H3C S9820-64H Switch

VXLAN Configuration Guide

New H3C Technologies Co., Ltd.
http://www.h3c.com

Software version: Release 6607 and later
Document version: 6W100-20200325

Preface

This configuration guide describes the VXLAN fundamentals and configuration procedures.

This preface includes the following topics about the documentation:

- Audience.
- Conventions.
- Documentation feedback.

Audience

This documentation is intended for:

- Network planners.
- Field technical support and servicing engineers.
- Network administrators working with the S9820-64H switch.

Conventions

The following information describes the conventions used in the documentation.

Command conventions

Convention	Description
Boldface	**Bold** text represents commands and keywords that you enter literally as shown.
Italic	*Italic* text represents arguments that you replace with actual values.
[ ]	Square brackets enclose syntax choices (keywords or arguments) that are optional.
{ x | y | ... }	Braces enclose a set of required syntax choices separated by vertical bars, from which you select one.
[x | y | ...]	Square brackets enclose a set of optional syntax choices separated by vertical bars, from which you select one or none.
{ x | y | ... } *	Asterisk marked braces enclose a set of required syntax choices separated by vertical bars, from which you select a minimum of one.
[x | y | ...] *	Asterisk marked square brackets enclose optional syntax choices separated by vertical bars, from which you select one choice, multiple choices, or none.
&<1-n>	The argument or keyword and argument combination before the ampersand (&) sign can be entered 1 to n times.
#	A line that starts with a pound (#) sign is comments.

GUI conventions

Convention	Description
Boldface	Window names, button names, field names, and menu items are in Boldface. For example, the **New User** window opens; click **OK**.
>	Multi-level menus are separated by angle brackets. For example, **File** > **Create** > **Folder**.

Symbols

Convention	Description
WARNING!	An alert that calls attention to important information that if not understood or followed can result in personal injury.
CAUTION:	An alert that calls attention to important information that if not understood or followed can result in data loss, data corruption, or damage to hardware or software.
IMPORTANT:	An alert that calls attention to essential information.
NOTE:	An alert that contains additional or supplementary information.
TIP:	An alert that provides helpful information.

Network topology icons

Convention	Description
	Represents a generic network device, such as a router, switch, or firewall.
	Represents a routing-capable device, such as a router or Layer 3 switch.
	Represents a generic switch, such as a Layer 2 or Layer 3 switch, or a router that supports Layer 2 forwarding and other Layer 2 features.
	Represents an access controller, a unified wired-WLAN module, or the access controller engine on a unified wired-WLAN switch.
	Represents an access point.
	Represents a wireless terminator unit.
	Represents a wireless terminator.
	Represents a mesh access point.
	Represents omnidirectional signals.
	Represents directional signals.
	Represents a security product, such as a firewall, UTM, multiservice security gateway, or load balancing device.
	Represents a security module, such as a firewall, load balancing, NetStream, SSL VPN, IPS, or ACG module.

Examples provided in this document

Examples in this document might use devices that differ from your device in hardware model, configuration, or software version. It is normal that the port numbers, sample output, screenshots, and other information in the examples differ from what you have on your device.

Documentation feedback

You can e-mail your comments about product documentation to info@h3c.com.

We appreciate your comments.

Contents

VXLAN overview

Virtual eXtensible LAN (VXLAN) is a MAC-in-UDP technology that provides Layer 2 connectivity between distant network sites across an IP network. VXLAN is typically used in data centers and the access layer of campus networks for multitenant services.

VXLAN benefits

VXLAN provides the following benefits:

- **Support for more virtual switched domains than VLANs**—Each VXLAN is uniquely identified by a 24-bit VXLAN ID. The total number of VXLANs can reach 16777216 (2^{24}). This specification makes VXLAN a better choice than 802.1Q VLAN to isolate traffic for user terminals.
- **Easy deployment and maintenance**—VXLAN requires deployment only on the edge devices of the transport network. Devices in the transport network perform typical Layer 3 forwarding.

VXLAN network model

As shown in Figure 1, a VXLAN is a virtual Layer 2 network (known as the overlay network) built on top of an existing physical Layer 3 network (known as the underlay network). The overlay network encapsulates inter-site Layer 2 frames into VXLAN packets and forwards the packets to the destination along the Layer 3 forwarding paths provided by the underlay network. The underlay network is transparent to tenants, and geographically dispersed sites of a tenant are merged into a Layer 2 network.

The transport edge devices assign user terminals to different VXLANs, and then forward traffic between sites for user terminals by using VXLAN tunnels. Supported user terminals include PCs, wireless terminals, and VMs on servers.

NOTE:

This document uses VMs as examples to describe the mechanisms of VXLAN. The mechanisms do not differ between different kinds of user terminals.

The transport edge devices are VXLAN tunnel endpoints (VTEP). The VTEP implementation of the device uses ACs, VSIs, and VXLAN tunnels to provide VXLAN services.

- **VSI**—A virtual switch instance is a virtual Layer 2 switched domain. Each VSI provides switching services only for one VXLAN. VSIs learn MAC addresses and forward frames independently of one another. VMs in different sites have Layer 2 connectivity if they are in the same VXLAN.
- **Attachment circuit (AC)**—An AC is a physical or virtual link that connects a VTEP to a local site. Typically, ACs are Ethernet service instances that are associated with the VSI of a VXLAN. Traffic received from an AC is assigned to the VSI associated with the AC. Ethernet service instances are created on site-facing Layer 2 interfaces. An Ethernet service instance matches a list of custom VLANs by using a frame match criterion.
- **VXLAN tunnel**—Logical point-to-point tunnels between VTEPs over the transport network. Each VXLAN tunnel can trunk multiple VXLANs.

VTEPs encapsulate VXLAN traffic in the VXLAN, outer UDP, and outer IP headers. The devices in the transport network forward VXLAN traffic only based on the outer IP header.

Figure 1 VXLAN network model

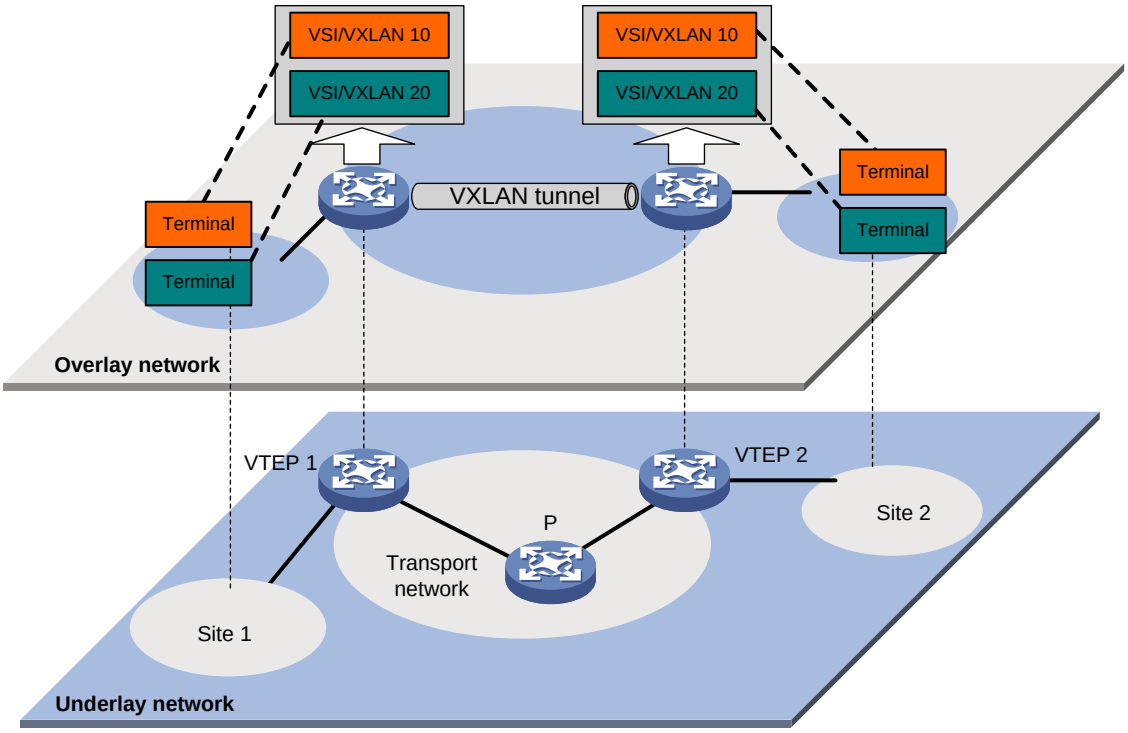

VXLAN packet format

As shown in Figure 2, a VTEP encapsulates a frame in the following headers:

- **8-byte VXLAN header**—VXLAN information for the frame.
  - **Flags**—If the I bit is 1, the VXLAN ID is valid. If the I bit is 0, the VXLAN ID is invalid. All other bits are reserved and set to 0.
  - **24-bit VXLAN ID**—Identifies the VXLAN of the frame. It is also called the virtual network identifier (VNI).
- **8-byte outer UDP header for VXLAN**—The default VXLAN destination UDP port number is 4789.
- **20-byte outer IP header**—Valid addresses of VTEPs or VXLAN multicast groups on the transport network. Devices in the transport network forward VXLAN packets based on the outer IP header.

Figure 2 VXLAN packet format

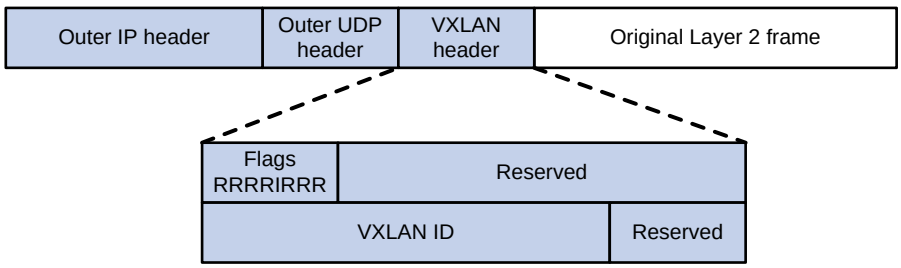

VXLAN working mechanisms

Generic VXLAN network establishment and forwarding process

The VTEP uses the following process to establish the VXLAN network and forward an inter-site frame:

1. Discovers remote VTEPs, establishes VXLAN tunnels, and assigns the VXLAN tunnels to VXLANs.
2. Assigns the frame to its matching VXLAN if the frame is sent between sites.
3. Performs MAC learning on the VXLAN's VSI.
4. Forwards the frame through VXLAN tunnels.

This section describes this process in detail. For intra-site frames in a VSI, the system performs typical Layer 2 forwarding, and it processes 802.1Q VLAN tags as described in "Access modes of VSIs."

VXLAN tunnel establishment and assignment

To provide Layer 2 connectivity for a VXLAN between two sites, you must create a VXLAN tunnel between the sites and assign the tunnel to the VXLAN.

VXLAN tunnel establishment

VXLAN supports manual and automatic VXLAN tunnel establishment.

- **Manual creation**—Manually create a VXLAN tunnel interface, and specify the tunnel source and destination IP addresses on the peer VTEPs.
- **Automatic creation**—Configure Ethernet Virtual Private Network (EVPN) to automatically discover VTEPs and set up VXLAN tunnels. For more information about EVPN, see *EVPN Configuration Guide*.

VXLAN tunnel assignment

VXLAN supports manual and automatic VXLAN tunnel assignment.

- **Manual assignment**—Manually assign VXLAN tunnels to VXLANs.
- **Automatic assignment**—Run EVPN to automatically assign VXLAN tunnels to VXLANs. For more information about EVPN, see *EVPN Configuration Guide*.

Assignment of traffic to VXLANs

Traffic from the local site to a remote site

The VTEP uses the following methods to assign customer frames to a VXLAN:

- **Ethernet service instance-to-VSI mapping**—This method uses the frame match criterion of an Ethernet service instance to match a list of VLANs on a site-facing Layer 2 interface. The frame match criterion specifies the characteristics of traffic from the VLANs, such as tagging status and VLAN IDs. The VTEP assigns customer traffic to a VXLAN by mapping the Ethernet service instance to a VSI.
- **VLAN-based VXLAN assignment**—This method maps a VLAN to a VXLAN. The VTEP assigns all frames of the VLAN to the VXLAN.

As shown in Figure 3, Ethernet service instance 1 matches VLAN 2 and is mapped to VSI A (VXLAN 10). When a frame from VLAN 2 arrives, the VTEP assigns the frame to VXLAN 10, and looks up VSI A's MAC address table for the outgoing interface.

Figure 3 Identifying traffic from the local site

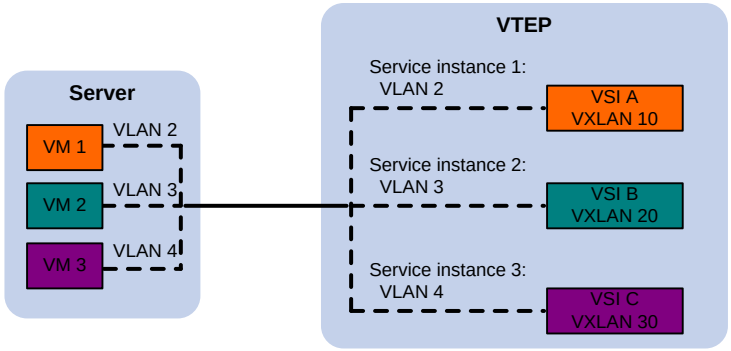

Traffic from a remote site to the local site

When a frame arrives at a VXLAN tunnel, the VTEP uses the VXLAN ID in the frame to identify its VXLAN.

MAC learning

The VTEP performs source MAC learning on the VSI as a Layer 2 switch.

- For traffic from the local site to the remote site, the VTEP learns the source MAC address before VXLAN encapsulation.
- For traffic from the remote site to the local site, the VTEP learns the source MAC address after removing the VXLAN header.

A VSI's MAC address table includes the following types of MAC address entries:

- **Local MAC**—MAC entries learned from the local site. The outgoing interfaces for the MAC address entries are site-facing interfaces.
  - **Static**—Manually added MAC entries.
  - **Dynamic**—Dynamically learned MAC entries.
- **Remote MAC**—MAC entries learned from a remote site, including static and dynamic MAC entries. The outgoing interfaces for the MAC addresses are VXLAN tunnel interfaces.
  - **Static**—Manually added MAC entries.
  - **Dynamic**—MAC entries learned in the data plane from incoming traffic on VXLAN tunnels. The learned MAC addresses are contained in the inner Ethernet header.
  - **BGP EVPN**—MAC entries advertised through BGP EVPN. For more information, see *EVPN Configuration Guide*.
  - **OpenFlow**—MAC entries issued by a remote controller through OpenFlow. For more information, see *OpenFlow Configuration Guide*.
  - **OVSDB**—MAC entries issued by a remote controller through OVSDB.

  The following shows the priority order of different types of remote MAC address entries:

  a. Static MAC address entries, and MAC address entries issued by a remote controller through OpenFlow or OVSDB. These types of entries have the same priority and overwrite each other.

  a. MAC address entries advertised through BGP EVPN.

  b. Dynamic MAC address entries.

Unicast forwarding

Intra-site unicast forwarding

The VTEP uses the following process to forward a known unicast frame within a site:

1. Identifies the VSI of the frame.
2. Looks up the destination MAC address in the VSI's MAC address table for the outgoing interface.
3. Sends the frame out of the matching outgoing interface.

As shown in Figure 4, VTEP 1 forwards a frame from VM 1 to VM 4 within the local site in VLAN 10 as follows:

1. Identifies that the frame belongs to VSI A when the frame arrives at Interface A.
2. Looks up the destination MAC address (MAC 4) in the MAC address table of VSI A for the outgoing interface.
3. Sends the frame out of the matching outgoing interface (Interface B) to VM 4 in VLAN 10.

