在Catalyst 9000交换机上将EVPN VxLAN迁移到IPv6底层

目录

简介

本文档介绍如何将EVPN VxLAN迁移到Catalyst 9000系列交换机上的IPv6底层。

先决条件

要求
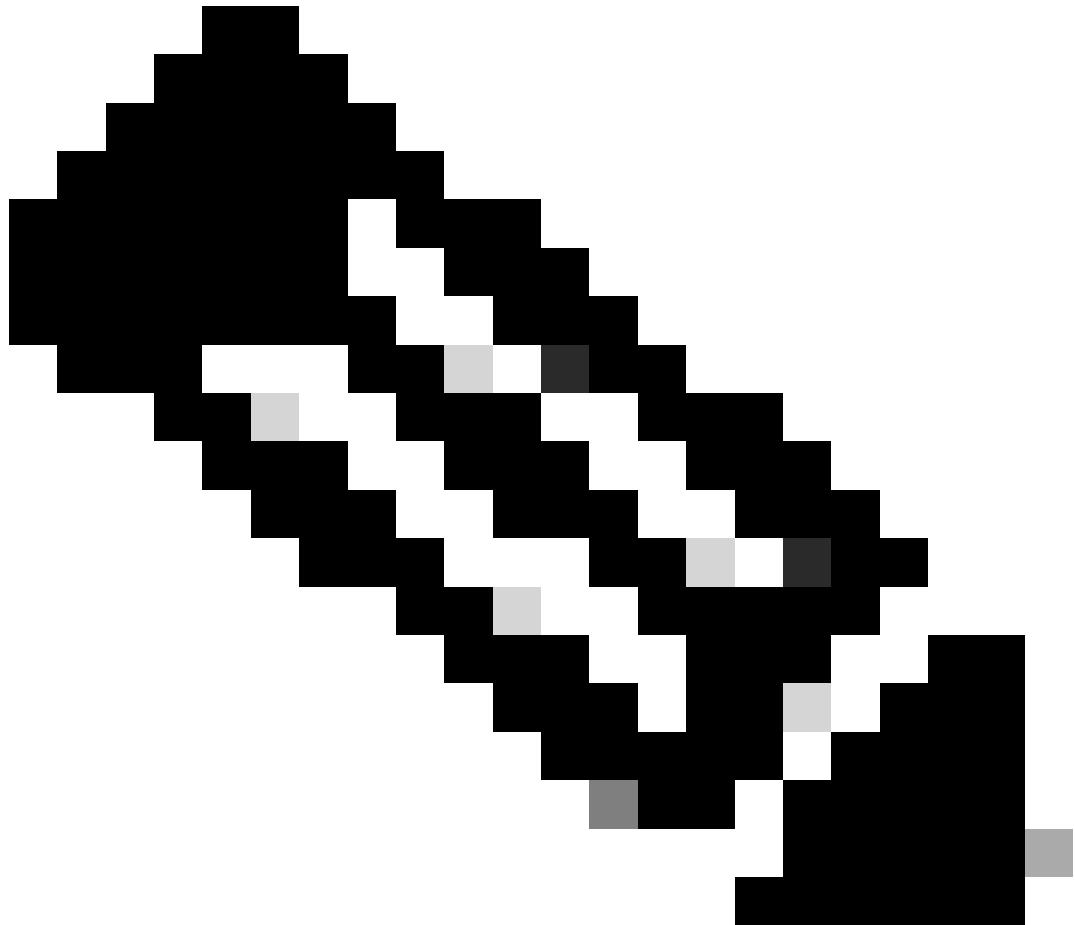
Cisco 建议您了解以下主题：

- 单播EVPN VxLAN功能、BGP和MVPN（组播虚拟专用网络）。
- IPv4和IPv6单播
- 组播概念和组播如何运行

使用的组件

本文档中的信息基于以下软件和硬件版本：

- Catalyst 9000 系列交换机

注意：9200、9500X和9600X不支持VXLANv6

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

背景信息

迁移到EVPN VXLANv6需要更改EVPN交换矩阵中的某些配置以启用IPv6基础。本文档详细介绍将现有EVPN VXLANv4部署迁移到绿地（仅VXLANv6）或棕地（双堆栈 — VXLANv4和VXLANv6）部署的相关配置更改和验证过程。

绿地EVPN VXLANv6部署需要：

- IPv6核心
- 将EVPN交换矩阵迁移到VXLANv6底层支持
- 将BGP EVPN邻居迁移到IPv6邻居对等

灰场EVPN VXLAN部署需要：

- IPv4 + IPv6核心
- 将EVPN交换矩阵无缝迁移到双堆栈(VXLANv4 + VXLANV6)底层
- BGP EVPN邻居对等从IPv4无缝迁移到IPv6邻居地址

术语

EVPN	以太网虚拟专用网络	允许BGP传输第2层MAC和第3层IP信息的扩展是EVPN，它使用多协议边界网关协议(MP-BGP)作为协议来分发属于VXLAN重叠网络的可达性信息。
VXLAN	虚拟可扩展LAN（局域网）	VXLAN旨在克服VLAN和STP的固有局限性。推荐的IETF标准[RFC 7348]提供与VLAN相同的以太网第2层网络服务，但灵活性更高。功能上，它是MAC-in-UDP封装协议，在第3层底层网络上作为虚拟重叠运行。
VTEP	虚拟隧道终端	这是执行封装和解封的设备
EVI	EVPN实例	EVPN实例(EVI)由虚拟网络标识符(VNI)表示。 EVI表示PE路由器上的VPN。它充当IP VPN路由和转发(VRF)的相同角色，并且EVI被分配了导入/导出路由目标(RT)
NVE	网络虚拟接口	进行封装和解封的逻辑接口
VNI	VXLAN网络标识符	唯一标识每个第2层子网或网段。VNI有两种类型： 对称(L2VNI):VTEP具有相同的VNI 非对称(L3VNI):VTEP没有相同的VNI并通过单个通用VNI路由。
BUM	广播、未知单播、组播	BUM流量通过与NVE配置下的VNI关联的Mcast组发送。
TRM	租户路由组播	基于BGP-EVPN的解决方案，可在VxLAN交换矩阵中连接到VTEPS的源和接收器之间启用组播路由[RFC7432]。L2TRM（第2层TRM）和L3TRM（第3层TRM）有两种类型
MDT	组播分布树	在VTEP之间构建的组播树，用于租户组播流量的封装和隧道。

