1. Configuring VXLAN

1 Configuring VXLAN

1.1 Overview

Virtual Extensible Local Area Network (VXLAN) is a virtual Ethernet based on the physical IP (overlay) network. It is a technology that encapsulates layer 2 (L2) Ethernet frames within layer 3 User Datagram Protocol (UDP) packets.

VXLAN has a 24-bit VXLAN network identifier (VNI). It allows users to create up to 16,000,000 isolated virtual networks to meet the requirements of multi-tenant environments and scale expansion, far surpassing the widely used Virtual Local Area Network (VLAN) technology that is limited to 4,000 isolated networks. VXLAN uses the IP multicast method to encapsulate multicast, broadcast, and unknown unicast packets, effectively controlling the broadcast domain in multi-tenant environments.

With the transformation of data centers, more and more virtual machines are deployed. In addition, as virtual machines must be migrated in L2 environments, scales of L2 networks increase. VXLAN can extend L2 networks over layer 3 (L3) networks, so that virtual machines can be moved to L3 networks interconnected to L2 networks without changing the IP addresses and MAC addresses, thereby ensuring service continuity.

Protocols and Standards

- RFC7348: Virtual Extensible Local Area Network (VXLAN) -- A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks

1.2 Applications

Application	Description
Virtual L2 Interconnection	Applicable to the virtual L2 interconnection on IP network using VXLAN.
Inter-Data Center Interconnection	Applicable to inter data center interconnection on IP core network using VXLAN.
VXLAN IP Gateway	Applicable to communication across VXLANs and with external networks through a VXLAN IP gateway.

1.2.1 Virtual L2 Interconnection

Scenario

Servers in a data center implement L2 interconnection over L3. As shown in the following figure, Servers B and C are not deployed on the same L2 network as Server D, but the three servers can implement L2 interconnection through a VXLAN. Logically, Servers B, C, and D are deployed on the same VLAN. Physically, they are interconnected over L3 network.

Figure 1-1

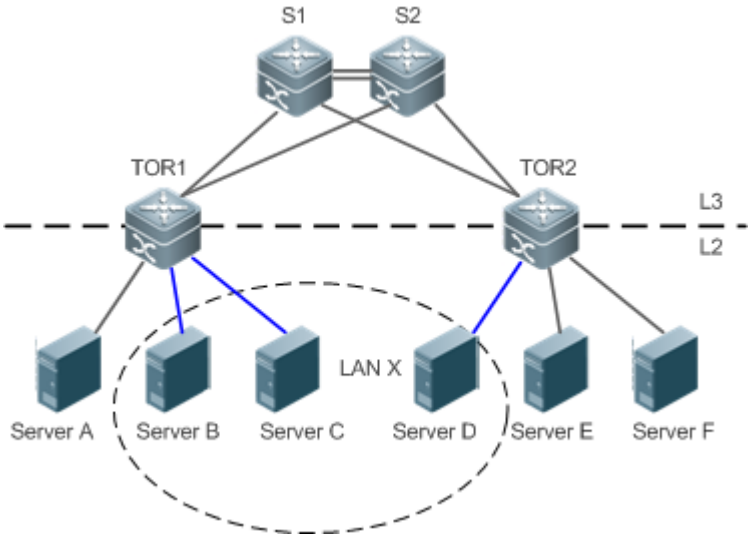

- Servers B, C, and D implement L2 interconnection over L3 using a VXLAN.
- Servers B, C, and D forward packets to each other on the VXLAN.

Remarks	S1 and S2 are core switches. TOR1 and TOR2 are access switches that support the VXLAN function. Servers A, B, C, D, E, and F are deployed in the same data center. Servers B, C, and D are located in the same broadcast domain.

Deployment

- Configure an IPv4 unicast routing protocol such as the OSPF protocol on S1, S2, TOR1, and TOR2 to ensure that unicast routes are reachable.
- Configure an IP multicast protocol (such as the PIM-DM protocol) on S1, S2, TOR1, and TOR2 to ensure that multicast routes are reachable.
- Configure a VXLAN on switches TOR1 and TOR2 in the data center to implement L2 interconnection via VLAN.

1.2.2 Inter-Data Center Interconnection

Scenario

Servers in data centers implement L2 interconnection across data centers using a VXLAN and the IP core network, as shown in the figure below.

Figure 1-2

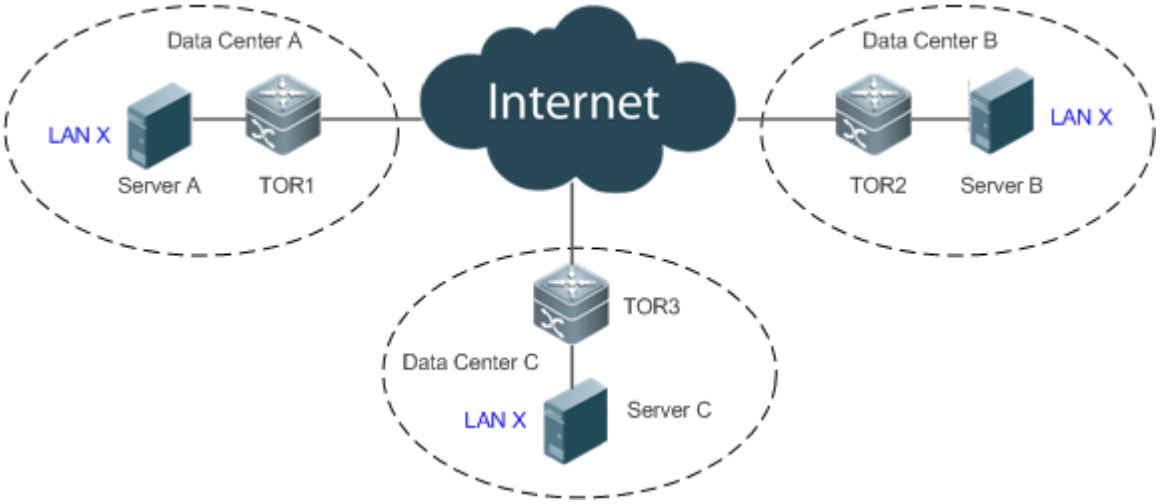

- Servers in Data Centers A, B, and C implement L2 interconnection on the IP core network by using a VXLAN.
- Servers in Data Centers A, B, and C forward packets to each other on the VXLAN.

Remarks	TOR1, TOR2, and TOR3 are access switches that support the VXLAN function in Data Centers A, B, and C, respectively. Servers A, B, and C are deployed in Data Centers A, B, and C, respectively, and they are interconnected at L2.

Deployment

- Configure an IPv4 unicast routing protocol such as the OSPF protocol on TOR1, TOR2, and TOR3 to ensure that unicast routes are reachable.
- Configure an IP multicast protocol (such as the PIM-DM protocol) on TOR1, TOR2, and TOR3 to ensure that multicast routes are reachable.
- Configure a VXLAN on switches TOR1, TOR2, and TOR3 in the data centers to implement L2 interconnection via VLAN.

1.2.3 VXLAN IP Gateway

Scenario

Servers implement communication across VXLANs and with external networks through a VXLAN IP gateway, as shown in the figure below.

Figure 1-3

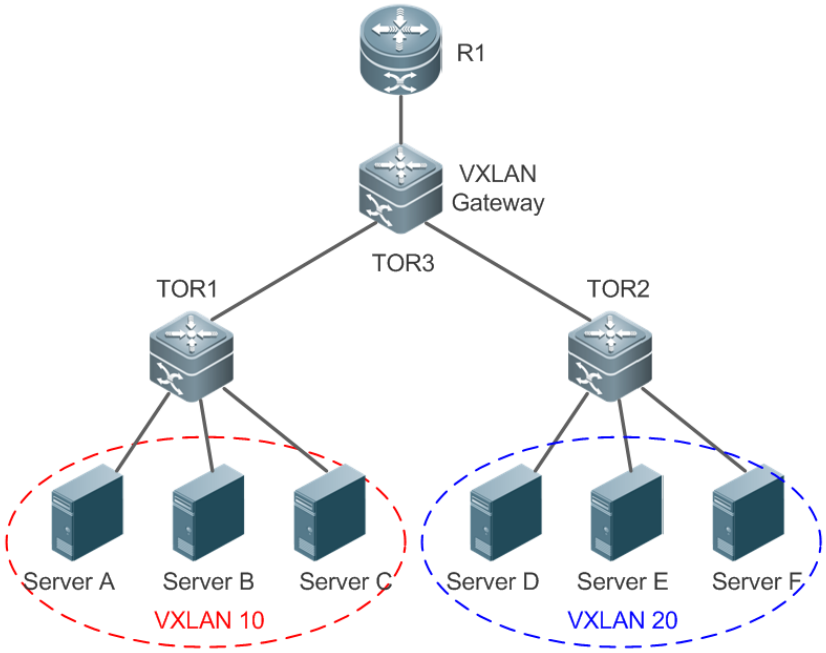

- Servers A, B, and C belong to VXLAN 10 and Servers D, E, and F belong to VXLAN 20.
- A server of VXLAN 10 can communicate with a server of VXLAN 20 through TOR3, for example, Server A can access Server D.
- External network devices can access servers of VXLAN 10 or VXLAN 20 through R1 and TOR3. Servers of VXLAN 10 and VXLAN 20 can access external networks through TOR3 and R1.

Remarks	Servers A, B, and C belong to VXLAN 10. They can be physical servers that implement VXLAN packet encapsulation via TOR1, or. virtual servers (virtual machines) that implement VXLAN packet encapsulation via Hypervisor. Servers D, E, and F belong to VXLAN 20. They can be physical servers that implement VXLAN packet encapsulation via TOR2, or virtual servers (virtual machines) that implement VXLAN packet encapsulation via Hypervisor. TOR3 provides the VXLAN routing function. It can implement routing between VXLANs and routing from a conventional network to a VXLAN. It serves as the IP gateway for VXLAN 10 and VXLAN 20.

Deployment

- Configure an IPv4 unicast routing protocol such as the OSPF protocol on TOR1, TOR2, and TOR3 to ensure that unicast routes are reachable.
- Configure a multicast routing protocol (such as the PIM-DM protocol) on TOR1, TOR2, and TOR3 to ensure that multicast routes are reachable.
- Configure the IP gateway for VXLAN 10 and VXLAN 20 on TOR3 to implement the communication between VXLAN 10 and VXLAN 20.

1.2.4 EVPN-based Multi-tenant Centralized Deployment

Scenario

VPN routing and forwarding (VRF) networks are usually allocated to different tenants to support the multi-tenant application in a data center. Multiple VXLANs can be assigned to each tenant. VXLANs of the same tenant can be mutually accessed through the L3 router, while VXLANs of different tenants cannot be mutually accessed, as shown in Figure 1-4.

Tenant A rents VRF-10, which includes VXLAN 10 and VXLAN 20. Servers HOST-1 and HOST-2 belong to VXLAN 10 and Servers HOST-3 and HOST-4 belong to VXLAN 20.

Tenant B rents VRF-20, which includes VXLAN 100. Servers HOST-5 and HOST-6 belong to VXLAN 100.

The networks of Tenant A and Tenant B are isolated from each other.

The entire network is formed by a Border Gateway Protocol (BGP) network and includes CORE and TOR switches. The BGP neighbor relationship is formed between every two devices and the BGP-EVPN protocol family is supported. All VXLAN gateways on the network are deployed in the core switches in a centralized manner.

