在Catalyst 9000交換機上將EVPN VxLAN遷移到IPv6底層

目錄

簡介

本文說明如何將EVPN VxLAN遷移到Catalyst 9000系列交換器上的IPv6底層。

必要條件

需求

思科建議您瞭解以下主題：

- 單播EVPN VxLAN功能、BGP和MVPN（組播虛擬專用網路）。
- IPv4和IPv6單播
- 組播概念和組播如何運行

採用元件

本文中的資訊係根據以下軟體和硬體版本：

- Catalyst 9000 系列交換器

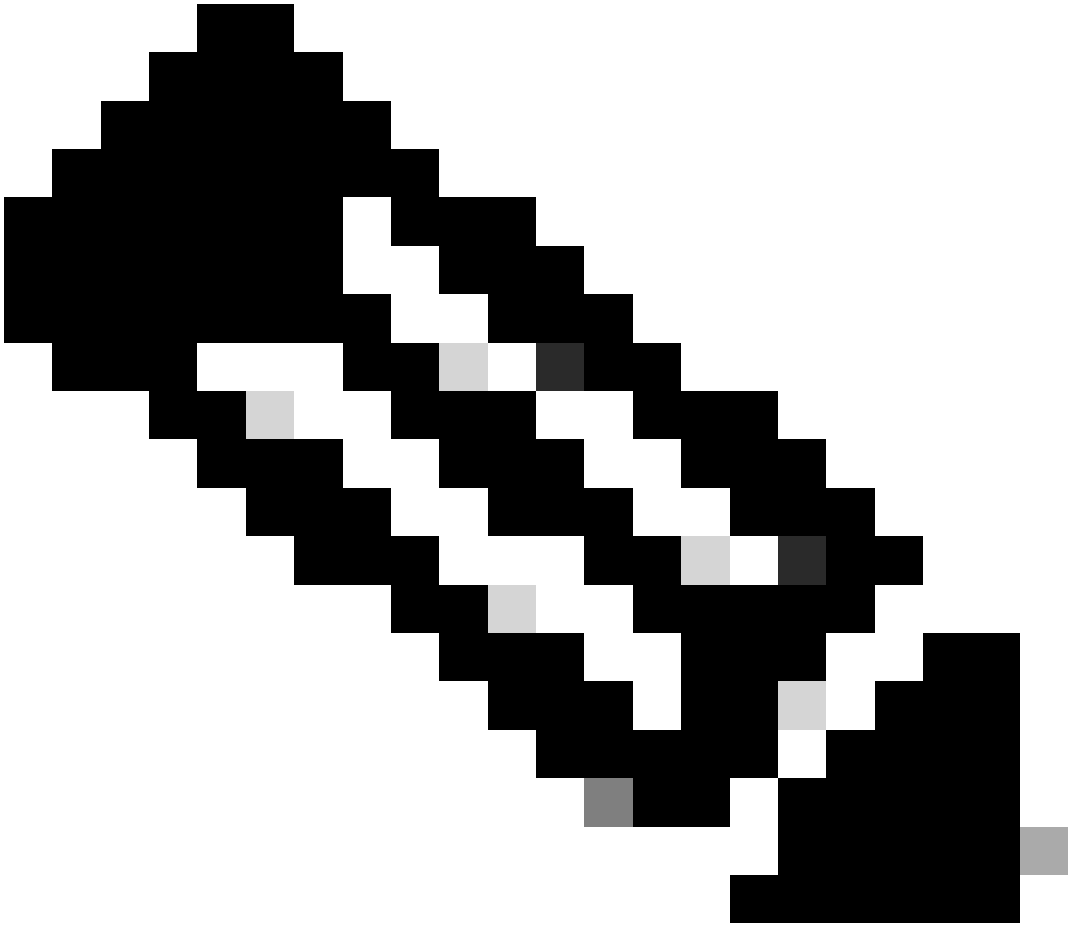

附註：9200、9500X和9600X不支援VXLANv6

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除（預設）的組態來啟動。如果您的網路運作中，請確保您瞭解任何指令可能造成的影響。

背景資訊

要遷移到EVPN VXLANv6，需要更改EVPN交換矩陣中的某些配置以啟用IPv6底層。本檔案將詳細介紹將現有EVPN VXLANv4部署遷移到綠場（僅限VXLANv6）或棕場（雙堆疊 — VXLANv4和VXLANv6）部署的相關配置更改和驗證程式。

綠地EVPN VXLANv6部署要求：

- IPv6核心
- 將EVPN交換矩陣遷移到VXLANv6底層支援
- 將BGP EVPN鄰居遷移到IPv6鄰居對等

Brownfield EVPN VXLAN部署需要：

- IPv4 + IPv6核心
- 將EVPN交換矩陣無縫遷移至雙堆疊(VXLANv4 + VXLANV6)襯底
- BGP EVPN鄰居對等從IPv4無縫遷移到IPv6鄰居地址

技術

EVPN	乙太網路虛擬私人網路	允許BGP傳輸第2層MAC和第3層IP資訊的擴展是EVPN，它使用多協定邊界網關協定(MP-BGP)作為協定來分發有關VXLAN重疊網路的可達性資訊。
VXLAN	虛擬可擴充LAN（區域網路）	VXLAN的用途是克服VLAN和STP的固有限制。建議採用的IETF標準[RFC 7348]提供與VLAN相同的乙太網第2層網路服務，但具有更高的靈活性。功能上，它是UDP內MAC封裝協定，在第3層底層網路上作為虛擬重疊運行。
VTEP	虛擬通道端點	這是負責執行封裝和解除封裝的裝置
EVI	EVPN例項	EVPN例項(EVI)由虛擬網路識別符號(VNI)表示。 EVI表示PE路由器上的VPN。它充當IP VPN路由和轉發(VRF)的相同角色，並且為EVI分配了匯入/匯出路由目標(RT)
NVE	網路虛擬介面	進行封裝和解除封裝的邏輯介面
VNI	VXLAN網路識別碼	唯一標識每個第2層子網或網段。VNI有兩種型別： 對稱(L2VNI):VTEP具有相同的VNI 非對稱(L3VNI):VTEP沒有相同的VNI，而是透過單一通用VNI路由。
BUM	廣播、未知單播、組播	在NVE設定下，BUM流量會透過連結到VNI的Mcast群組傳送。
TRM	租戶路由多點傳送	基於BGP-EVPN的解決方案，可在VxLAN交換矩陣中通過VTEPS連線的源和接收器之間啟用組播路由[RFC7432]。有兩種型別：L2TRM（第2層TRM）和L3TRM（第3層TRM）

MDT	多點傳送發佈樹	在VTEP之間構建的組播樹，用於租戶組播流量的封裝和通道化。
PVLAN	專用VLAN	將VLAN的乙太網路廣播網域分割為子網域，這樣您就可以將交換器上的連線埠彼此隔離。
MIB	管理資訊庫	A 簡單網路管理協定(SNMP)監控器對象
PIM-BIDIR	通訊協定無關多點傳送雙向	一種PIM，其中的流量僅沿著共用樹轉發 該位置位於組的集合點(RP)。
VFI	虛擬轉送例項	虛擬網橋埠，能夠根據目標MAC地址、源MAC地址學習和老化等執行本地橋接功能，如轉發。
IRB	整合式路由和橋接	啟用第2層VPN和第3層VPN覆蓋，允許覆蓋範圍內的終端主機在同一子網內和VPN內的不同子網之間相互通訊。
IMET	包含型組播乙太網標籤	也稱為BGP路由型別3(RT3)，用於自動發現遠端對等點，以便透過VXLAN建立BUM通道。IMET路由傳送從遠端對等點通告的遠端（輸出）VNI，這可能與本地VNI不同。這些遠端VNI稱為下游指派的VNI。
DAG	分散式任播閘道	所有VTEP上的預設網關功能。相同的網關IP存在於所有VTEP上，並允許交換矩陣中的移動性。

限制

- 只有Cat9k交換機支援無縫遷移
- 僅考慮一個NVE介面和全域性遷移

這些EVPN功能不支援VXLANv6底層

- 集中式網關
- 多宿主支援
- L3多點傳送(TRM)
- 含輸入複製的L2TRM
- 使用預設MDT的L2TRM（多點傳送複製）
- 帶有預設MDT的L3TRM
- 含資料MDT的L3TRM
- 邊界網關（多站點）
- 接入VFI

- PVLAN
- MIB
- 適用於多點傳送底層的PIM-BIDIR

無縫遷移概念概述

棕色EVPN VXLAN部署需要將網路從VXLANv4逐步遷移到VXLANv6襯底。為實現此EVPN VXLAN網路，需要逐步從IPv4遷移到IPv6底層並允許部分EVPN網路遷移到IPv6底層以及網路的其他部分繼續使用IPv4底層但網路中的所有節點仍然可以連線。

為了對單播和BUM（廣播、未知單播和多播）入口複製實現這種無縫遷移，EVPN節點需要支持雙堆疊VTEP。雙堆疊VTEP節點具有與同一VNI（VXLAN網路識別碼）相關聯的兩個VTEP位址（IPv4和IPv6）。 在底層遷移期間，這兩個VTEP IP地址在單個BGP EVPN更新(BGP EVPN Dual-Next-hop update)中通告給對等體，並為接收節點提供選擇底層中任一層用於流量轉發的選項。