PVLAN	专用VLAN	将VLAN的以太网广播域划分为子域，这样您就可以将交换机上的端口彼此隔离。
MIB	管理信息库	A 简单网络管理协议(SNMP)监控器对象
PIM-BIDIR	协议无关组播双向	一种PIM，其中流量仅沿着共享树转发 根植于组的交汇点(RP)。
VFI	虚拟转发实例	虚拟网桥端口，能够根据目的MAC地址、源MAC地址学习和老化等执行本地桥接功能，例如转发。
IRB	集成路由和桥接	启用第2层VPN和第3层VPN重叠，允许重叠中的终端主机在同一子网内和VPN内的不同子网之间相互通信。
IMET	包含组播以太网标记	也称为BGP路由类型3(RT3)，用于自动发现远程对等体，以便在VXLAN上设置BUM隧道。IMET路由承载从远程对等体通告的远程（出口）VNI，这可能与本地VNI不同。这些远程VNI称为下游分配的VNI。
DAG	分布式任播网关	所有VTEP的默认网关功能。相同的网关IP存在于所有VTEP上，并允许交换矩阵中的移动性。

限制

- 仅Cat9k交换机支持无缝迁移
- 仅考虑一个NVE接口和全局迁移

这些EVPN功能不支持VXLANv6基础

- 集中网关
- 多宿主支持
- L3组播(TRM)
- 带有入口复制的L2TRM
- 使用默认MDT的L2TRM（组播复制）
- 使用默认MDT的L3TRM
- L3TRM与数据MDT
- 边界网关（多站点）
- 接入VFI
- PVLAN
- MIB
- 用于组播底层的PIM-BIDIR

无缝迁移概念概述

灰场EVPN VXLAN部署需要逐步将网络从VXLANv4迁移到VXLANv6底层。要实现此EVPN，VXLAN网络需要逐步从IPv4迁移到IPv6底层并允许部分EVPN网络迁移到IPv6底层而网络的其他部分继续使用IPv4底层但网络中的所有节点仍然可以连接。

要实现单播和BUM（广播、未知单播和组播）入口复制的无缝迁移，EVPN节点需要支持双栈VTEP。双堆栈VTEP节点有两个VTEP地址（IPv4和IPv6）与同一VNI（VXLAN网络标识符）关联。在底层迁移期间，这两个VTEP IP地址在单个BGP EVPN更新(BGP EVPN Dual-Next-hop更新)中通告给对等体，并为接收节点提供选项以选择底层中的任何一个进行流量转发。

BGP EVPN双下一跳更新通告

BGP双下一跳更新传输两个下一跳：

- MP_REACH_NLRI(EVPN Routetype-2/Routetype-5)/PMSI-tunnel(EVPN Routetype-3)属性中的主下一跳（现有底层）
- BGP隧道封装属性中的辅助下一跳（迁移基础）(23)

作为主要和辅助节点传送的VTEP IP取决于EVPN节点的迁移模式。

此表详细说明了双下一跳更新中携带的主要/辅助VTEP IP

迁移模式	主要下一跳	辅助下一跳
VXLANv4到VXLANv6	IPv4 VTEP	IPv6 VTEP
VXLANv6到VXLANv4	IPv6 VTEP	IPv4 VTEP

BGP枝叶/边缘EVPN双下一跳更新处理

接收此BGP EVPN双下一跳更新的枝叶/边缘/边界节点使用其中一个收到的下一跳作为远程VTEP进行转发。用于衬底的下一跳取决于设备上配置的这些迁移策略。

- 本地VTEP地址
- 本地衬底偏好

此表详细说明了本地配置的策略如何决定使用哪个底层来转发数据包

已收到BGP更新	本地VTEP 地址	本地衬底偏好	VXLAN基础 单播/BUM-IR

双下一跳(IPv4 + IPv6)	仅IPv4 VTEP	不适用	VXLANv4
双下一跳(IPv4 + IPv6)	仅IPv6 VTEP	不适用	VXLANv6
双下一跳(IPv4 + IPv6)	双堆栈(IPv4 + IPv6 VTEP IP)	IPv4	VXLANv4
双下一跳(IPv4 + IPv6)	双堆栈(IPv4 + IPv6 VTEP IP)	IPv6	VXLANv6
单一IPv4下一跳	仅IPV4 VTEP	不适用	VXLANv4
单一IPv4下一跳	仅IPV6 VTEP	不适用	无VXLAN基础
单一IPv4下一跳	双堆栈(IPv4 + IPv6 VTEP IP)	不适用	VXLANv4
单一IPv6下一跳	仅IPV4 VTEP	不适用	无VXLAN基础
单一IPv6下一跳	仅IPV6 VTEP	不适用	VXLANv6
单一IPv6下一跳	双堆栈(IPv4 + IPv6 VTEP IP)	不适用	VXLANv6