Figure 4 Intra-site unicast

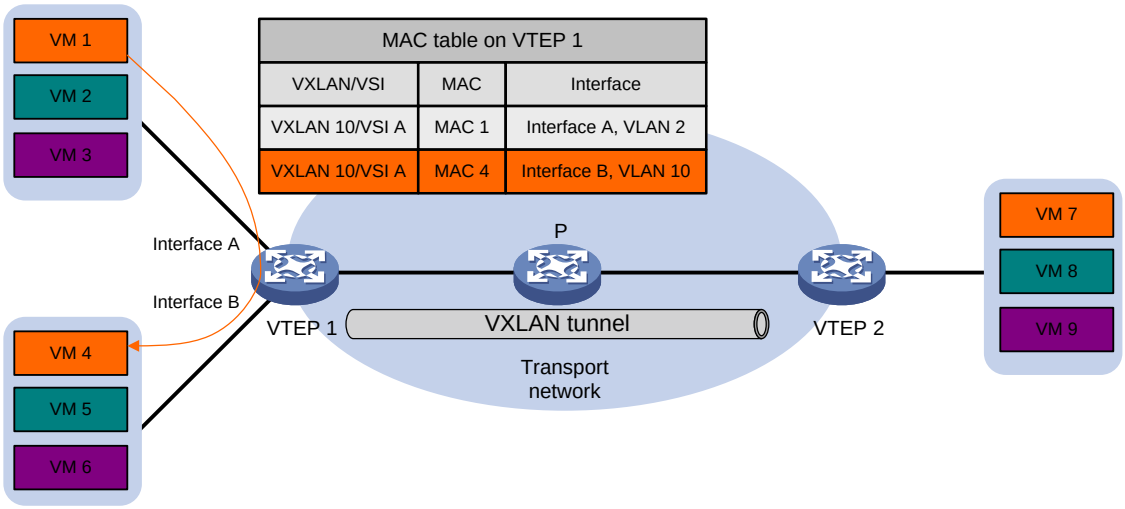

Inter-site unicast forwarding

The following process (see Figure 5) applies to a known unicast frame between sites:

1. The source VTEP encapsulates the Ethernet frame in the VXLAN/UDP/IP header.

   In the outer IP header, the source IP address is the source VTEP's VXLAN tunnel source IP address. The destination IP address is the VXLAN tunnel destination IP address.
2. The source VTEP forwards the encapsulated packet out of the outgoing VXLAN tunnel interface found in the VSI's MAC address table.
3. The intermediate transport devices (P devices) forward the frame to the destination VTEP by using the outer IP header.
4. The destination VTEP removes the headers on top of the inner Ethernet frame. It then performs MAC address table lookup in the VXLAN's VSI to forward the frame out of the matching outgoing interface.

Figure 5 Inter-site unicast

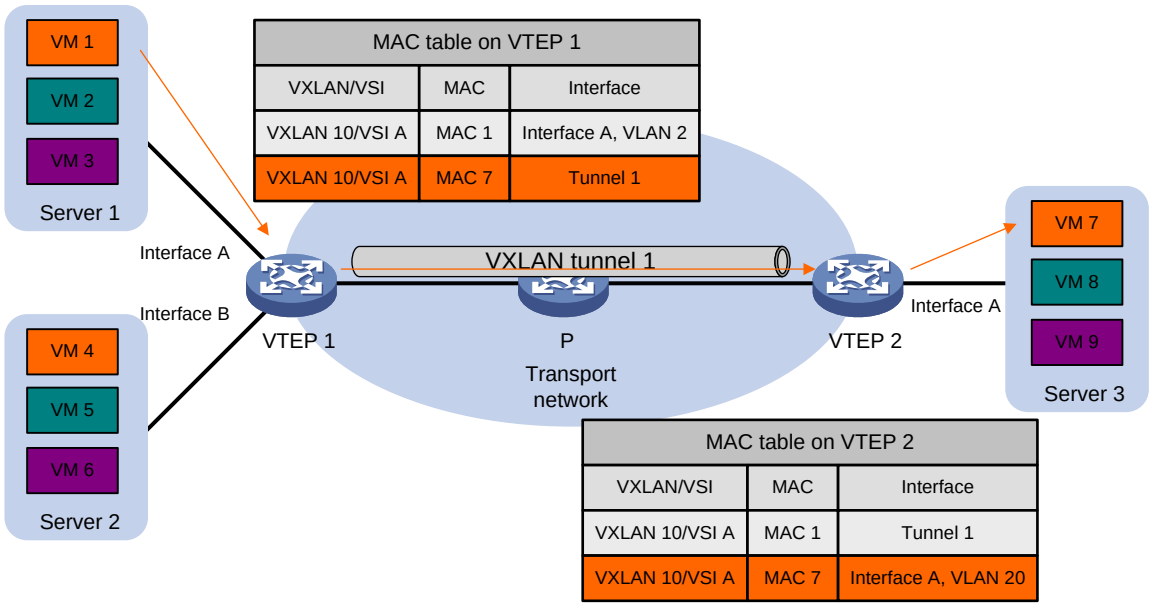

Flood

The source VTEP floods a broadcast, multicast, or unknown unicast frame to all site-facing interfaces and VXLAN tunnels in the VXLAN, except for the incoming interface. Each destination VTEP floods the inner Ethernet frame to all site-facing interfaces in the VXLAN. To avoid loops, the destination VTEPs do not flood the frame back to VXLAN tunnels.

VXLAN supports unicast mode (also called head-end replication), multicast mode (also called tandem replication), and flood proxy mode for flood traffic.

Unicast mode (head-end replication)

As shown in Figure 6, the source VTEP replicates the flood frame, and then sends one replica to the destination IP address of each VXLAN tunnel in the VXLAN.

Figure 6 Unicast mode

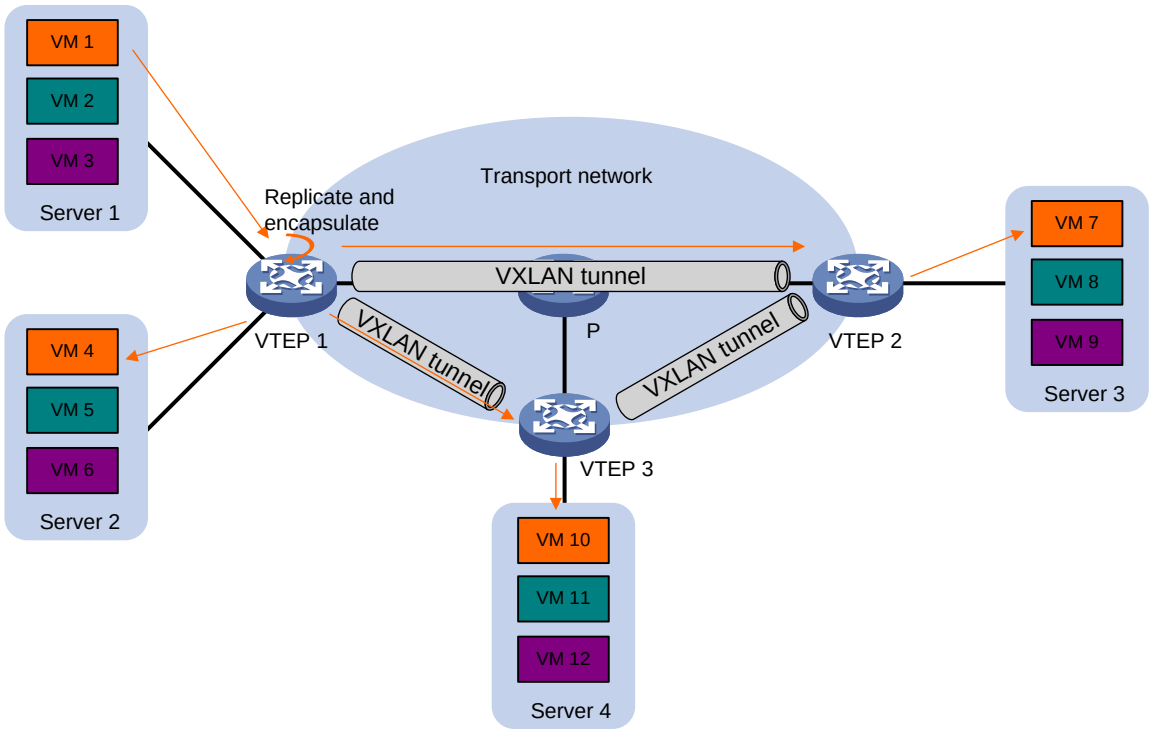

Multicast mode (tandem replication)

As shown in Figure 7, the source VTEP sends the flood frame in a multicast VXLAN packet destined for a multicast group address. Transport network devices replicate and forward the packet to remote VTEPs based on their multicast forwarding entries.

Figure 7 Multicast mode

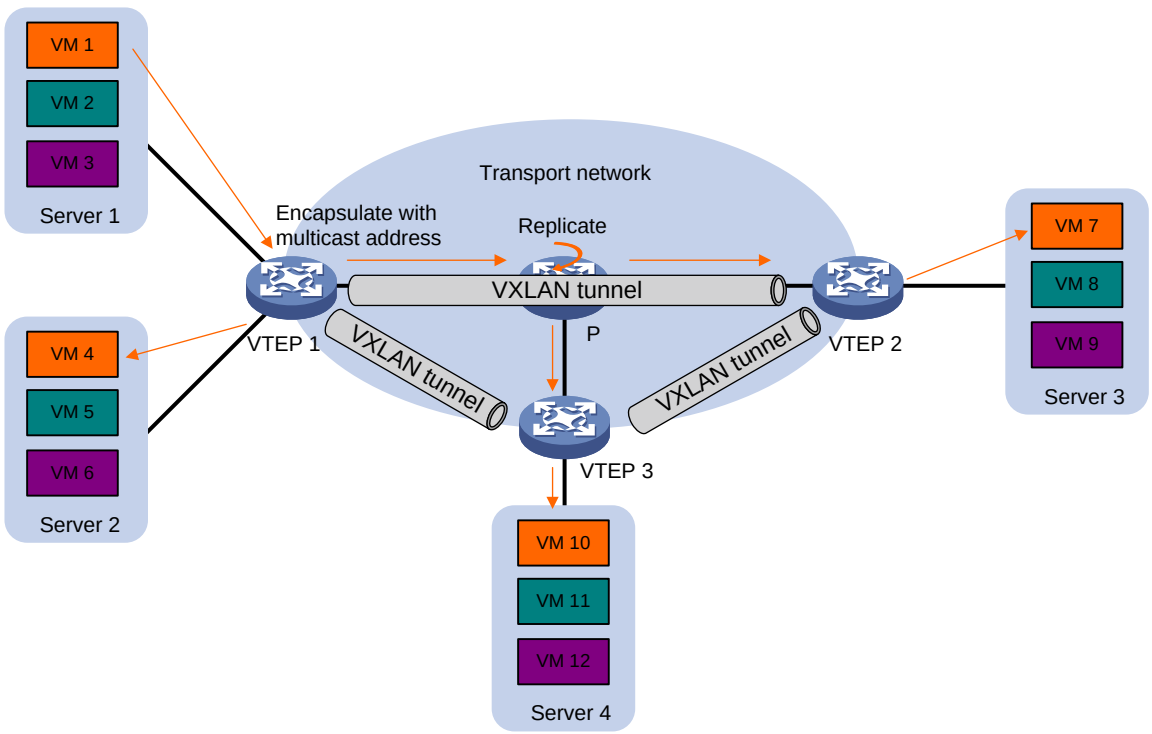

Flood proxy mode (proxy server replication)

As shown in Figure 8, the source VTEP sends the flood frame in a VXLAN packet over a VXLAN tunnel to a flood proxy server. The flood proxy server replicates and forwards the packet to each remote VTEP through its VXLAN tunnels.

The flood proxy mode applies to VXLANs that have many sites. This mode reduces flood traffic in the transport network without using a multicast protocol. To use a flood proxy server, you must set up a VXLAN tunnel to the server on each VTEP.

Figure 8 Flood proxy mode

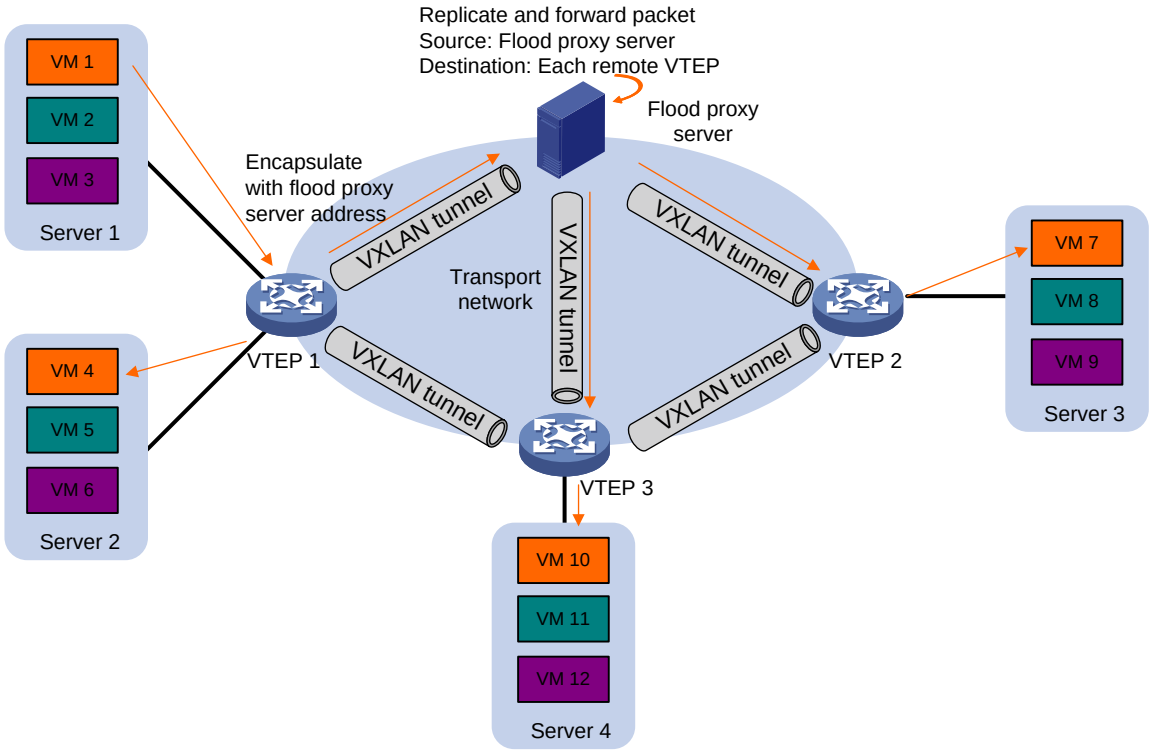

The flood proxy mode is typically used in SDN transport networks that have a flood proxy server. For VTEPs to forward packets based on the MAC address table issued by an SDN controller, you must perform the following tasks on the VTEPs:

- Disable remote-MAC address learning by using the **`vxlan tunnel mac-learning disable`** command.
- Disable source MAC check on all transport-facing interfaces by using the **`undo mac-address static source-check enable`** command. If the VTEP is an IRF fabric, you must also disable the feature on all IRF ports.

Access modes of VSIs

The access mode of a VSI determines how the VTEP processes the 802.1Q VLAN tags in the Ethernet frames.

VLAN access mode

In this mode, Ethernet frames received from or sent to the local site must contain 802.1Q VLAN tags.

- For an Ethernet frame received from the local site, the VTEP removes all its 802.1Q VLAN tags before forwarding the frame.

- For an Ethernet frame destined for the local site, the VTEP adds 802.1Q VLAN tags to the frame before forwarding the frame.

In VLAN access mode, VXLAN packets sent between sites do not contain 802.1Q VLAN tags. You can use different 802.1Q VLANs to provide the same service in different sites.

Ethernet access mode

The VTEP does not process the 802.1Q VLAN tags of Ethernet frames received from or sent to the local site.

- For an Ethernet frame received from the local site, the VTEP forwards the frame with the 802.1Q VLAN tags intact.
- For an Ethernet frame destined for the local site, the VTEP forwards the frame without adding 802.1Q VLAN tags.