BGP EVPN雙下一躍點更新通告

BGP雙下一跳更新攜帶兩個下一跳：

- MP_REACH_NLRI(EVPN Routetype-2/Routetype-5)/PMSI-tunnel(EVPN Routetype-3)屬性中的主要下一躍點（現有底層）
- BGP通道封裝屬性(23)中的次要下一躍點（遷移底層）

作為主要和次要承載的VTEP IP取決於EVPN節點的遷移模式。

此表詳細說明了雙下一跳更新中攜帶的主要/輔助VTEP IP

遷移模式	主要下一跳	次要Nexthop
VXLANv4到VXLANv6	IPv4 VTEP	IPv6 VTEP
VXLANv6到VXLANv4	IPv6 VTEP	IPv4 VTEP

BGP分葉/邊緣EVPN雙下一躍點更新處理

接收此BGP EVPN雙下一跳更新的枝葉/邊緣/邊界節點使用其中一個接收的下一跳作為遠端VTEP進行轉發。用於襯底的下一跳取決於裝置上配置的這些遷移策略。

- 本地VTEP地址
- 本地底層優先使用程度

此表詳細說明了本地配置的策略如何決定使用哪個底層來轉發資料包

已收到BGP更新	本地VTEP 地址	本地底層優先使用程度	VXLAN底層 單點傳播/BUM-IR
雙下一跳(IPv4 + IPv6)	僅限IPv4 VTEP	不適用	VXLANv4
雙下一跳(IPv4 + IPv6)	僅IPv6 VTEP	不適用	VXLANv6
雙下一跳(IPv4 + IPv6)	雙堆疊(IPv4 + IPv6 VTEP IP)	IPv4	VXLANv4
雙下一跳(IPv4 + IPv6)	雙堆疊(IPv4 + IPv6 VTEP IP)	IPv6	VXLANv6
單個IPv4下一跳	僅限IPV4 VTEP	不適用	VXLANv4
單個IPv4下一跳	僅IPV6 VTEP	不適用	無VXLAN底層
單個IPv4下一跳	雙堆疊(IPv4 + IPv6 VTEP IP)	不適用	VXLANv4
單個IPv6下一跳	僅限IPV4 VTEP	不適用	無VXLAN底層
單個IPv6下一跳	僅IPV6 VTEP	不適用	VXLANv6
單個IPv6下一跳	雙堆疊(IPv4 + IPv6 VTEP IP)	不適用	VXLANv6

配置（VXLAN底層遷移模式）

「interface nve」配置下的新cli命令可用於設定VXLAN底層遷移模式以及單播和組播的底層首選項。

用於單播和BUM-Ingress複製的遷移模式CLI

<#root>

interface nve 1

```
     vxlan encapsulation ?
         dual-stack  Encapsulation type dual-stack
         ipv4        Encapsulation type IPv4
         ipv6        Encapsulation type IPv6
     vxlan encapsulation dual-stack ?
         prefer-ipv4  Dual-stack underlay with ipv4 preference
         prefer-ipv6  Dual-stack underlay with ipv6 preference
```

下表詳細說明了單播和BUM-IR遷移模式的CLI配置

CLI組態	本地VTEP IP和 單點傳播/BUM-IR底層
int nve 1 vxlan封裝ipv4 （這是選用的，因為預設vxlan封裝是ipv4）	IPv4（VXLANv4底層）
int nve 1 vxlan封裝ipv6	IPv6（VXLANv6底層）
int nve 1 vxlan封裝雙堆疊優先使用 — ipv4	雙堆疊(IPv4 + IPv6)(優選VXLANv4襯底)
int nve 1 vxlan封裝雙堆疊優先 — ipv6	雙堆疊(IPv4 + IPv6)(優選VXLANv6底層)

靜態組播複製的遷移模式CLI

<#root>

interface nve 1

```
     vxlan encapsulation ?
         dual-stack  Encapsulation type dual-stack
         ipv4        Encapsulation type IPv4
      ipv6        Encapsulation type IPv6
```

```
 vxlan encapsulation dual-stack ?
prefer-ipv4  Dual-stack underlay with ipv4 preference
 prefer-ipv6  Dual-stack underlay with ipv6 preference
 vxlan encapsulation dual-stack prefer-ipv4 underlay-mcast ?
     ipv4  Select IPv4 multicast underlay
     ipv6  Select IPv6 multicast underlay
 vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ?
     ipv4  Select IPv4 multicast underlay
     ipv6  Select IPv6 multicast underlay
```

CLI配置	靜態多點傳送底層
int nve 1 member vni <L2VNI> mcast-group <v4-mcast-group> vxlan封裝ipv4 （這是選用的，因為預設vxlan封裝是ipv4）	在為L2VNI配置的IPv4底層組播組上傳送和接收組播流量
int nve 1 member vni <L2VNI> mcast-group <v6-mcast-group> vxlan封裝ipv6	在為L2VNI配置的IPv6底層組播組上傳送和接收組播流量
int nve 1 member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group> vxlan封裝雙堆疊優先 — ipv6	雙堆疊(IPv4 +IPv6) 在為L2VNI配置的IPv4和IPv6底層組播組上接收組播流量 僅在為L2VNI配置的IPv4底層組播組上傳送組播流量
int nve 1 member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group> vxlan封裝雙堆疊優先使用 — ipv4	雙堆疊(IPv4 +IPv6) 在為L2VNI配置的IPv4和IPv6底層組播組上接收組播流量 僅在為L2VNI配置的IPv6底層組播組上傳送組播流量
int nve 1	雙堆疊(IPv4 +IPv6)

member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group> vxlan封裝雙堆疊優先 — ipv6 底層組播ipv4	在為L2VNI配置的IPv4和IPv6底層組播組上接收組播流量 僅在為L2VNI配置的IPv4底層組播組上傳送組播流量
int nve 1 member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group> vxlan封裝雙堆疊優先 — ipv4底層 — 多點傳送ipv6	雙堆疊(IPv4 +IPv6) 在為L2VNI配置的IPv4和IPv6底層組播組上接收組播流量 僅在為L2VNI配置的IPv6底層組播組上傳送組播流量
int nve 1 member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group> vxlan封裝雙堆疊優先 — ipv6 底層組播ipv6	雙堆疊(IPv4 +IPv6) 在為L2VNI配置的IPv4和IPv6底層組播組上接收組播流量 僅在為L2VNI配置的IPv6底層組播組上傳送組播流量
int nve 1 member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group> vxlan封裝雙堆疊優先使用ipv4底層 — 多點傳送ipv4	雙堆疊(IPv4 +IPv6) 在為L2VNI配置的IPv4和IPv6底層組播組上接收組播流量 僅在為L2VNI配置的IPv4底層組播組上傳送組播流量