Figure 1-4

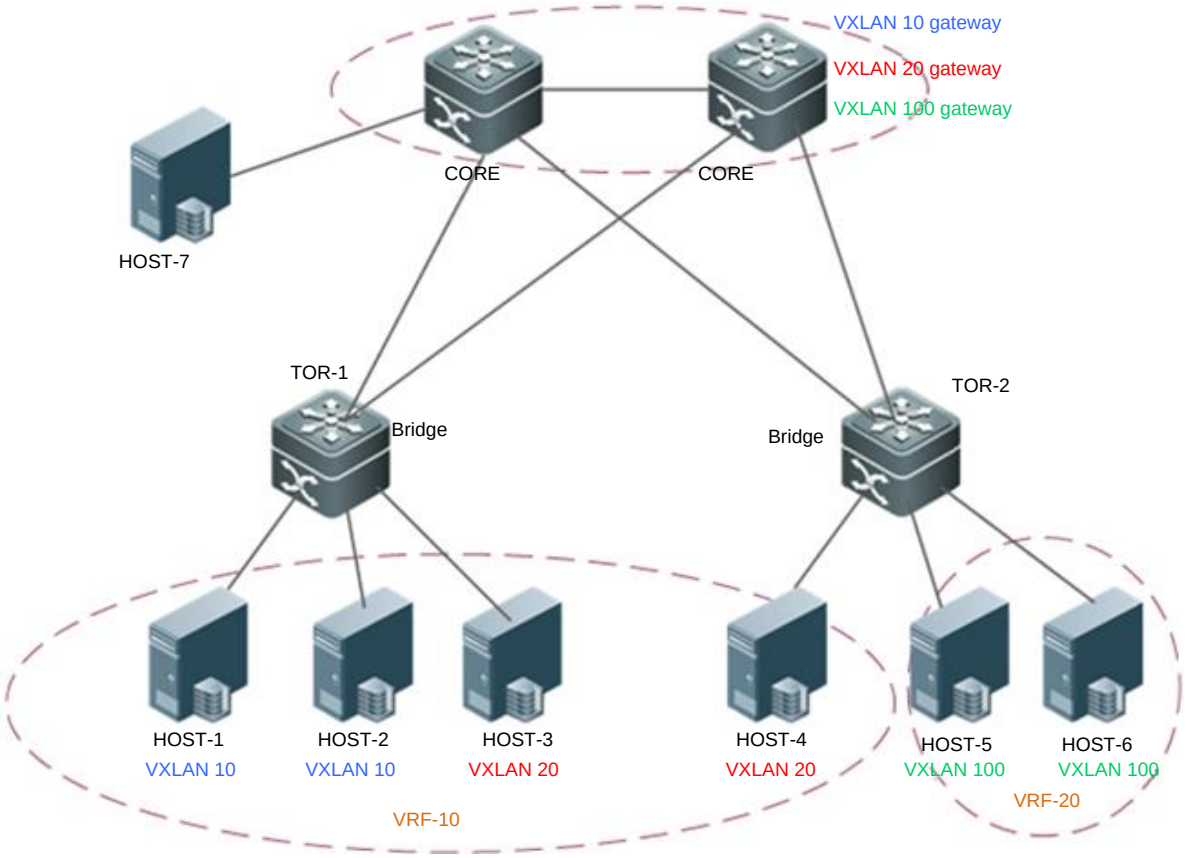

- Packets between HOST-1 and HOST-2 are forwarded through TOR-1 at L2 within the VXLAN.
- Packets between HOST-3 and HOST-4 are forwarded through TOR-1 > CORE > TOR-2 at L2 within the VXLAN.
- Packets between HOST-5 and HOST-6 are forwarded through TOR-2 at L2 within the VXLAN.

- Packets between VXLAN 10 and VXLAN 20 are forwarded through TOR-1 > CORE > TOR-2 at L3 across the VXLANs.
- VRF-10 and VRF-20 cannot communicate with each other.

Remarks:	CORE indicates a core switch that supports the VXLAN function. TOR1 and TOR2 are access switches that support the VXLAN function. HOST-1, HOST-2, HOST-3, HOST-4, HOST-5, and HOST-6 are servers in the data center.

Deployment

- Configure an Internet Protocol version 4 (IPv4) unicast routing protocol, for example, the Open Shortest Path First (OSPF) protocol, on the switches to ensure that unicast routes are reachable.
- Configure the BGP routing protocol (supporting EVPN) on the switches to establish neighbor relationships between each other.
- Deploy the VXLAN gateway on the core switches.
- Deploy the VXLAN bridge on the TOR switches.

1.2.5 EVPN-based Multi-tenant Distributed Deployment

Scenario

The EVPN-based multi-tenant distributed deployment applies to data center networks that support multiple tenants. The difference between this deployment and the EVPN-based multi-tenant centralized deployment described in section 1.2.4 lies in that: On the distributed deployment network, gateways are deployed on the TOR switches, as shown in Figure 1-5.

Tenant A rents VRF-10, which includes VXLAN 10 and VXLAN 20.

Tenant B rents VRF-20, which includes VXLAN 100.

The networks of Tenant A and Tenant B are isolated from each other.

The entire network is formed by a BGP network and includes CORE and TOR switches. The BGP neighbor relationship is formed between every two devices and the BGP-EVPN protocol family is supported.

VXLAN gateways are deployed on TOR switches on the network. Anycast gateways can be deployed so that the IP addresses and MAC addresses of all gateways on the network are kept consistent. In this way, the gateway configuration does not need to be modified no matter which TOR switch a virtual machine of a customer is migrated to.

VXLANs are unnecessarily deployed on the core switches.

Figure 1-5

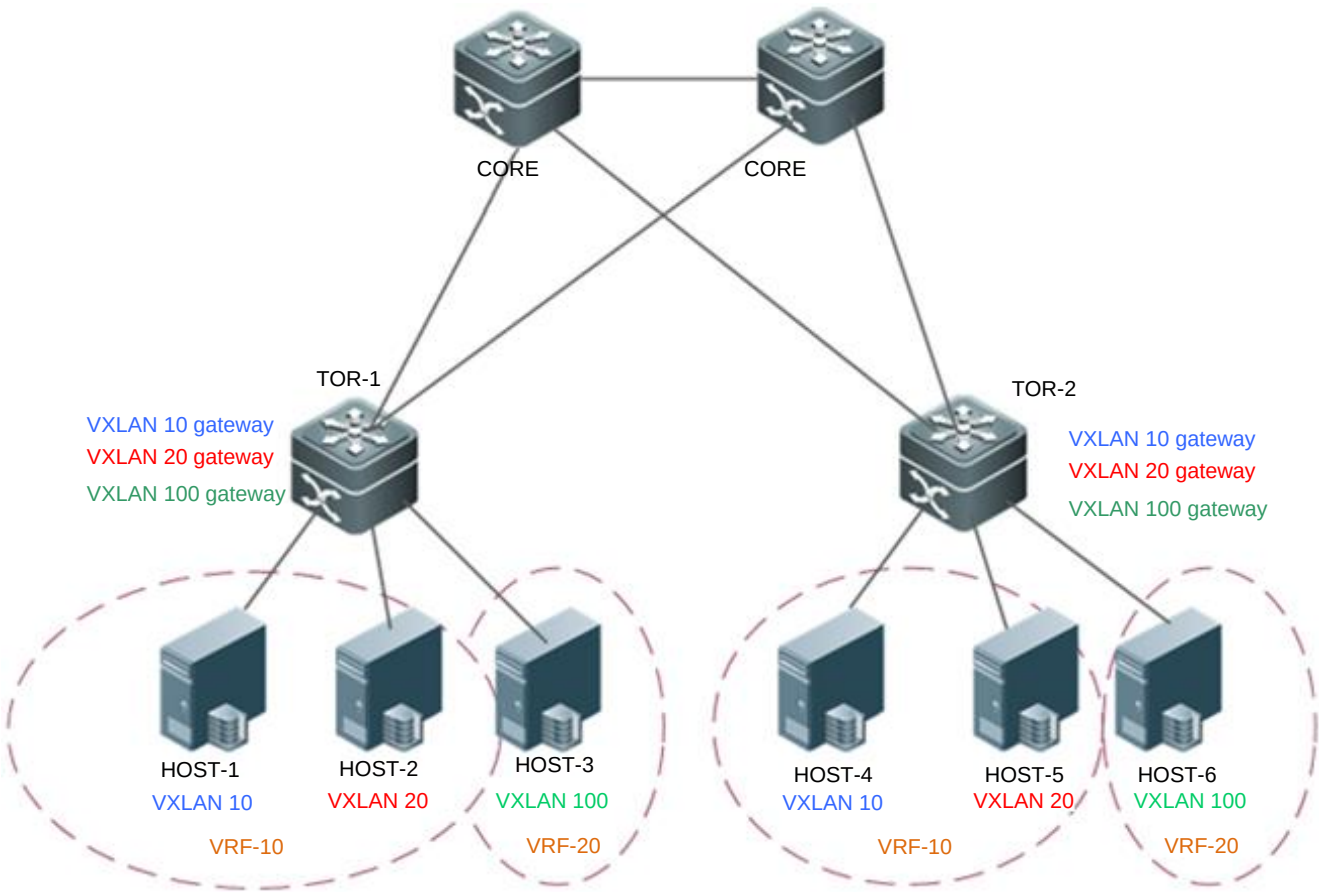

- Packets between HOST-1 and HOST-4 are forwarded through TOR-1 > TOR-2 at L2 within the VXLAN.
- Packets between HOST-1 and HOST-2 are forwarded through TOR-1 at L3 across the VXLANs.
- Packets between HOST-1 and HOST-5 are forwarded through TOR-1 > TOR-2 at L3 across the VXLANs.
- VRF-10 and VRF-20 cannot communicate with each other.

Remarks:	CORE indicates a core switch that supports the BGP-EVPN function. TOR1 and TOR2 are access switches that support the VXLAN function. HOST-1, HOST-2, HOST-3, HOST-4, HOST-5, and HOST-6 are servers in the data center.

Deployment

- Configure an IPv4 unicast routing protocol, for example, the OSPF protocol, on the switches to ensure that unicast routes are reachable.
- Configure the BGP routing protocol (supporting EVPN) on the switches to establish neighbor relationships between each other.
- Deploy the VXLAN bridge on the core switches if required.
- Deploy the VXLAN gateway on the TOR switches.

1.2.6 EVPN-based Single-tenant VXLAN Routing Deployment

Scenario

Single-tenant VXLAN route deployment is shown in Figure 1-6.

In this scenario, only the VRF-10 is deployed, which includes VXLAN 10 and VXLAN 20.

The border devices are connected to the external network. These devices are deployed in VRF-10 (including VXLAN 90) and interconnect with the external network at L3 via the overlay router interface.

The entire network is formed by a BGP network and includes TOR and border devices. The BGP neighbor relationship is formed between every two devices (except between Border-1 and Border-2) and the BGP-EVPN protocol family is supported.

The TOR and border devices must use a symmetric VXLAN (VXLAN 100) for interconnection with each other. The border devices import network routes to the TOR switches through the symmetric VXLAN.

VXLAN gateways are deployed on TOR switches on the network. Anycast gateways can be deployed so that the IP addresses and MAC addresses of all gateways on the network are kept consistent. In this way, the gateway configuration does not need to be modified no matter which TOR switch a virtual machine of a customer is migrated to.