配置（VXLAN底层迁移模式）

"interface nve"配置下的新cli命令可用于设置VXLAN底层迁移模式以及单播和组播的底层首选项。

单播和BUM — 入口复制的迁移模式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
```

下表详细说明了单播和BUM-IR迁移模式的CLI配置

CLI 配置	本地VTEP IP和 单播/BUM-IR衬底
int nve 1 vxlan encapsulation ipv4 （这是可选的，因为默认vxlan封装是ipv4）	IPv4（VXLANv4底层）
int nve 1 vxlan encapsulation ipv6	IPv6（VXLANv6底层）
int nve 1 vxlan encapsulation dual-stack prefer-ipv4	双堆栈(IPv4 + IPv6)(首选VXLANv4底层)
int nve 1 vxlan encapsulation dual-stack prefer-ipv6	双堆栈(IPv4 + IPv6)(首选VXLANv6底层)

静态组播复制的迁移模式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	静态组播底层

int nve 1
 	
    member vni \<L2VNI> mcast-group \<v4-mcast-group>
 	
    vxlan encapsulation ipv4
 	
（这是可选的，因为默认vxlan封装是ipv4）
	在为L2VNI配置的IPv4底层组播组上发送和接收组播流量

int nve 1
 	
    member vni \<L2VNI> mcast-group \<v6-mcast-group>
 	
    vxlan encapsulation ipv6
	在为L2VNI配置的IPv6底层组播组上发送和接收组播流量

int nve 1
 	
    member vni \<L2VNI> mcast-group \<v4-mcast-group> \<v6-mcast-group>
 	
    vxlan encapsulation dual-stack prefer-ipv6
	双堆栈(IPv4 +IPv6)
	
	在为L2VNI配置的IPv4和IPv6底层组播组上接收组播流量
	
	仅在为L2VNI配置的IPv4底层组播组上发送组播流量

int nve 1
 	
    member vni \<L2VNI> mcast-group \<v4-mcast-group> \<v6-mcast-group>
 	
    vxlan encapsulation dual-stack prefer-ipv4
	双堆栈(IPv4 +IPv6)
	
	在为L2VNI配置的IPv4和IPv6底层组播组上接收组播流量
	
	仅在为L2VNI配置的IPv6底层组播组上发送组播流量

int nve 1 member vni \<L2VNI> mcast-group \<v4-	双堆栈(IPv4 +IPv6) 在为L2VNI配置的IPv4和IPv6底层组播组上接收组播

mcast-group> <v6-mcast-group>
 	
  vxlan encapsulation dual-stack prefer-ipv6
 	
  底层组播IPV4
	流量
	
	仅在为L2VNI配置的IPv4底层组播组上发送组播流量

int nve 1
 	
  member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group>
 	
  vxlan encapsulation dual-stack prefer-ipv4 underlay-mcast ipv6
	双堆栈(IPv4 +IPv6)
	
	在为L2VNI配置的IPv4和IPv6底层组播组上接收组播流量
	
	仅在为L2VNI配置的IPv6底层组播组上发送组播流量

int nve 1
 	
  member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group>
 	
  vxlan encapsulation dual-stack prefer-ipv6
 	
  底层组播IPV6
	双堆栈(IPv4 +IPv6)
	
	在为L2VNI配置的IPv4和IPv6底层组播组上接收组播流量
	
	仅在为L2VNI配置的IPv6底层组播组上发送组播流量

int nve 1
 	
  member vni <L2VNI> mcast-group <v4-mcast-group> <v6-mcast-group>
 	
  vxlan encapsulation dual-stack prefer-ipv4 underlay-mcast ipv4
	双堆栈(IPv4 +IPv6)
	
	在为L2VNI配置的IPv4和IPv6底层组播组上接收组播流量
	
	仅在为L2VNI配置的IPv4底层组播组上发送组播流量

基础迁移程序
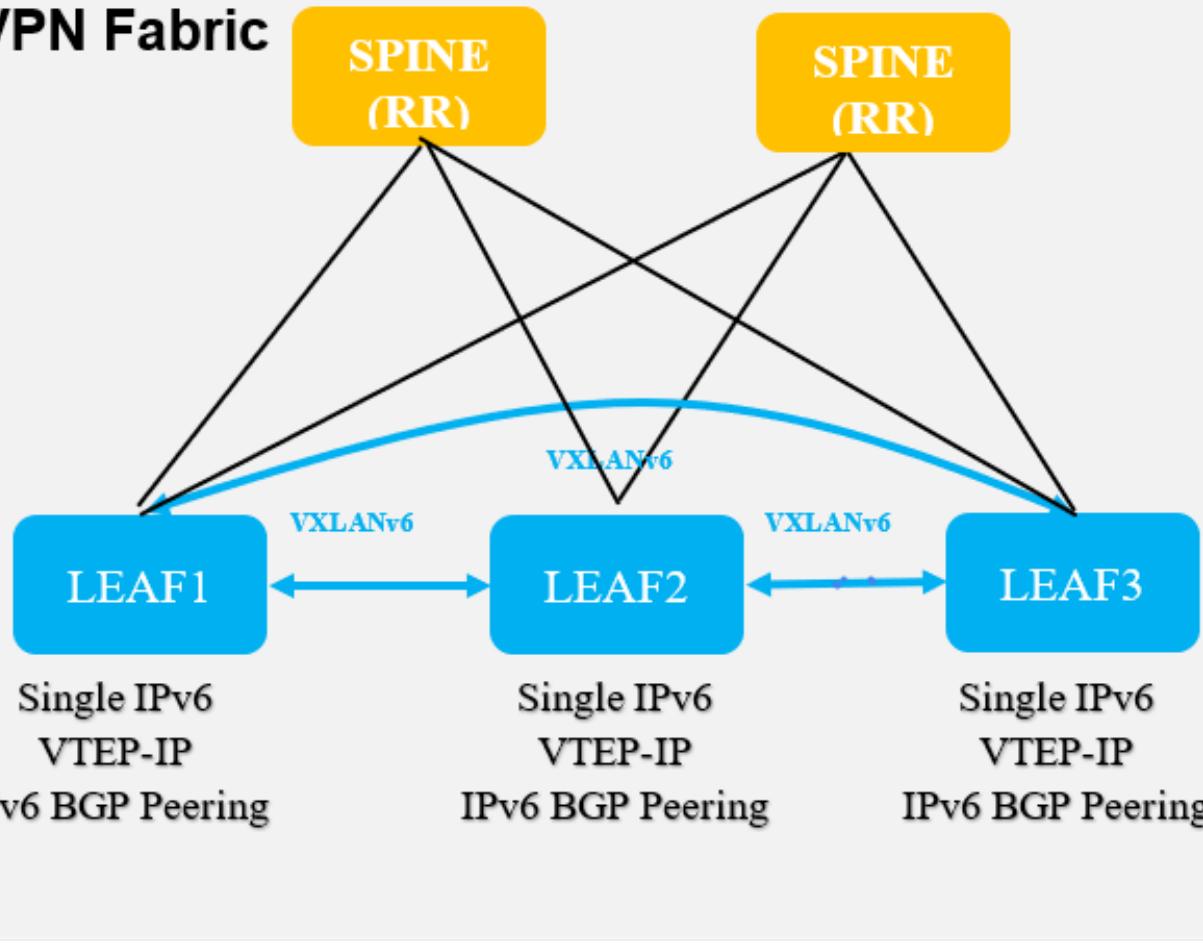
EVPN L2网关和EVPN IRB（分布式任播网关）部署的底层迁移步骤相同

VXLANv4至VXLANv6迁移

VXLANv6部署在下层中有一个IPv6传输。VXLAN隧道和BGP邻居关系均基于IPv6。

网络图

EVPN Fabric

VxLANv4到VxLANv6的单播迁移

下表详细说明了单播流量的VxLANv4到VXLANv6底层迁移所需的配置更改示例。

迁移步骤	VXLANv4底层	VXLANv6底层	描述
	EVPN路由器ID配置		
1		l2vpn router-id 10.1.1.1	配置l2vpn router-id以用作EVPN路由器ID
	VXLAN VTEP IP配置		
2	interface Loopback1 ip address 10.2.2.2 255.255.255.255	interface Loopback1 ipv6地址 2001:DB8:2::2/128	与使用IPv6地址配置的VXLAN关联的环回接口。此IPV6地址用于VXLAN的本地IPv6 VTEP。