底層遷移程式

EVPN L2網關和EVPN IRB（分散式任播網關）部署的底層遷移步驟相同

VXLANv4到VXLANv6的遷移

VXLANv6部署在下層具有單個IPv6傳輸。VXLAN隧道和BGP鄰居關係都基於IPv6。

網路圖表

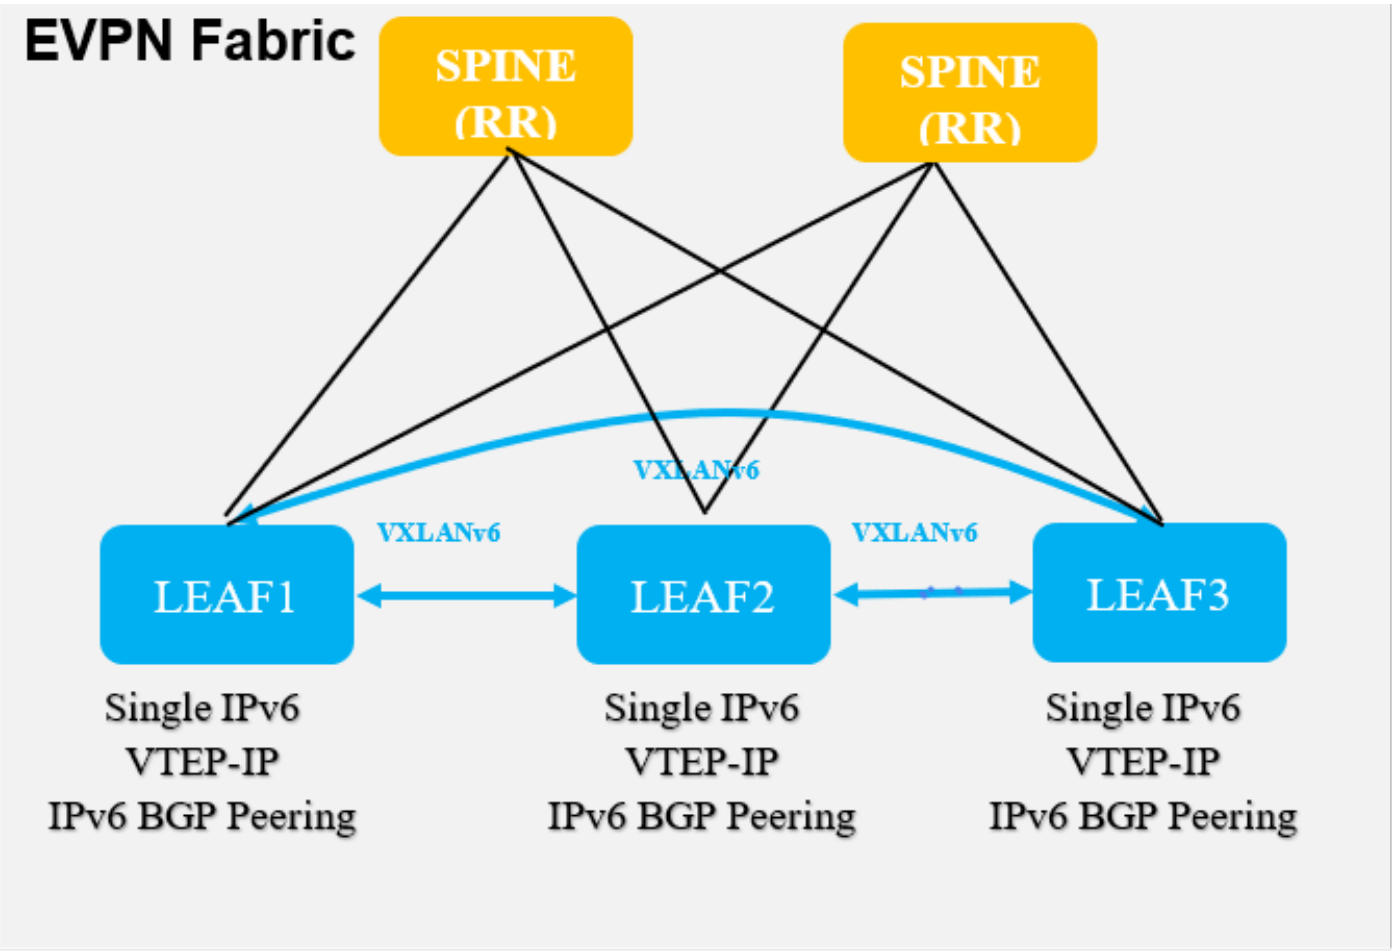

單播VxLANv4到VxLANv6遷移

下表詳細說明了單播流量的VxLANv4到VXLANv6底層遷移所需的配置更改示例。

遷移步驟	VXLANv4底層	VXLANv6底層	說明
	EVPN路由器ID配置		
1		l2vpn router-id 10.1.1.1	配置l2vpn router-id以用作EVPN router-id
	VXLAN VTEP IP組態		
2	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255	interface Loopback1 ipv6 address 2001:DB8:2::2/128	與使用IPv6地址配置的VXLAN關聯的環回介面。此IPV6地址用於VXLAN的本機IPv6 VTEP。

	interface nve1 source-interface Loopback1	interface nve1 source-interface Loopback1	
3	interface Loopback1 ip ospf 1區域0 interface nve1 source-interface Loopback1	interface Loopback1 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	為介面的IPv6地址啟用IGP（如OSPF）
	底層遷移模式配置		
4		interface nve1 vxlan封裝ipv6	必須使用「vxlan encapsulation ipv6」配置VXLANv6底層配置VXLAN NVE介面
	單播路由配置		
5		ipv6單播路由	啟用IPv6路由
	IGP配置		
6	router ospf 1	ipv6 router ospf 1 router-id 10.1.1.1	啟用用於IPv6的OSPF
	BGP組態		
7		router bgp 100 bgp router-id 10.2.2.1	設定BGP路由器ID
8	router bgp 100 neighbor 10.99.99 remote-as 100	router bgp 100 neighbor 2001:DB8:99::99	BGP EVPN對等已移至IPv6鄰居地址

	neighbor 10.99.99.99 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.99.99.99啟用 neighbor 10.99.99 send-community both exit-address-family ! exit-address-family	remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 鄰居 2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	

BUM輸入複製VxLANv4到VxLANv6的遷移

下表詳細說明了從BUM-IR進行VxLANv4到VXLANv6底層遷移所需的配置更改示例

遷移步驟	VXLANv4底層	VXLANv6底層	說明
	EVPN路由器ID配置		
1		l2vpn router-id 10.1.1.1	將l2vpn router-id配置為用作EVPN router-id
	VXLAN VTEP IP組態		
2	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 interface nve1 source-interface	interface Loopback1 ipv6 address 2001:DB8:2::2/128 interface nve1 source-interface	與使用IPv6地址配置的VXLAN關聯的環回介面。此IPV6地址用於VXLAN的本地IPv6 VTEP

	Loopback1	Loopback1	
3	interface Loopback1 ip ospf 1區域0 interface nve1 source-interface Loopback1	interface Loopback1 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	為介面的IPv6地址啟用IGP（如OSPF）
	底層遷移模式配置		
4		interface nve1 vxlan封裝ipv6	必須使用「vxlan encapsulation ipv6」配置VXLANv6底層配置VXLAN NVE介面
	單播路由配置		
5		ipv6單播路由	啟用IPv6路由
	IGP配置		
6	router ospf 1	ipv6 router ospf 1 router-id 10.1.1.1	啟用用於IPv6的OSPF
	BGP組態		
7		router bgp 100 bgp router-id 10.2.2.1	設定BGP路由器ID
8	router bgp 100 neighbor 10.9.9 remote-as 100	router bgp 100 neighbor 2001:DB8:99::99 remote-	BGP EVPN對等已移至IPv6鄰居地址

	neighbor 10.9.9.9 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.9.9.9啟用 neighbor 10.9.9 send-community both exit-address-family ! exit-address-family	as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 鄰居2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	

靜態組播複製VxLANv4到VxLANv6遷移

下表詳細說明了靜態組播複製的VxLANv4到VXLANv6底層遷移所需的配置更改示例

遷移步驟	VXLANv4底層	VXLANv6底層	說明
	靜態組播複製配置		
1	interface nve1 member vni 20011 mcast-group 226.1.1.1	interface nve1 member vni 20011 mcast — 組FF05::1	配置靜態IPv6複製組播地址
	底層遷移模式配置		
2		interface nve1 vxlan封裝ipv6	必須使用「vxlan encapsulation ipv6」配置VXLANv6底層配置VXLAN NVE介面
	單播路由配置		

3		ipv6單播路由	啟用IPv6路由
	組播路由配置		
4	ip multicast-routing	ipv6多點傳送路由	啟用IPv6組播路由
5	ip pim rp-address 10.9.9.9	ipv6 pim rp-address 2001:DB8::99:99	將PIM RP地址遷移到IPv6

Brownfield - VXLANv4和VXLANv6無縫遷移

棕色現場部署在底層具有可傳遞的雙IPv4/IPv6傳輸，可實現無縫遷移。VXLAN隧道和BGP鄰居最初基於IPv4，然後無縫遷移到基於IPv6的IPv4（遷移後，可以選擇從底層刪除IPv4）。 換句話說，單個VTEP可以遷移到雙IPv4和IPv6，而其他的VTEP繼續使用IPv4運行。一旦交換矩陣內的所有VTEP都支援雙IPv4和IPv6，單個VTEP現在可以遷移到IPv6。