Figure 1-6

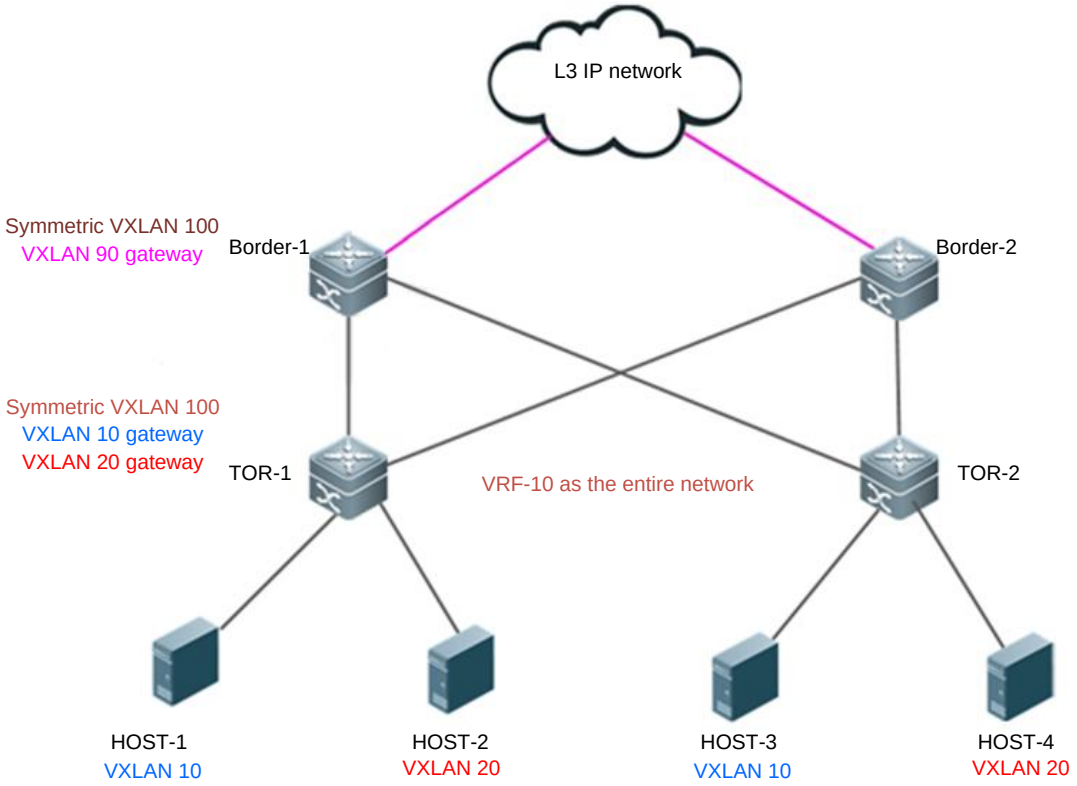

- Packets between HOST-1 and HOST-3 are forwarded through TOR-1 > TOR-2 at L2 within the VXLAN.
- Packets between HOST-1 and HOST-2 are forwarded through TOR-1 at L3 across the VXLANs.

- To access the external network, HOST-1 forwards packets to the border device through TOR1 at L3 across the VXLANs, and then the border device forwards the packets to the external network at L3.

Deployment

- Configure an IPv4 unicast routing protocol, for example, the OSPF protocol, on the switches to ensure that unicast routes are reachable.
- Configure the BGP routing protocol (supporting EVPN) on the switches to establish neighbor relationships between each other (except between the border devices).
- Deploy the VXLAN on the border devices for L3 interconnection with the external network.
- Deploy the VXLAN gateway on the TOR switches.

1.2.7 EVPN-based Multi-tenant VXLAN Route Deployment

Scenario

VRF networks are usually allocated to different tenants to support the multi-tenant application in a data center. Multiple VXLANs can be assigned to each tenant. VXLANs of the same tenant can be mutually accessed through the L3 router, while VXLANs of different tenants cannot be mutually accessed, as shown in Figure 1-7.

Tenant A rents VRF-10, which includes VXLAN 10 and VXLAN 20.

Tenant B rents VRF-20, which includes VXLAN 30.

The border devices are connected to the external network. These devices are deployed in VRF-30 (including VXLAN 90) and interconnect with the external network at L3 via the overlay router interface.

The networks of Tenant A and Tenant B are isolated from each other.

The entire network is formed by a BGP network and includes TOR and border devices. The BGP neighbor relationship is formed between every two devices (except between Border-1 and Border-2) and the BGP-EVPN protocol family is supported

The TOR and border devices must use a symmetric VXLAN (VXLAN 100 and VXLAN200) for interconnection with each other. The border devices import network routes to the TOR switches through the symmetric VXLAN.

VXLAN gateways are deployed on TOR switches on the network. Anycast gateways can be deployed so that the IP addresses and MAC addresses of all gateways on the network are kept consistent. In this way, the gateway configuration does not need to be modified no matter which TOR switch a virtual machine of a customer is migrated to.

To improve high availability (HA) of the network, multiple VXLAN gateways can be deployed on the core switch to form centralized all-active anycast gateways to ensure redundant backup of gateways.

Figure 1-7

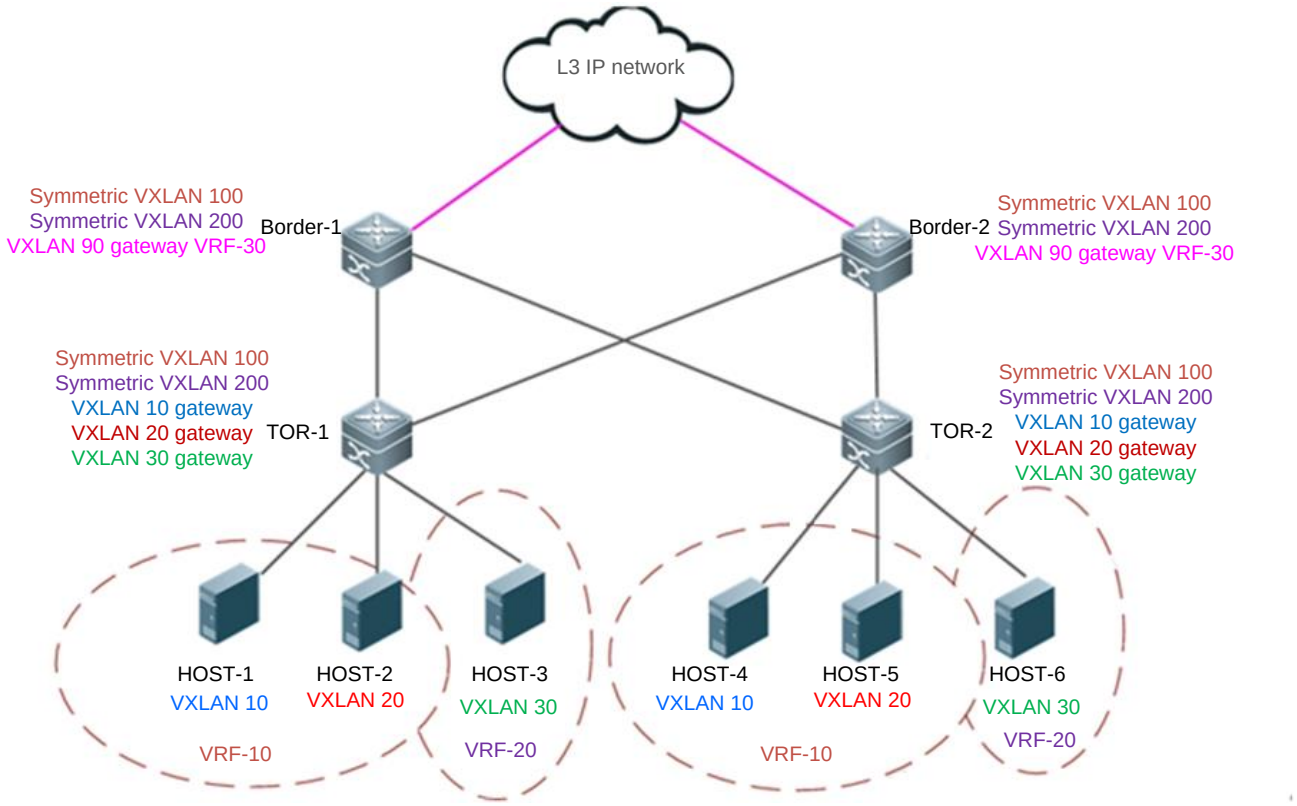

- Packets between HOST-1 and HOST-4 are forwarded through TOR-1 > TOR-2 at L2 within the VXLAN.
- Packets between HOST-1 and HOST-2 are forwarded through TOR-1 at L3 across the VXLANs.
- To access the external network, HOST-1 forwards packets to the border device through TOR1 at L3 across the VXLANs, and then the border device forwards the packets to the external network at L3.

Deployment

- Configure an IPv4 unicast routing protocol, for example, the OSPF protocol, on the switches to ensure that unicast routes are reachable.
- Configure the BGP routing protocol (supporting EVPN) on the switches to establish neighbor relationships between each other (except between the border devices).
- Deploy the VXLAN on the border devices for L3 interconnection with the external network.
- Deploy the VXLAN gateway on the TOR switches.

1.3 Features

Basic Concepts

VXLAN Packet Format

A VXLAN encapsulates the Ethernet frames into UDP packets and transmits them on the IP core network.

The VXLAN defines a VTEP entity, which encapsulates the data generated by the virtual machine into the UDP headers, and sends the data out. After the encapsulation, the MAC address and VLAN information of the virtual machine no longer serves as the basis for data forwarding.

The VTEP entity can be software, a hardware server, or other device. If the VTEP function is directly integrated into a hypervisor (also called virtual machine monitor), all virtual machine traffic is marked with new VXLAN tags and UDP headers before entering the switch. This is equivalent to creating a tunnel between any two virtual machines.

As the VLAN information of the virtual machine is externally invisible, a new VXLAN label (VNI) is added. VNIs replace VLANs to represent different VXLAN segments. Same as the forwarding behavior of VLANs, only the virtual machines with the same VNI in the same VXLAN segment can communicate with each other.

The new UDP header and VNI form a new frame structure. After receiving the data frame sent from the virtual machine, a VTEP encapsulates four elements (which are the VXLAN header, outer UDP header, outer IPv4 header, and outer Ethernet frame header from inside out) to form a new frame header. In the new frame header, the original source and destination MAC addresses, inner VLAN tag, and Ethernet type that are carried by the inner data frame remain the same.

The format of an encapsulated VXLAN frame is as follows:

Figure 1-8

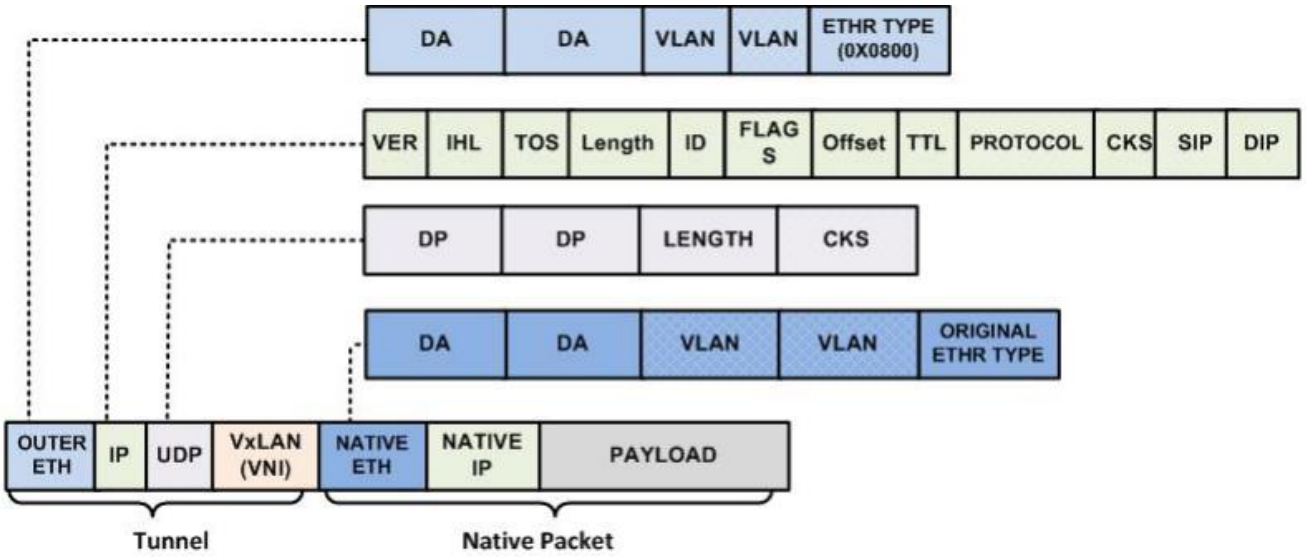

1.3.1 Packet Format

VXLAN Header Information

Figure 1-9

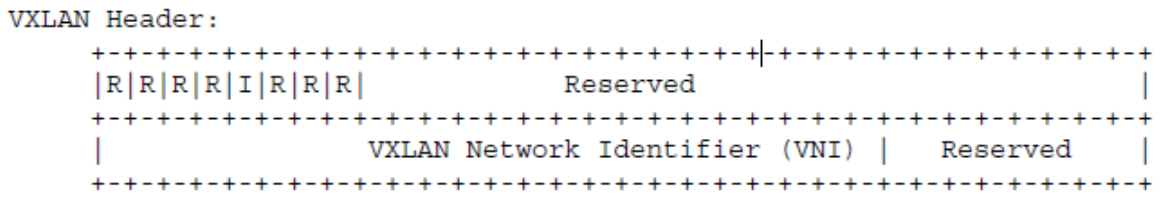

A VXLAN header has 64 bits. In the design of the current protocol version, the sole purpose of a VXLAN header is to carry the 24-bit VNI assigned by the VTEP.