	interface nve1
	
	  source-interface Loopback1
		interface nve1
		
		  source-interface Loopback1	

3	interface Loopback1
	
	ip ospf 1 area 0
	
	interface nve1
	
	  source-interface Loopback1
		interface Loopback1
		
		  ipv6 ospf 1区域0
		
		interface nve1
		
		  source-interface Loopback1
			为接口的IPv6地址启用类似IGP的OSPF

	底层迁移模式配置		
4		interface nve1 vxlan encapsulation ipv6	必须使用"vxlan encapsulation ipv6"配置VXLANv6底层配置VXLAN NVE接口
	单播路由配置		
5		ipv6 unicast-routing	启用IPv6路由
	IGP配置		
6	router ospf 1	ipv6 router ospf 1 router-id 10.1.1.1	启用OSPF for IPv6
	BGP配置		
7		router bgp 100 bgp router-id 10.2.2.1	配置BGP路由器ID
8	router bgp 100 neighbor 10.99.99.99	router bgp 100 neighbor	BGP EVPN对等已移至IPv6邻居地址

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

从BUM入口复制VxLANv4迁移到VxLANv6

下表详细列出了BUM-IR的VxLANv4到VXLANv6底层迁移所需的配置更改示例

迁移步骤	VXLANv4底层	VXLANv6底层	描述
	EVPN路由器ID配置		
1		l2vpn router-id 10.1.1.1	将l2vpn router-id配置为EVPN路由器id
	VXLAN VTEP IP配置		
2	interface Loopback1 ip address 10.2.2.2 255.255.255.255 interface nve1	interface Loopback1 ipv6地址 2001:DB8:2::2/128 interface nve1	与使用IPv6地址配置的VXLAN关联的环回接口。此IPV6地址用于VXLAN的本地IPv6 VTEP

	source-interface Loopback1	source-interface Loopback1	
3	interface Loopback1 ip ospf 1 area 0 interface nve1 source-interface Loopback1	interface Loopback1 ipv6 ospf 1区域0 interface nve1 source-interface Loopback1	为接口的IPv6地址启用类似IGP的OSPF
	底层迁移模式配置		
4		interface nve1 vxlan encapsulation ipv6	必须使用"vxlan encapsulation ipv6"配置VXLANv6底层配置VXLAN NVE接口
	单播路由配置		
5		ipv6 unicast-routing	启用IPv6路由
	IGP配置		
6	router ospf 1	ipv6 router ospf 1 router-id 10.1.1.1	启用OSPF for IPv6
	BGP配置		
7		router bgp 100 bgp router-id 10.2.2.1	配置BGP路由器ID
8	router bgp 100 neighbor 10.9.9.9	router bgp 100 neighbor	BGP EVPN对等已移至IPv6邻居地址

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

静态组播复制VxLANv4到VxLANv6迁移

下表详细说明了静态组播复制的VxLANv4到VXLANv6底层迁移所需的配置更改示例

迁移步骤	VXLANv4底层	VXLANv6底层	描述
	静态组播复制配置		
1	interface nve1 member vni 20011 mcast-group 226.1.1.1	interface nve1 member vni 20011 mcast — 组FF05::1	配置静态IPv6复制组播地址
	底层迁移模式配置		
2		interface nve1 vxlan encapsulation ipv6	必须使用"vxlan encapsulation ipv6"配置VXLANv6底层配置VXLAN NVE接口
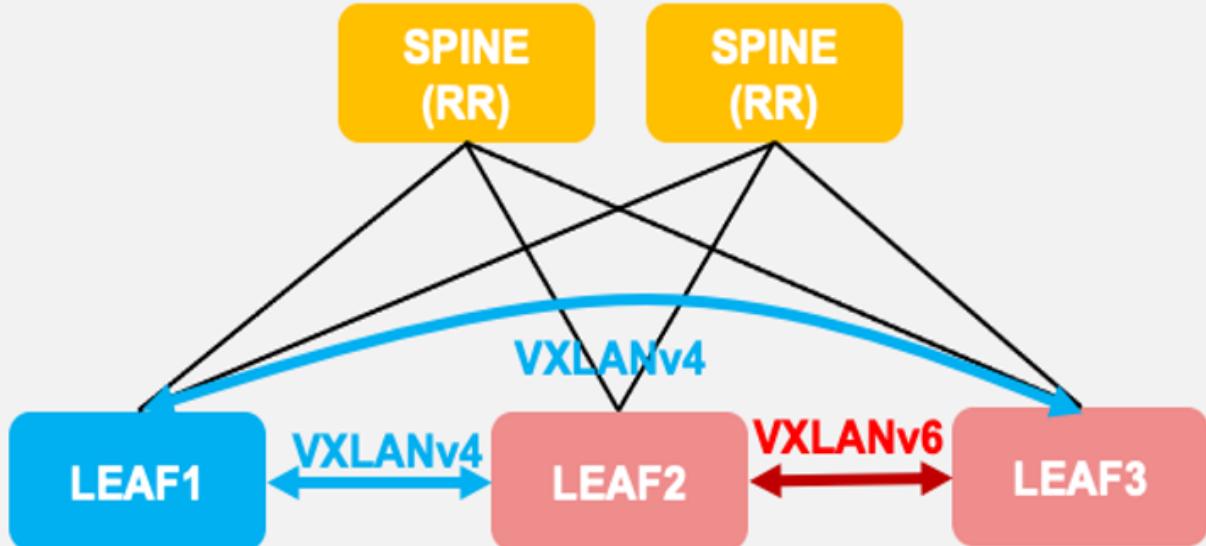
	单播路由配置		
3		ipv6 unicast-routing	启用IPv6路由
	组播路由配置		
4	ip multicast-routing	ipv6组播路由	启用IPv6组播路由
5	ip pim rp-address 10.9.9.9	ipv6 pim rp-address 2001:DB8::99:99	将PIM RP地址迁移到IPv6

Brownfield - VXLANv4和VXLANv6无缝迁移

灰场部署具有基础上的可传递的双IPv4/IPv6传输，可实现无缝迁移。VXLAN隧道和BGP邻居最初基于IPv4，然后无缝迁移至基于IPv6（迁移后，可以选择从底层删除IPv4）。 换句话说，单个VTEP可以迁移到双IPv4和IPv6，而其他的VTEP继续使用IPv4。一旦交换矩阵内的所有VTEP都支持双IPv4和IPv6，单个VTEP现在就可以迁移到IPv6。

网络图

灰场单播VxLANv4到双协议栈迁移

下表详细说明了单播流量的Brownfield VxLANv4到双栈底层迁移所需的配置更改示例