網路圖表

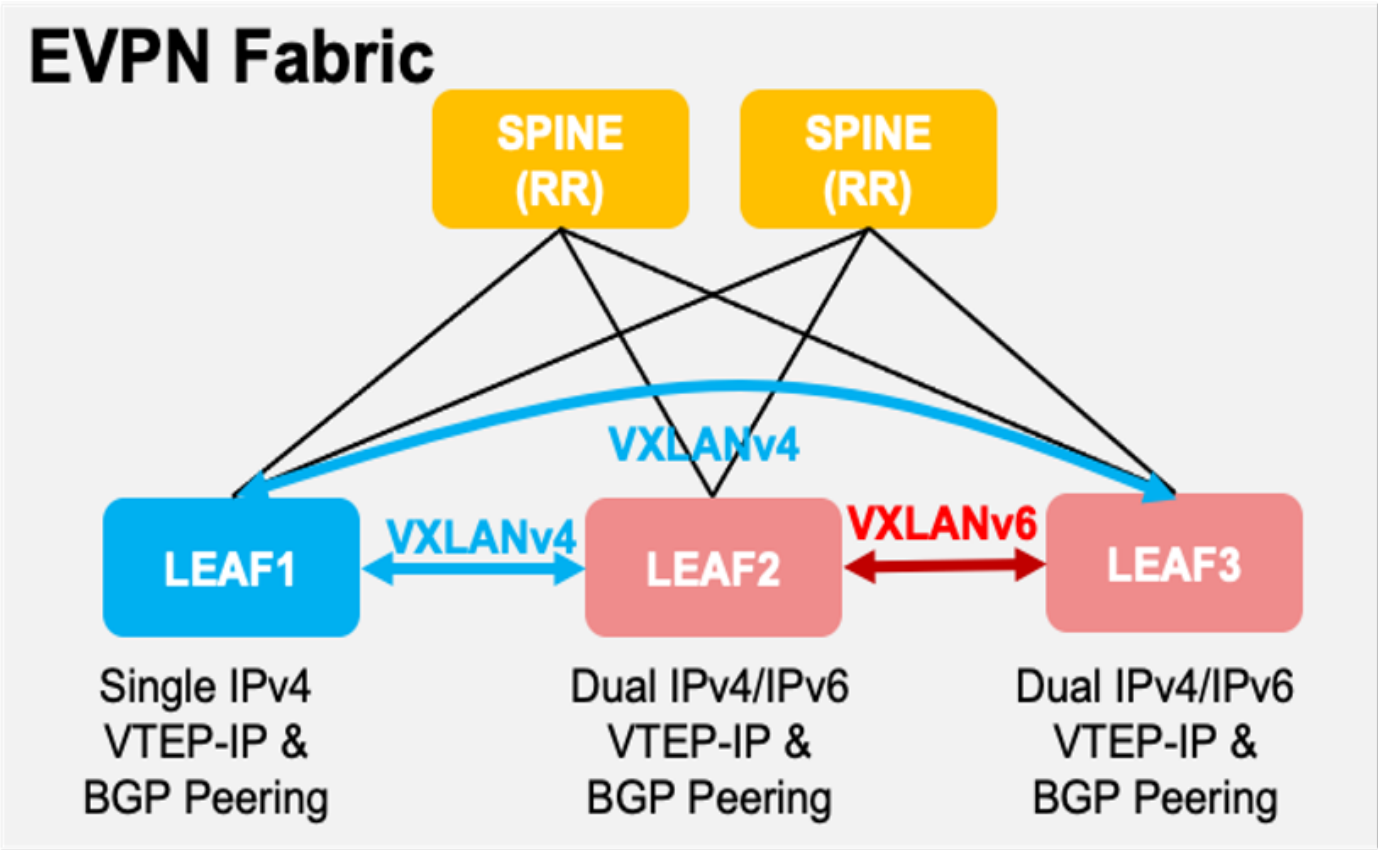

Brownfield單播VxLANv4到雙堆疊的遷移

下表詳細說明了單播流量從Brownfield VxLANv4到雙堆疊底層遷移所需的配置更改示例

遷移步驟	VXLANv4底層	雙堆疊（首選VxLANv6底層）	說明
	L2VPN Router-ID Configuration		
1		l2vpn router-id 10.2.2.3	將l2vpn router-id配置為用作EVPN router-id
	VXLAN VTEP IP組態		
2	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 interface nve1 source-interface Loopback1	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 ipv6 address 2001:DB8:2::2/128 interface nve1 source-interface Loopback1	與使用IPv4和IPv4地址配置的VXLAN關聯的環回介面。
3	interface Loopback1 ip ospf 1區域0 interface nve1 source-interface Loopback1	interface Loopback1 ip ospf 1區域0 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	為介面的IPv4和IPv6地址啟用IGP（如OSPF）
	底層遷移模式配置		
4		interface nve1 vxlan封裝雙堆疊優先 —	必須使用「vxlan encapsulation dual-stack prefer-ipv6」為雙堆疊配置VXLAN NVE介面，但必須使用「

		ipv6	vxlan encapsulation dual-stack prefer-ipv6」為雙堆疊配置VXLANv6底層
	單播路由配置		
6		ipv6單播路由	啟用IPv6路由
	IGP配置		
7	router ospf 1	router ospf 1 ! ipv6 router ospf 1 router-id 10.1.1.1	為IPv4和IPv6啟用OSPF
	BGP組態		
8		router bgp 100 bgp router-id 10.2.2.1	設定BGP路由器ID
9	router bgp 100 neighbor 10.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.9.9.9啟用 neighbor 10.9.9 send-community both	router bgp 100 neighbor 10.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0 neighbor 2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.9.9.9啟用	使用IPv4和IPv6鄰居位址的BGP EVPN對等

	exit-address-family ! exit-address-family	neighbor 10.9.9 send-community both 鄰居2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	

Brownfield BUM輸入複製VxLANv4到雙堆疊遷移

此表詳細說明了從Brownfield VxLANv4遷移到BUM-IR的雙堆疊底層所需的配置更改示例

遷移步驟	VXLANv4底層	雙堆疊（首選VxLANv6底層）	說明
	L2VPN Router-ID Configuration		
1		l2vpn router-id 10.2.2.3	將l2vpn router-id配置為用作EVPN router-id
	VXLAN VTEP IP組態		
2	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 interface nve1 source-interface Loopback1	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 ipv6 address 2001:DB8:2::2/128 interface nve1 source-interface Loopback1	與VXLAN關聯的環回介面，配置了IPv4和IPv6地址。
3	interface Loopback1 ip ospf 1區域0 interface nve1	interface Loopback1 ip ospf 1區域0 ipv6 ospf 1區域0	為介面的IPv4和IPv6地址啟用IGP（如OSPF）

	source-interface Loopback1	interface nve1 source-interface Loopback1	
	底層遷移模式配置		
4		interface nve1 vxlan封裝雙堆疊優先 — ipv6	必須使用「vxlan encapsulation dual-stack prefer-ipv6」為雙堆疊配置VXLAN NVE介面，但必須使用「vxlan encapsulation dual-stack prefer-ipv6」為雙堆疊配置VXLANV6襯底
	單播路由配置		
5		ipv6單播路由	啟用IPv6路由
	IGP配置		
6	router ospf 1	router ospf 1 ipv6 router ospf 1 router-id 10.1.1.1	為IPv4和IPv6啟用OSPF
	BGP組態		
7		router bgp 100 bgp router-id 10.2.2.1	設定BGP路由器ID
8	router bgp 100 neighbor 10.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0	router bgp 100 neighbor 10.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0	使用IPv4和IPv6鄰居位址的BGP EVPN對等

	! address-family l2vpn evpn 鄰居10.9.9.9啟用 neighbor 10.9.9 send-community both exit-address-family ! exit-address-family	neighbor 2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.9.9.9啟用 neighbor 10.9.9 send-community both 鄰居2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	

棕色靜態多點傳送複製VxLANv4到雙堆疊遷移

下表詳細說明了靜態組播複製的Brownfield VxLANv4到雙堆疊底層遷移所需的配置更改示例

遷移步驟	VXLANv4底層	雙堆疊（VxLANv4多點傳送底層）	說明
	靜態組播複製配置		
1	interface nve1 member vni 20011 mcast-group 226.1.1.1	interface nve1 member vni 20011 mcast-group 226.1.1.1 FF05::1	配置靜態IPv4和靜態IPv6複製組播地址
	底層遷移模式配置		
2		interface nve1 vxlan封裝雙堆疊優先 — ipv6底層多點傳送ipv4	必須使用「vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4」配置VXLAN NVE介面

	單播路由配置		
3		ipv6單播路由	啟用IPv6路由
	IPv6組播路由配置		
4	ip multicast-routing	ip multicast-routing ! ipv6多點傳送路由	啟用IPV4和IPv6組播路由
5	ip pim rp-address 10.9.9.9	ip pim rp-address 10.9.9.9 ! ipv6 pim rp-address2001:DB8::99:99	配置IPV4和IPv6 PIM RP

Brownfield雙協定棧到VXLANv6的無縫遷移

在將所有網路遷移到雙堆疊後，網路只能遷移到VXLANv6底層網路。要實現此目的，需要在裝置上完成此配置。

從單播雙堆疊遷移到VXLANv6

下表詳細說明了從Brownfield雙協定棧到VxLANv6僅底層遷移單播流量所需的配置更改示例

遷移步驟	雙堆疊（首選VxLANv6底層）	VXLANv6底層	說明
	VXLAN VTEP IP組態		
1	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 ipv6 address 2001:DB8:2::2/128	interface Loopback1 ipv6 address 2001:DB8:2::2/128 interface nve1 source-interface	與VXLAN關聯的環回介面，僅配置了IPv6地址

	interface nve1 source-interface Loopback1	Loopback1	
2	interface Loopback1 ip ospf 1區域0 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	interface Loopback1 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	類似OSPF的IGP僅對介面的IPv6地址啟用
	底層遷移模式配置		
3	interface nve1 vxlan封裝雙堆疊優先— ipv6	interface nve1 vxlan封裝pv6	必須為VXLANv6底層配置VXLAN NVE介面「vxlan封裝ipv6」
	IGP配置		
4	router ospf 1 ! ipv6 router ospf 1 router-id 10.1.1.1	ipv6 router ospf 1 router-id 10.1.1.1	僅對和IPv6啟用OSPF
	BGP組態		
5	router bgp 100 neighbor 10.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0	router bgp 100 neighbor 2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0	僅使用IPv6鄰居地址的BGP EVPN對等

	neighbor 2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.9.9.9啟用 neighbor 10.9.9 send-community both 鄰居 2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	! address-family l2vpn evpn 鄰居2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	