In Ethernet access mode, VXLAN packets sent between VXLAN sites contain 802.1Q VLAN tags. You must use the same VLAN to provide the same service between sites.

ARP flood suppression

ARP flood suppression reduces ARP request broadcasts by enabling the VTEP to reply to ARP requests on behalf of VMs.

As shown in Figure 9, this feature snoops ARP packets to populate the ARP flood suppression table with local and remote MAC addresses. If an ARP request has a matching entry, the VTEP replies to the request on behalf of the VM. If no match is found, the VTEP floods the request to both local and remote sites.

Figure 9 ARP flood suppression

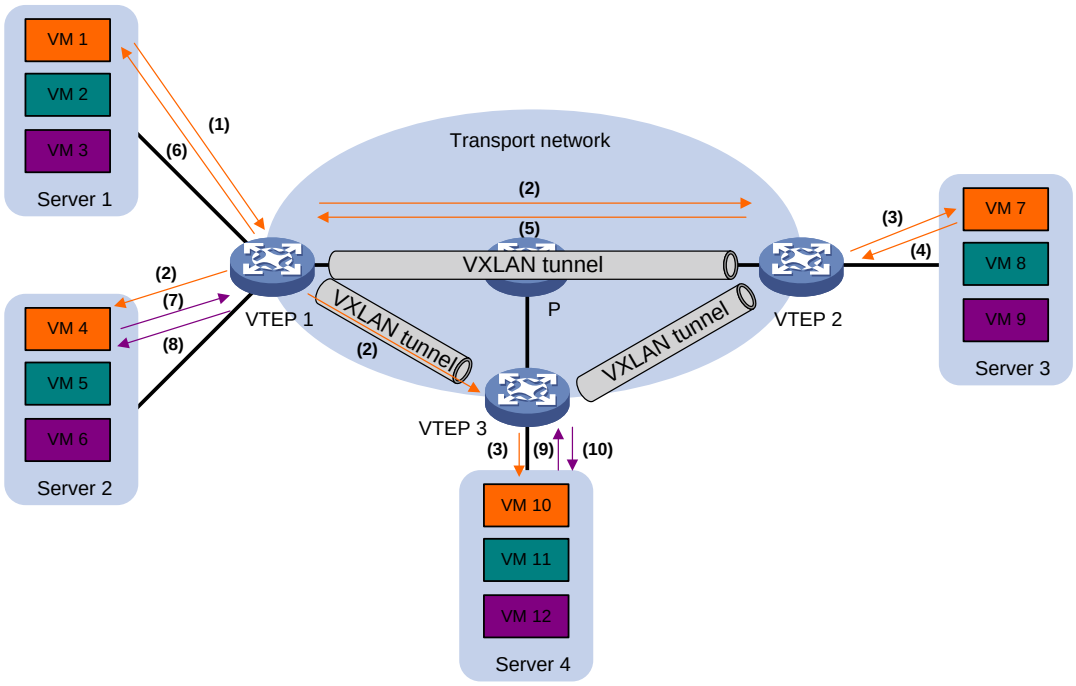

ARP flood suppression uses the following workflow:

1. VM 1 sends an ARP request to obtain the MAC address of VM 7.
2. VTEP 1 creates a suppression entry for VM 1, and floods the ARP request in the VXLAN.
3. VTEP 2 and VTEP 3 de-encapsulate the ARP request. The VTEPs create a suppression entry for VM 1, and broadcast the request in the local site.

4. VM 7 sends an ARP reply.
5. VTEP 2 creates a suppression entry for VM 7 and forwards the ARP reply to VTEP 1.
6. VTEP 1 de-encapsulates the ARP reply, creates a suppression entry for VM 7, and forwards the ARP reply to VM 1.
7. VM 4 sends an ARP request to obtain the MAC address of VM 1 or VM 7.
8. VTEP 1 creates a suppression entry for VM 4 and replies to the ARP request.
9. VM 10 sends an ARP request to obtain the MAC address of VM 1.
10. VTEP 3 creates a suppression entry for VM 10 and replies to the ARP request.

Protocols and standards

RFC 7348, *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*

Configuring basic VXLAN features

Restrictions: Loop prevention restriction

Do not enable the spanning tree feature on VTEPs. To prevent loops on site-facing interfaces, enable BPDU guard.

VXLAN tasks at a glance

To configure basic VXLAN settings, perform the following tasks on VTEPs:

1. Creating a VXLAN on a VSI
2. Configuring a VXLAN tunnel
3. Manually assigning VXLAN tunnels to a VXLAN
4. Assigning customer frames to a VSI
5. (Optional.) Managing MAC address entries
   - Configuring static MAC address entries
   - Disabling local-MAC address learning
   - Disabling remote-MAC address learning
   - Enabling local-MAC logging
   - Setting the MAC learning priority of an Ethernet service instance
   - Configuring VXLAN over VXLAN
6. Configuring a multicast-mode VXLAN

   If the network is multicast dense, configure the VTEP to flood VXLAN traffic in multicast mode.
7. (Optional.) Configuring VXLAN packet parameters
   - Setting the destination UDP port number of VXLAN packets
   - Configuring VXLAN packet check
   - Enabling default VXLAN decapsulation
8. (Optional.) Reducing flood traffic in the transport network
   - Disabling flooding for a VSI
   - Confining the flood traffic of an Ethernet service instance
   - Enabling ARP flood suppression
9. Maintaining VXLAN networks
   - Enabling VXLAN packet statistics
   - Testing the reachability of a remote VM

Prerequisites for VXLAN

Configure a routing protocol on the devices in the transport network to make sure the VTEPs can reach one another.

Creating a VXLAN on a VSI

Restrictions and guidelines

If you use both the **restrain** and **bandwidth** commands on a VSI, the following rules apply:

- If the restraint bandwidth is 0, the **restrain** command takes effect.
- If the restraint bandwidth is not 0, the **bandwidth** command takes effect.

Procedure

1. Enter system view.
   system-view
2. Enable L2VPN.
   l2vpn enable
   By default, L2VPN is disabled.
3. Create a VSI and enter VSI view.
   vsi *vsi-name*
4. Enable the VSI.
   undo shutdown
   By default, a VSI is enabled.
5. Create a VXLAN and enter VXLAN view.
   vxlan *vxlan-id*
   You can create only one VXLAN on a VSI.
   The VXLAN ID must be unique for each VSI.
6. (Optional.) Configure VSI parameters:
   a. Return to VSI view.
      quit
   b. Configure a VSI description.
      description *text*
      By default, a VSI does not have a description.
   c. Set the MTU for the VSI.
      mtu *size*
      The default MTU for a VSI is 1500 bytes.
   d. Set the maximum bandwidth for the VSI.
      bandwidth *bandwidth*
      By default, the maximum bandwidth is not limited for a VSI.
   e. Set the broadcast, unknown multicast, or unknown unicast restraint bandwidth for the VSI.
      restrain { **broadcast** | **multicast** | **unknown-unicast** } *bandwidth*
      By default, a VSI's broadcast restraint bandwidth, unknown multicast restraint bandwidth, and unknown unicast restraint bandwidth are not set.
   f. Enable MAC address learning for the VSI.
      mac-learning enable
      By default, MAC address learning is enabled for a VSI.

Configuring a VXLAN tunnel

Manually creating a VXLAN tunnel

About this task

When you manually create a VXLAN tunnel, specify addresses on the local VTEP and the remote VTEP as the tunnel source and destination addresses, respectively.

Restrictions and guidelines

As a best practice, do not configure multiple VXLAN tunnels to use the same source and destination IP addresses.

Make sure the following VXLAN tunnels are not associated with the same VXLAN when they have the same tunnel destination IP address:

- A VXLAN tunnel automatically created by EVPN.
- A manually created VXLAN tunnel.

For more information about EVPN, see *EVPN Configuration Guide*.

This task provides basic VXLAN tunnel configuration. For more information about tunnel configuration and commands, see *Layer 3—IP Services Configuration Guide* and *Layer 3—IP Services Command Reference*.

Procedure

1. Enter system view.

   system-view

2. (Optional.) Specify a global source IP address for VXLAN tunnels.

   tunnel global source-address *ipv4-address*

   By default, no global source IP address is specified for VXLAN tunnels.

   A VXLAN tunnel uses the global source address if you do not specify a source interface or source address for the tunnel.

3. Create a VXLAN tunnel interface and enter tunnel interface view.

   interface tunnel *tunnel-number* **mode vxlan**

   The endpoints of a tunnel must use the same tunnel mode.

4. Specify a source address for the tunnel. Choose one of the following methods:

   - Specify a source IP address for the tunnel.

     source *ipv4-address*

     The specified IP address is used in the outer IP header of tunneled VXLAN packets.

   - Specify a source interface for the tunnel.

     source *interface-type interface-number*

     The IP address of the specified interface is used in the outer IP header of tunneled VXLAN packets.

   By default, no source IP address or source interface is specified for a tunnel.

   Do not perform this step if you are using OVSDB for VXLAN tunnel management.

   For a multicast-mode VXLAN, the source IP address cannot be a loopback interface's address, and the source interface cannot be a loopback interface.

5. Specify a destination IP address for the tunnel.

   destination *ipv4-address*

   By default, no destination IP address is specified for a tunnel.

Specify the remote VTEP's IP address. This IP address will be the destination IP address in the outer IP header of tunneled VXLAN packets.

Enabling BFD on a VXLAN tunnel

About this task

Enable BFD on both ends of a VXLAN tunnel for quick link connectivity detection. The VTEPs periodically send BFD single-hop control packets to each other through the VXLAN tunnel. A VTEP sets the tunnel state to Defect if it has not received control packets from the remote end for 5 seconds. In this situation, the tunnel interface state is still Up. The tunnel state will change from Defect to Up if the VTEP can receive BFD control packets again.

Restrictions and guidelines

You must enable BFD on both ends of a VXLAN tunnel.

Procedure

1. Enter system view.

   system-view
2. Specify the reserved VXLAN.

   reserved vxlan *vxlan-id*

   By default, no VXLAN has been reserved.

   For BFD sessions to come up, you must reserve a VXLAN.

   You can specify only one reserved VXLAN on the VTEP. The reserved VXLAN cannot be the VXLAN created on any VSI.
3. Enter VXLAN tunnel interface view.

   interface tunnel *tunnel-number*
4. Enable BFD on the tunnel.

   tunnel bfd enable destination-mac *mac-address*

   By default, BFD is disabled on a tunnel.

Manually assigning VXLAN tunnels to a VXLAN

About this task

To provide Layer 2 connectivity for a VXLAN between two sites, you must assign the VXLAN tunnel between the sites to the VXLAN.

You can assign multiple VXLAN tunnels to a VXLAN, and configure a VXLAN tunnel to trunk multiple VXLANs. For a unicast-mode VXLAN, the system floods unknown unicast, multicast, and broadcast traffic to each tunnel associated with the VXLAN. If a flood proxy server is used, the VTEP sends flood traffic to the server through the flood proxy tunnel. The flood proxy server replicates and forwards flood traffic to remote VTEPs.

Restrictions and guidelines

For full Layer 2 connectivity in the VXLAN, make sure the VXLAN contains the VXLAN tunnel between each pair of sites in the VXLAN.

Procedure

1. Enter system view.

   system-view
2. Enter VSI view.

vsi *vsi-name*

3. Enter VXLAN view.
   vxlan *vxlan-id*
4. Assign VXLAN tunnels to the VXLAN.
   tunnel { *tunnel-number* [**backup-tunnel** *tunnel-number* | **flooding-proxy** }] | **all** }
   By default, a VXLAN does not contain any VXLAN tunnels.

Parameter	Description
backup-tunnel *tunnel-number*	Specifies a backup tunnel. When the primary VXLAN tunnel is operating correctly, the backup VXLAN tunnel does not forward traffic. When the primary VXLAN tunnel goes down, traffic is switched to the backup VXLAN tunnel.
flooding-proxy	Enables flood proxy on a tunnel for it to send flood traffic to the flood proxy server. The flood proxy server replicates and forwards flood traffic to remote VTEPs.

Assigning customer frames to a VSI

Restrictions and guidelines for configuring traffic assignment methods

VLAN-based VXLAN assignment is mutually exclusive with the manually created Ethernet service instances and the Ethernet service instances automatically created for 802.1X or MAC authentication VSI manipulation. To create these Ethernet service instances, you must first disable VLAN-based VXLAN assignment by using the **undo vxlan vlan-based** command. To enable VLAN-based VXLAN assignment, you must first delete all Ethernet service instances.

Mapping a static Ethernet service instance to a VSI

About this task

A static Ethernet service instance matches a list of VLANs on a site-facing interface. The VTEP assigns customer traffic from the VLANs to a VXLAN by mapping the Ethernet service instance to a VSI.

Restrictions and guidelines

Ethernet service instance bindings of VSIs are mutually exclusive with port bridging, QinQ, and VLAN mapping on a Layer 2 Ethernet interface or Layer 2 aggregate interface. Do not configure these features simultaneously on the same interface. Otherwise, the features cannot take effect.

Make sure the VLANs that each Ethernet service instance matches have been created and the site-facing interface where the Ethernet service instance is configured has been assigned to the VLANs.

For information about the frame match criterion configuration restrictions and guidelines of Ethernet service instances, see *VXLAN Command Reference*.

Procedure

1. Enter system view.

system-view

2. Enter interface view.
   - Enter Layer 2 Ethernet interface view.
     interface *interface-type interface-number*
   - Enter Layer 2 aggregate interface view.
     interface bridge-aggregation *interface-number*
3. Create an Ethernet service instance and enter Ethernet service instance view.
   service-instance *instance-id*
4. Configure a frame match criterion. Choose one of the following options:
   - Match frames tagged with the specified outer 802.1Q VLAN IDs.
     encapsulation s-vid *vlan-id* [**only-tagged**]
     encapsulation s-vid *vlan-id-list*
   - Match frames tagged with the specified outer and inner 802.1Q VLAN IDs.
     encapsulation s-vid *vlan-id-list* **c-vid** *vlan-id*
     encapsulation s-vid *vlan-id* **c-vid** { *vlan-id* | **all** }
   - Match any 802.1Q tagged or untagged frames.
     encapsulation { **tagged** | **untagged** }
   - Match frames that do not match any other service instance on the interface.
     encapsulation default

     An interface can contain only one Ethernet service instance that uses the **encapsulation default** criterion.

     An Ethernet service instance that uses the **encapsulation default** criterion matches any frames if it is the only instance on the interface.

   By default, an Ethernet service instance does not contain a frame match criterion.
5. (Optional.) Set the bandwidth limit for the Ethernet service instance.
   bandwidth *bandwidth*

   By default, no bandwidth limit is set for an Ethernet service instance.
6. Map the Ethernet service instance to a VSI.
   xconnect vsi *vsi-name* [**access-mode** { **ethernet** | **vlan** }] [**track** *track-entry-number*&<1-3>]

   By default, an Ethernet service instance is not mapped to any VSI.

Mapping dynamic Ethernet service instances to VSIs

About this task

The 802.1X or MAC authentication feature can use the authorization VSI, the guest VSI, the Auth-Fail VSI, and the critical VSI to control the access of users to network resources. When assigning a user to a VSI, 802.1X or MAC authentication sends the VXLAN feature the VSI information and the user's access information, including access interface, VLAN, and MAC address. Then the VXLAN feature creates a dynamic Ethernet service instance for the user and maps it to the VSI. For more information about 802.1X authentication and MAC authentication, see *Security Configuration Guide*.

A dynamic Ethernet service instance supports the following traffic match modes:

- **VLAN-based mode**—Matches frames by VLAN ID.
- **MAC-based mode**—Matches frames by VLAN ID and source MAC address.

By default, dynamic Ethernet service instances use VLAN-based traffic match mode. To use MAC-based traffic match mode for dynamic Ethernet service instances, you must enable MAC authentication or 802.1X authentication that uses MAC-based access control.