迁移步骤	VXLANv4底层	双堆栈（首选VxLANv6底层）	描述
	L2VPN路由器ID配置		
1		l2vpn router-id 10.2.2.3	将l2vpn router-id配置为EVPN路由器id
	VXLAN VTEP IP配置		
2	interface Loopback1 ip address 10.2.2.2 255.255.255.255 interface nve1 source-interface Loopback1	interface Loopback1 ip address 10.2.2.2 255.255.255.255 ipv6地址2001:DB8:2::2/128 interface nve1 source-interface Loopback1	与VXLAN关联的环回接口，使用IPv4和IPv4地址进行配置。
3	interface Loopback1 ip ospf 1 area 0 interface nve1 source-interface Loopback1	interface Loopback1 ip ospf 1 area 0 ipv6 ospf 1区域0 interface nve1 source-interface Loopback1	为接口的IPv4和IPv6地址启用类似IGP的OSPF
	底层迁移模式配置		
4		interface nve1	VXLAN NVE接口必须为双堆栈配置"vxlan encapsulation dual-stack

		vxlan encapsulation dual-stack prefer-ipv6	prefer-ipv6", 但必须优先配置VXLANv6底层
	单播路由配置		
6		ipv6 unicast-routing	启用IPv6路由
	IGP配置		
7	router ospf 1	router ospf 1 ! ipv6 router ospf 1 router-id 10.1.1.1	为IPv4和IPv6启用OSPF
	BGP配置		
8		router bgp 100 bgp router-id 10.2.2.1	配置BGP路由器ID

9	router bgp 100
	 neighbor 10.9.9.9 remote-as 100
	 neighbor 10.9.9.9 update-source Loopback0
	!
	 address-family l2vpn evpn
	  邻居10.9.9.9激活
	 neighbor 10.9.9.9 send-community both
	exit-address-family
		router bgp 100
		 neighbor 10.9.9.9 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
		  邻居10.9.9.9激活
			使用IPv4和IPv6邻居地址的BGP EVPN对等

! exit-address-family	neighbor 10.9.9.9 send-community both 邻居2001:DB8:99::99激活 neighbor 2001:DB8:99::99 send-community both exit-address-family	

Brownfield BUM入口复制VxLANv4到双堆栈迁移

下表详细说明了从Brownfield VxLANv4迁移到BUM-IR的双堆栈底层所需的配置更改示例

迁移步骤	VXLANv4底层	双堆栈（首选VxLANv6底层）	描述
	L2VPN路由器ID配置		
1		l2vpn router-id 10.2.2.3	将l2vpn router-id配置为EVPN路由器id
	VXLAN VTEP IP配置		
2	interface Loopback1 ip address 10.2.2.2 255.255.255.255 interface nve1 source-interface Loopback1	interface Loopback1 ip address 10.2.2.2 255.255.255.255 ipv6地址2001:DB8:2::2/128 interface nve1 source-interface Loopback1	与VXLAN关联的环回接口，使用IPv4和IPv6地址进行配置。
3	interface Loopback1 ip ospf 1 area 0 interface nve1	interface Loopback1 ip ospf 1 area 0 ipv6 ospf 1区域0	为接口的IPv4和IPv6地址启用类似IGP的OSPF

		source-interface Loopback1	interface nve1 source-interface Loopback1	
	底层迁移模式配置			
4			interface nve1 vxlan encapsulation dual-stack prefer-ipv6	VXLAN NVE接口必须为双堆栈配置"vxlan encapsulation dual-stack prefer-ipv6"，但必须配置prefer VXLANV6底层
	单播路由配置			
5			ipv6 unicast-routing	启用IPv6路由
	IGP配置			
6	router ospf 1		router ospf 1 ipv6 router ospf 1 router-id 10.1.1.1	为IPv4和IPv6启用OSPF
	BGP配置			
7			router bgp 100 bgp router-id 10.2.2.1	配置BGP路由器ID

8	router bgp 100
	
	neighbor 10.9.9.9 remote-as 100
	
	  neighbor 10.9.9.9 update-source Loopback0
	
	!
			router bgp 100
			
			  neighbor 10.9.9.9 remote-as 100
			