- Flag (8 bits): The I bit must be set to 1 to indicate a valid VNI, and the R bit must be set to 0.
- VXLAN segment ID/VNI: Includes 24 bits and indicates the VXLAN network identifier. Only the virtual machines that belong to the same VXLAN can communicate with each other.
- Reserved: The 24th bit and 8th bit are reserved, and are set to 0.

Outer UDP Header

Figure 1-10

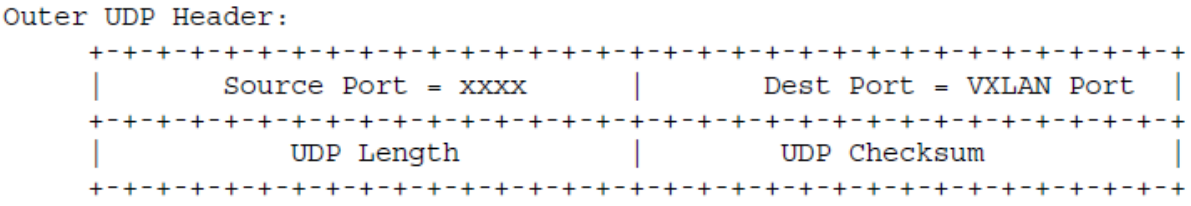

The definitions of the fields of the UDP header are as follows:

- Source Port: Indicates the source port ID of the UDP packet. Assigned by the VTEP, the source port ID is the result of the hash operation of the L2 header of the data frame. This hash result can serve as the basis for traffic load balancing.
- Dest Port: Indicates the destination port ID. The port ID assigned by the Internet Assigned Numbers Authority (IANA) is 4789.
- UDP Length: Indicates the length of the UDP header.
- UDP Checksum: Indicates the UDP checksum, which is set to 0 for transmission.

Outer IP Header

Figure 1-11

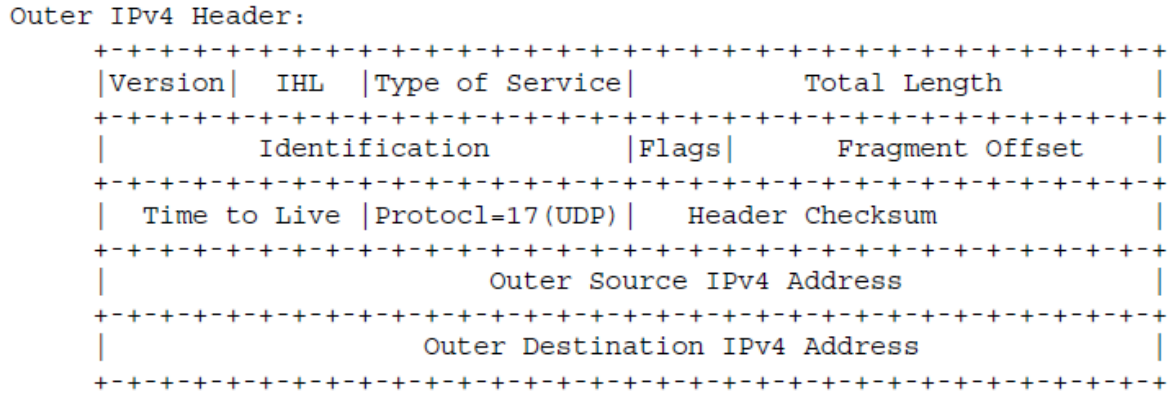

The definitions of the fields of the outer IP header are as follows:

- Source IPv4 Address: Identifies the IP address of the VTEP that corresponds to the virtual machine.
- Destination IPv4 Address: Indicates the unicast or multicast IP address. If it is a unicast IP address, it indicates the IP address of the VTEP corresponding to the virtual machine to be communicated with.

The IP address of the outer IP header is no longer the address of the virtual machines of both communication parties, but the address of the VTEPs at both ends of the tunnel. If the hypervisor directly takes over the work of the VTEP, the IP address is the IP address of the NIC of the server that runs the hypervisor. If the VTEP is an access switch, the IP address is the IP address of an egress interface or the IP address of an L3 switch virtual interface (SVI).

↘ Outer Ethernet Header

Figure 1-12

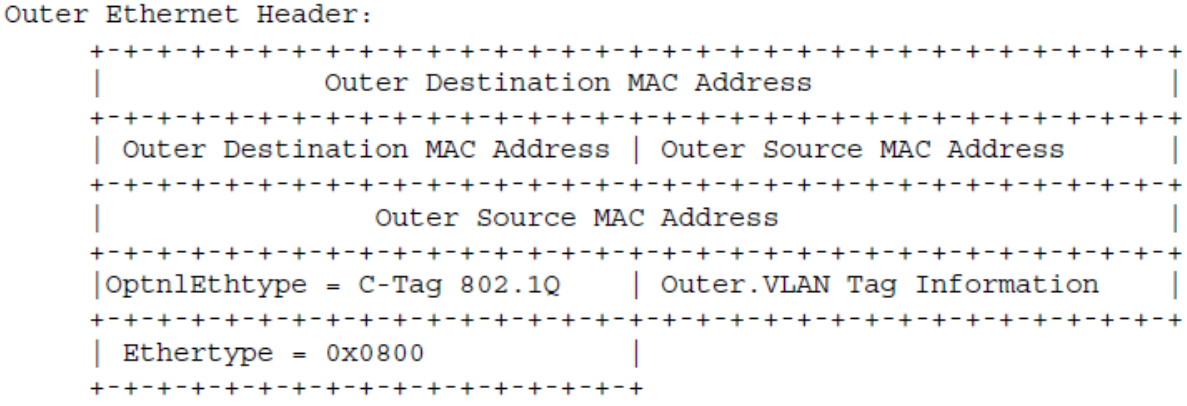

The definitions of the fields of the outer Ethernet header are as follows:

- Destination MAC Address: Indicates the MAC address of the destination VTEP or the MAC address of the L3 gateway. If the outer Ethernet header is encapsulated into a multicast packet, the destination MAC address indicates the multicast MAC address.
- VLAN tag: Optional.

1.3.2 Forwarding Model

VXLAN Bridging Principle

VXLAN encapsulates Ethernet packets within UDP packets to transmit them on the IP network. On the receiver, the VXLAN packets are decapsulated into Ethernet packets and then forwarded, as shown in Figure 1-13.

Figure 1-13

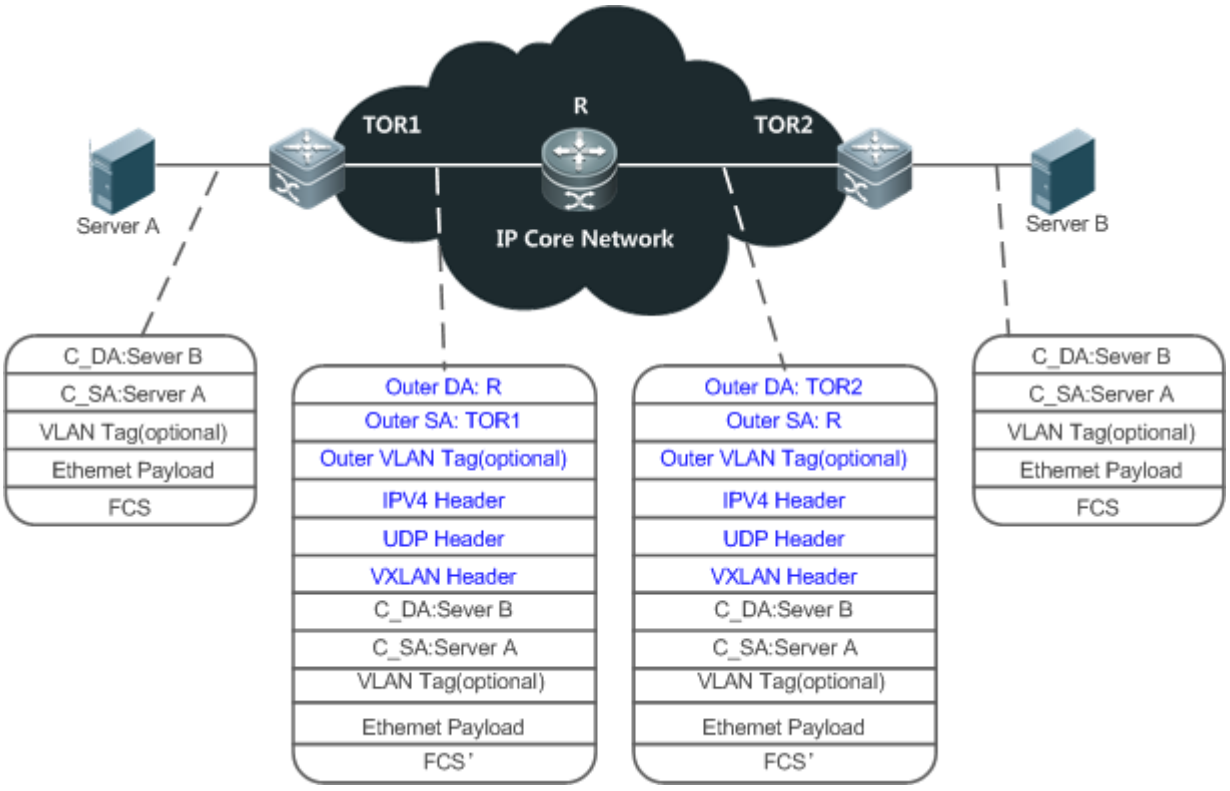

- Switch TOR1 receives the Ethernet packet from LAN X, and then encapsulates the packet into a VXLAN packet.
- The VXLAN packet is forwarded in the IP core network. As shown in Figure 1-13, R forwards the VXLAN packet.
- Switch TOR2 receives the VXLAN packet, and then decapsulates and forwards it at L2 of the LAN.

Overview

Feature	Description
VXLAN Bridging and Forwarding	Encapsulates broadcast, multicast, and unknown unicast packets into IP multicast packets to realize flooding. The well-known unicast packets are encapsulated and forwarded by searching the VXLAN address table for the MAC address and IP address.