從BUM-Ingress複製雙堆疊到VXLANv6的遷移

下表詳細說明了從Brownfield Dual-Stack遷移到VxLANv6（僅支援BUM-IR遷移）所需的配置更改示例

遷移步驟	雙堆疊（首選VxLANv6底層）	VXLANv6底層	說明
1	interface Loopback1 IP 網址 10.2.2.2 255.255.255.255 ipv6 address 2001:DB8:2::2/128 interface nve1	interface Loopback1 ipv6 address 2001:DB8:2::2/128 interface nve1 source-interface Loopback1	與VXLAN關聯的環回介面，僅配置了IPv6地址

	source-interface Loopback1		
2	interface Loopback1 ip ospf 1區域0 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	interface Loopback1 ipv6 ospf 1區域0 interface nve1 source-interface Loopback1	類似OSPF的IGP僅對介面的IPv6地址啟用
	底層遷移模式配置		
3	interface nve1 vxlan封裝雙堆疊優先— ipv6	interface nve1 vxlan封裝pv6	必須為VXLANv6底層配置VXLAN NVE介面「vxlan封裝ipv6」
	IGP配置		
4	router ospf 1 ! ipv6 router ospf 1 router-id 10.1.1.1	ipv6 router ospf 1 router-id 10.1.1.1	僅為IPv6啟用OSPF
	BGP組態		
5	router bgp 100 neighbor 10.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0 neighbor	router bgp 100 neighbor 2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 !	僅使用IPv6鄰居地址的BGP EVPN對等

	2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.9.9.9啟用 neighbor 10.9.9 send-community both 鄰居 2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	address-family l2vpn evpn 鄰居2001:DB8:99::99啟用 neighbor 2001:DB8:99::99 send-community both exit-address-family	

靜態組播複製雙堆疊到VXLANv6的遷移

下表詳細說明了使用多播IPv4底層的Brownfield雙協定棧到使用多播IPv6底層的Brownfield雙協定棧進行靜態組播複製所需的配置更改示例

遷移步驟	雙堆疊（多點傳送VxLANv4底層）	雙堆疊（組播VxLANv6底層）	說明
	底層遷移模式配置		
1	interface nve1 vxlan封裝雙堆疊優先 — ipv6底層多點傳送ipv4	interface nve1 vxlan封裝雙棧優先ipv6底層注入ipv6	必須使用「vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv6」配置VXLAN NVE介面，以便仍然在V4和V6上接收組播流量，但僅在V6底層上傳送

靜態組播複製雙棧IPv6組播到IPv6組播底層遷移

下表詳細說明了將具有組播IPv6底層的Brownfield雙堆疊更改為僅支援靜態組播複製的VXLANv6底層所需的配置更改示例

遷移步驟	雙堆疊（含多點傳送VxLANv6底層）	VXLANv6底層	說明
	靜態組播複製配置		
1	interface nve1 member vni 20011 mcast-group 226.1.1.1 FF05::1	interface nve1 member vni 20011 mcast — 組FF05::1	僅配置靜態IPv6複製組播地址
	底層遷移模式配置		
2	interface nve1 vxlan封裝雙堆疊優先 — ipv6底層多點傳送ipv4	interface nve1 vxlan封裝ipv6	必須使用「vxlan encapsulation ipv6」配置VXLAN NVE介面
	IPv6組播路由配置		
3	ip multicast-routing ! ipv6多點傳送路由	ipv6多點傳送路由	僅啟用IPv6組播路由
4	ip pim rp-address 10.9.9.9 ! ipv6 pim rp-address2001:DB8::99:99	ipv6 pim rp-address2001:DB8::99:99	僅配置IPv6 PIM RP

骨幹/路由反射器遷移

路由反射器可以反射雙下一跳更新，即使不升級到17.9.2版，因為輔助下一跳地址已編碼在可選BGP傳遞隧道封裝屬性中（現有BGP實現已支援接收和反射傳遞隧道封裝屬性）。

尚未遷移到17.9.2的路由反射器/主幹能夠：

- 僅在可到達主下一跳時才反映雙下一跳更新
- 僅通過IPv4對等建立BGP鄰居關係

遷移到17.9.2的路由反射器/主幹能夠：

- 如果可訪問主要或輔助下一跳或同時訪問兩者，則反映雙下一跳更新
- 通過IPv4和IPv6對等建立BGP鄰居關係

主幹/路由反射器V4到V6 EVPN交換矩陣遷移

此表詳細說明了從V4核心遷移到V6核心所需的骨幹/RR配置更改示例

遷移步驟	V4 EVPN光纖	V6 EVPN光纖	說明
	單播路由配置		
1	ip routing	ipv6單播路由	啟用IPv6路由
	BGP組態		
2		router bgp 100 bgp router-id 10.3.3.3	設定BGP路由器ID
3	router bgp 100 neighbor 10.1.1.1 remote-as 100 neighbor 10.1.1.1 update-source Loopback0 ! address-family l2vpn evpn	router bgp 100 neighbor 2001:DB8:1::1 remote-as 100 neighbor 2001:DB8:1::1 update-source Loopback0 ! address-family l2vpn evpn 啟用鄰居2001:DB8:1::1	BGP EVPN對等已移至IPv6鄰居地址。

	鄰居10.1.1.1啟用 neighbor 10.1.1.1 send-community both exit-address-family	neighbor 2001:DB8:1::1 send-community both exit-address-family	

Brownfield主幹/路由反射器V4到V4+V6 EVPN交換矩陣遷移

此表詳細說明了從V4核心遷移到V4+V6核心所需的骨幹/RR配置更改示例

遷移步驟	V4 EVPN光纖	V4+V6 EVPN交換矩陣	說明
	單播路由配置		
1	ip routing	ip routing ipv6單播路由	啟用IPv6路由
	BGP組態		
2		router bgp 100 bgp router-id 10.3.3.3	設定BGP路由器ID
3	router bgp 100 neighbor 10.1.1.1 remote-as 100 neighbor 10.1.1.1 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.1.1.1啟用 neighbor 10.1.1.1	router bgp 100 neighbor 10.1.1.1 remote-as 100 neighbor 10.1.1.1 update-source Loopback0 neighbor 2001:DB8:1::1 remote-as 100 neighbor 2001:DB8:1::1 update-source Loopback0 ! address-family l2vpn evpn	使用IPv6和IPv6鄰居地址的BGP EVPN對等。

	send-community both exit-address-family	鄰居10.1.1.1啟用 neighbor 10.1.1.1 send-community both 啟用鄰居2001:DB8:1::1 neighbor 2001:DB8:1::1 send-community both exit-address-family	

主幹/路由反射器V4+V6到V6 EVPN交換矩陣遷移

此表詳細說明了從V4+V6核心遷移到V6核心所需的骨幹/RR配置更改示例

遷移步驟	V4+V6 EVPN交換矩陣	V6 EVPN光纖	說明
	BGP組態		
1	router bgp 100 neighbor 10.1.1.1 remote-as 100 neighbor 10.1.1.1 update-source Loopback0 neighbor 2001:DB8:1::1 remote-as 100 neighbor 2001:DB8:1::1 update-source Loopback0 ! address-family l2vpn evpn 鄰居10.1.1.1啟用 neighbor 10.1.1.1	router bgp 100 neighbor 2001:DB8:1::1 remote-as 100 neighbor 2001:DB8:1::1 update-source Loopback0 ! address-family l2vpn evpn 啟用鄰居2001:DB8:1::1 neighbor 2001:DB8:1::1 send-community both exit-address-family !	使用IPv6鄰居位址的BGP EVPN對等。

	send-community both 啟用鄰居 2001:DB8:1::1 neighbor 2001:DB8:1::1 send-community both exit-address-family		