Configuring the VLAN-based traffic match mode

To use the VLAN-based traffic match mode, configure 802.1X authentication or MAC authentication and perform one of the following tasks:

- Configure the guest VSI, Auth-Fail VSI, or critical VSI on the 802.1X- or MAC authentication-enabled interface.
- Issue an authorization VSI to an 802.1X or MAC authentication user from a remote AAA server.

Then, the device will automatically create a dynamic Ethernet service instance for the 802.1X or MAC authentication user and map the Ethernet service instance to a VSI.

For more information about configuring 802.1X authentication and MAC authentication, see *Security Configuration Guide*.

Configuring the MAC-based traffic match mode

1. Enter system view.

   system-view
2. Enter interface view.
   - Enter Layer 2 Ethernet interface view.

     interface *interface-type interface-number*
   - Enter Layer 2 aggregate interface view.

     interface bridge-aggregation *interface-number*
3. Enable MAC-based traffic match mode for dynamic Ethernet service instances on the interface.

   mac-based ac

   By default, MAC-based traffic match mode is disabled for dynamic Ethernet service instances. Dynamic Ethernet service instances use VLAN-based traffic match mode.

   For more information about this command, see *VXLAN Command Reference*.
4. Enable MAC authentication or 802.1X authentication that uses MAC-based access control.

   To use the MAC-based traffic match mode, configure MAC authentication or 802.1X authentication that uses MAC-based access control and perform one of the following tasks:
   - Configure the guest VSI, Auth-Fail VSI, or critical VSI on the 802.1X- or MAC authentication-enabled interface.
   - Issue an authorization VSI to an 802.1X or MAC authentication user from a remote AAA server.

   Then, the device will automatically create a dynamic Ethernet service instance for the 802.1X or MAC authentication user and map the Ethernet service instance to a VSI.

   For more information about configuring 802.1X authentication and MAC authentication, see *Security Configuration Guide*.

Configuring VLAN-based VXLAN assignment

About this task

VLAN-based VXLAN assignment enables the device to assign all traffic of a VLAN to a VXLAN. If you enable this feature and map a VLAN to a VXLAN, the device automatically performs the following operations:

1. Creates an Ethernet service instance that uses the VLAN ID as its instance ID on each interface in the VLAN. The matching outer VLAN ID of the Ethernet service instances is the VLAN ID.
2. Maps the Ethernet service instances to the VSI of the VXLAN.

Restrictions and guidelines

Do not configure this feature together with EVPN distributed relay. For information about EVPN distributed relay, see *EVPN Configuration Guide*.

If you map a VLAN to a VXLAN, the VTEP cannot perform non-VXLAN Layer 2 forwarding in the VLAN. Also, the VLAN interface of the VLAN cannot perform Layer 3 forwarding.

The Ethernet service instance creation or deletion time is affected by the number of VLANs mapped to a VXLAN and the number of trunk ports assigned to the VLANs. The larger the numbers, the longer the time. During AC creation or deletion, other operations are queued.

Prerequisites

Use the **vxlan** command to create the VXLAN to which a VLAN is mapped.

Procedure

1. Enter system view.
   system-view
2. Enable VLAN-based VXLAN assignment.
   vxlan vlan-based
   By default, VLAN-based VXLAN assignment is disabled.
3. Create a VLAN and enter VLAN view.
   vlan *vlan-id*
   Do not specify VLAN 1 for VLAN-based VXLAN assignment.
4. Map the VLAN to a VXLAN.
   vxlan vni *vxlan-id*
   By default, a VLAN is not mapped to a VXLAN.

Managing MAC address entries

About MAC address entry management

Local-MAC address entries can be manually added or dynamically learned. You can log local MAC addresses and local-MAC changes.

Remote-MAC address entries have a variety of types, including manually added entries and dynamically learned entries.

Configuring static MAC address entries

Restrictions and guidelines

Do not configure static remote-MAC entries for VXLAN tunnels that are automatically established by using EVPN.

- EVPN re-establishes VXLAN tunnels if the transport-facing interface goes down and then comes up. If you have configured static remote-MAC entries, the entries are deleted when the tunnels are re-established.
- EVPN re-establishes VXLAN tunnels if you perform configuration rollback. If the tunnel IDs change during tunnel re-establishment, configuration rollback fails, and static remote-MAC entries on the tunnels cannot be restored.

  For more information about EVPN, see *EVPN Configuration Guide*.

Procedure

1. Enter system view.
   system-view
2. Add a static local-MAC address entry.
   mac-address static *mac-address* **interface** *interface-type interface-number* **service-instance** *instance-id* **vsi** *vsi-name*
   For successful configuration, make sure the Ethernet service instance has been mapped to the VSI.
3. Add a static remote-MAC address entry.
   mac-address static *mac-address* **interface tunnel** *tunnel-number* **vsi** *vsi-name*
   For the setting to take effect, make sure the VSI's VXLAN has been specified on the VXLAN tunnel.

Disabling local-MAC address learning

Restrictions and guidelines

When MAC address learning is disabled for Ethernet service instances, you can only configure static local-MAC address entries by using the **mac-address static** command.

Prerequisites

Before you enable MAC address learning for an Ethernet service instance, you must use the **mac-learning enable** command to enable MAC address learning for the associated VSI.

Procedure

1. Enter system view.
   system-view
2. Enter interface view.
   - Enter Layer 2 Ethernet interface view**.**
     interface *interface-type interface-number*
   - Enter Layer 2 aggregate interface view.
     interface bridge-aggregation *interface-number*
3. Enter Ethernet service instance view.
   service-instance *instance-id*
4. Disable MAC address learning for the Ethernet service instance.
   learning mode disable
   By default, MAC address learning is enabled for Ethernet service instances.

Disabling remote-MAC address learning

About this task

When network attacks occur, disable remote-MAC address learning to prevent the device from learning incorrect remote MAC addresses. You can manually add static remote-MAC address entries.

Procedure

1. Enter system view.
   system-view

2. Disable remote-MAC address learning.

   vxlan tunnel mac-learning disable

   By default, remote-MAC address learning is enabled.

Enabling local-MAC logging

About this task

When the local-MAC logging feature is enabled, the VXLAN module immediately sends a log message with its local MAC addresses to the information center. When a local MAC address is added or removed, a log message is also sent to the information center to report the local-MAC change.

With the information center, you can set log message filtering and output rules, including output destinations. For more information about configuring the information center, see *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.

   system-view

2. Enable local-MAC logging.

   vxlan local-mac report

   By default, local-MAC logging is disabled.

Setting the MAC learning priority of an Ethernet service instance

About this task

A VSI uses the MAC learning priority to control MAC address learning of its Ethernet service instances. An Ethernet service instance with high MAC learning priority takes precedence over an Ethernet service instance with low MAC learning priority when they learn the same MAC address. For example:

- A MAC address entry of a high-priority Ethernet service instance can be overwritten only when the MAC address is learned on another high-priority Ethernet service instance.
- A MAC address entry of a low-priority Ethernet service instance is overwritten when the MAC address is learned on a high-priority Ethernet service instance or another low-priority Ethernet service instance.

Procedure

1. Enter system view.

   system-view

2. Enter interface view.
   - Enter Layer 2 Ethernet interface view**.**

     interface *interface-type interface-number*
   - Enter Layer 2 aggregate interface view.

     interface bridge-aggregation *interface-number*
3. Enter Ethernet service instance view.

   service-instance *instance-id*
4. Set the MAC learning priority of the Ethernet service instance.

   mac-address mac-learning priority { **high** | **low** }

By default, the MAC learning priority of an Ethernet service instance is low.

This setting takes effect only after the Ethernet service instance is mapped to a VSI.

Configuring VXLAN over VXLAN

About this task

By default, the device de-encapsulates an incoming VXLAN packet if the packet's destination UDP port number is the VXLAN destination UDP port number (configured by using **vxlan udp-port**). For VXLAN packets received from a non-transport-facing interface on the device to traverse the VXLAN network through VXLAN tunnels, perform the following tasks on the interface:

- Enable VXLAN over VXLAN.
- Configure Ethernet service instance and VSI settings for matching the VXLAN packets.

When receiving VXLAN packets on the interface, the device adds a second layer of VXLAN encapsulation to the packets and forwards them over VXLAN tunnels.

Restrictions and guidelines

An interface enabled with VXLAN over VXLAN does not de-encapsulate incoming VXLAN packets. Do not enable this feature on a transport-facing interface.

For an aggregate interface, you do not need to enable VXLAN over VXLAN on its member ports if this feature is already enabled on that aggregate interface.

Procedure

1. Enter system view.

   system-view
2. Enter interface view.
   - Enter Layer 2 Ethernet interface view.

     interface *interface-type interface-number*
   - Enter Layer 2 aggregate interface view.

     interface bridge-aggregation *interface-number*
3. Enable VXLAN over VXLAN.

   vxlan-over-vxlan enable

   By default, VXLAN over VXLAN is disabled on an interface.

Configuring a multicast-mode VXLAN

About multicast methods for multicast-mode VXLANs

A multicast-mode VXLAN supports the following multicast methods:

- **PIM**—VTEPs and transport network devices run PIM to generate multicast forwarding entries. To forward multicast traffic correctly, you must use the source IP address of an up VXLAN tunnel as the source IP address of multicast VXLAN packets. As a best practice, use the source IP address of a VXLAN tunnel that uses the IP address of a loopback interface. If the VTEP has multiple transport-facing interfaces, PIM dynamically selects the outgoing interfaces for multicast VXLAN packets.
- **IGMP host**—VTEPs and transport network devices run PIM and IGMP to generate multicast forwarding entries.
  - Transport-facing interfaces of VTEPs act as IGMP hosts.
  - Transport network devices connected to a VTEP run IGMP.

- All transport network devices run PIM.

  On a VTEP, you must use the IP address of the transport-facing interface as the source IP address for multicast VXLAN packets. If the VTEP has multiple transport-facing interfaces, multicast VXLAN packets are sent to the transport network through the interface that provides the source IP address for multicast VXLAN packets.

VTEPs in a multicast-mode VXLAN can use different multicast methods.

Prerequisites for multicast-mode VXLANs

For a multicast-mode VXLAN to flood traffic, you must perform the following tasks in addition to multicast-mode configuration:

- Enable IP multicast routing on all VTEPs and transport network devices.
- Configure a multicast routing protocol on transport network devices. A VTEP can be both a multicast source and multicast group member. As a best practice, use BIDIR-PIM.
- Enable IGMP on transport network devices that are connected to an IGMP host-enabled VTEP.

Configuring a multicast-mode VXLAN that uses the PIM method

1. Enter system view.

   system-view
2. Enter VSI view.

   vsi *vsi-name*
3. Enter VXLAN view.

   vxlan *vxlan-id*
4. Assign a multicast group address for flood traffic, and specify a source IP address for multicast VXLAN packets.

   group *group-address* **source** *source-address*

   By default, a VXLAN uses unicast mode for flood traffic. No multicast group address or source IP address is specified for multicast VXLAN packets.

   You must assign all VTEPs in a multicast-mode VXLAN to the same multicast group.
5. Enter interface view.

   interface *interface-type interface-number*

   Enable PIM on the loopback interface and all transport-facing interfaces.
6. Enable PIM. Choose one of the following modes:
   - Enable PIM-SM.

     pim sm
   - Enable PIM-DM.

     pim dm

   By default, PIM is disabled on an interface.

Configuring a multicast-mode VXLAN that uses the IGMP host method

1. Enter system view.

   system-view

2. Enter VSI view.

   vsi *vsi-name*

3. Enter VXLAN view.

   vxlan *vxlan-id*

4. Assign a multicast group address for flood traffic, and specify a source IP address for multicast VXLAN packets.

   group *group-address* **source** *source-address*

   By default, a VXLAN uses unicast mode for flood traffic. No multicast group address or source IP address is specified for multicast VXLAN packets.

   You must assign all VTEPs in a multicast-mode VXLAN to the same multicast group.

5. Enter the view of the transport-facing interface.

   interface *interface-type interface-number*

6. Enable the IGMP host feature.

   igmp host enable

   By default, the IGMP host feature is disabled on an interface.

   The IGMP host feature enables the interface to send IGMP reports in response to IGMP queries before it can receive traffic from the multicast group.

Setting the destination UDP port number of VXLAN packets

1. Enter system view.

   system-view

2. Set a destination UDP port for VXLAN packets.

   vxlan udp-port *port-number*

   By default, the destination UDP port number is 4789 for VXLAN packets.

   You must configure the same destination UDP port number on all VTEPs in a VXLAN.

Configuring VXLAN packet check

About this task

The device can check the UDP checksum and 802.1Q VLAN tags of each received VXLAN packet.

- **UDP checksum check**—The device always sets the UDP checksum of VXLAN packets to zero. For compatibility with third-party devices, a VXLAN packet can pass the check if its UDP checksum is zero or correct. If its UDP checksum is incorrect, the VXLAN packet fails the check and is dropped.
- **VLAN tag check**—The device checks the inner Ethernet header of each VXLAN packet for 802.1Q VLAN tags. If the header contains 802.1Q VLAN tags, the device drops the packet.

Restrictions and guidelines

If a remote VTEP uses the Ethernet access mode, its VXLAN packets might contain 802.1Q VLAN tags. To prevent the local VTEP from dropping the VXLAN packets, do not execute the **vxlan invalid-vlan-tag discard** command on the local VTEP.

The access mode is configurable by using the **xconnect vsi** command.

Procedure

1. Enter system view.
   ```
   system-view
   ```
2. Enable the VTEP to drop VXLAN packets that fail UDP checksum check.
   ```
   vxlan invalid-udp-checksum discard
   ```
   By default, the VTEP does not check the UDP checksum of VXLAN packets.
3. Enable the VTEP to drop VXLAN packets that have 802.1Q VLAN tags in the inner Ethernet header.
   ```
   vxlan invalid-vlan-tag discard
   ```
   By default, the VTEP does not check the inner Ethernet header for 802.1Q VLAN tags.

Enabling default VXLAN decapsulation

About this task

If a VXLAN tunnel is configured on only one VTEP of a pair of VTEPs, the VXLAN tunnel is a unidirectional tunnel to the VTEP not configured with the tunnel. In this situation, that VTEP drops the VXLAN packets received from the unidirectional VXLAN tunnel. For a VTEP to receive VXLAN packets from a unidirectional VXLAN tunnel, enable default VXLAN decapsulation on the interface whose IP address is the tunnel destination address. The VTEP will decapsulate all the VXLAN packets destined for the IP address of that interface.

Restrictions and guidelines

This feature takes effect only when the specified interface has an IP address.

Default VXLAN decapsulation does not take effect on bidirectional VXLAN tunnels. If you remove the one-way communication issue for a VXLAN tunnel by configuring the tunnel on both the local and remote VTEPs, this feature no longer takes effect on that tunnel.

Procedure

1. Enter system view.
   ```
   system-view
   ```
2. Enable default VXLAN decapsulation.
   vxlan default-decapsulation source interface *interface-type interface-number*

   By default, default VXLAN decapsulation is disabled.

Disabling flooding for a VSI

About this task

By default, the VTEP floods broadcast, unknown unicast, and unknown multicast frames received from the local site to the following interfaces in the frame's VXLAN:

- All site-facing interfaces except for the incoming interface.
- All VXLAN tunnel interfaces.

When receiving broadcast, unknown unicast, and unknown multicast frames on VXLAN tunnel interfaces, the device floods the frames to all site-facing interfaces in the frames' VXLAN.

To confine a kind of flood traffic, disable flooding for that kind of flood traffic on the VSI bound to the VXLAN.

To exclude a remote MAC address from the remote flood suppression done by using this feature, enable selective flood for the MAC address. The VTEP will flood the frames destined for the MAC address to remote sites.

Procedure

1. Enter system view.
   system-view
2. Enter VSI view.
   vsi *vsi-name*
3. Disable flooding for the VSI.
   flooding disable { **all** | { **broadcast** | **unknown-multicast** | **unknown-unicast** } * } [**all-direction**]
   By default, flooding is enabled for a VSI.
4. (Optional.) Enable selective flood for a MAC address.
   selective-flooding mac-address *mac-address*

Confining the flood traffic of an Ethernet service instance

About this task

By default, an Ethernet service instance sends flood traffic to the other Ethernet service instances of the same VSI. To prevent broadcast storms, you can disable an Ethernet service instance from flooding traffic to the other Ethernet service instances of the same VSI on the local port.