VXLAN Routing Principle

VXLANs interconnect with each other through the VXLAN IP gateway, as shown in Figure 1-14.

Figure 1-14

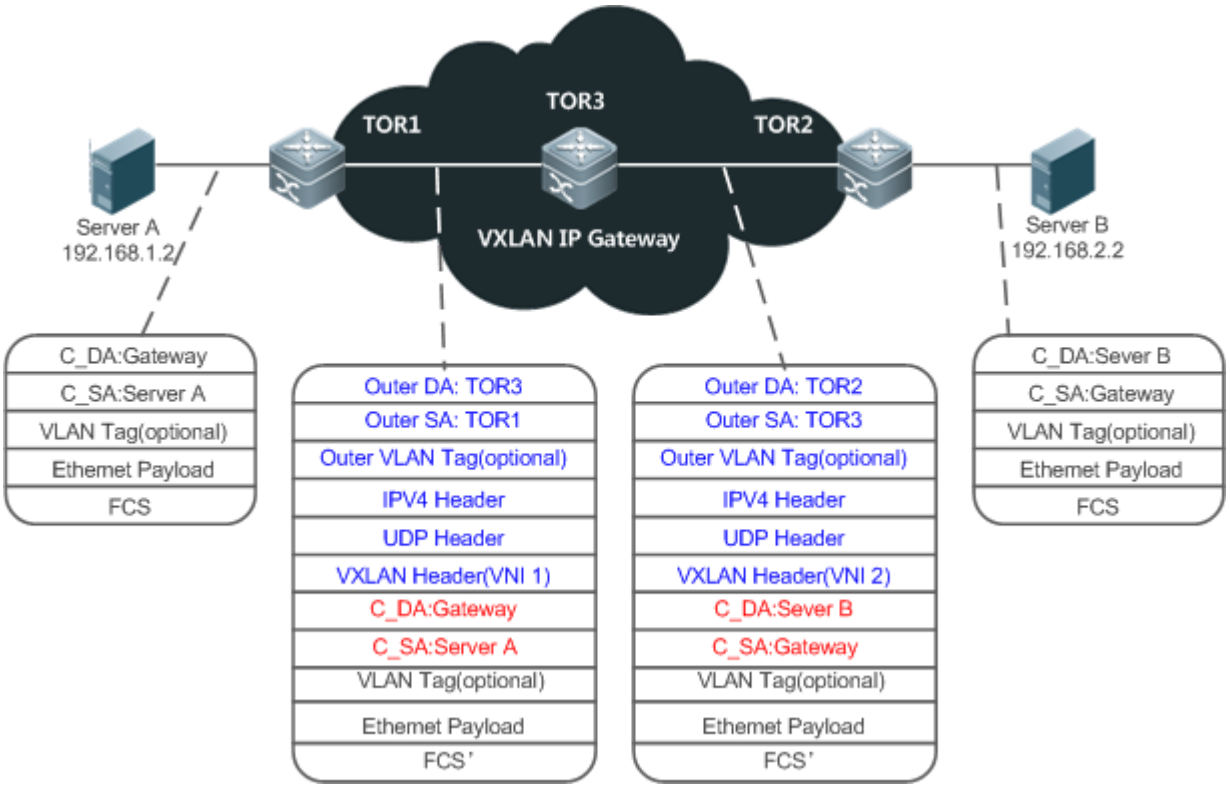

- To implement cross-VXLAN communication, Server A first sends a packet to the IP gateway, which is deployed on TOR3.
- The packet sent by Server A is encapsulated by TOR1 into a VXLAN packet and then sent to TOR3.
- After receiving the VXLAN packet, TOR3 finds that the destination MAC address is the local MAC address and sends the packet to TOR2 after VXLAN routing.
- After receiving the packet from TOR3, TOR2 decapsulates the packet and sends it to Server B.

Overview

Feature	Description
VXLAN Routing and Forwarding	Implements cross-VXLAN communication and supports communication between a conventional IP network and a VXLAN. A VXLAN router can serve as a VXLAN IP gateway.

1.3.3 Forwarding Process

Working Principle

AS shown in Figure 1-15, three servers use a VXLAN to achieve L2 interconnection on the IP network. The VXLAN VNI is 100.

Figure 1-15

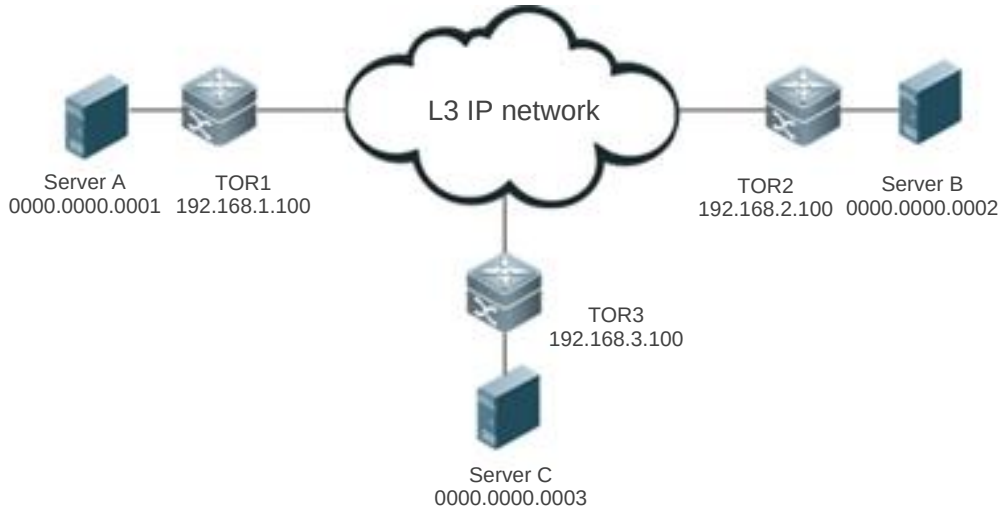

The VXLAN packet forwarding process is described by using an example in which Server A sends an Address Resolution Protocol (ARP) request to Server B and Server B returns an ARP response.

Figure 1-16

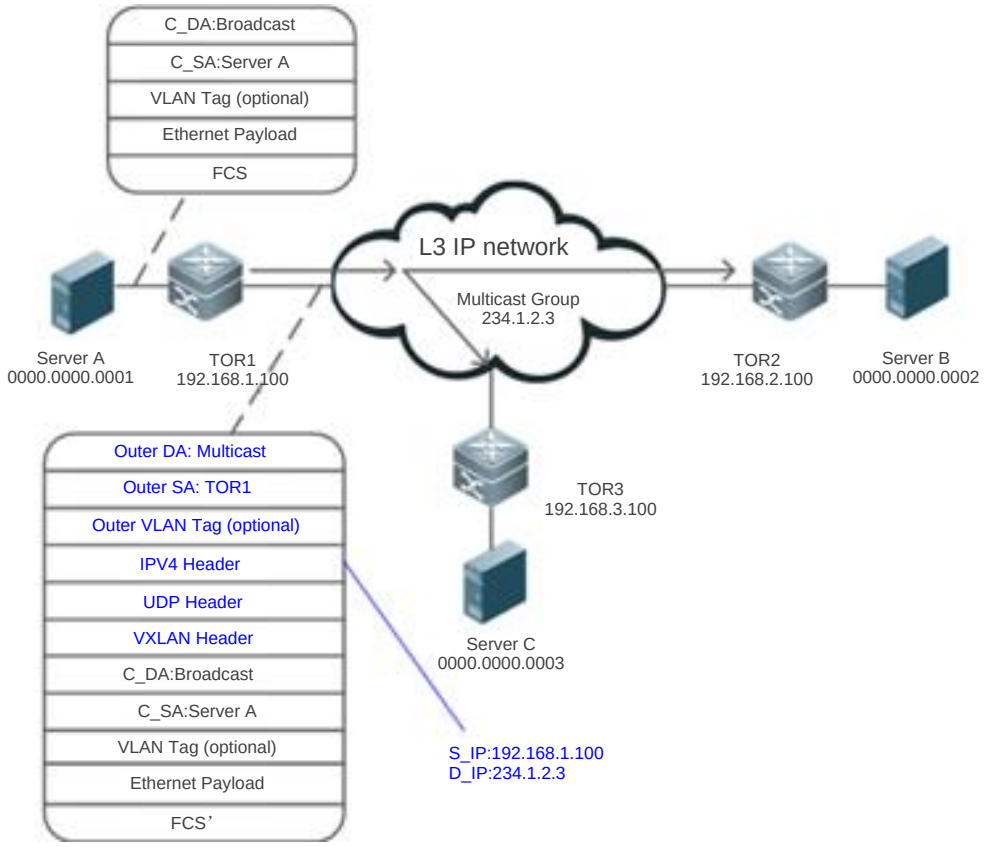

1. The same multicast group (234.1.2.3) is configured for TOR1, TOR2, and TOR3. Server A sends an ARP request packet to TOR1. As the packet is a broadcast packet, it is flooded in IP multicast mode. The destination IP address is 234.1.2.3 and the source IP address is 192.168.1.100.

2. The IP core network forwards the multicast VXLAN packet.

Figure 1-17

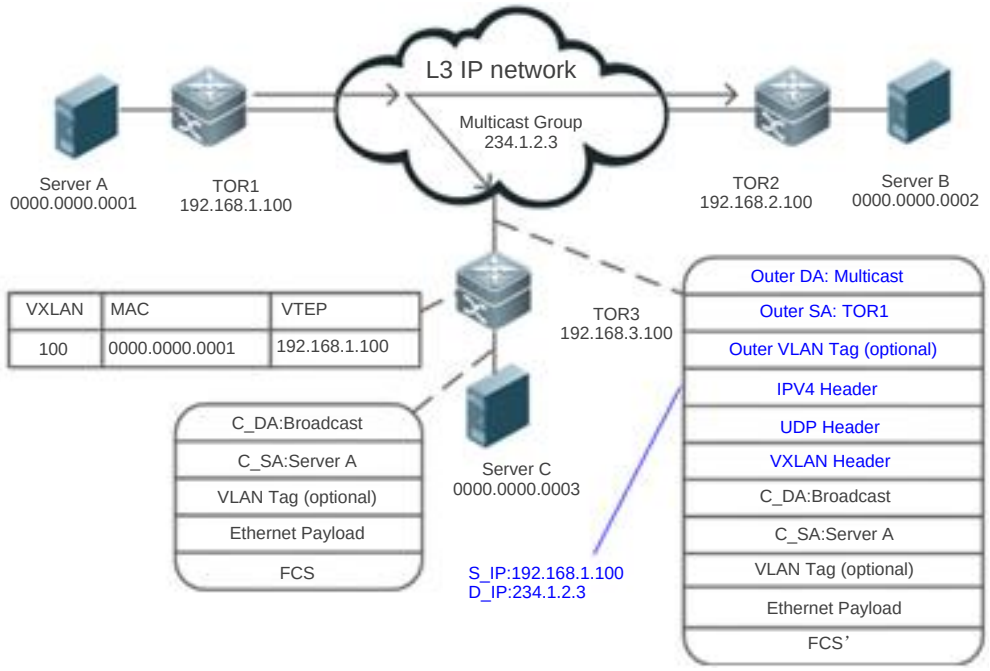

3. After receiving the VXLAN packet, TOR3 decapsulates the packet into an Ethernet packet and implements VXLAN address learning (the VXLAN ID is 100, the MAC address is 0000.0000.0001, and the IP address is 192.168.1.100).

