세그먼트 라우팅 MPLS를 통한 Nexus L2 EVPN 구현

목차

소개

이 문서에서는 Cisco Nexus 9000 Series 스위치에 Layer2 EVPN over Segment Routing MPLS를 구축/구성하는 방법에 대해 설명합니다.

사전 요구 사항

요구 사항

BGP, OSPF, MPLS, LDP, RSVP, EVPN, SR(Segment Routing)에 대한 지식 필요

사용되는 구성 요소

9.3.(3)과 함께 실행되는 Cisco Nexus 스위치 93360YC-FX2

9.3.(3)과 함께 실행되는 Cisco Nexus 스위치 93240YC-FX2

이 문서의 정보는 특정 랩 환경의 디바이스를 토대로 작성되었습니다.이 문서에 사용된 모든 디바이스는 초기화된(기본) 컨피그레이션으로 시작되었습니다.현재 네트워크가 작동 중인 경우, 모든 명령어의 잠재적인 영향을 미리 숙지하시기 바랍니다.

배경

레이어 2 VPN을 정의하고, VPLS/L2-EVPN은 IP/MPLS 네트워크를 통해 단일 논리적 스위치드 아키텍처에서 고객의 여러 지점을 연결하는 Multipoint-to-Multipoint Layer 2 VPN 서비스입니다.

레이어 2 EVPN-MPLS SR:

- EVPN(RFC 7432)은 가상화된 데이터 센터 네트워크에서 차세대 이더넷 서비스에 사용된 BGP MPLS 기반 솔루션입니다.

- EVPN은 기존 MPLS 기술의 'RD', 'RT', 'VRF'와 같은 여러 구성 요소를 사용합니다.
- EVPN은 코어에서 컨트롤 플레인 기반 MAC 학습을 활성화하여 기존 VPLS와 대조적으로 작동합니다.
- EVPN에서 EVPN 인스턴스에 참여하는 PE는 MP-BGP 프로토콜을 사용하여 컨트롤 플레인에서 고객 MAC 경로를 학습합니다.
- 컨트롤 플레인 MAC 학습은 EVPN이 VPLS 단점을 해결할 수 있도록 하는 다양한 이점을 제공합니다. 여기에는 플로우당 로드 밸런싱을 통한 멀티홈(multihoming) 지원 등이 포함됩니다.
- SR L2 EVPN은 NXOS 9.3(1)에서 제공되는 새로운 기능이며 Nexus 9300 FX2 시리즈 플랫폼에서 지원됩니다.

SR MPLS를 통한 L2 EVPN의 제한:

- 세그먼트 라우팅 레이어 2 EVPN 플러딩은 인그레스 복제 메커니즘을 기반으로 함
- BUM 트래픽에 EVPN Type 3 경로를 사용합니다.
- MPLS 코어는 멀티캐스트를 지원하지 않습니다.
- ARP 억제는 지원되지 않습니다.
- VPC의 일관성 검사는 지원되지 않습니다.
- 동일한 L2 EVI와 L3 EVI를 함께 구성할 수 없습니다.