Procedure

1. Enter system view.
   system-view
2. Enter interface view.
   - Enter Layer 2 Ethernet interface view.
     interface *interface-type interface-number*
   - Enter Layer 2 aggregate interface view.
     interface bridge-aggregation *interface-number*
3. Enter Ethernet service instance view.
   service-instance *instance-id*
4. Confine the flood traffic of the Ethernet service instance.
   flooding disable source-port
   By default, an Ethernet service instance sends flood traffic to the other Ethernet service instances of the same VSI.

Enabling ARP flood suppression

Restrictions and guidelines

The aging timer is fixed at 25 minutes for ARP flood suppression entries. If the suppression table is full, the VTEP stops learning new entries. For the VTEP to learn new entries, you must wait for old entries to age out, or use the **reset arp suppression vsi** command to clear the table.

If the **flooding disable** command is configured, set the MAC aging timer to a higher value than the aging timer for ARP flood suppression entries on all VTEPs. This setting prevents the traffic blackhole that occurs when a MAC address entry ages out before its ARP flood suppression entry ages out. To set the MAC aging timer, use the **mac-address timer** command.

When remote ARP learning is disabled for VXLANs, the device does not use ARP flood suppression entries to respond to ARP requests received on VXLAN tunnels.

Procedure

1. Enter system view.
   system-view
2. Enter VSI view.
   vsi *vsi-name*
3. Enable ARP flood suppression.
   arp suppression enable
   By default, ARP flood suppression is disabled.

Enabling VXLAN packet statistics

Enabling packet statistics for a VSI

Restrictions and guidelines

To display the packet statistics for a VSI, use the **display l2vpn vsi verbose** command in any view.

To clear the packet statistics for a VSI, use the **reset l2vpn statistics vsi** command in user view.

Procedure

1. Enter system view.
   system-view
2. Enter VSI view.
   vsi *vsi-name*
3. Enable packet statistics for the VSI.
   statistics enable
   By default, the packet statistics feature is disabled for all VSIs.

Enabling packet statistics for an AC

Restrictions and guidelines

For the **statistics enable** command to take effect on an Ethernet service instance, you must configure a frame match criterion for the Ethernet service instance and map it to a VSI. When you modify the frame match criterion or VSI mapping, the packet statistics of the instance are cleared.

Enabling packet statistics for an Ethernet service instance

1. Enter system view.
   system-view
2. Enter interface view.
   - Enter Layer 2 Ethernet interface view.

interface *interface-type interface-number*

- Enter Layer 2 aggregate interface view.
  interface bridge-aggregation *interface-number*

3. Enter Ethernet service instance view.
   service-instance *instance-id*
4. Enable packet statistics for the Ethernet service instance.
   statistics enable

   By default, the packet statistics feature is disabled for all Ethernet service instances.

Enabling packet statistics for Ethernet service instances of a VLAN

1. Enter system view.
   system-view
2. Enter VLAN view.
   vlan *vlan-id*
3. Enable packet statistics for Ethernet service instances of the VLAN.
   ac statistics enable

   By default, packet statistics are disabled for Ethernet service instances of a VLAN.

   This feature enables packet statistics for the Ethernet service instances that are automatically created for VLAN-based VXLAN assignment. Before you enable this feature, you must use the **vxlan vlan-based** command to enable VLAN-based VXLAN assignment.

Enabling packet statistics for VXLAN tunnels

About this task

VXLAN tunnels can be manually or automatically created. For manually created VXLAN tunnels, you can enable packet statistics on a per-tunnel interface basis. For automatically created VXLAN tunnels, you can enable packet statistics globally in system view.

To display the packet statistics for a VXLAN tunnel, use the **display interface tunnel** command in any view.

To clear the packet statistics for a VXLAN tunnel, use the **reset counters interface tunnel** command in user view.

Enabling packet statistics for a manually created VXLAN tunnel

1. Enter system view.
   system-view
2. Enter VXLAN tunnel interface view.
   interface tunnel *tunnel-number* [**mode vxlan**]
3. Enable packet statistics for the tunnel.
   statistics enable

   By default, the packet statistics feature is disabled for manually created VXLAN tunnels.

Enabling packet statistics for automatically created VXLAN tunnels

1. Enter system view.
   system-view
2. Enable packet statistics for automatically created VXLAN tunnels.
   tunnel statistics vxlan auto [**destination** *ipv4-address*]

   By default, the packet statistics feature is disabled for automatically created VXLAN tunnels.

This command enables the device to collect packet statistics for VXLAN tunnels that are automatically created by EVPN or OVSDB. For more information about EVPN, see *EVPN Configuration Guide*. For more information about OVSDB, see "Configuring the VTEP as an OVSDB VTEP."

Testing the reachability of a remote VM

About this task

This feature enables the device to test the reachability of a remote VM by simulating a local VM to send ICMP echo requests. The requests are encapsulated in Layer 2 data frames and then sent to the remote VM in the specified VXLAN. The device determines the reachability of the remote VM based on the response time and the number of received ICMP echo replies.

Restrictions and guidelines

This feature is not supported if EVPN distributed relay is configured on the device. For more information about EVPN distributed relay, see *EVPN Configuration Guide*.

Procedure

Execute the following command in any view to test the reachability of a remote VM:

emulate-ping vxlan [**-c** *count* | **-m** *interval* | **-s** *packet-size* | **-t** *time-out*] * **vxlan-id** *vxlan-id* **source-mac** *mac-address* **destination-mac** *mac-address*

Display and maintenance commands for VXLANs

Execute **display** commands in any view and **reset** commands in user view.

Task	Command
Display ARP flood suppression entries on VSIs.	**display arp suppression vsi** [**name** *vsi-name*] [**slot** *slot-number*] [**count**]
Display information about the multicast groups that contain IGMP host-enabled interfaces.	**display igmp host group** [*group-address* | **interface** *interface-type interface-number*] [**verbose**]
Display information about tunnel interfaces.	**display interface** [**tunnel** [*number*]] [**brief** [**description** | **down**]]
Display MAC address entries for VSIs.	**display l2vpn mac-address** [**vsi** *vsi-name*] [**dynamic**] [**count** | **verbose**]
Display information about Ethernet service instances.	**display l2vpn service-instance** [**interface** *interface-type interface-number* [**service-instance** *instance-id*]] [**verbose**]
Display information about VSIs.	**display l2vpn vsi** [**name** *vsi-name*] [**verbose**]
Display VXLAN tunnel information for VXLANs.	**display vxlan tunnel** [**vxlan** *vxlan-id*]
Clear ARP flood suppression entries on VSIs.	**reset arp suppression vsi** [**name** *vsi-name*]
Clear dynamic MAC address entries on VSIs.	**reset l2vpn mac-address** [**vsi** *vsi-name*]
Clear packet statistics on ACs.	**reset l2vpn statistics ac** [**interface**

Task	Command
	interface-type interface-number **service-instance** *instance-id*]
Clear packet statistics on VSIs.	**reset l2vpn statistics vsi** [**name** *vsi-name*]

NOTE:

For more information about the **display interface tunnel** command, see tunneling commands in *Layer 3—IP Services Command Reference*.

VXLAN configuration examples

Example: Configuring a unicast-mode VXLAN

Network configuration

As shown in Figure 10:

- Configure VXLAN 10 as a unicast-mode VXLAN on Switch A, Switch B, and Switch C to provide Layer 2 connectivity for the VMs across the network sites.
- Manually establish VXLAN tunnels and assign the tunnels to VXLAN 10.
- Enable remote-MAC address learning.

Figure 10 Network diagram

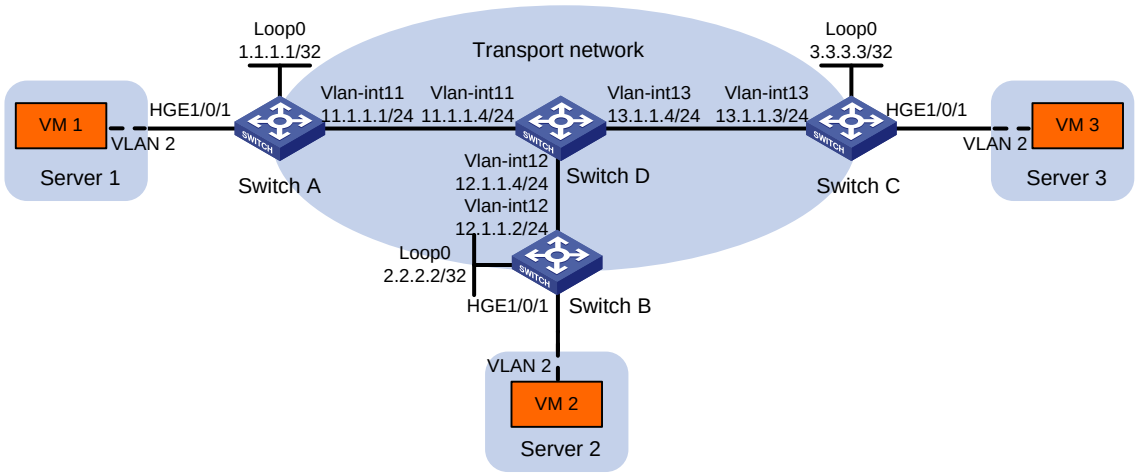

Procedure

1. Create VLANs and VLAN interfaces on all devices. (Details not shown.)
2. Configure IP addresses and unicast routing settings:

   # Assign IP addresses to interfaces, as shown in Figure 10. (Details not shown.)

   # Configure OSPF on all transport network switches (Switches A through D). (Details not shown.)
3. Configure Switch A:

   # Enable L2VPN.

   ```
   <SwitchA> system-view
   [SwitchA] l2vpn enable
   ```

Create VSI **vpna** and VXLAN 10.

```
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] vxlan 10
[SwitchA-vsi-vpna-vxlan-10] quit
[SwitchA-vsi-vpna] quit
```

Assign an IP address to Loopback 0. The IP address will be used as the source IP address of the VXLAN tunnels to Switch B and Switch C.

```
[SwitchA] interface loopback 0
[SwitchA-Loopback0] ip address 1.1.1.1 255.255.255.255
[SwitchA-Loopback0] quit
```

Create a VXLAN tunnel to Switch B. The tunnel interface name is **Tunnel 1**.

```
[SwitchA] interface tunnel 1 mode vxlan
[SwitchA-Tunnel1] source 1.1.1.1
[SwitchA-Tunnel1] destination 2.2.2.2
[SwitchA-Tunnel1] quit
```

Create a VXLAN tunnel to Switch C. The tunnel interface name is **Tunnel 2**.

```
[SwitchA] interface tunnel 2 mode vxlan
[SwitchA-Tunnel2] source 1.1.1.1
[SwitchA-Tunnel2] destination 3.3.3.3
[SwitchA-Tunnel2] quit
```

Assign Tunnel 1 and Tunnel 2 to VXLAN 10.

```
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] vxlan 10
[SwitchA-vsi-vpna-vxlan-10] tunnel 1
[SwitchA-vsi-vpna-vxlan-10] tunnel 2
[SwitchA-vsi-vpna-vxlan-10] quit
[SwitchA-vsi-vpna] quit
```

On HundredGigE 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchA] interface hundredgige 1/0/1
[SwitchA-HundredGigE1/0/1] port link-type trunk
[SwitchA-HundredGigE1/0/1] port trunk permit vlan 2
[SwitchA-HundredGigE1/0/1] service-instance 1000
[SwitchA-HundredGigE1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI **vpna**.

```
[SwitchA-HundredGigE1/0/1-srv1000] xconnect vsi vpna
[SwitchA-HundredGigE1/0/1-srv1000] quit
[SwitchA-HundredGigE1/0/1] quit
```

4. Configure Switch B:

Enable L2VPN.

```
<SwitchB> system-view
[SwitchB] l2vpn enable
```

Create VSI **vpna** and VXLAN 10.

```
[SwitchB] vsi vpna
[SwitchB-vsi-vpna] vxlan 10
[SwitchB-vsi-vpna-vxlan-10] quit
[SwitchB-vsi-vpna] quit
```

Assign an IP address to Loopback 0. The IP address will be used as the source IP address of the VXLAN tunnels to Switch A and Switch C.

```
[SwitchB] interface loopback 0
[SwitchB-Loopback0] ip address 2.2.2.2 255.255.255.255
[SwitchB-Loopback0] quit
```

Create a VXLAN tunnel to Switch A. The tunnel interface name is **Tunnel 2**.

```
[SwitchB] interface tunnel 2 mode vxlan
[SwitchB-Tunnel2] source 2.2.2.2
[SwitchB-Tunnel2] destination 1.1.1.1
[SwitchB-Tunnel2] quit
```

Create a VXLAN tunnel to Switch C. The tunnel interface name is **Tunnel 3**.

```
[SwitchB] interface tunnel 3 mode vxlan
[SwitchB-Tunnel3] source 2.2.2.2
[SwitchB-Tunnel3] destination 3.3.3.3
[SwitchB-Tunnel3] quit
```

Assign Tunnel 2 and Tunnel 3 to VXLAN 10.

```
[SwitchB] vsi vpna
[SwitchB-vsi-vpna] vxlan 10
[SwitchB-vsi-vpna-vxlan-10] tunnel 2
[SwitchB-vsi-vpna-vxlan-10] tunnel 3
[SwitchB-vsi-vpna-vxlan-10] quit
[SwitchB-vsi-vpna] quit
```

On HundredGigE 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchB] interface hundredgige 1/0/1
[SwitchB-HundredGigE1/0/1] port link-type trunk
[SwitchB-HundredGigE1/0/1] port trunk permit vlan 2
[SwitchB-HundredGigE1/0/1] service-instance 1000
[SwitchB-HundredGigE1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI **vpna**.

```
[SwitchB-HundredGigE1/0/1-srv1000] xconnect vsi vpna
[SwitchB-HundredGigE1/0/1-srv1000] quit
[SwitchB-HundredGigE1/0/1] quit
```

5. Configure Switch C:

Enable L2VPN.

```
<SwitchC> system-view
[SwitchC] l2vpn enable
```

Create VSI **vpna** and VXLAN 10.

```
[SwitchC] vsi vpna
[SwitchC-vsi-vpna] vxlan 10
[SwitchC-vsi-vpna-vxlan-10] quit
[SwitchC-vsi-vpna] quit
```

Assign an IP address to Loopback 0. The IP address will be used as the source IP address of the VXLAN tunnels to Switch A and Switch B.

```
[SwitchC] interface loopback 0
[SwitchC-Loopback0] ip address 3.3.3.3 255.255.255.255
[SwitchC-Loopback0] quit
```

Create a VXLAN tunnel to Switch A. The tunnel interface name is **Tunnel 1**.

```
[SwitchC] interface tunnel 1 mode vxlan
[SwitchC-Tunnel1] source 3.3.3.3
[SwitchC-Tunnel1] destination 1.1.1.1
```

```
[SwitchC-Tunnel1] quit
```

Create a VXLAN tunnel to Switch B. The tunnel interface name is **Tunnel 3**.

```
[SwitchC] interface tunnel 3 mode vxlan
[SwitchC-Tunnel3] source 3.3.3.3
[SwitchC-Tunnel3] destination 2.2.2.2
[SwitchC-Tunnel3] quit
```

Assign Tunnel 1 and Tunnel 3 to VXLAN 10.

```
[SwitchC] vsi vpna
[SwitchC-vsi-vpna] vxlan 10
[SwitchC-vsi-vpna-vxlan-10] tunnel 1
[SwitchC-vsi-vpna-vxlan-10] tunnel 3
[SwitchC-vsi-vpna-vxlan-10] quit
[SwitchC-vsi-vpna] quit
```

