

Implementazione di Nexus L2 EVPN su MPLS routing segmento

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Sfondo](#)

[Configurazione](#)

[Esempio di rete](#)

[Verifica](#)

[Riferimenti](#)

Introduzione

In questo documento viene descritto come distribuire/configurare VPN di layer 2 su MPLS di routing dei segmenti sugli switch Cisco Nexus serie 9000.

Prerequisiti

Requisiti

Necessità di conoscere BGP, OSPF, MPLS, LDP, RSVP, EVPN, Segment Routing(SR)

Componenti usati

Cisco Nexus switch 93360YC-FX2 in esecuzione con 9.3(3)

Cisco Nexus switch 93240YC-FX2 in esecuzione con 9.3(3)

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.

Sfondo

Definizione di VPN di livello 2, VPLS/L2-EVPN è un servizio VPN di livello 2 da punto a punto multiplo che connette più filiali di un cliente in un'unica architettura a commutazione logica su una rete IP/MPLS.

EVPN-MPLS SR Layer 2:


- EVPN (RFC 7432) è una soluzione basata su BGP MPLS che è stata utilizzata per i servizi Ethernet di nuova generazione nella rete virtualizzata del centro dati
- EVPN utilizza diversi elementi di base, quali 'RD', 'RT' e 'VRF' delle tecnologie MPLS esistenti
- EVPN opera in contrasto con il VPLS esistente consentendo l'apprendimento MAC basato sul control plane nel core
- In EVPN, i PE che partecipano alle istanze EVPN apprendono le route MAC del cliente nel control-plane utilizzando il protocollo MP-BGP
- L'apprendimento MAC Control-Plane fornisce una serie di vantaggi che consentono a EVPN di risolvere le carenze VPLS, incluso il supporto del multihoming con bilanciamento del carico per flusso
- SR L2 EVPN è una nuova funzionalità disponibile in NXOS 9.3(1) ed è supportata sulle piattaforme Nexus serie 9300 FX2

Limitazioni per EVPN L2 su SR MPLS:

- Il flusso EVPN di livello 2 del routing dei segmenti è basato sul meccanismo di replica in entrata
- Utilizza un router EVPN tipo 3 per il traffico BUM
- Il core MPLS non supporta il multicast
- Eliminazione ARP non supportata
- Verifica della coerenza su VPC non supportata
- Non è possibile configurare insieme le stesse EVI L2 e EVI L3

Configurazione

Esempio di rete



Fasi della configurazione di alto livello:

- Funzionalità di installazione
- Configure Ip address - Underlay
- Configurare IGP -OSPF
- Configurazione di MP-BGP
- Configurazione della sovrapposizione VLAN ed EVPN
- Configurazione host finale per layer 2

SPINE -1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
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 192.168.1.1/32 index 211 route-map label-index-spine1 permit 10 set label-index 211	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	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

SPINE -2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
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	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
<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	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 192.168.1.4/32 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</pre>	<pre>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</pre>	<pre>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</pre>

Verifica

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, l-injecte

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.fafd]:[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.fafd]:[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.fafd]:[0]:[0.0.0.0]/216	192.168.1.4	100	0	i	
* i	192.168.1.4	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[32]:[100.0.0.100]/248	192.168.1.4	100	0	i	
* i	192.168.1.4	100	0	i	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	0	i	
* i	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, l-injecte

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	
* i	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	
* i	192.168.1.3	100	0	i	
*>[3]:[0]:[32]:[192.168.1.3]/88	192.168.1.3	100	0	i	
* i	192.168.1.3	100	0	i	
Route Distinguisher: 192.168.1.4:37864 (L2VNI 1001)					
*>[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[0]:[0.0.0.0]/216	192.168.1.4	100	32768	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.fafd]:[32]:[100.0.0.100]/248	192.168.1.4	100	32768	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	
*>[3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	100	32768	i	

Riferimenti

[White paper sul routing dei segmenti sugli switch con piattaforma Cisco Nexus 9500, 9300, 9200, 3200 e 3100](#)

[Configurazione di EVPN di layer 2 su MPLS routing segmento](#)