Figure 1-18

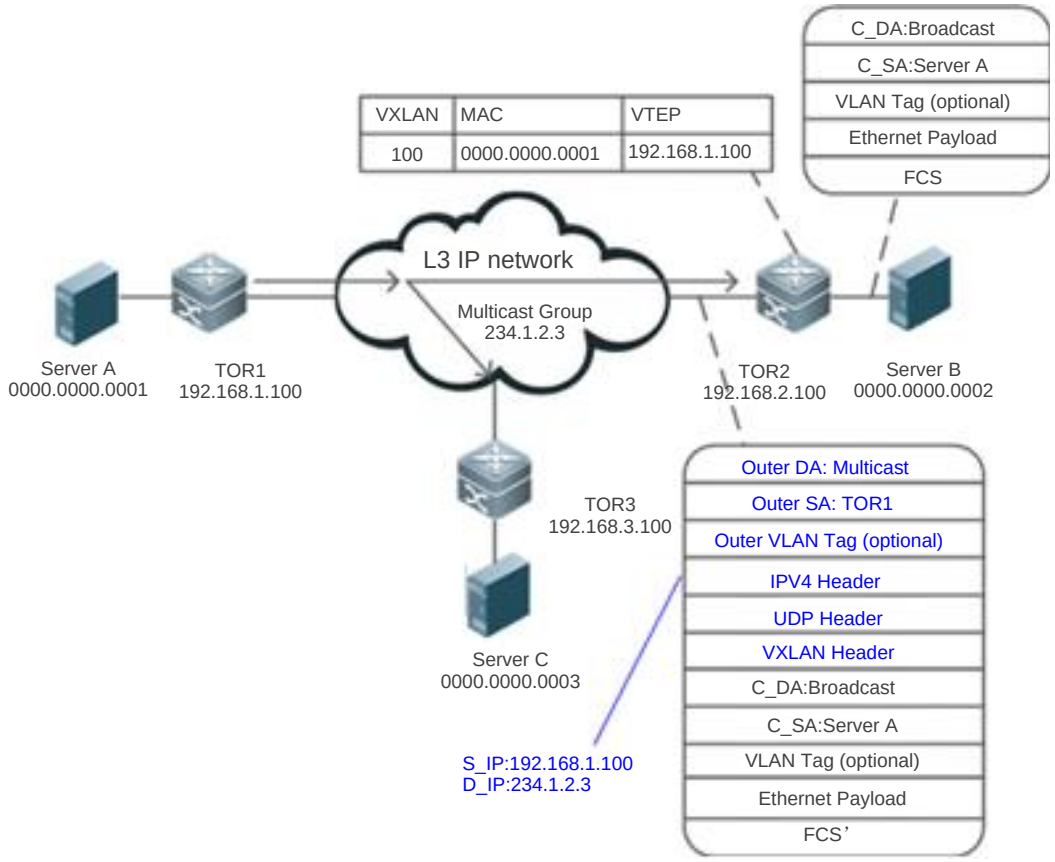

4. After receiving the VXLAN packet, TOR2 decapsulates the packet into an Ethernet packet, implements address learning (the VXLAN ID is 100, the MAC address is 0000.0000.0001, and the IP address is 192.168.1.100) and forwards the packet. Then, Server B receives the ARP request packet and returns a response packet.

Figure 1-19

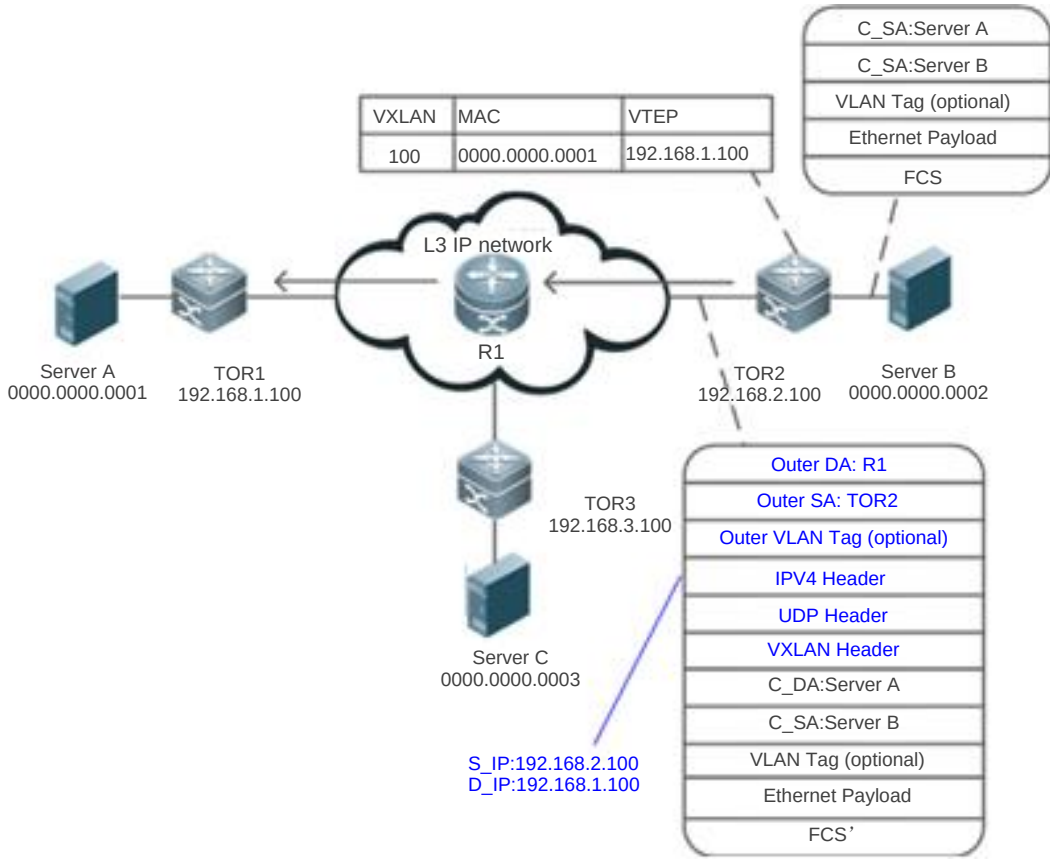

5. After receiving the ARP response packet from Server B, TOR2 searches the address table and finds that the destination IP address is 192.168.1.100. Then, TOR2 encapsulates the packet into a unicast VXLAN packet (the outer source IP address is 192.168.2.100) destined for the switch at 192.168.1.100.

Figure 1-20

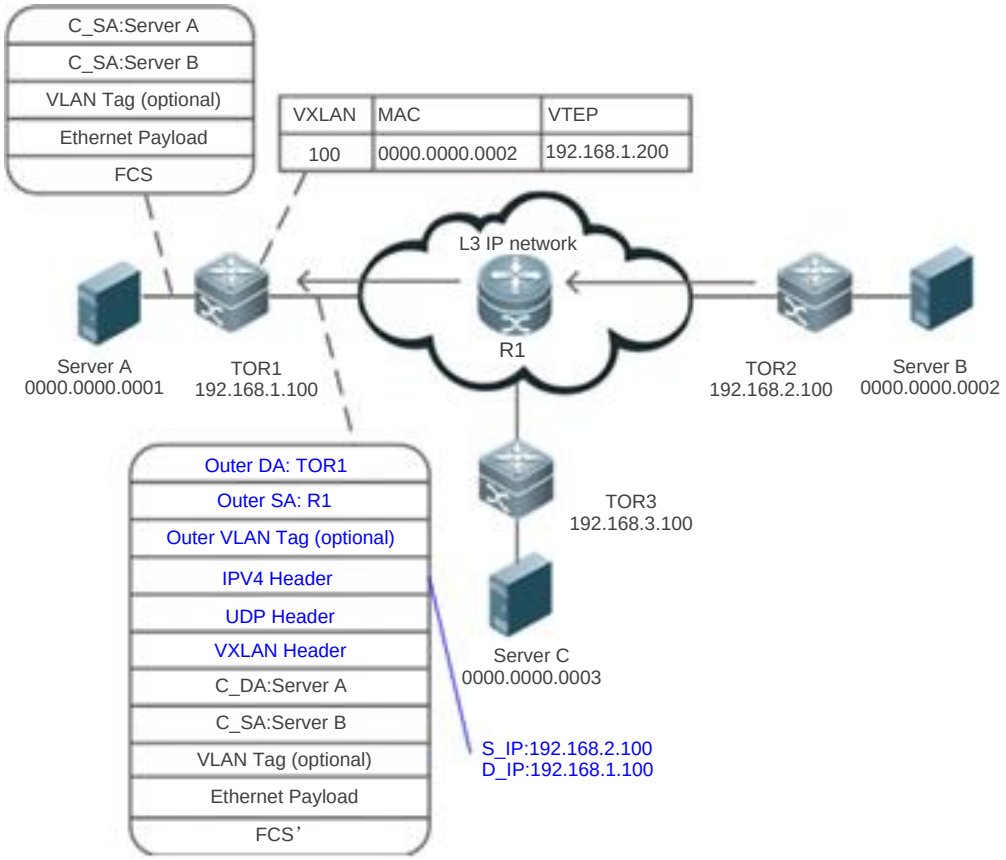

6. The IP core network forwards the VXLAN packet.
7. TOR1 receives the ARP response packet encapsulated in the VXLAN, decapsulates the packet into an Ethernet packet, implements VXLAN address learning (the VXLAN ID is 100, the MAC address is 0000.0000.0002, and the IP address is 192.168.2.100), and forwards the packet. Then, Server A receives the ARP response packet.

Multicast VXLAN Packet Flooding

A VXLAN uses multicast packets to flood broadcast, multicast, and unknown unicast packets. After receiving an ARP request packet, TOR1 encapsulates the packet into a multicast VXLAN packet and sends it to TOR2 and TOR3, as shown in Figure 1-21.

Figure 1-21

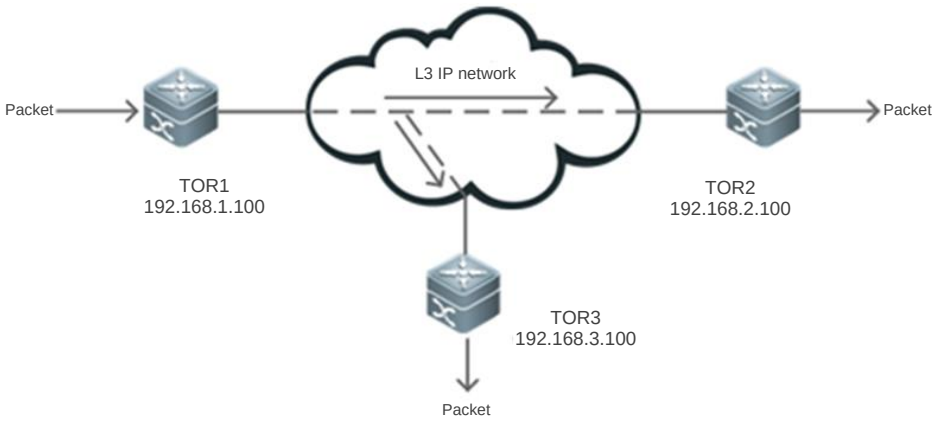

VTEP Address Learning

As shown in Figure 1-21, in the process of using multicast packets to flood broadcast, multicast, and unknown unicast packets, TOR2 and TOR3 learn the VTEP information during decapsulation, and therefore establish neighbor relationships.

Related Configuration

Configuring VXLAN Type Instance

No VXLAN instance is configured on the switches by default.

Run the **vxlan** *vni-number* command to create a VXLAN instance.

Configuring the Loopback Port on Local End

In VTEP configuration mode, a loopback port is required to a VTEP which should be configured with a unique VTEP IP as the source IP address of the VXLAN network.