On HundredGigE 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchC] interface hundredgige 1/0/1
[SwitchC-HundredGigE1/0/1] port link-type trunk
[SwitchC-HundredGigE1/0/1] port trunk permit vlan 2
[SwitchC-HundredGigE1/0/1] service-instance 1000
[SwitchC-HundredGigE1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI **vpna**.

```
[SwitchC-HundredGigE1/0/1-srv1000] xconnect vsi vpna
[SwitchC-HundredGigE1/0/1-srv1000] quit
[SwitchC-HundredGigE1/0/1] quit
```

Verifying the configuration

1. Verify the VXLAN settings on the VTEPs. This example uses Switch A.

Verify that the VXLAN tunnel interfaces on the VTEP are up.

```
[SwitchA] display interface tunnel 1
Tunnel1
Current state: UP
Line protocol state: UP
Description: Tunnel1 Interface
Bandwidth: 64 kbps
Maximum transmission unit: 1464
Internet protocol processing: Disabled
Last clearing of counters: Never
Tunnel source 1.1.1.1, destination 2.2.2.2
Tunnel protocol/transport UDP_VXLAN/IP
Last 300 seconds input rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Last 300 seconds output rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Input: 0 packets, 0 bytes, 0 drops
Output: 0 packets, 0 bytes, 0 drops
```

Verify that the VXLAN tunnels have been assigned to the VXLAN.

```
[SwitchA] display l2vpn vsi verbose
VSI Name: vpna
  VSI Index               : 0
  VSI State               : Up
  MTU                     : 1500
```

```
  Bandwidth               : Unlimited
  Broadcast Restrain      : Unlimited
  Multicast Restrain      : Unlimited
  Unknown Unicast Restrain: Unlimited
  MAC Learning            : Enabled
  MAC Table Limit         : -
  MAC Learning rate       : -
  Drop Unknown            : -
  Flooding                : Enabled
  Statistics              : Disabled
  VXLAN ID                : 10
  Tunnels:
    Tunnel Name          Link ID    State  Type        Flood proxy
    Tunnel1              0x5000001  Up     Manual      Disabled
    Tunnel2              0x5000002  Up     Manual      Disabled
  ACs:
    AC                              Link ID    State    Type
    HGE1/0/1 srv1000                0          Up       Manual
```

Verify that the VTEP has learned the MAC addresses of remote VMs.

```
<SwitchA> display l2vpn mac-address
MAC Address : cc3e-5f9c-6cdb
VSI Name    : vpna
State       : Dynamic
Link ID/Name      Aging
Tunnel1           Aging

MAC Address : cc3e-5f9c-23dc
VSI Name    : vpna
State       : Dynamic
Link ID/Name      Aging
Tunnel2           Aging
--- 2 mac address(es) found  ---
```

2. Verify that VM 1, VM 2, and VM 3 can ping each other. (Details not shown.)

Example: Configuring a multicast-mode VXLAN

Network configuration

As shown in Figure 11:

- Configure VXLAN 10 as a multicast-mode VXLAN on Switch A, Switch B, and Switch C to provide Layer 2 connectivity for the VMs across the network sites.
- Manually establish VXLAN tunnels and assign the tunnels to VXLAN 10.
- Enable remote-MAC address learning.

Figure 11 Network diagram

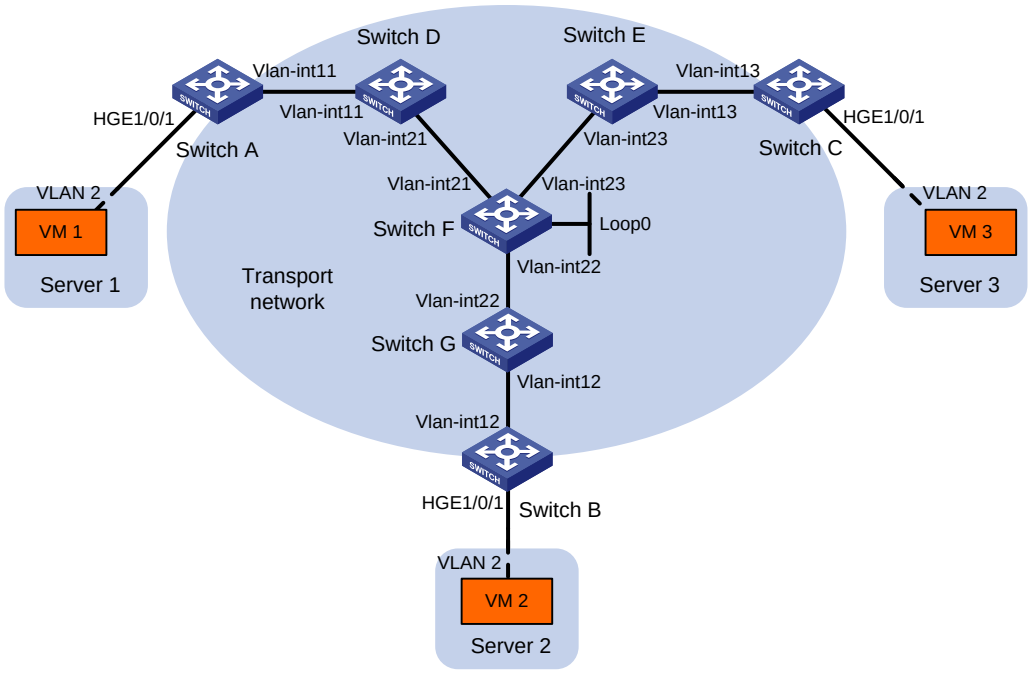

Table 1 IP address assignment

Device	Interface	IP address	Device	Interface	IP address
Switch A:			**Switch C:**		
	VLAN-interface 11	11.1.1.1/24		VLAN-interface 13	13.1.1.3/24
Switch D:			**Switch E:**		
	VLAN-interface 11	11.1.1.4/24		VLAN-interface 13	13.1.1.5/24
	VLAN-interface 21	21.1.1.4/24		VLAN-interface 23	23.1.1.5/24
Switch F:			**Switch G:**		
	VLAN-interface 21	21.1.1.6/24		VLAN-interface 12	12.1.1.7/24
	VLAN-interface 22	22.1.1.6/24		VLAN-interface 22	22.1.1.7/24
	VLAN-interface 23	23.1.1.6/24	**Switch B:**		
	Loop 0	6.6.6.6/32		VLAN-interface 12	12.1.1.2/24

Procedure

1. Create VLANs and VLAN interfaces on all devices. (Details not shown.)
2. Configure IP addresses and unicast routing settings:

   # Assign IP addresses to interfaces, as shown in Figure 11. (Details not shown.)

   # Configure OSPF on all transport network switches (Switches A through G). (Details not shown.)
3. Configure Switch A:

   # Enable L2VPN.

   ```
   <SwitchA> system-view
   [SwitchA] l2vpn enable
   ```

   # Enable IP multicast routing.

   ```
   [SwitchA] multicast routing
   [SwitchA-mrib] quit
   ```

Create VSI **vpna** and VXLAN 10.

```
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] vxlan 10
[SwitchA-vsi-vpna-vxlan-10] quit
[SwitchA-vsi-vpna] quit
```

Assign an IP address to VLAN-interface 11, and enable the IGMP host feature on the interface. This interface's IP address will be the source IP address of VXLAN packets sent by the VTEP.

```
[SwitchA] interface vlan-interface 11
[SwitchA-Vlan-interface11] ip address 11.1.1.1 24
[SwitchA-Vlan-interface11] igmp host enable
[SwitchA-Vlan-interface11] quit
```

Create a VXLAN tunnel to Switch B. The tunnel interface name is **Tunnel 1**.

```
[SwitchA] interface tunnel 1 mode vxlan
[SwitchA-Tunnel1] source 11.1.1.1
[SwitchA-Tunnel1] destination 12.1.1.2
[SwitchA-Tunnel1] quit
```

Create a VXLAN tunnel to Switch C. The tunnel interface name is **Tunnel 2**.

```
[SwitchA] interface tunnel 2 mode vxlan
[SwitchA-Tunnel2] source 11.1.1.1
[SwitchA-Tunnel2] destination 13.1.1.3
[SwitchA-Tunnel2] quit
```

Assign Tunnel 1 and Tunnel 2 to VXLAN 10.

```
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] vxlan 10
[SwitchA-vsi-vpna-vxlan-10] tunnel 1
[SwitchA-vsi-vpna-vxlan-10] tunnel 2
```

Configure the multicast group address and source IP address for multicast VXLAN packets.

```
[SwitchA-vsi-vpna-vxlan-10] group 225.1.1.1 source 11.1.1.1
[SwitchA-vsi-vpna-vxlan-10] quit
[SwitchA-vsi-vpna] quit
```

On HundredGigE 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchA] interface hundredgige 1/0/1
[SwitchA-HundredGigE1/0/1] port link-type trunk
[SwitchA-HundredGigE1/0/1] port trunk permit vlan 2
[SwitchA-HundredGigE1/0/1] service-instance 1000
[SwitchA-HundredGigE1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI **vpna**.

```
[SwitchA-HundredGigE1/0/1-srv1000] xconnect vsi vpna
[SwitchA-HundredGigE1/0/1-srv1000] quit
[SwitchA-HundredGigE1/0/1] quit
```

4. Configure Switch B:

Enable L2VPN.

```
<SwitchB> system-view
[SwitchB] l2vpn enable
```

Enable IP multicast routing.

```
[SwitchB] multicast routing
[SwitchB-mrib] quit
```

Create VSI **vpna** and VXLAN 10.

```
[SwitchB] vsi vpna
[SwitchB-vsi-vpna] vxlan 10
[SwitchB-vsi-vpna-vxlan-10] quit
[SwitchB-vsi-vpna] quit
```

Assign an IP address to VLAN-interface 12, and enable the IGMP host feature on the interface. This interface's IP address will be the source IP address of VXLAN packets sent by the VTEP.

```
[SwitchB] interface vlan-interface 12
[SwitchB-Vlan-interface12] ip address 12.1.1.2 24
[SwitchB-Vlan-interface12] igmp host enable
[SwitchB-Vlan-interface12] quit
```

Create a VXLAN tunnel to Switch A. The tunnel interface name is **Tunnel 2**.

```
[SwitchB] interface tunnel 2 mode vxlan
[SwitchB-Tunnel2] source 12.1.1.2
[SwitchB-Tunnel2] destination 11.1.1.1
[SwitchB-Tunnel2] quit
```

Create a VXLAN tunnel to Switch C. The tunnel interface name is **Tunnel 3**.

```
[SwitchB] interface tunnel 3 mode vxlan
[SwitchB-Tunnel3] source 12.1.1.2
[SwitchB-Tunnel3] destination 13.1.1.3
[SwitchB-Tunnel3] quit
```

Assign Tunnel 2 and Tunnel 3 to VXLAN 10.

```
[SwitchB] vsi vpna
[SwitchB-vsi-vpna] vxlan 10
[SwitchB-vsi-vpna-vxlan-10] tunnel 2
[SwitchB-vsi-vpna-vxlan-10] tunnel 3
```

Configure the VXLAN multicast group address and the source IP address for VXLAN packets.

```
[SwitchB-vsi-vpna-vxlan-10] group 225.1.1.1 source 12.1.1.2
[SwitchB-vsi-vpna-vxlan-10] quit
[SwitchB-vsi-vpna] quit
```

On HundredGigE 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchB] interface hundredgige 1/0/1
[SwitchB-HundredGigE1/0/1] port link-type trunk
[SwitchB-HundredGigE1/0/1] port trunk permit vlan 2
[SwitchB-HundredGigE1/0/1] service-instance 1000
[SwitchB-HundredGigE1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI **vpna**.

```
[SwitchB-HundredGigE1/0/1-srv1000] xconnect vsi vpna
[SwitchB-HundredGigE1/0/1-srv1000] quit
[SwitchB-HundredGigE1/0/1] quit
```

5. Configure Switch C:

Enable L2VPN.

```
<SwitchC> system-view
[SwitchC] l2vpn enable
```

Enable IP multicast routing.

```
[SwitchC] multicast routing
[SwitchC-mrib] quit
```

Create VSI **vpna** and VXLAN 10.

```
[SwitchC] vsi vpna
[SwitchC-vsi-vpna] vxlan 10
[SwitchC-vsi-vpna-vxlan-10] quit
[SwitchC-vsi-vpna] quit
```

Assign an IP address to VLAN-interface 13, and enable the IGMP host feature on the interface. This interface's IP address will be the source IP address of VXLAN packets sent by the VTEP.

```
[SwitchC] interface vlan-interface 13
[SwitchC-Vlan-interface13] ip address 13.1.1.3 24
[SwitchC-Vlan-interface13] igmp host enable
[SwitchC-Vlan-interface13] quit
```

Create a VXLAN tunnel to Switch A. The tunnel interface name is **Tunnel 1**.

```
[SwitchC] interface tunnel 1 mode vxlan
[SwitchC-Tunnel1] source 13.1.1.3
[SwitchC-Tunnel1] destination 11.1.1.1
[SwitchC-Tunnel1] quit
```

Create a VXLAN tunnel to Switch B. The tunnel interface name is **Tunnel 3**.

```
[SwitchC] interface tunnel 3 mode vxlan
[SwitchC-Tunnel3] source 13.1.1.3
[SwitchC-Tunnel3] destination 12.1.1.2
[SwitchC-Tunnel3] quit
```

Assign Tunnel 1 and Tunnel 3 to VXLAN 10.

```
[SwitchC] vsi vpna
[SwitchC-vsi-vpna] vxlan 10
[SwitchC-vsi-vpna-vxlan-10] tunnel 1
[SwitchC-vsi-vpna-vxlan-10] tunnel 3
```

Configure the multicast group address and source IP address for VXLAN multicast packets.

```
[SwitchC-vsi-vpna-vxlan-10] group 225.1.1.1 source 13.1.1.3
[SwitchC-vsi-vpna-vxlan-10] quit
[SwitchC-vsi-vpna] quit
```

On HundredGigE 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchC] interface hundredgige 1/0/1
[SwitchC-HundredGigE1/0/1] port link-type trunk
[SwitchC-HundredGigE1/0/1] port trunk permit vlan 2
[SwitchC-HundredGigE1/0/1] service-instance 1000
[SwitchC-HundredGigE1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI **vpna**.

```
[SwitchC-HundredGigE1/0/1-srv1000] xconnect vsi vpna
[SwitchC-HundredGigE1/0/1-srv1000] quit
[SwitchC-HundredGigE1/0/1] quit
```

6. Configure Switch D:

Enable IP multicast routing.

```
<SwitchD> system-view
[SwitchD] multicast routing
[SwitchD-mrib] quit
```

Enable IGMP and PIM-SM on VLAN-interface 11.

```
[SwitchD] interface vlan-interface 11
```

```
[SwitchD-Vlan-interface11] igmp enable
[SwitchD-Vlan-interface11] pim sm
[SwitchD-Vlan-interface11] quit
```

Enable PIM-SM on VLAN-interface 21.

```
[SwitchD] interface vlan-interface 21
[SwitchD-Vlan-interface21] pim sm
[SwitchD-Vlan-interface21] quit
```

Enable BIDIR-PIM.

```
[SwitchD] pim
[SwitchD-pim] bidir-pim enable
[SwitchD-pim] quit
```

7. Configure Switch E:

Enable IP multicast routing.

```
<SwitchE> system-view
[SwitchE] multicast routing
[SwitchE-mrib] quit
```

Enable IGMP and PIM-SM on VLAN-interface 13.

```
[SwitchE] interface vlan-interface 13
[SwitchE-Vlan-interface13] igmp enable
[SwitchE-Vlan-interface13] pim sm
[SwitchE-Vlan-interface13] quit
```

Enable PIM-SM on VLAN-interface 23.

```
[SwitchE] interface vlan-interface 23
[SwitchE-Vlan-interface23] pim sm
[SwitchE-Vlan-interface23] quit
```

Enable BIDIR-PIM.

```
[SwitchE] pim
[SwitchE-pim] bidir-pim enable
[SwitchE-pim] quit
```

