

Implemente EVPN de Capa 3 sobre SR MPLS [Ospf / iBGP] [PE-CE es OSPF] en Nexus 9300

Contenido

[Introducción](#)

[Prerequisites](#)

[Requirements](#)

[Componentes Utilizados](#)

[Antecedentes](#)

[MPLS L3VPN Recap](#)

[Descripción General de EVPN con L3VPN \(MPLS SR\)](#)

[Diagrama de la red](#)

[Configuración](#)

[Verificación](#)

[Información Relacionada](#)

Introducción

Este documento describe cómo implementar/configurar el switching de etiquetas multiprotocolo (MPLS) EVPN sobre routing de segmentos (SR) de capa 3 en productos Nexus 9300 con el protocolo PE-CE como OSPF (Open Shortest Path First).

Prerequisites

Requirements

Cisco recomienda que tenga conocimiento sobre estos temas:

- Border Gateway Protocol (BGP)
- Abrir primero la ruta más corta (OSPF)
- L3VPN
- EVPN
- Routing de segmentos (SR)

Componentes Utilizados

La información que contiene este documento se basa en las siguientes versiones de software y hardware.

- Hardware SPINE - 93360YC-FX2 ejecutándose con 9.3.(3)
- Hardware LEAF - 93240YC-FX2 ejecutándose con 9.3.(3)
- CLIENTE: 93216TC-FX2 (host-1), Catalyst-3750 (host-2)

The information in this document was created from the devices in a specific lab environment. All of

the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Antecedentes

MPLS L3VPN Recap

Una VPN es:

- Una red basada en IP que ofrece servicios de red privada a través de una infraestructura pública.
- Un conjunto de sitios que pueden comunicarse entre sí de forma privada a través de Internet u otras redes públicas o privadas.

Las VPN convencionales se crean mediante la configuración de una malla completa de túneles o circuitos virtuales permanentes (PVC) a todos los sitios de una VPN. Este tipo de VPN no es fácil de mantener o ampliar, ya que para agregar un nuevo sitio es necesario cambiar cada dispositivo de borde en la VPN.

Las VPN basadas en MPLS se crean en la Capa 3 y se basan en el modelo de peer. El modelo de peer permite al proveedor de servicios y al cliente intercambiar información de ruteo de Capa 3. El proveedor de servicios transmite los datos entre las instalaciones del cliente sin la participación del cliente.

Las VPN MPLS son más fáciles de administrar y ampliar que las VPN convencionales. Cuando se agrega un nuevo sitio a una VPN MPLS, sólo se debe actualizar el router de borde del proveedor de servicios que proporciona servicios al sitio del cliente.

Estos son los componentes de MPLS VPN:

- Router del proveedor (P): router en el núcleo de la red del proveedor. Los routers PE ejecutan el switching MPLS y no adjuntan etiquetas VPN a los paquetes enrutados. Las etiquetas VPN se utilizan para dirigir los paquetes de datos a la red privada correcta o al router de borde del cliente.
- Router PE: router que asocia la etiqueta VPN a los paquetes entrantes según la interfaz o subinterfaz en la que se reciben, y también conecta las etiquetas de núcleo MPLS. Un router PE se conecta directamente a un router CE.
- Router (C) del cliente: router en el proveedor de servicios de Internet (ISP) o en la red empresarial.
- Router de borde del cliente (CE): router de borde en la red del ISP que se conecta al router PE de la red. Un router CE debe interactuar con un router PE.

Descripción General de EVPN con L3VPN (MPLS SR)

Las implementaciones de Data Center (DC) han adoptado VXLAN EVPN (o) MPLS EVPN por sus ventajas, como el aprendizaje del plano de control EVPN, la multiempresa, la movilidad fluida, la

redundancia y las incorporaciones POD más sencillas. De forma similar, el CORE es una red MPLS L3VPN basada en el protocolo de distribución de etiquetas (LDP) o una transición de la capa subyacente tradicional basada en LDP de L3VPN MPLS a una solución más sofisticada como el routing de segmentos (SR).

El routing de segmentos se adopta por sus ventajas como:

- Planos de control de MPLS e IGP unificados
- Métodos de ingeniería de tráfico más sencillos
- Configuración más sencilla
- adopción de SDN
- EVPN (RFC 7432) es una solución basada en MPLS de BGP que se ha utilizado para los servicios Ethernet de última generación en una red de Data Center virtualizada.
- EVPN utiliza varios bloques de creación como RD, RT y VRF de las tecnologías MPLS que existen.
- La EVPN L3 sobre SR que se introdujo en la versión NXOS 7.0(3)I6(1) utiliza la ruta EVPN Tipo 5 con encapsulación MPLS.
- La EVPN L3 sobre SR ofrece varios arrendatarios, escalabilidad y alto rendimiento para servicios de Data Center evolucionados.

Nota: En DC, el plano de datos puede ser VXLAN o MPLS.

VPN MPLS L3 tradicional

Bloques de generación principales: RD, RT y VRF

Capa subyacente para transporte: IGP, LDP y RSVP-TE

Capa de superposición para servicio: VPNv4 y VPNv6


VPN L3 MPLS sobre SR

Bloques de generación principales: RD, RT y VRF

Capa subyacente para transporte: IGP/BGP-LU y SR-TE

Capa de superposición para servicio: EVPN

Diagrama de la red



Configuración

SPINE-1 Configuration		
Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
<pre>feature-set mpls feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam mpls label range 5000 45000 segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.1/32 index 211 route-map label-index-spine1 permit 20 set label-index 211</pre>	<pre>interface Ethernet1/33 ip address 172.16.0.1/24 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 172.16.1.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown interface loopback0 ip address 192.168.1.1/32 ip router ospf 100 area 0.0.0.0 router ospf 100 segment-routing mpls router-id 192.168.1.1</pre>	<pre>router bgp 65001 router-id 192.168.1.1 address-family ipv4 unicast network 192.168.1.1/32 route-map label-index-spine1 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 172.16.0.2 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</pre>

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
```

Verificación

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 Label	Prefix Table Id	FEC (Prefix/Tunnel id)	Next-Hop	Interface	Out Label	
16321	0x1	192.168.1.4/32	172.16.0.1	Eth1/51	16321	SWAP
"	0x1	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	

Información Relacionada

- [VPN MPLS BGP Multiprotocolo](#)
- [Informe técnico de routing de segmentos en switches de plataforma Cisco Nexus 9500, 9300, 9200, 3200 y 3100](#)
- [Configuración de EVPN de Capa 3 y VPN de Capa 3 sobre MPLS de Ruteo de Segmentos](#)