驗證

以下各節詳細說明了用於驗證基本遷移功能的show命令。

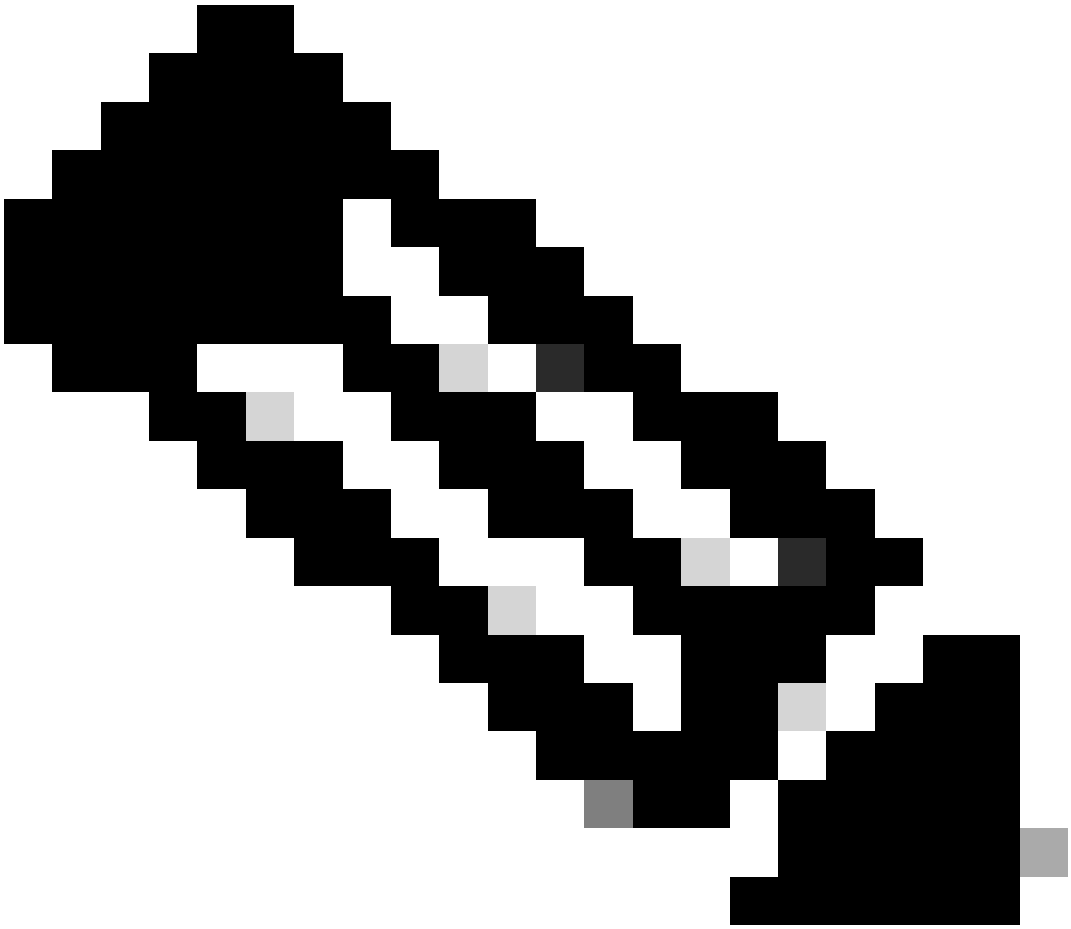

附註：有關詳細的驗證和故障排除過程，請參閱BGP VXLANv6遷移故障排除指南。（即將推出）

本地VTEP配置

綠地VXLANv6

<#root>

#show nve interface nve1 detail

Interface: nve1, State: Admin Up, Oper Up

Encapsulation: Vxlan IPv6

Multicast BUM encapsulation: Vxlan IPv6

BGP host reachability: Enabled, VxLAN dport: 4789
VNI number: L3CP 1 L2CP 6 L2DP 0

source-interface: Loopback1 (primary: 2001:DB8:1::2 vrf: 0)

tunnel interface: Tunnel0

Pkts In Bytes In Pkts Out Bytes Out
0 0 0 0

雙堆疊（首選IPv6）

<#root>

#show nve interface nve1 detail

Interface: nve1, State: Admin Up, Oper Up

Encapsulation: Vxlan dual stack prefer IPv6

Multicast BUM encapsulation: Vxlan IPv4

BGP host reachability: Enabled, VxLAN dport: 4789
VNI number: L3CP 1 L2CP 6 L2DP 0

source-interface: Loopback1 (primary: 10.1.1.2 2001:DB8:1::2 vrf: 0)

tunnel interface: Tunnel0 Tunnel1

Pkts In Bytes In Pkts Out Bytes Out
0 0 0 0

L3功能

L3 VRF VTEP

<#root>

#

show bgp l2vpn evpn local-vtep vrf red

```
   Local VTEP vrf red:
   Protocol: IPv4
     RMAC Address: AABB.CC81.F500

  VTEP-IP:10.1.1.2

     SEC-VTEP-IP:2001:DB8:1::2

     VNI: 30000
     BDI:Vlan3
   Protocol: IPv6
     RMAC Address: AABB.CC81.F500

 VTEP-IP:10.1.1.2

     SEC-VTEP-IP:2001:DB8:1::2

     VNI: 30000
     BDI:Vlan3
```

BGP EVPN路由型别5路由

源路由

<#root>

#show bgp l2vpn evpn route-type 5

```
    BGP routing table entry for [5][100:101][0][24][192.168.11.0]/17, version 127
    Paths: (1 available, best #1, table EVPN-BGP-Table)
      Advertised to update-groups:
        1
      Refresh Epoch 1
      Local, imported path from base
       0.0.0.0 (via vrf red) from 0.0.0.0 (10.1.1.1)
```

```
        Origin incomplete, metric 0, localpref 100, weight 32768, valid, external, best
        EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, local vtep: 0.0.0.0, VNI Label 30000
        Extended Community: RT:100:100 ENCAP:8 Router MAC:AABB.CC81.F500
        Tunnel Encapsulation Attribute:
         Encap type: 8
          Secondary nexthop address 2001:DB8:1::2
        rx pathid: 0, tx pathid: 0x0
        Updated on Apr 22 2022 09:28:45 PST
```

遠端路由

```
<#root>
#
show bgp l2vpn evpn route-type 5
    BGP routing table entry for [5][100:102][0][24][192.168.11.0]/17, version 164
    Paths: (1 available, best #1, table EVPN-BGP-Table)
     Not advertised to any peer
     Refresh Epoch 2
     Local
10.2.2.2
 (metric 21) (via default) from 10.9.9.9 (10.99.99.99)
--> Primary Nexthop
        Origin incomplete, metric 0, localpref 100, valid, internal, best
        EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 30000, MPLS VPN Label 0
        Extended Community: RT:100:100 ENCAP:8 Router MAC:AABB.CC81.F600
        Originator: 10.2.2.1, Cluster list: 10.9.9.9
        Tunnel Encapsulation Attribute:
         Encap type: 8
          Secondary nexthop address 2001:DB8:2::2(active)                          --> Secondary Nexthop
         rx pathid: 0, tx pathid: 0x0
         Updated on Apr 22 2022 13:02:02 PST
```

BGP L3VPN路由

來源為L3 VRF的路由

```
<#root>
```

```
#show bgp vpnv4 unicast all 192.168.11.0


     Local
        0.0.0.0 (via vrf red) from 0.0.0.0 (10.1.1.1)
          Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
          Extended Community: RT:100:100

          Local vxlan vtep:


           vrf:red, vni:30000
           local router mac:AABB.CC81.F500
           encap:4

           vtep-ip:10.2.1.2



           sec-vtep-ip:2001:DB8:2::2


           bdi:Vlan3
          mpls labels in/out 18/nolabel(red)
          rx pathid: 0, tx pathid: 0x0
          Updated on Apr 21 2022 07:43:07 PST
```

L3VRF遠端（從EVPN匯入）路由

<#root>

```
#sh bgp vpnv4 uni all 192.168.11.0


    BGP routing table entry for 100:101:192.168.11.0/24, version 24
    Paths: (3 available, best #3, table red)
     Not advertised to any peer
     Refresh Epoch 2
     Local, imported path from [5][100:102][0][24][192.168.11.0]/17 (global)


2001:DB8:2::2

 (metric 20) (via default) from 10.9.9.9 (10.99.99.99)
       Origin incomplete, metric 0, localpref 100, valid, internal
       Extended Community: RT:100:100 ENCAP:8 Router MAC:AABB.CC81.F600
       Originator: 10.2.2.1, Cluster list: 10.9.9.9
       Tunnel Encapsulation Attribute:
        Encap type: 8
          Secondary nexthop address 2001:DB8:2::2


 Local vxlan vtep:


        vrf:red, vni:30000
        local router mac:AABB.CC81.F500
        encap:4


 vtep-ip:10.1.1.2
```

```
        sec-vtep-ip:2001:DB8:1::2


        bdi:Vlan3
      Remote VxLAN:


        Topoid 0x1(vrf red)
        Remote Router MAC:AABB.CC81.F600
        Encap 8
        Egress VNI 30000
        RTEP 2001:DB8:2::2


      mpls labels in/out 18/nolabel
      rx pathid: 0, tx pathid: 0
      Updated on Apr 22 2022 13:02:02 PST
```

L3RIB IP路由

<#root>

```
#show ip route vrf red 192.168.2.0


   Routing Table: red
   Routing entry for 192.168.2.0/32, 1 known subnets
   B     192.168.2.2 [200/0]
via 2001:DB8:2::2 (red:ipv6)
, 01:08:20, Vlan3
```

<#root>

```
#show ipv6 route vrf red2001:DB8:10::/128


   Routing entry for2001:DB8:10::/128
    Known via "bgp 100", distance 200, metric 0
    Tag 10, type internal
    Route count is 1/1, share count 0
    Routing paths:


2001:DB8:3::2%
default, Vlan3%default
      Route metric is 0, traffic share count is 1
      MPLS label: nolabel
      From 2001:DB8:6363:6363::
```

```
      opaque_ptr 0x7F6945444B78
      Last updated 04:44:10 ago
```

L3FIB/CEF路由

<#root>

#

show ip cef vrf red 192.168.2.2

```
   192.168.2.2/32
```

nexthop 2001:DB8:2::2 Vlan3

#show ipv6 cef vrf red2001:DB8:10::/128

2001:10::/128

** nexthop 2001:DB8:3::2 Vlan3**

VXLANv6 L3流量轉送

<#root>

#

show ip cef vrf red 192.168.2.2

```
    192.168.2.2/32
```

**   nexthop 2001:DB8:2::2 Vlan3**

#show ipv6 cef vrf red2001:DB8:10::/128

```
     2001:10::/128
```

**   nexthop 2001:DB8:3::2 Vlan3**

#show ip interface Vlan3 stats

```
    Vlan3
```

5 minutes input rate 0 bits/sec, 0 packet/sec,

```
     5 minutes output rate 0 bits/sec, 0 packet/sec,

     0 packets input, 0 bytes,

     0 packets output, 0 bytes.
```