			  neighbor 10.9.9.9 update-source Loopback0
			
			neighbor 2001:DB8:99::99 remote-as 100
				使用IPv4和IPv6邻居地址的BGP EVPN对等

address-family l2vpn evpn 邻居10.9.9.9激活 neighbor 10.9.9.9 send-community both exit-address-family ! exit-address-family	neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 邻居10.9.9.9激活 neighbor 10.9.9.9 send-community both 邻居2001:DB8:99::99激活 neighbor 2001:DB8:99::99 send-community both exit-address-family	

灰场静态组播复制VxLANv4到双堆栈迁移

下表详细说明了将Brownfield VxLANv4迁移到静态组播复制的双堆栈底层所需的配置更改示例

迁移步骤	VXLANv4底层	双堆栈（VxLANv4组播底层）	描述
	静态组播复制配置		
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	配置静态IPv4和静态IPv6复制组播地址
	底层迁移模式配置		
2		interface nve1 vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4	必须使用"vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4"配置VXLAN NVE接口
	单播路由配置		

3		ipv6 unicast-routing	启用IPv6路由
	IPv6组播路由配置		
4	ip multicast-routing	ip multicast-routing ! ipv6组播路由	启用IPV4和IPv6组播路由
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双协议栈到VXLANv6的无缝迁移

只有在所有网络都迁移到双协议栈之后，才能将网络迁移到VXLANv6底层网络。为实现此目的，需要在设备上完成此配置。

单播双协议栈到VXLANv6的迁移

下表详细说明了仅支持单播流量迁移的Brownfield双协议栈到VxLANv6所需的配置更改示例

迁移步骤	双堆栈（首选VxLANv6底层）	VXLANv6底层	描述
	VXLAN VTEP IP配置		
1	interface Loopback1 ip address 10.2.2.2 255.255.255.255 ipv6地址 2001:DB8:2::2/128 interface nve1 source-interface	interface Loopback1 ipv6地址2001:DB8:2::2/128 interface nve1 source-interface Loopback1	与VXLAN关联的环回接口，仅配置了IPv6地址

	Loopback1		

2	interface Loopback1
	
	 ip ospf 1 area 0
	
	 ipv6 ospf 1区域0
	
	interface nve1
	
	  source-interface Loopback1
		interface Loopback1
		
		  ipv6 ospf 1区域0
		
		interface nve1
		
		  source-interface Loopback1
			类似IGP的OSPF仅对接口的IPv6地址启用

	底层迁移模式配置		

3	interface nve1
	
	  vxlan encapsulation dual-stack prefer-ipv6
		interface nve1
		
		  vxlan encapsulation pv6
			VXLAN NVE接口必须配置为"vxlan encapsulation ipv6"，适用于VXLANv6基础

	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
			仅对和IPv6启用OSPF

	BGP配置		

5	router bgp 100
	
	neighbor 10.9.9.9 remote-as 100
	
	  neighbor 10.9.9.9 update-source Loopback0
	
	  neighbor 2001:DB8:99::99
		router bgp 100
		
		neighbor 2001:DB8:99::99 remote-as 100
		
		  neighbor 2001:DB8:99::99 update-source Loopback0
		
		!
		
		  address-family l2vpn evpn
			仅使用IPv6邻居地址的BGP EVPN对等

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

从BUM-Ingress复制双协议栈到VXLANv6的迁移

下表详细说明了从Brownfield Dual-Stack到VxLANv6（仅支持BUM-IR迁移）所需的配置更改示例

迁移步骤	双堆栈（首选 VxLANv6底层）	VXLANv6底层	描述
1	interface Loopback1 ip address 10.2.2.2 255.255.255.255 ipv6地址 2001:DB8:2::2/128 interface nve1 source-interface Loopback1	interface Loopback1 ipv6地址2001:DB8:2::2/128 interface nve1 source-interface Loopback1	与VXLAN关联的环回接口，仅配置了IPv6地址