구성

네트워크 다이어그램

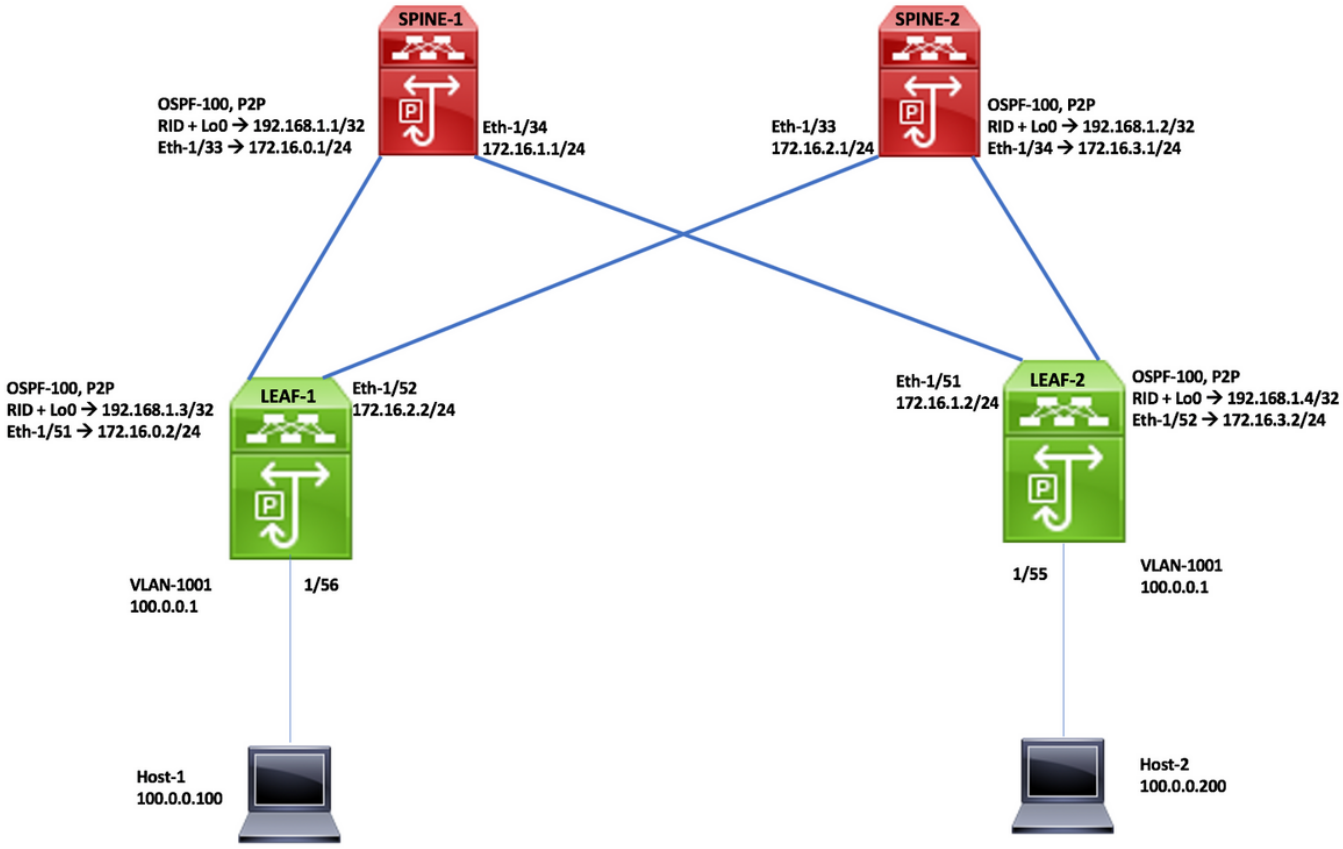

상위 레벨 구성 단계:

- 기능 설치
- IP 주소 구성 -언더레이

- IGP -OSPF 구성
- MP-BGP 구성
- VLAN 및 EVPN 오버레이 구성
- 레이어 2에 대한 엔드 호스트 구성

SPINE -1 Configuration

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

```
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
```

OSPF Configuration

```
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
```

BGP/EVPN Configuration

```
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

```
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
192.168.1.2/32index 221

route-map label-index-spine2 permit 10
set label-index 221
```

OSPF Configuration

```
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
```

BGP/EVPN Configuration

```
router bgp 65001
router-id 192.168.1.2
address-family ipv4 unicast
network 192.168.1.2/32route-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

```
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
```

OSPF Configuration

```
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
```

BGP/EVPN Configuration

```
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
```