L2功能

L2 EVI VTEP

<#root>

```
#show l2vpn evpn evi 1 detail

   EVPN instance:       1 (VLAN Based)
   RD:                  10.1.1.3:1 (auto)
   Import-RTs:          100:1
   Export-RTs:          100:1
   Per-EVI Label:       none
   State:               Established
   Replication Type:    Ingress
   Encapsulation:       vxlan
   IP Local Learn:      Enabled (global)
   Adv. Def. Gateway:   Enabled (global)
   Re-originate RT5:    Disabled
   Adv. Multicast:      Enabled (global)
   Vlan:                11
    Protected:          False
    Ethernet-Tag:       0
    State:              Established
    Flood Suppress:     Attached
    Core If:            Vlan3
    Access If:          Vlan11
    NVE If:             nve1
    RMAC:               aabb.cc81.f500
    Core Vlan:          3
    L2 VNI:             20011
    L3 VNI:             30000

VTEP IP:         10.1.1.2

    Sec. VTEP IP:    2001:DB8:1::2

    VRF:             red
    IPv4 IRB:        Enabled
    IPv6 IRB:        Enabled
    Pseudoports:
      Ethernet0/1 service instance 11
        Routes: 1 MAC, 1 MAC/IP
```

```
    Peers:

      10.2.2.2
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD
      2001:DB8:3::2
        Routes: 1 MAC, 3 MAC/IP, 1 IMET, 0 EAD
```

BGP EVPN路由型别2路由

源路由

<#root>

```
#show bgp l2vpn evpn route-type 2


   BGP routing table entry for [2][10.1.1.3:1][0][48][001100110011][32][192.168.11.254]/24, version 132
   Paths: (3 available, best #1, table evi_1)
     Advertised to update-groups:
      1
   Refresh Epoch 1
   Local
    :: (via default) from 0.0.0.0 (10.1.1.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, multipath, best
      EVPN ESI: 00000000000000000000, Label1 20011
      Extended Community: RT:100:1 RT:100:100 ENCAP:8 EVPN DEF GW:0:0
        Router MAC:AABB.CC81.F500

      Tunnel Encapsulation Attribute:


        Encap type: 8

          Secondary nexthop address 2001:DB8:1::2(active)


      Local irb vxlan vtep:
        vrf:red, l3-vni:30000
        local router mac:AABB.CC81.F500
        core-irb interface:Vlan3

        vtep-ip:10.1.1.2



        sec-vtep-ip:2001:DB8:1::2


      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 22 2022 09:28:34 PST
    Refresh Epoch 2
```

遠端路由

<#root>

#show bgp l2vpn evpn route-type 2

```
    BGP routing table entry for [2][2.2.2.3:1][0][48][001100110011][32][192.168.11.254]/24, version 140
    Paths: (1 available, best #1, table EVPN-BGP-Table)
     Flag: 0x100
     Not advertised to any peer
     Refresh Epoch 2
     Local
      10.2.2.2 (metric 21) (via default) from 10.9.9.9 (10.99.99.99)
     <--
```

Primary Nexthop

```
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 20011
      Extended Community: RT:100:1 RT:100:100 ENCAP:8 EVPN DEF GW:0:0
       Router MAC:AABB.CC81.F600
      Originator: 10.2.2.1, Cluster list: 10.9.9.9
      Tunnel Encapsulation Attribute:
        Encap type: 8
         Secondary nexthop address 2001:DB8:2::2(active)
                         <--
```

Secondary Nexthop

```
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 22 2022 13:01:53 PST
```

L2RIB EVPN MAC路由

<#root>

#show l2route evpn mac ip

```
   EVI       ETag  Prod    Mac Address                                Host IP
----- ---------- ----- -------------- ---------------------------------------- -------------------------
    1          0   BGP 0011.0011.0011                             192.168.11.254

    1          0  L2VPN 0011.0011.0011                             192.168.11.254
```

#show l2route evpn mac ip detail

```
   EVPN Instance:            1
   Ethernet Tag:             0
   Producer Name:            BGP
   MAC Address:              0011.0011.0011
   Host IP:                  192.168.11.254
   Sequence Number:          0
   Label 2:                  0
   ESI:                      0000.0000.0000.0000.0000
```

```
   MAC Route Flags:          BInt(Brm)Dgr
   Next Hop(s):              V:20011 2001:DB8:2::2
```

```
#show l2route evpn mac mac-address 0011.0011.0011 detail
```

```
   EVPN Instance:            1
   Ethernet Tag:             0
   Producer Name:            BGP
   MAC Address:              0011.0011.0011
   Num of MAC IP Route(s):   2
   Sequence Number:          0
   ESI:                      0000.0000.0000.0000.0000
   Flags:                    BInt(Brm)
   Num of Default Gateways:  2
   Next Hop(s):              V:20011 10.1.1.2
```

L2FIB單播路由

<#root>

```
#show l2fib bridge-domain 11 detail
```

```
   Bridge Domain : 11
   Reference Count : 12
   Replication ports count : 3
   Unicast Address table size : 2
   IP Multicast Prefix table size : 1

   Flood List Information :
   Olist: 1035, Ports: 3

   Port Information :
   BD_PORT Gi1/0/1:11

   VXLAN_REP PL:22(1) T:VXLAN_REP [IR]20011:2001:DB8:2::2

   VXLAN_REP PL:18(1) T:VXLAN_REP [IR]20011:2001:DB8:3::2

   Unicast Address table information :
   aabb.0000.0021 VXLAN_UC PL:21(1) T:VXLAN_UC [MAC]20011:2001:DB8:2::2

   aabb.0000.0031 VXLAN_UC PL:17(1) T:VXLAN_UC [MAC]20011:2001:DB8:3::2
```

```
   IP Multicast Prefix table information :
   Source: *, Group: 239.21.21.21, IIF: Null, Adjacency: Olist: 6160, Ports: 1
```

```
#show l2fib path-list 17 detail

   VXLAN_UC Pathlist 17: topo 11, 1 paths, none
   ESI: 0000.0000.0000.0000.0000
   path 2001:DB8:3::2, type VXLAN, evni 20011, vni 20011, source MAC
   oce type: vxlan_header, sw_handle 0x7FA98894B318
   forwarding oce 0x7FA988AAE538 type adjacency, IPV6 midchain out of Tunnel0, addr 2001:DB8:3::2, cid:
   output chain:
   oce type: evpn_vxlan_encap, sw_handle 0x7FA988938728
   oce type: vxlan_header, sw_handle 0x7FA98894B380
   forwarding oce 0x7FA988AAE538 type adjacency,
IPV6 midchain out of Tunnel0, addr 2001:DB8:3::2,
cid: 1
```

VXLANv6 L2流量轉送

<#root>

```
#show interface Tunnel1

     Tunnel1 is up, line protocol is up
     Hardware is Tunnel
     MTU 9216 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
     Encapsulation TUNNEL, loopback not set
     Keepalive not set
     Tunnel linestate evaluation up
     Tunnel source 2001:DB8:1::2

     Tunnel protocol/transport MUDP/IPV6                <-- VXLANv6 tunnel

     TEID 0x0, sequencing disabled
     Checksumming of packets disabled
     source_port:4789, destination_port:0
     Tunnel TTL 255
     Tunnel transport MTU 9216 bytes
     Tunnel transmit bandwidth 8000 (kbps)
     Tunnel receive bandwidth 8000 (kbps)
     Last input never, output never, output hang never
     Last clearing of "show interface" counters 02:38:42
     Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 8
     Queueing strategy: fifo
```

```
     Output queue: 0/0 (size/max)
     5 minute input rate 0 bits/sec, 0 packets/sec
     5 minute output rate 0 bits/sec, 0 packets/sec


 0 packets input, 0 bytes

, 0 no buffer

     Received 0 broadcasts (0 IP multicasts)


     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort


0 packets output, 0 bytes

, 0 underruns


 Output 0 broadcasts (0 IP multicasts)


     0 output errors, 0 collisions, 0 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
```

多點傳送功能

適用於BUM-IR的BGP EVPN路由型別3路由

源路由

<#root>

#

show bgp l2vpn evpn route-type 3

```
    BGP routing table entry for [3][10.1.1.3:1][0][32][10.1.1.3]/17, version 116
    Paths: (1 available, best #1, table evi_1)
     Advertised to update-groups:
       1
     Refresh Epoch 1
     Local
      :: (via default) from 0.0.0.0 (10.1.1.1)
         Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
         Extended Community: RT:100:1 ENCAP:8 EVPN Mcast Flags:1

         Tunnel Encapsulation Attribute:


          Encap type: 8

           Secondary nexthop address 2001:DB8:1::2(active)


          PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, vni:20011 tunnel identifier: 0000 0000
```

```
          Local irb vxlan vtep:
           vrf:red, l3-vni:30000
           local router mac:AABB.CC81.F500
           core-irb interface:Vlan3
           vtep-ip:10.1.1.2
           sec-vtep-ip:2001:DB8:1::2
         rx pathid: 0, tx pathid: 0x0
         Updated on Apr 22 2022 09:28:34 PST
```