Configuring VLAN Associated with VXLAN Instance

Run the **extend-vlan** *vlan-id* command in VXLAN instance configuration mode to configure the associated VLAN.

1.4 Configuration

1.4.1 Configuring VXLAN Bridging

Configuration Effect

- Create a VXLAN instance and provide IP core network-based L2 virtual network services.

Notes

- The VXLAN instances require support from existing unicast routes on the network. Therefore, an IPv4 unicast routing protocol, for example, the OSPF protocol, must be configured on the network devices.

Configuration Steps

Creating VXLAN Instances

- Mandatory.

Configuring Loopback Interface Associated with Local End

- Mandatory.

Associating VXLAN Instances with VLANs

- Mandatory.
- Only after a VLAN is associated with a VXLAN instance, packets of the VLAN can be encapsulated into VXLAN packets and then forwarded.
- After a VLAN is associated with a VXLAN, all packets of the VLAN will be encapsulated into VXLAN packets. Therefore, an SVI cannot be used as the VLAN IP gateway on the device.

Configuring VXLAN UDP Destination Port

- Optional. As the VXLAN UDP destination port used by earlier devices may not be Port 4789, you can run this command to achieve compatibility. In addition, you can also run this command to customize the VXLAN UDP destination port.
- The VXLAN UDP destination port 4789 designated by IANA is used by default.

Verification

Verify that packets of a VLAN associated with a VXLAN instance are forwarded as VXLAN packets.

- Verify that the local and remote devices of a VXLAN can receive and send packets of a VLAN associated with the VXLAN.
- Run the **show vxlan** *vni-number* command to check whether the local and remote VXLAN devices can learn mutual VTEP neighbor relationships.
- Run the **show vxlan mac** command to check whether the VXLAN MAC address is learned.
- Run the **show vxlan udp-port** command to display the VXLAN UDP destination port.

Related Commands

Creating or Entering VXLAN Instances

Command	**vxlan** *vni-number*
Parameter Description	*vni-number*: Indicates the VNI. The value ranges from 1 to 16,777,215.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuring Loopback Interface Associated with Local End

Command	**source** loopback *loopback-port-id*
Parameter Description	*loopback-port-id*: Indicates the loopback port ID.

Command Mode	VTEP configuration mode
Usage Guide	The local VTEP IP address is the configured loopback interface IP address.

Allowing VXLAN Instances to Forward VLAN Packets

Command	**extend-vlan** *vlan-id*
Parameter Description	*vlan-id*: Indicates the VXLAN instance that can forward VLAN packets.
Command Mode	VXLAN configuration mode
Usage Guide	Instances cannot use the same extended VLAN ID.

Configuring VXLAN UDP Destination Port

Command	**vxlan udp-port** *port-number*
Parameter Description	*port-number*: Indicates the UDP destination port ID. The value ranges from **0 to 65535** and the default value is **4789**.
Command Mode	Global configuration mode
Usage Guide	Note that the UDP destination port cannot be same as commonly used UDP ports.

Configuration Example

Configuring an VXLAN Instance

Scenario Figure1-22	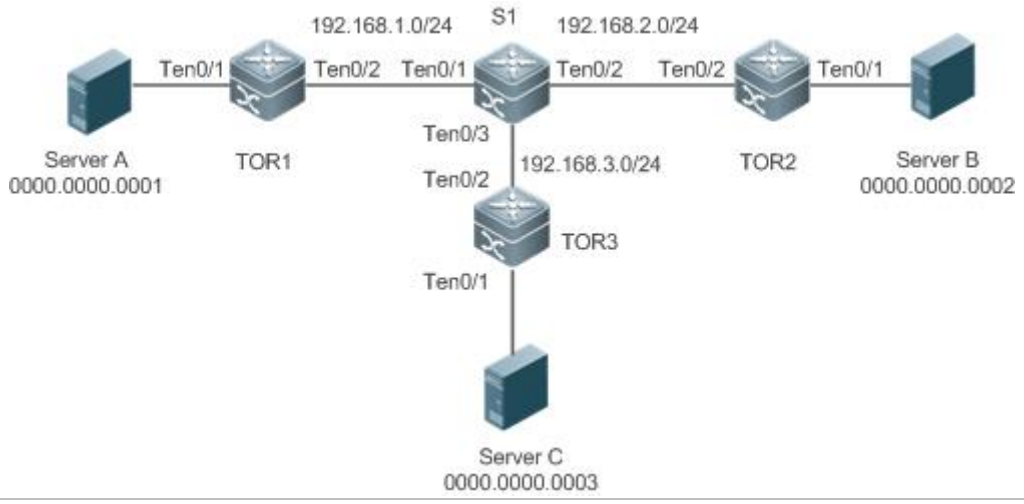
Configuration Steps	● Configure an IPv4 unicast routing protocol such as the OSPF protocol on TOR1, TOR2, TOR3, and S1 to ensure that unicast routes are reachable. ● Configure a VLAN (such as VLAN 100) on TOR1, TOR2, and TOR3.

	• Create a VXLAN instance on TOR1, TOR2, and TOR3. • Configure the loopback interface associated with local end on TOR1, TOR2, and TOR3. • Associate the VXLAN instance with the VLAN on TOR1, TOR2, and TOR3.
S1	S1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. S1(config)# interface TenGigabitEthernet 0/1 S1(config-if-TenGigabitEthernet 0/1)# no switchport S1(config-if-TenGigabitEthernet 0/1)# ip address 192.168.1.200 255.255.255.0 S1(config-if-TenGigabitEthernet 0/1)# exit S1(config)# interface TenGigabitEthernet 0/2 S1(config-if-TenGigabitEthernet 0/2)# no switchport S1(config-if-TenGigabitEthernet 0/2)# ip address 192.168.2.200 255.255.255.0 S1(config-if-TenGigabitEthernet 0/2)# exit S1(config)# interface TenGigabitEthernet 0/3 S1(config-if-TenGigabitEthernet 0/3)# no switchport S1(config-if-TenGigabitEthernet 0/3)# ip address 192.168.3.200 255.255.255.0 S1(config-if-TenGigabitEthernet 0/3)# exit S1(config)# interface Loopback 0 S1(config-if-Loopback 0)# ip address 10.10.10.10 255.255.255.255 S1(config)# router ospf 1 S1(config-router)# network 10.10.10.10 0.0.0.0 area 1 S1(config-router)# network 192.168.1.200 0.0.0.0 area 1 S1(config-router)# network 192.168.2.200 0.0.0.0 area 1 S1(config-router)# network 192.168.3.200 0.0.0.0 area 1 S1(config-router)# exit S1(config)# end S1(config)#
TOR1	TOR1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR1(config)# vlan 100 TOR1(config-vlan)# exit TOR1(config)# interface TenGigabitEthernet 0/1 TOR1(config-if-TenGigabitEthernet 0/1)# switchport mode access TOR1(config-if-TenGigabitEthernet 0/1)# switchport access vlan 100 TOR1(config-if-TenGigabitEthernet 0/1)# exit TOR1(config)# interface TenGigabitEthernet0/2 TOR1(config-if-TenGigabitEthernet 0/2)# ip address 192.168.1.100 225.255.255.0 TOR1(config-if-TenGigabitEthernet 0/2)# exit TOR1(config)# interface Loopback 0 TOR1(config-if-Loopback 0)# ip address 1.1.1.1 255.255.255.255 TOR1(config)# router ospf 1

	TOR1(config-router)# network 1.1.1.1 0.0.0.0 area 1 TOR1(config-router)# network 192.168.1.100 0.0.0.0 area 1 TOR1 (config-router)# exit TOR1(config)# vtep TOR1(config-vtep)# source loopback 0 TOR1(config-vtep)# exit TOR1(config)# vxlan 100 TOR1(config-vxlan)# extend-vlan 100 TOR1(config-vxlan)# end TOR1(config)#
TOR2	TOR2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR2(config)# vlan 100 TOR2(config-vlan)# exit TOR2(config)# interface TenGigabitEthernet 0/1 TOR2(config-if-TenGigabitEthernet 0/1)# switchport mode access TOR2(config-if-TenGigabitEthernet 0/1)# switchport access vlan 100 TOR2(config-if-TenGigabitEthernet 0/1)# exit TOR2(config)# interface TenGigabitEthernet0/2 TOR2(config-if-TenGigabitEthernet 0/2)# ip address 192.168.2.100 255.255.255.0 TOR2(config-if-TenGigabitEthernet 0/2)# exit TOR2 (config)# interface Loopback 0 TOR2 (config-if-Loopback 0)# ip address 2.2.2.2 255.255.255.255 TOR2 (config)# router ospf 1 TOR2 (config-router)# network 2.2.2.2 0.0.0.0 area 1 TOR2 (config-router)# network 192.168.2.100 0.0.0.0 area 1 TOR2 (config-router)# exit TOR2(config)# vtep TOR2(config-vtep)# source loopback 0 TOR2(config-vtep)# exit TOR2(config)# vxlan 100 TOR2(config-vxlan)# extend-vlan 100 TOR2(config-vxlan)# end TOR2(config)#
TOR3	TOR3# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR3(config)# vlan 100 TOR3(config-vlan)# exit TOR3(config)# interface TenGigabitEthernet 0/1 TOR3(config-if-TenGigabitEthernet 0/1)# switchport mode access TOR3(config-if-TenGigabitEthernet 0/1)# switchport access vlan 100

```
TOR3(config-if-TenGigabitEthernet 0/1)# exit
TOR3(config)# interface TenGigabitEthernet 0/2
TOR3(config-if-TenGigabitEthernet 0/2)# ip address 192.168.3.100 255.255.255.0
TOR3(config-if-TenGigabitEthernet 0/2)# exit
TOR3 (config)# interface Loopback 0
TOR3 (config-if-Loopback 0)# ip address 3.3.3.3 255.255.255.255
TOR3 (config)# router ospf 1
TOR3 (config-router)# network 3.3.3.3 0.0.0.0 area 1
TOR3 (config-router)# network 192.168.3.100 0.0.0.0 area 1
TOR3 (config-router)# exit
TOR3(config)# vtep
TOR3(config-vtep)# source loopback 0
TOR3(config-vtep)# exit
TOR3(config)# vxlan 100
TOR3(config-vxlan)# extend-vlan 100
TOR3(config-vxlan)# end
TOR3(config)#
```

Verification

- Verify that HOST-A, Host-B, and HOST-C can normally ping each other.

TOR1

```
TOR1# show vxlan 100
VXLAN 100
   Symmetric property  : FALSE
   Source Address     : 1.1.1.1
   Multicast Group    : -
   Extend VLAN        : 100
   VTEP Adjacency Count: 2
   VTEP Adjacency List :
   Interface           Source IP      Destination IP   Type
   ---------------------- --------------- --------------- -------
   OverlayTunnel 4097    1.1.1.1        2.2.2.2        dynamic
   OverlayTunnel 4098    1.1.1.1        3.3.3.3        dynamic
```

TOR2

```
TOR2# show vxlan 100
VXLAN 100
   Symmetric property  : FALSE
   Source Address     : 2.2.2.2
   Multicast Group    : -
   Extend VLAN        : 100
   VTEP Adjacency Count: 2
   VTEP Adjacency List :
   Interface           Source IP      Destination IP   Type
   ---------------------- --------------- --------------- -------
```

	OverlayTunnel 4097 2.2.2.2 1.1.1.1 dynamic OverlayTunnel 4098 2.2.2.2 3.3.3.3 dynamic
TOR3	TOR3# show vxlan 100 VXLAN 100 Symmetric property : FALSE Source Address : 3.3.3.3 Multicast Group : - Extend VLAN : 100 VTEP Adjacency Count: 2 VTEP Adjacency List : Interface Source IP Destination IP Type ---------- ---------- ---------- ------- OverlayTunnel 4097 3.3.3.3 1.1.1.1 dynamic OverlayTunnel 4098 3.3.3.3 2.2.2.2 dynamic

1.4.2 Configuring VXLAN Routing

Configuration Effect

- Create a VXLAN instance and associate it with the overlay router interface. Provide the VXLAN routing (IP gateway) function to achieve cross-VXLAN communication.