2	interface Loopback1 ip ospf 1 area 0 ipv6 ospf 1区域0 interface nve1 source-interface Loopback1	interface Loopback1 ipv6 ospf 1区域0 interface nve1 source-interface Loopback1	类似IGP的OSPF仅对接口的IPv6地址启用
	底层迁移模式配置		
3	interface nve1 vxlan encapsulation dual-stack prefer-ipv6	interface nve1 vxlan encapsulation pv6	VXLAN NVE接口必须配置为"vxlan encapsulation ipv6"，适用于VXLANv6基础
	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	仅启用IPv6的OSPF
	BGP配置		
5	router bgp 100 neighbor 10.9.9.9 remote-as 100 neighbor 10.9.9.9 update-source Loopback0 neighbor 2001:DB8:99::99 remote-as 100 neighbor	router bgp 100 neighbor 2001:DB8:99::99 remote-as 100 neighbor 2001:DB8:99::99 update-source Loopback0 ! address-family l2vpn evpn 邻居2001:DB8:99::99激活	仅使用IPv6邻居地址的BGP EVPN对等

	2001:DB8:99::99 update-source Loopback0
	
	 !
	
	 address-family l2vpn evpn
	
	  邻居10.9.9.9激活
	
	  neighbor 10.9.9.9 send-community both
	
	  邻居 2001:DB8:99::99激活
	
	  neighbor 2001:DB8:99::99 send-community both
	
	  exit-address-family
		neighbor 2001:DB8:99::99 send-community both
		
		 exit-address-family	

静态组播复制双协议栈到VXLANv6的迁移

下表详细说明了具有组播IPv4底层的灰场双协议栈到具有组播IPv6底层的灰场双协议栈进行静态组播复制所需的配置更改示例

迁移步骤	双堆栈（组播VxLANv4底层）	双堆栈（组播VxLANv6底层）	描述
	底层迁移模式配置		
1	interface nve1 vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4	interface nve1 vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv6	必须使用"vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv6"配置VXLAN NVE接口，才能在V4和V6上仍接收组播流量，但仅在V6下发组播流量

静态组播复制双栈IPv6组播到IPv6组播底层迁移

下表详细说明了具有组播IPv6基础的Brownfield双协议栈到仅支持VXLANv6的静态组播复制所需的配置更改示例

迁移步骤	双堆栈（带组播 VxLANv6底层）	VXLANv6底层	描述
	静态组播复制配置		
1	interface nve1 member vni 20011 mcast-group 226.1.1.1 FF05::1	interface nve1 member vni 20011 mcast — 组FF05::1	仅配置静态IPv6复制组播地址
	底层迁移模式配置		
2	interface nve1 vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4	interface nve1 vxlan encapsulation ipv6	必须使用"vxlan encapsulation ipv6"配置VXLAN NVE接口
	IPv6组播路由配置		
3	ip multicast-routing ! ipv6组播路由	ipv6组播路由	仅启用IPv6组播路由
4	ip pim rp-address 10.9.9.9 ! ipv6 pim rp-address2001:DB8::99:99	ipv6 pim rp-address2001:DB8::99:99	仅配置IPv6 PIM RP

主干/路由反射器迁移

路由反射器可以反映双下一跳更新，即使不升级到17.9.2版本，因为辅助下一跳地址已编码在可选BGP传递隧道封装属性中（现有BGP实现已支持接收并反映传递隧道封装属性）。

尚未迁移到17.9.2的路由反射器/主干能够：

- 仅在可到达主下一跳时反映双下一跳更新
- 只有通过IPv4对等的BGP邻居关系

迁移到17.9.2的路由反射器/主干能够：

- 如果可以到达主或辅助下一跳或两者，则反映双下一跳更新
- 通过IPv4和IPv6对等建立BGP邻居关系

主干/路由反射器V4至V6 EVPN交换矩阵迁移

下表详细说明了从V4核心迁移到V6核心所需的主干/路由反射器配置更改示例

迁移步骤	V4 EVPN交换矩阵	V6 EVPN交换矩阵	描述
	单播路由配置		
1	ip routing	ipv6 unicast-routing	启用IPv6路由
	BGP配置		
2		router bgp 100 bgp router-id 10.3.3.3	配置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 邻居10.1.1.1激活 neighbor 10.1.1.1 send-community both	router bgp 100 neighbor 2001:DB8:1::1 remote-as 100 neighbor 2001:DB8:1::1 update-source Loopback0 ! address-family l2vpn evpn 激活邻居2001:DB8:1::1 neighbor 2001:DB8:1::1 send-community both exit-address-family	BGP EVPN对等已移至IPv6邻居地址。

