168.1.3:37864					
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272	192.168.1.3	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272	192.168.1.3	100	0	i	
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	0	i	
Route Distinguisher: 192.168.1.4:37864 (L2VNI 1001)					
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[0]:[0.0.0.0]/216	192.168.1.4	100	32768	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[0]:[0.0.0.0]/216	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248	192.168.1.4	100	32768	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248	192.168.1.4	100	0	i	
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	0	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	32768	i	

Referenties

[Segment-routing op Cisco Nexus 9500, 9300, 9200, 3200 en 3100 platform switchingmodule in witboek](#)

[Layer 2 EVPN configureren via segment routing MPLS](#)