遠端路由

<#root>

```
#show bgp l2vpn evpn route-type 3
```

```
   BGP routing table entry for [3][10.2.2.3:2][0][32][10.2.2.3]/17, version 151
   Paths: (1 available, best #1, table EVPN-BGP-Table)
    Flag: 0x100
    Not advertised to any peer
    Refresh Epoch 2
    Local
```

10.2.2.2

```
 (metric 21) (via default) from 10.9.9.9 (10.99.99.99)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:2 ENCAP:8 EVPN Mcast Flags:1
      Originator: 10.2.2.1, Cluster list: 10.9.9.9
      Tunnel Encapsulation Attribute:
        Encap type: 8
          Secondary nexthop address 2001:DB8:2::2(active)
      PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, vni:20012 tunnel identifier: < Tunnel Endpoi
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 22 2022 13:01:53 PST
```

適用於BUM-IR的L2RIB EVPN IMET路由

<#root>

```
#sh l2route evpn imet  detail
```

```
   EVPN Instance:            1
```

```
   Ethernet Tag:             0
   Producer Name:            BGP
   Router IP Addr:           10.3.3.3
   Route Ethernet Tag:       0
   Tunnel Flags:             0
   Tunnel Type:              Ingress Replication
   Tunnel Labels:            20011
```

Tunnel ID: **2001:DB8:3::2**

```
   Multicast Proxy:          IGMP
   Next Hop(s):              V:0 2001:DB8:3::2
```

靜態多點傳送複製路由

<#root>

#show ipv6 mroute ff05::1

```
    Multicast Routing Table
    Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
   Timers: Uptime/Expires
   Interface state: Interface, State
```

On All VTEPS

```
   (*, FF05::1), 00:11:31/never, RP2001:DB8::99:99, flags: SCJ
    Incoming interface: TenGigabitEthernet1/1/1
    RPF nbr: FE80::822D:BFFF:FE7B:1DC8
    Immediate Outgoing interface list:
```

Tunnel0, Forward, 00:11:31/never

On Sender VTEP

(2000::1:1, FF05::1)

, 00:10:59/00:00:41, flags: SFJT

Incoming interface:

Loopback0

```
    RPF nbr: FE80::822D:BFFF:FE9B:8480
    Immediate Outgoing interface list:
     TenGigabitEthernet1/1/1, Forward, 00:10:24/00:03:08
    Inherited Outgoing interface list:
     Tunnel0, Forward, 00:11:31/never
```

On Receiver VTEP

```
   (2000::2:2, FF05::1), 00:10:34/00:00:49, flags: SJT
    Incoming interface: TenGigabitEthernet1/1/1
    RPF nbr: FE80::822D:BFFF:FE7B:1DC8
    Inherited Outgoing interface list:
```

Tunnel0,

```
 Forward, 00:11:31/never
```

VXLANv6多點傳送轉送

<#root>

#show ipv6 mfib ff05::1

```
    Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                    ET - Data Rate Exceeds Threshold, K - Keepalive
                    DDE - Data Driven Event, HW - Hardware Installed
                    ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                    MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                    MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                    e   - Encap helper tunnel flag.
    I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                    NS - Negate Signalling, SP - Signal Present,
                    A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                    MA - MFIB Accept, A2 - Accept backup,
                    RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

    Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
    Other counts:      Total/RPF failed/Other drops
    I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
    Default
```

On All VTEPS

```
   (*,FF05::1) Flags: C HW
      SW Forwarding: 0/0/0/0, Other: 0/0/0
      HW Forwarding:   1/0/277/0, Other: 0/0/0
      TenGigabitEthernet1/1/1 Flags: A NS
```

Tunnel0

, VXLAN v6 Decap Flags: F NS
        Pkts: 0/0/0   Rate: 0 pps

On Sender VTEP

   (2000::1:1,FF05::1) Flags: HW
     SW Forwarding: 2/0/257/0, Other: 0/0/0

 HW Forwarding: 698/1/174/1

, Other: 0/0/0

 Null0 Flags: A

     TenGigabitEthernet1/1/1 Flags: F NS
      Pkts: 0/0/0   Rate: 0 pps

On Receiver VTEP

   (2000::2:2,FF05::1) Flags: HW
     SW Forwarding: 1/0/259/0, Other: 0/0/0

HW Forwarding: 259/1/184/1

, Other: 0/0/0
     TenGigabitEthernet1/1/1 Flags: A

Tunnel0, VXLAN v6 Decap Flags: F NS

       Pkts: 0/0/1   Rate: 0 pps

配置示例

EVPN L2網關VXLANv4部署

```
l2vpn evpn instance 1 vlan-based
 encapsulation vxlan
 replication-type ingress
!
l2vpn evpn instance 2 vlan-based
  encapsulation vxlan
  replication-type ingress
!
l2vpn
 router-id 10.1.1.3
```

```
!
spanning-tree mode rapid-pvst
spanning-tree extend system-id
!
vlan configuration 11
 member evpn-instance 1 vni 20011
vlan configuration 12
 member evpn-instance 2 vni 20012
vlan internal allocation policy ascending
!
vlan 3,11-12
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.255
 ip ospf 1 area 0
!
interface Loopback1
 ip address 10.1.1.2 255.255.255.255
 ip ospf 1 area 0
!
interface Ethernet1/0
 no switchport
 ip address 10.0.1.2 255.255.255.252
 ip ospf network point-to-point
 ip ospf 1 area 0
!
interface nve1
 no ip address
 source-interface Loopback1
 host-reachability protocol bgp
 member vni 20011 ingress-replication
 member vni 20012 ingress-replication
!
router ospf 1
 redistribute connected
!
router bgp 100
 bgp router-id 10.1.1.1
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor 10.9.9.9 remote-as 100
 neighbor 10.9.9.9 update-source Loopback0
 !
 address-family l2vpn evpn
  neighbor 10.9.9.9 activate
  neighbor 10.9.9.9 send-community both
 exit-address-family
```

EVPN DAG（分散式任播閘道）IRB VXLANv4部署

```
vrf definition red
  rd 100:101
!
address-family ipv4
  route-target export 100:100
  route-target import 100:100
  route-target export 100:100 stitching
```

```
  route-target import 100:100 stitching
 exit-address-family
!
 address-family ipv6
  route-target export 100:200
  route-target import 100:200
  route-target export 100:200 stitching
  route-target import 100:200 stitching
 exit-address-family
!
l2vpn evpn
 default-gateway advertise
!
l2vpn evpn instance 1 vlan-based
 encapsulation vxlan
 replication-type ingress
!
l2vpn evpn instance 2 vlan-based
 encapsulation vxlan
 replication-type ingress
!
l2vpn
 router-id 10.1.1.3
!
spanning-tree mode rapid-pvst
spanning-tree extend system-id
!
vlan configuration 3
 member vni 30000
vlan configuration 11
 member evpn-instance 1 vni 20011
vlan configuration 12
 member evpn-instance 2 vni 20012
vlan internal allocation policy ascending
!
vlan 3,11-12
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.255
 ip ospf 1 area 0
!
interface Loopback1
 ip address 10.1.1.2 255.255.255.255
 ip ospf 1 area 0
!
interface Loopback192
 vrf forwarding red
 ip address 192.168.1.1 255.255.255.255
 ip pim sparse-mode
 !
interface Ethernet1/0
 no switchport
 ip address 10.0.1.2 255.255.255.252
 ip pim sparse-mode
 ip ospf network point-to-point
 ip ospf 1 area 0
!
interface nve1
 no ip address
 source-interface Loopback1
 host-reachability protocol bgp
 member vni 30000 vrf red
```

```
 member vni 20011 ingress-replication
 member vni 20012 ingress-replication
!
router ospf 1
 redistribute connected
!
router bgp 100
 bgp router-id 10.1.1.1
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor 10.9.9.9 remote-as 100
 neighbor 10.9.9.9 update-source Loopback0
 !
 address-family l2vpn evpn
  neighbor 10.9.9.9 activate
  neighbor 10.9.9.9 send-community both
 exit-address-family
 !
 address-family ipv4 vrf red
  advertise l2vpn evpn
  redistribute connected
  redistribute static
 exit-address-family
 !
 address-family ipv6 vrf red
  redistribute connected
  advertise l2vpn evpn
 exit-address-family
```

相關資訊

- BGP EVPN VXLAN設定指南
- BGP通道封裝屬性(rfc9012)
- BGP VXLANv6遷移故障排除指南，瞭解詳細的驗證和故障排除過程。（即將推出）
- 技術支援與文件 - Cisco Systems

關於此翻譯

思科已使用電腦和人工技術翻譯本文件，讓全世界的使用者能夠以自己的語言理解支援內容。請注意，即使是最佳機器翻譯，也不如專業譯者翻譯的內容準確。Cisco Systems, Inc. 對這些翻譯的準確度概不負責，並建議一律查看原始英文文件（提供連結）。