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
```

확인

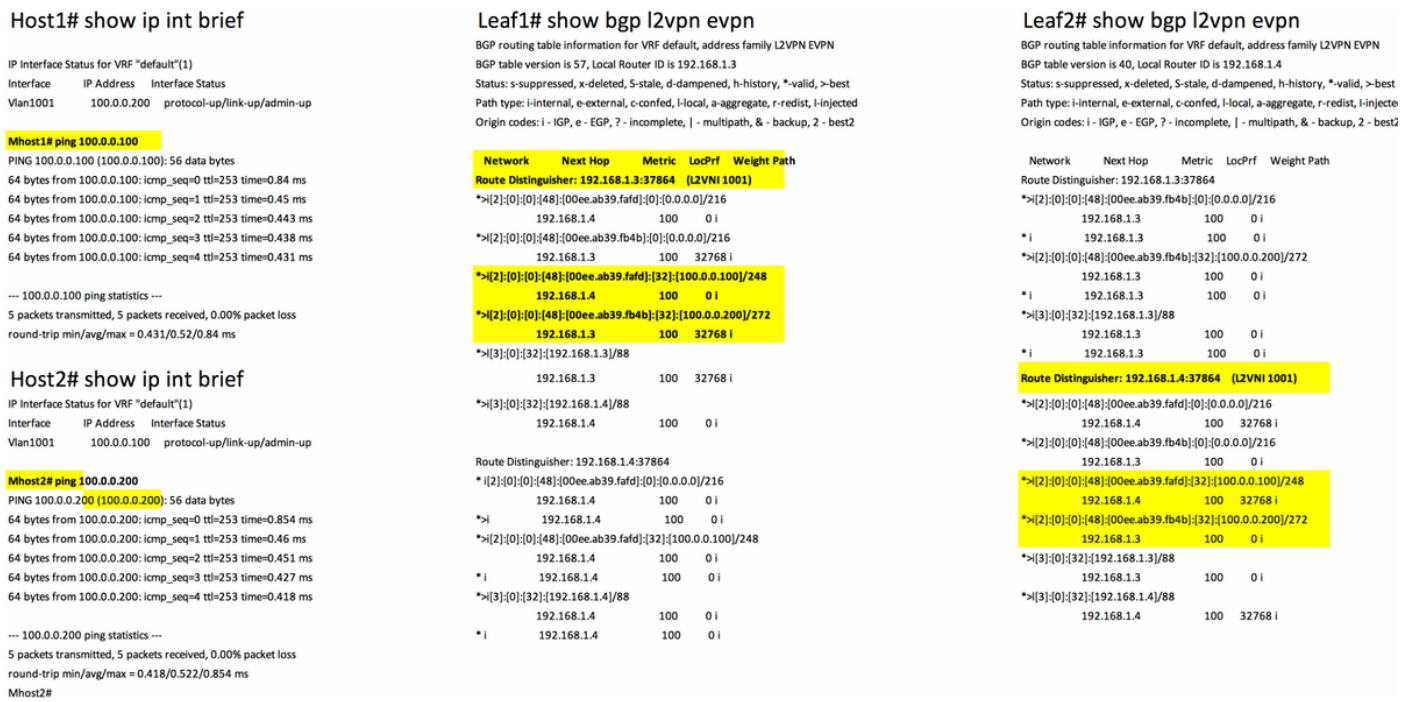

Host1# show ip int brief

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-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)
*>i[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[0]:[0.0.0.0]/216
192.168.1.4 100 0 i
*>l[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216
192.168.1.3 100 32768 i
*>i[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[32]:[100.0.0.100]/248
192.168.1.4 100 0 i
*>l[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272
192.168.1.3 100 32768 i
*>l[3]:[0]:[32]:[192.168.1.3]/88
192.168.1.3 100 32768 i
*>i[3]:[0]:[32]:[192.168.1.4]/88
192.168.1.4 100 0 i

Route Distinguisher: 192.168.1.4:37864
* i[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
*>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
*>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
*>l[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
*>l[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
*>l[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)
*>i[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[0]:[0.0.0.0]/216
192.168.1.4 100 32768 i
*>l[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[0]:[0.0.0.0]/216
192.168.1.3 100 0 i
*>i[2]:[0]:[0]:[48]:[00ee.ab39.fafd]:[32]:[100.0.0.100]/248
192.168.1.4 100 32768 i
*>l[2]:[0]:[0]:[48]:[00ee.ab39.fb4b]:[32]:[100.0.0.200]/272
192.168.1.3 100 0 i
*>l[3]:[0]:[32]:[192.168.1.3]/88
192.168.1.3 100 0 i
*>i[3]:[0]:[32]:[192.168.1.4]/88
192.168.1.4 100 32768 i

참조

[Cisco Nexus 9500, 9300, 9200, 3200 및 3100 플랫폼 스위치의 세그먼트 라우팅 백서]

[세그먼트 라우팅 MPLS를 통한 Layer2 EVPN 구성]