8. Configure Switch F:

Enable IP multicast routing.

```
<SwitchF> system-view
[SwitchF] multicast routing
[SwitchF-mrib] quit
```

Enable PIM-SM on VLAN-interface 21, VLAN-interface 22, VLAN-interface 23, and Loopback 0.

```
[SwitchF] interface vlan-interface 21
[SwitchF-Vlan-interface21] pim sm
[SwitchF-Vlan-interface21] quit
[SwitchF] interface vlan-interface 22
[SwitchF-Vlan-interface22] pim sm
[SwitchF-Vlan-interface22] quit
[SwitchF] interface vlan-interface 23
[SwitchF-Vlan-interface23] pim sm
[SwitchF-Vlan-interface23] quit
[SwitchF] interface loopback 0
[SwitchF-LoopBack0] pim sm
[SwitchF-LoopBack0] quit
```

Enable BIDIR-PIM.

```
[SwitchF] pim
[SwitchF-pim] bidir-pim enable
```

Configure VLAN-interface 22 as a candidate-BSR, and configure Loopback 0 as a candidate-RP for BIDIR-PIM.

```
[SwitchF-pim] c-bsr 22.1.1.6
[SwitchF-pim] c-rp 6.6.6.6 bidir
[SwitchF-pim] quit
```

9. Configure Switch G:

Enable IP multicast routing.

```
<SwitchG> system-view
[SwitchG] multicast routing
[SwitchG-mrib] quit
```

Enable IGMP and PIM-SM on VLAN-interface 12.

```
[SwitchG] interface vlan-interface 12
[SwitchG-Vlan-interface12] igmp enable
[SwitchG-Vlan-interface12] pim sm
[SwitchG-Vlan-interface12] quit
```

Enable PIM-SM on VLAN-interface 22.

```
[SwitchG] interface vlan-interface 22
[SwitchG-Vlan-interface22] pim sm
[SwitchG-Vlan-interface22] quit
```

Enable BIDIR-PIM.

```
[SwitchG] pim
[SwitchG-pim] bidir-pim enable
[SwitchG-pim] quit
```

Verifying the configuration

1. Verify the VXLAN settings on the VTEPs. This example uses Switch A.

Verify that the VXLAN tunnel interfaces on the VTEP are up.

```
[SwitchA] display interface tunnel 1
Tunnel1
Current state: UP
Line protocol state: UP
Description: Tunnel1 Interface
Bandwidth: 64 kbps
Maximum transmission unit: 1464
Internet protocol processing: Disabled
Last clearing of counters: Never
Tunnel source 11.1.1.1, destination 12.1.1.2
Tunnel protocol/transport UDP_VXLAN/IP
Last 300 seconds input rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Last 300 seconds output rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Input: 0 packets, 0 bytes, 0 drops
Output: 0 packets, 0 bytes, 0 drops
```

Verify that the VXLAN tunnels have been assigned to the VXLAN.

```
[SwitchA] display l2vpn vsi verbose
VSI Name: vpna
```

```
  VSI Index               : 0
  VSI State               : Up
  MTU                     : 1500
  Bandwidth               : Unlimited
  Broadcast Restrain      : Unlimited
  Multicast Restrain      : Unlimited
  Unknown Unicast Restrain: Unlimited
  MAC Learning            : Enabled
  MAC Table Limit         : -
  MAC Learning rate       : -
  Drop Unknown            : -
  Flooding                : Enabled
  Statistics              : Disabled
  VXLAN ID                : 10
  Tunnels:
    Tunnel Name           Link ID    State  Type        Flood proxy
    Tunnel1               0x5000001  Up     Manual      Disabled
    Tunnel2               0x5000002  Up     Manual      Disabled
    MTunnel0              0x6000000  Up     Auto        Disabled
  ACs:
    AC                                Link ID    State    Type
    HGE1/0/1 srv1000                  0          Up       Manual
```

Verify that the VTEP has learned the MAC addresses of remote VMs.

```
<SwitchA> display l2vpn mac-address
MAC Address : cc3e-5f9c-6cdb
VSI Name    : vpna
State       : Dynamic
Link ID/Name      Aging
Tunnel1           Aging

MAC Address : cc3e-5f9c-23dc
VSI Name    : vpna
State       : Dynamic
Link ID/Name      Aging
Tunnel2           Aging
--- 2 mac address(es) found  ---
```

Verify that the VTEP has joined the VXLAN multicast group on VLAN-interface 11.

```
<SwitchA> display igmp host group
IGMP host groups in total: 1
 Vlan-interface11(11.1.1.1):
  IGMP host groups in total: 1
   Group address      Member state      Expires
   225.1.1.1          Idle              Off
```

2. Verify that VM 1, VM 2, and VM 3 can ping each other. (Details not shown.)

Configuring the VTEP as an OVSDB VTEP

About OVSDB VTEP

An H3C network virtualization controller can use the Open vSwitch Database (OVSDB) management protocol to deploy and manage VXLANs on VTEPs. To work with a controller, you must configure the VTEP as an OVSDB VTEP.

Working mechanisms

As shown in Figure 12, an OVSDB VTEP stores all of its VXLAN settings in the form of entries in an OVSDB database. The OVSDB database, OVSDB VTEP service, and the controller interact through the OVSDB server. The controller communicates with the OVSDB server through the OVSDB protocol to manage the OVSDB database. The OVSDB VTEP service reads and writes data in the OVSDB database through the OVSDB server.

The OVSDB VTEP service performs the following operations to manage the VXLAN settings on the VTEP:

- Converts data in the OVSDB database into VXLAN configuration and deploys the configuration to the VTEP. For example, create or remove a VXLAN or VXLAN tunnel.
- Adds site-facing interface information and the global source address of VXLAN tunnels to the OVSDB database. The information is reported to the controller by the OVSDB server.

Figure 12 OVSDB network model

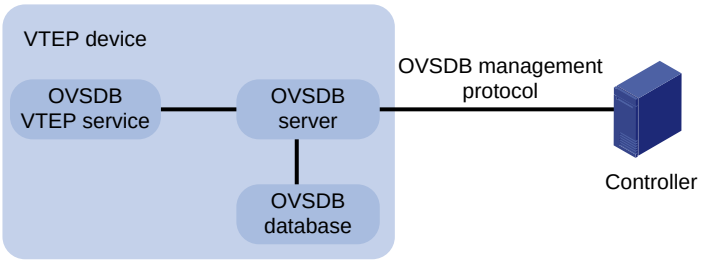

Protocols and standards

RFC 7047, *The Open vSwitch Database Management Protocol*

Restrictions and guidelines: OVSDB VTEP configuration

You can configure a VTEP both at the CLI and through a controller. As a best practice, do not manually remove the VXLAN configuration issued by the controller.

OVSDB VTEP tasks at a glance

To configure OVSDB VTEPs, perform the following tasks:

1. Setting up an OVSDB connection to a controller

    - Configuring active SSL connection settings
    - Configuring passive SSL connection settings
    - Configuring active TCP connection settings
    - Configuring passive TCP connection settings
2. Enabling the OVSDB server
3. Enabling the OVSDB VTEP service
4. Specifying a global source address for VXLAN tunnels
5. Specifying a VTEP access port
6. Enabling flood proxy on multicast VXLAN tunnels

   If you use a flood proxy server, you must enable flood proxy globally on multicast tunnels.
7. (Optional.) Disabling the ACLs issued by the OVSDB controller

Prerequisites for OVSDB VTEP configuration

Before you configure the VTEP as an OVSDB VTEP, enable L2VPN by using the **`l2vpn enable`** command.

Before you set up SSL connections to controllers, you must configure SSL as described in *Security Configuration Guide*.

Setting up an OVSDB connection to a controller

About OVSDB connection types

The OVSDB server supports the following types of OVSDB connections:

- **Active SSL connection**—The OVSDB server initiates an SSL connection to the controller.
- **Passive SSL connection**—The OVSDB server accepts the SSL connection from the controller.
- **Active TCP connection**—The OVSDB server initiates a TCP connection to the controller.
- **Passive TCP connection**—The OVSDB server accepts the TCP connection from the controller.

Restrictions and guidelines for OVSDB controller connection setup

When you set up OVSDB connections, follow these restrictions and guidelines:

- You can set up multiple OVSDB connections. For the device to establish the connections, you must enable the OVSDB server. You must disable and then re-enable the OVSDB server if it has been enabled.
- You must specify the same PKI domain and CA certificate file for all active and passive SSL connections.

Prerequisites for OVSDB controller connection setup

Make sure you have configured a PKI domain before specifying it for SSL. For more information about configuring a PKI domain, see *Security Configuration Guide*.

Configuring active SSL connection settings

1. Enter system view.
   system-view
2. Specify a PKI domain for SSL.
   ovsdb server pki domain *domain-name*
   By default, no PKI domain is specified for SSL.
3. (Optional.) Specify a CA certificate file for SSL.
   ovsdb server bootstrap ca-certificate *ca-filename*
   By default, SSL uses the CA certificate file in the PKI domain.
   If the specified CA certificate file does not exist, the device obtains a self-signed certificate from the controller. The obtained file uses the name specified for the *ca-filename* argument.
4. Set up an active SSL connection.
   ovsdb server ssl ip *ip-address* **port** *port-number*
   By default, the device does not have active OVSDB SSL connections.
   You can set up a maximum of eight OVSDB SSL connections.

Configuring passive SSL connection settings

1. Enter system view.
   system-view
2. Specify a PKI domain for SSL.
   ovsdb server pki domain *domain-name*
   By default, no PKI domain is specified for SSL.
3. (Optional.) Specify a CA certificate file for SSL.
   ovsdb server bootstrap ca-certificate *ca-filename*
   By default, SSL uses the CA certificate file in the PKI domain.
   If the specified CA certificate file does not exist, the device obtains a self-signed certificate from the controller. The obtained file uses the name specified for the *ca-filename* argument.
4. Enable the device to listen for SSL connection requests.
   ovsdb server pssl [**port** *port-number*]
   By default, the device does not listen for SSL connection requests.
   You can specify only one port to listen for OVSDB SSL connection requests.

Configuring active TCP connection settings

1. Enter system view.
   system-view
2. Set up an active TCP connection.
   ovsdb server tcp ip *ip-address* **port** *port-number*
   By default, the device does not have active OVSDB TCP connections.
   You can set up a maximum of eight active OVSDB TCP connections.

Configuring passive TCP connection settings

1. Enter system view.
   system-view
2. Enable the device to listen for TCP connection requests.
   ovsdb server ptcp [**port** *port-number*] [**acl** *acl-number*]
   By default, the device does not listen for TCP connection requests.
   You can specify only one port to listen for OVSDB TCP connection requests.

Enabling the OVSDB server

Prerequisites

Make sure you have complete OVSDB connection setup before you enable the OVSDB server. If you change OVSDB connection settings after the OVSDB server is enabled, you must disable and then re-enable the OVSDB server for the change to take effect.

Procedure

1. Enter system view.
   system-view
2. Enable the OVSDB server.
   ovsdb server enable
   By default, the OVSDB server is disabled.

Enabling the OVSDB VTEP service

1. Enter system view.
   system-view
2. Enable the OVSDB VTEP service.
   vtep enable
   By default, the OVSDB VTEP service is disabled.

Specifying a global source address for VXLAN tunnels

About this task

The VTEP reports the global VXLAN tunnel source address to the controller for VXLAN tunnel setup.

Restrictions and guidelines

For correct VXLAN deployment and VTEP management, do not manually specify tunnel-specific source addresses for VXLAN tunnels if OVSDB is used.

Procedure

1. Enter system view.
   system-view
2. Specify a global source address for VXLAN tunnels.

tunnel global source-address { *ipv4-address* | **ipv6** *ipv6-address* }

By default, no global source address is specified for VXLAN tunnels.

Specifying a VTEP access port

About this task

For the controller to manage a site-facing interface, you must specify the interface as a VTEP access port.

Procedure

1. Enter system view.

   system-view
2. Enter interface view.

   interface *interface-type interface-number*
3. Specify the interface as a VTEP access port.

   vtep access port

   By default, an interface is not a VTEP access port.

Enabling flood proxy on multicast VXLAN tunnels

About this task

If you use a flood proxy server, you must enable flood proxy globally on multicast tunnels. Then the multicast tunnels are converted into flood proxy tunnels. The VTEP sends broadcast, multicast, and unknown unicast traffic for a VXLAN to the flood proxy server through the tunnels. The flood proxy server then replicates and forwards flood traffic to remote VTEPs.

Restrictions and guidelines

Flood proxy is supported on multicast VXLAN tunnels only when the OVSDB controller is a NSX controller from VMware.

After you enable flood proxy on multicast VXLAN tunnels, if the controller issues VSI configuration, the system automatically disables ARP flood suppression on all VSIs issued by the controller. If the controller does not issue VSI configuration, the system does not automatically change the state of ARP flood suppression.

If you do not enable flood proxy on multicast VXLAN tunnels, the system does not automatically change the state of ARP flood suppression regardless of whether the controller issues VSI configuration.

Procedure

1. Enter system view.

   system-view
2. Enable flood proxy on multicast VXLAN tunnels.

   vxlan tunnel flooding-proxy

   By default, flood proxy is disabled on multicast VXLAN tunnels.

Disabling the ACLs issued by the OVSDB controller

About this task

Perform this task on a VTEP to disable all the ACLs issued by the OVSDB controller in order to save ACL resources on the VTEP.

Prerequisites

Before you perform this task, you must enable the OVSDB VTEP service by using the **vtep enable** command.

Procedure

1. Enter system view.
   ```
   system-view
   ```
2. Disable the ACLs issued by the OVSDB controller.
   ```
   vtep acl disable
   ```
   By default, the ACLs issued by the OVSDB controller are enabled on a VTEP.

OVSDB VTEP configuration examples

Example: Configuring a unicast-mode VXLAN

Network configuration

As shown in Figure 13, configure the controller cluster to deploy unicast-mode VXLAN 10 to Switch A, Switch B, and Switch C to provide Layer 2 connectivity for the VMs across the network sites.

Figure 13 Network diagram

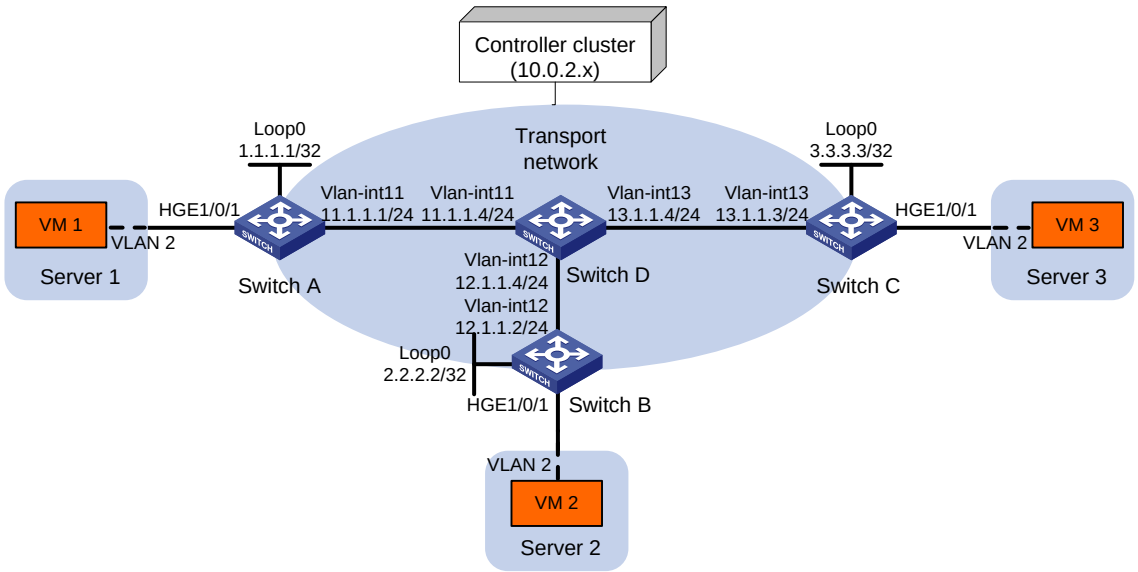

Procedure

1. Create VLANs and VLAN interfaces on all devices. (Details not shown.)
2. Configure IP addresses and unicast routing settings:

Assign IP addresses to interfaces, as shown in Figure 13. (Details not shown.)

