

セグメントルーティングMPLSでのNexus L2 EVPNの実装

内容

[概要](#)

[前提条件](#)

[要件](#)

[使用するコンポーネント](#)

[背景](#)

[設定](#)

[ネットワーク図](#)

[確認](#)

[参考資料](#)

概要

このドキュメントでは、Cisco Nexus 9000シリーズスイッチにレイヤ2 EVPN over Segment Routing MPLSを導入/設定する方法について説明します。

前提条件

要件

BGP、OSPF、MPLS、LDP、RSVP、EVPN、セグメントルーティング(SR)に関する知識が必要

使用するコンポーネント

9.3.(3)で動作するCisco Nexusスイッチ93360YC-FX2

9.3.(3)で動作するCisco Nexusスイッチ93240YC-FX2

このドキュメントの情報は、特定のラボ環境にあるデバイスに基づいて作成されました。このドキュメントで使用するすべてのデバイスは、初期（デフォルト）設定の状態から起動しています。対象のネットワークが実稼働中である場合には、どのようなコマンドについても、その潜在的な影響について確実に理解しておく必要があります。

背景

レイヤ2 VPNを定義します。VPLS/L2-EVPNは、IP/MPLSネットワーク上の単一の論理スイッチアーキテクチャで、顧客の複数のブランチを接続するマルチポイントレイヤ2 VPNサービスです。

レイヤ2 EVPN-MPLS SR:


- EVPN(RFC 7432)はBGP MPLSベースのソリューションで、仮想化データセンターネットワークの次世代イーサネットサービスに使用されています
- EVPNは、既存のMPLSテクノロジーから「RD」、「RT」、「VRF」などの複数の構成要素を使用します
- EVPNは、コアでコントロールプレーンベースのMACラーニングを有効にすることで、既存のVPLSとは対照的に動作します
- EVPNでは、EVPNインスタンスに参加しているPEは、MP-BGPプロトコルを使用してコントロールプレーン内のカスタマーMACルートを学習します
- コントロールプレーンのMACラーニングは、フロー単位のロードバランシングによるマルチホーミングのサポートなど、EVPNがVPLSの欠点に対処できるようにする多くの利点を提供します
- SR L2 EVPNはNXOS 9.3(1)で利用可能な新機能で、Nexus 9300 FX2シリーズプラットフォームでサポートされています

L2 EVPN over SR MPLSの制限事項：

- セグメントルーティングレイヤ2 EVPNフラッディングは、入力レプリケーションメカニズムに基づいています
- BUMトラフィックにEVPNタイプ3ルートを使用する
- MPLSコアはマルチキャストをサポートしない
- ARP抑制はサポートされていません
- VPCでの整合性チェックはサポートされていません
- 同じL2 EVIとL3 EVIを一緒に設定することはできません

設定

ネットワーク図



高度な設定手順：

- インストール機能
- IPアドレスの設定 – アンダーレイ
- IGP-OSPFの設定
- MP-BGPの設定
- VLANとEVPNオーバーレイの設定
- レイヤ2のエンドホストの設定

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
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	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	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	OSPF Configuration	BGP/EVPN Configuration
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	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	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
fabric forwarding anycast-gateway-mac 0000.0000.1111 mpls label range 5000 450000	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	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
vlan 1,1001 segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.4/32 index 321	interface Ethernet1/55 switchport switchport mode trunk switchport trunk allowed vlan 1001 no shutdown	evpn encapsulation mpls source-interface loopback0
vlan 1001 evi auto	interface loopback0 ip address 192.168.1.4/32 ip router ospf 100 area 0.0.0.0	
route-map label-index-Leaf2 permit 10 set label-index 321	router ospf 100 segment-routing mpls router-id 192.168.1.4	
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		

確認

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					
*>[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	
* [3]:[0]:[32]:[192.168.1.4]/88	192.168.1.4	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.fb4b]:[0]:[0.0.0.0]/216	192.168.1.3	100	0	i	
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248					
*>[2]:[0]:[0]:[48]:[00ee.ab39.faf4]:[32]:[100.0.0.100]/248	192.168.1.4	100	32768	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	

参考資料

[『Segment Routing on Cisco Nexus 9500, 9300, 9200, 3200, and 3100 Platform Switches White paper』](#)

セグメントルーティングMPLSでのレイヤ2 EVPNの設定