	exit-address-family		

棕地主干/路由反射器V4至V4+V6 EVPN交换矩阵迁移

下表详细说明了从V4核心迁移到V4+V6核心所需的主干/路由反射器配置更改示例

迁移步骤	V4 EVPN交换矩阵	V4+V6 EVPN交换矩阵	描述
	单播路由配置		
1	ip routing	ip routing ipv6 unicast-routing	启用IPv6路由
	BGP配置		
2		router bgp 100 bgp router-id 10.3.3.3	配置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 邻居10.1.1.1激活 neighbor 10.1.1.1 send-community both exit-address-family	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 邻居10.1.1.1激活 neighbor 10.1.1.1 send-community both	使用IPv6和IPv6邻居地址的BGP EVPN对等。

	激活邻居2001:DB8:1::1 neighbor 2001:DB8:1::1 send-community both exit-address-family		

主干/路由反射器V4+V6至V6 EVPN交换矩阵迁移
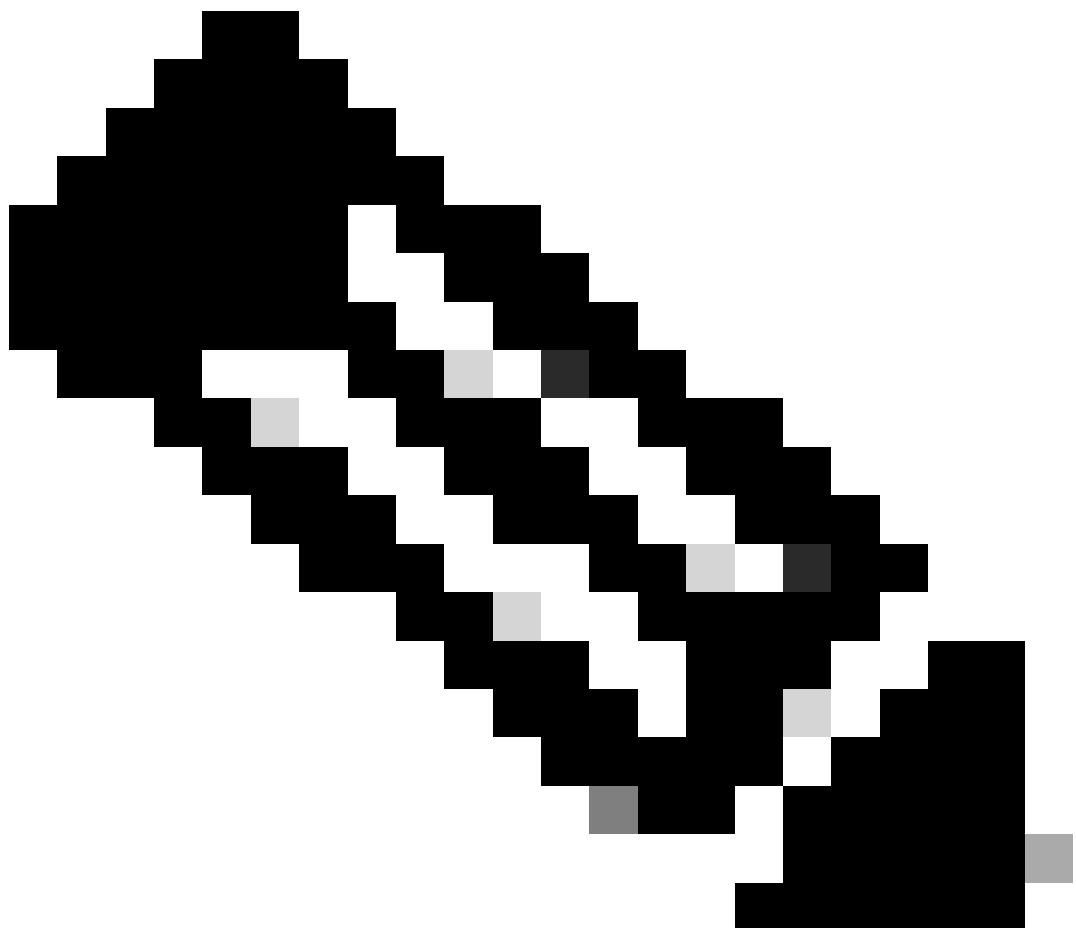
下表详细说明了从V4+V6核心迁移到V6核心所需的主干/路由反射器配置更改示例

迁移步骤	V4+V6 EVPN交换矩阵	V6 EVPN交换矩阵	描述
	BGP配置		
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 邻居10.1.1.1激活 neighbor 10.1.1.1 send-community both 激活邻居 2001:DB8:1::1	router bgp 100 neighbor 2001:DB8:1::1 remote-as 100 neighbor 2001:DB8:1::1 update-source Loopback0 ! address-family l2vpn evpn 激活邻居2001:DB8:1::1 neighbor 2001:DB8:1::1 send-community both exit-address-family !	使用IPv6邻居地址的BGP EVPN对等。

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

验证

以下各节详细介绍用于验证基本迁移功能的show命令。

> 注意：有关详细的验证和故障排除过程，请参阅BGP VXLANv6迁移故障排除指南。(不久就到)

本地VTEP配置

绿地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

双堆栈（首选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路由类型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
```

远程路由

<#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

源自第3层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远程（从EVPN导入）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流量转发

<#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路由类型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
```

远程路由

<#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单播路由

<#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流量转发

<#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

组播功能

BGP EVPN路由用于BUM-IR的第3类路由

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
```

远程路由

<#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
```

用于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
```

静态组播复制路由

<#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组播转发

<#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

示例配置

EVPN L2网关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（分布式任播网关）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
```

相关信息

- [BGP EVPN VXLAN配置指南](#)
- [BGP隧道封装属性(rfc9012)](#)
- BGP VXLANv6迁移故障排除指南，了解详细的验证和故障排除过程。(不久就到)
- [技术支持和文档 - Cisco Systems](#)