Configure OSPF on all transport network switches (Switches A through D). (Details not shown.)

3. Configure Switch A:

   # Enable L2VPN.

   ```
   <SwitchA> system-view
   [SwitchA] l2vpn enable
   ```

   # Configure active TCP connection settings.

   ```
   [SwitchA] ovsdb server tcp ip 10.0.2.15 port 6632
   ```

   # Enable the OVSDB server.

   ```
   [SwitchA] ovsdb server enable
   ```

   # Enable the OVSDB VTEP service.

   ```
   [SwitchA] vtep enable
   ```

   # Assign an IP address to Loopback 0. Specify the IP address as the global source address for VXLAN tunnels.

   ```
   [SwitchA] interface loopback 0
   [SwitchA-LoopBack0] ip address 1.1.1.1 255.255.255.255
   [SwitchA-LoopBack0] quit
   [SwitchA] tunnel global source-address 1.1.1.1
   ```

   # Specify site-facing interface HundredGigE 1/0/1 as a VTEP access port.

   ```
   [SwitchA] interface hundredgige 1/0/1
   [SwitchA-HundredGigE1/0/1] vtep access port
   [SwitchA-HundredGigE1/0/1] quit
   ```

4. Configure Switch B:

   # Enable L2VPN.

   ```
   <SwitchB> system-view
   [SwitchB] l2vpn enable
   ```

   # Configure active TCP connection settings.

   ```
   [SwitchB] ovsdb server tcp 10.0.2.15 port 6632
   ```

   # Enable the OVSDB server.

   ```
   [SwitchB] ovsdb server enable
   ```

   # Enable the OVSDB VTEP service.

   ```
   [SwitchB] vtep enable
   ```

   # Assign an IP address to Loopback 0. Specify the IP address as the global source address for VXLAN tunnels.

   ```
   [SwitchB] interface loopback 0
   [SwitchB-LoopBack0] ip address 2.2.2.2 255.255.255.255
   [SwitchB-LoopBack0] quit
   [SwitchB] tunnel global source-address 2.2.2.2
   ```

   # Specify site-facing interface HundredGigE 1/0/1 as a VTEP access port.

   ```
   [SwitchB] interface hundredgige 1/0/1
   [SwitchB-HundredGigE1/0/1] vtep access port
   [SwitchB-HundredGigE1/0/1] quit
   ```

5. Configure Switch C:

   # Enable L2VPN.

   ```
   <SwitchC> system-view
   [SwitchC] l2vpn enable
   ```

Configure active TCP connection settings.

```
[SwitchC] ovsdb server tcp ip 10.0.2.15 port 6632
```

Enable the OVSDB server.

```
[SwitchC] ovsdb server enable
```

Enable the OVSDB VTEP service.

```
[SwitchC] vtep enable
```

Assign an IP address to Loopback 0. Specify the IP address as the global source address for VXLAN tunnels.

```
[SwitchC] interface loopback 0
[SwitchC-LoopBack0] ip address 3.3.3.3 255.255.255.255
[SwitchC-LoopBack0] quit
[SwitchC] tunnel global source-address 3.3.3.3
```

Specify site-facing interface HundredGigE 1/0/1 as a VTEP access port.

```
[SwitchC] interface hundredgige 1/0/1
[SwitchC-HundredGigE1/0/1] vtep access port
[SwitchC-HundredGigE1/0/1] quit
```

6. Configure VXLAN settings on the controller. (Details not shown.)

Verifying the configuration

1. Verify the VXLAN settings on the VTEPs. This example uses Switch A.

Verify that the VXLAN tunnel interfaces on the VTEP are up.

```
[SwitchA] display interface tunnel 1
Tunnel1
Current state: UP
Line protocol state: UP
Description: Tunnel1 Interface
Bandwidth: 64 kbps
Maximum transmission unit: 1464
Internet protocol processing: Disabled
Last clearing of counters: Never
Tunnel source 1.1.1.1, destination 2.2.2.2
Tunnel protocol/transport UDP_VXLAN/IP
Last 300 seconds input rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Last 300 seconds output rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Input: 0 packets, 0 bytes, 0 drops
Output: 0 packets, 0 bytes, 0 drops
```

Verify that the VXLAN tunnels have been assigned to the VXLAN.

```
[SwitchA] display l2vpn vsi verbose
VSI Name: evpn2014
  VSI Index               : 0
  VSI State               : Up
  MTU                     : 1500
  Bandwidth               : Unlimited
  Broadcast Restrain      : Unlimited
  Multicast Restrain      : Unlimited
  Unknown Unicast Restrain: Unlimited
  MAC Learning            : Enabled
  MAC Table Limit         : -
```

```
  MAC Learning rate       : -
  Drop Unknown            : -
  Flooding                : Enabled
  Statistics              : Disabled
  VXLAN ID                : 10
  Tunnels:
    Tunnel Name           Link ID    State  Type        Flood proxy
    Tunnel1               0x5000001  Up     Manual      Disabled
    Tunnel2               0x5000002  Up     Manual      Disabled
  ACs:
    AC                              Link ID    State    Type
    HGE1/0/1 srv2                   0          Up       Manual
```

Verify that the VTEP has learned the MAC addresses of remote VMs.

```
<SwitchA> display l2vpn mac-address
MAC Address : cc3e-5f9c-6cdb
VSI Name    : evpn2014
State       : Dynamic
Link ID/Name      Aging
Tunnel1           Aging

MAC Address : cc3e-5f9c-23dc
VSI Name    : evpn2014
State       : Dynamic
Link ID/Name      Aging
Tunnel2           Aging
--- 2 mac address(es) found  ---
```

2. Verify that VM 1, VM 2, and VM 3 can ping each other. (Details not shown.)

Example: Configuring flood proxy for a VXLAN

Network configuration

As shown in Figure 14:

- Configure the controller cluster to deploy VXLAN 10 to Switch A, Switch B, and Switch C to provide Layer 2 connectivity for the VMs across the network sites.
- Enable flood proxy for VXLAN 10.
- Use the MAC address entries issued by the controller to direct traffic forwarding on Switch A, Switch B, and Switch C.

Figure 14 Network diagram

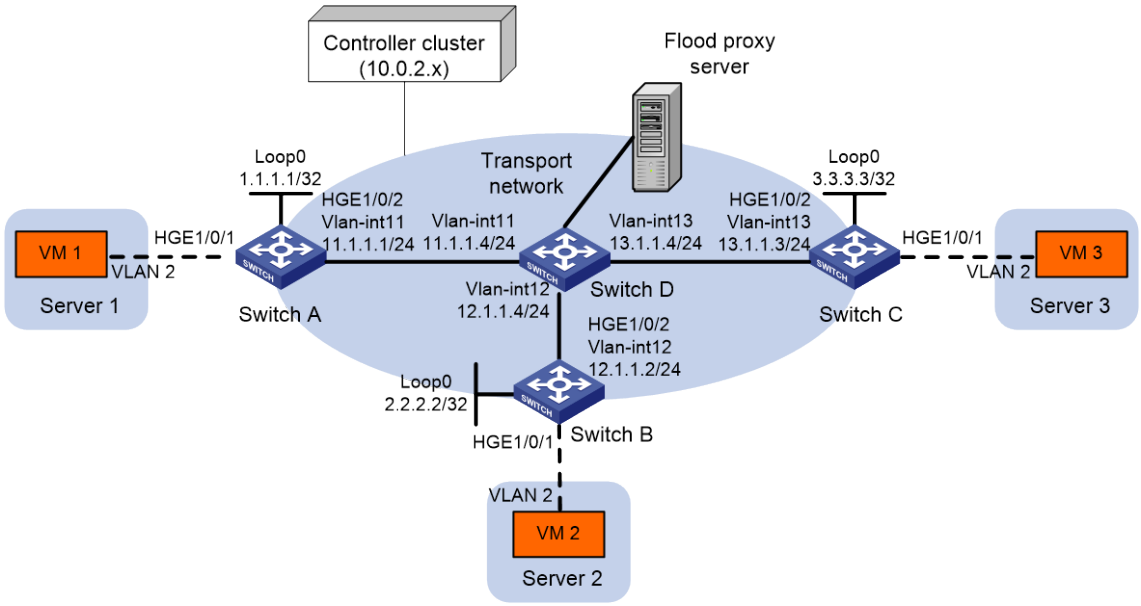

Procedure

1. Create VLANs and VLAN interfaces on all devices. (Details not shown.)
2. Configure IP addresses and unicast routing settings:

   # Assign IP addresses to interfaces, as shown in Figure 14. (Details not shown.)

   # Configure OSPF on all transport network switches (Switches A through D). (Details not shown.)
3. Configure Switch A:

   # Enable L2VPN.

   ```
   <SwitchA> system-view
   [SwitchA] l2vpn enable
   ```

   # Configure active TCP connection settings.

   ```
   [SwitchA] ovsdb server tcp ip 10.0.2.15 port 6632
   ```

   # Enable the OVSDB server.

   ```
   [SwitchA] ovsdb server enable
   ```

   # Enable the OVSDB VTEP service.

   ```
   [SwitchA] vtep enable
   ```

   # Assign an IP address to Loopback 0.

   ```
   [SwitchA] interface loopback 0
   [SwitchA-LoopBack0] ip address 1.1.1.1 255.255.255.255
   [SwitchA-LoopBack0] quit
   ```

   # Specify the IP address of Loopback 0 as the global source address for VXLAN tunnels.

   ```
   [SwitchA] tunnel global source-address 1.1.1.1
   ```

   # Specify site-facing interface HundredGigE 1/0/1 as a VTEP access port.

   ```
   [SwitchA] interface hundredgige 1/0/1
   [SwitchA-HundredGigE1/0/1] vtep access port
   [SwitchA-HundredGigE1/0/1] quit
   ```

   # Disable source MAC check on transport-facing interface HundredGigE 1/0/2.

   ```
   [SwitchA] interface hundredgige 1/0/2
   ```

```
[SwitchA-HundredGigE1/0/2] undo mac-address static source-check enable
[SwitchA-HundredGigE1/0/2] quit
```

Disable remote-MAC address learning.

```
[SwitchA] vxlan tunnel mac-learning disable
```

Enable flood proxy on multicast VXLAN tunnels.

```
[SwitchA] vxlan tunnel flooding-proxy
```

4. Configure Switch B:

Enable L2VPN.

```
<SwitchB> system-view
[SwitchB] l2vpn enable
```

Configure active TCP connection settings.

```
[SwitchB] ovsdb server tcp ip 10.0.2.15 port 6632
```

Enable the OVSDB server.

```
[SwitchB] ovsdb server enable
```

Enable the OVSDB VTEP service.

```
[SwitchB] vtep enable
```

Assign an IP address to Loopback 0.

```
[SwitchB] interface loopback 0
[SwitchB-LoopBack0] ip address 2.2.2.2 255.255.255.255
[SwitchB-LoopBack0] quit
```

Specify the IP address of Loopback 0 as the global source address for VXLAN tunnels.

```
[SwitchB] tunnel global source-address 2.2.2.2
```

Specify site-facing interface HundredGigE 1/0/1 as a VTEP access port.

```
[SwitchB] interface hundredgige 1/0/1
[SwitchB-HundredGigE1/0/1] vtep access port
[SwitchB-HundredGigE1/0/1] quit
```

Disable source MAC check on transport-facing interface HundredGigE 1/0/2.

```
[SwitchB] interface hundredgige 1/0/2
[SwitchB-HundredGigE1/0/2] undo mac-address static source-check enable
[SwitchB-HundredGigE1/0/2] quit
```

Disable remote-MAC address learning.

```
[SwitchB] vxlan tunnel mac-learning disable
```

Enable flood proxy on multicast VXLAN tunnels.

```
[SwitchB] vxlan tunnel flooding-proxy
```

5. Configure Switch C:

Enable L2VPN.

```
<SwitchC> system-view
[SwitchC] l2vpn enable
```

Configure active TCP connection settings.

```
[SwitchC] ovsdb server tcp 10.0.2.15 port 6632
```

Enable the OVSDB server.

```
[SwitchC] ovsdb server enable
```

Enable the OVSDB VTEP service.

```
[SwitchC] vtep enable
```

Assign an IP address to Loopback 0.

```
[SwitchC] interface loopback 0
```

```
[SwitchC-LoopBack0] ip address 3.3.3.3 255.255.255.255
[SwitchC-LoopBack0] quit
```

Specify the IP address of Loopback 0 as the global source address for VXLAN tunnels.

```
[SwitchC] tunnel global source-address 3.3.3.3
```

Specify site-facing interface HundredGigE 1/0/1 as a VTEP access port.

```
[SwitchC] interface hundredgige 1/0/1
[SwitchC-HundredGigE1/0/1] vtep access port
[SwitchC-HundredGigE1/0/1] quit
```

Disable source MAC check on transport-facing interface HundredGigE 1/0/2.

```
[SwitchC] interface hundredgige 1/0/2
[SwitchC-HundredGigE1/0/2] undo mac-address static source-check enable
[SwitchC-HundredGigE1/0/2] quit
```

Disable remote-MAC address learning.

```
[SwitchC] vxlan tunnel mac-learning disable
```

Enable flood proxy on multicast VXLAN tunnels.

```
[SwitchC] vxlan tunnel flooding-proxy
```

6. Configure VXLAN settings on the controller, and configure the flood proxy server. (Details not shown.)

Verifying the configuration

1. Verify the VXLAN settings on the VTEPs. This example uses Switch A.

Verify that the VXLAN tunnel interfaces on the VTEP are up.

```
[SwitchA] display interface tunnel
Tunnel1
Current state: UP
Line protocol state: UP
Description: Tunnel1 Interface
Bandwidth: 64 kbps
Maximum transmit unit: 1464
Internet protocol processing: disabled
Last clearing of counters: Never
Tunnel source 1.1.1.1, destination 2.2.2.2
Tunnel protocol/transport UDP_VXLAN/IP
Last 300 seconds input rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Last 300 seconds output rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Input: 0 packets, 0 bytes, 0 drops
Output: 0 packets, 0 bytes, 0 drops
```

Verify that the VXLAN tunnels have been assigned to the VXLAN, and flood proxy has been enabled on the multicast VXLAN tunnel.

```
[SwitchA] display l2vpn vsi verbose
VSI Name: evpn2014
  VSI Index               : 0
  VSI State               : Up
  MTU                     : 1500
  Bandwidth               : Unlimited
  Broadcast Restrain      : Unlimited
  Multicast Restrain      : Unlimited
  Unknown Unicast Restrain: Unlimited
```

```
  MAC Learning            : Enabled
  MAC Table Limit         : -
  MAC Learning rate       : -
  Drop Unknown            : -
  Flooding                : Enabled
  Statistics              : Disabled
  VXLAN ID                : 10
  Tunnels:
    Tunnel Name          Link ID    State  Type        Flood proxy
    Tunnel1              0x5000001  Up     Manual      Disabled
    Tunnel2              0x5000002  Up     Manual      Disabled
    Tunnel3              0x5000003  Up     Manual      Enabled
  ACs:
    AC                               Link ID    State    Type
    HGE1/0/1 srv2                    0          Up       Manual
```

Verify that the VTEP has obtained the MAC addresses of remote VMs from the controller.

```
<SwitchA> display l2vpn mac-address
MAC Address : cc3e-5f9c-6cdb
VSI Name    : evpn2014
State       : OVSDB
Link ID/Name      Aging
Tunnel1           NotAging

MAC Address : cc3e-5f9c-23dc
VSI Name    : evpn2014
State       : OVSDB
Link ID/Name      Aging
Tunnel2           NotAging
--- 2 mac address(es) found  ---
```

2. Verify that VM 1, VM 2, and VM 3 can ping each other. (Details not shown.)