Notes

- The VXLAN instances require support from existing unicast routes on the network. Therefore, an IPv4 unicast routing protocol, for example, the OSPF protocol, must be configured on the network devices.

Configuration Steps

Creating VXLAN Instances

- Mandatory.

Configuring Loopback Interface Associated with Local End

- Mandatory.

Associating VXLAN Instance with Overlay Router Interface

- Mandatory.
- Only after the VXLAN is associated with the overlay router interface, the device can provide the VXLAN routing function and serve as a VXLAN IP gateway.

Configuring VXLAN UDP Destination Port

- Optional. As the VXLAN UDP destination port used by earlier devices may not be Port 4789, you can run this command to achieve compatibility. In addition, you can also run this command to customize the VXLAN UDP destination port.

- The VXLAN UDP destination port 4789 designated by IANA is used by default.

Verification

After the VXLAN routing function is configured, communication between VXLANs and that between a conventional IP network and a VXLAN are implemented.

- Run the **show vxlan** *vni-number* command to check whether the local and remote VXLAN devices can learn mutual VTEP neighbor relationships.
- Run the **show vxlan mac** command to check whether the VXLAN MAC address is learned.
- Run the **show arp** command to check whether the ARP entry of the VXLAN IP gateway is learned.
- Run the **show ip route** command to check whether the routes of VXLAN IP gateways are learned.
- Run the **show vxlan udp-port** command to display the VXLAN UDP destination port.

Related Commands

Creating or Entering VXLAN Instances

Command	**vxlan** *vni-number*
Parameter Description	*vni-number*: Indicates the VNI. The value ranges from 1 to 16,777,215.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuring Loopback Interface Associated with Local End

Command	**source** loopback *loopback-port-id*
Parameter Description	*loopback-port-id*: Indicates the loopback port ID.
Command Mode	VTEP configuration mode
Usage Guide	The local VTEP IP address is the configured loopback interface IP address.

Creating Overlay Router Interfaces

Command	**interface overlayrouter** *port-id*
Parameter Description	*port-id*: Indicates the ID of the overlay router interface.
Command Mode	Global configuration mode
Usage Guide	Similar to the SVI in a VLAN, this interface serves as the VXLAN IP gateway in the VXLAN routing environment.

Configuring IP Address for Overlay Router Interface

Command	**ip address** *ip-address mask*
Parameter Description	*ip-address*: Indicates the IP address of the overlay router interface. *mask*: Indicates the subnet mask.
Command Mode	Interface configuration mode
Usage Guide	Similar to the IP address of the SVI in a VLAN, this IP address serves as the address of the VXLAN IP gateway in the VXLAN routing environment.

Associating VXLAN Instance with Overlay Router Interface

Command	**router-interface** *interface-name*
Parameter Description	*interface-name*: Indicates the name of the overlay router interface.
Command Mode	VXLAN configuration mode
Usage Guide	Different VXLANs cannot be associated with the same overlay router interface.

Configuring VXLAN UDP Destination Port

Command	**vxlan udp-port** *port-number*
Parameter Description	*port-number*: Indicates the UDP destination port ID. The value ranges from **0** to **65535** and the default value is **4789**.
Command Mode	Global configuration mode
Usage Guide	Note that the UDP destination port cannot be same as commonly used UDP ports.

Configuration Example

Configuring an VXLAN Instance

Scenario	
Figure1-23	

<table>
<tr><td></td><td> 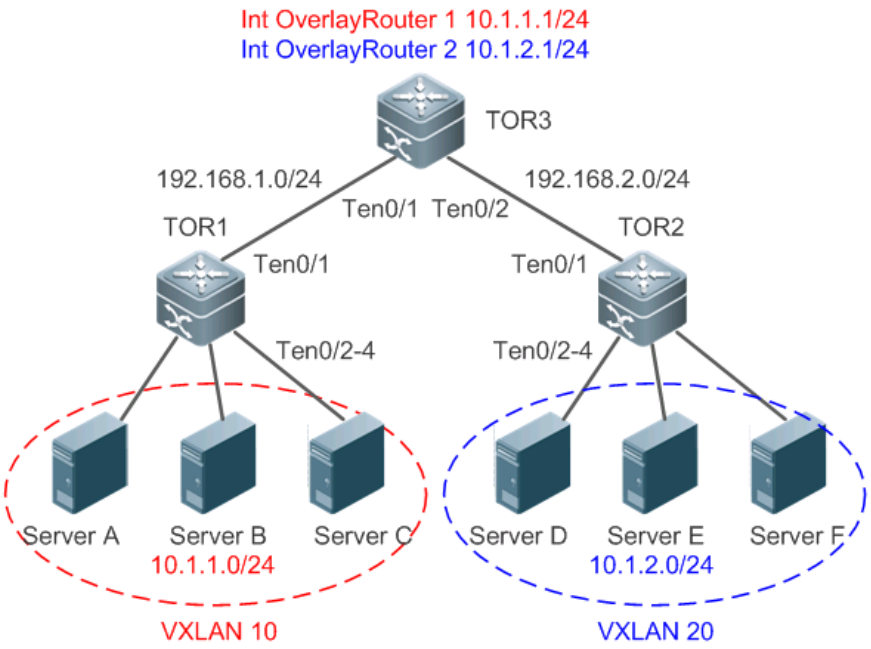
</td></tr>
<tr><td></td><td></td></tr>
<tr><td>Configuration Steps</td><td>

Configure an IPv4 unicast routing protocol such as the OSPF protocol on TOR1, TOR2, and TOR3 to ensure that unicast routes are reachable.
Configure a VLAN on TOR1 and TOR2 to implement VXLAN bridging.
Create a VXLAN instance on TOR1, TOR2, and TOR3.
Configure the loopback interface associated with local end on TOR1, TOR2, and TOR3.
Associate the VXLAN instance with the VLAN on TOR1 and TOR2.
Create an overlay router interface and configure the VXLAN gateway IP address on TOR3.
Associate the VXLAN instance with the overlay router interface on TOR3 to realize VXLAN routing.

</td></tr>
<tr><td>Server A</td><td>Configure the host IP address, and configure the gateway IP address as 10.1.1.1 on Server A. Configure Server B and Server C same as Server A. Ensure that Servers A, B, and C all belong to VXLAN 10.</td></tr>
<tr><td>Server D</td><td>Configure the host IP address, and configure the gateway IP address as 10.1.2.1 on Server D. Configure Server E and Server F same as Server D. Ensure that Servers D, E, and F all belong to VXLAN 20.</td></tr>
<tr><td>TOR1</td><td>
<pre>
TOR1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
TOR1(config)# interface TenGigabitEthernet 0/1
TOR1(config-if-TenGigabitEthernet 0/1)# ip address 192.168.1.100 255.255.255.0
TOR1(config-if-TenGigabitEthernet 0/1)# exit
TOR1(config)# interface Loopback 0
TOR1(config-if-Loopback 0)# ip address 1.1.1.1 255.255.255.255
TOR1(config-if-Loopback 0)# exit
TOR1(config)# router ospf 1
TOR1(config-router)# network 1.1.1.1 0.0.0.0 area 1
TOR1(config-router)# network 192.168.1.100 0.0.0.0 area 1
TOR1(config-router)# exit
</pre>
</td></tr>
</table>

	TOR1(config)# vlan 100 TOR1(config-vlan)# exit TOR1(config)# interface range TenGigabitEthernet 0/2-4 TOR1(config-if-range)# switchport mode access TOR1(config-if-range)# switchport access vlan 100 TOR1(config-if-range)# exit TOR1(config)# vtep TOR1(config-vtep)# source loopback 0 TOR1(config-vtep)# exit TOR1(config)# vxlan 10 TOR1(config-vxlan)# extend-vlan 100 TOR1(config-vxlan)# end TOR1(config)#
TOR2	TOR2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR2(config)# interface TenGigabitEthernet 0/1 TOR2(config-if-TenGigabitEthernet 0/1)# ip address 192.168.2.100 255.255.255.0 TOR2(config-if-TenGigabitEthernet 0/1)# exit TOR2(config)# interface Loopback 0 TOR2(config-if-Loopback 0)# ip address 2.2.2.2 255.255.255.255 TOR2(config-if-Loopback 0)# exit TOR2(config)# router ospf 1 TOR2(config-router)# network 2.2.2.2 0.0.0.0 area 1 TOR2(config-router)# network 192.168.2.100 0.0.0.0 area 1 TOR2(config-router)# exit TOR2(config)# vlan 200 TOR2(config-vlan)# exit TOR2(config)# interface range TenGigabitEthernet 0/2-4 TOR2(config-if-range)# switchport mode access TOR2(config-if-range)# switchport access vlan 200 TOR2(config-if-range)# exit TOR2(config)# vtep TOR2(config-vtep)# source loopback 0 TOR2(config-vtep)# exit TOR2(config)# vxlan 20

	TOR2(config-vxlan)# extend-vlan 200 TOR2(config-vxlan)# end TOR2(config)#
TOR3	TOR3# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR3(config)# interface TenGigabitEthernet 0/1 TOR3(config-if-TenGigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0 TOR3(config-if-TenGigabitEthernet 0/1)# exit TOR3(config)# interface TenGigabitEthernet 0/2 TOR3(config-if-TenGigabitEthernet 0/2)# ip address 192.168.2.1 255.255.255.0 TOR3(config-if-TenGigabitEthernet 0/2)# exit TOR3(config)# interface Loopback 0 TOR3(config-if-Loopback 0)# ip address 3.3.3.3 255.255.255.255 TOR3(config-if-Loopback 0)# exit TOR3(config)# router ospf 1 TOR3(config-router)# network 3.3.3.3 0.0.0.0 area 1 TOR3(config-router)# network 192.168.1.1 0.0.0.0 area 1 TOR3(config-router)# network 192.168.2.1 0.0.0.0 area 1 TOR3(config-router)# exit TOR3(config)# interface OverlayRouter 1 OR3(config-if-OverlayRouter 1)# overlay mode vxlan TOR3(config-if-OverlayRouter 1)# ip address 10.1.1.1 255.255.255.0 TOR3(config-if-OverlayRouter 1)# exit TOR3(config)# interface OverlayRouter 2 TOR3(config-if-OverlayRouter 2)# overlay mode vxlan TOR3(config-if-OverlayRouter 2)# ip address 10.1.2.1 255.255.255.0 TOR3(config-if-OverlayRouter 2)# exit TOR3(config)# vtep TOR3(config-vtep)# source loopback 0 TOR3(config-vtep)# exit TOR3(config)# vxlan 10 TOR3(config-vxlan)# router-interface OverlayRouter 1 TOR3(config-vxlan)# exit TOR3(config)# vxlan 20 TOR3(config-vxlan)# router-interface OverlayRouter 2 TOR3(config-vxlan)# end TOR3(config)#
Verification	● Verify that Servers A, B, and D can normally ping each other.
TOR1	TOR1# show vxlan 10

	VXLAN 10 Symmetric property : FALSE Source Address : 1.1.1.1 Multicast Group : - Extend VLAN : 100 VTEP Adjacency Count: 1 VTEP Adjacency List : Interface Source IP Destinaton IP Type ---------- ---------- ---------- ------- OverlayTunnel 4097 1.1.1.1 3.3.3.3 dynamic
TOR2	TOR2# show vxlan 20 VXLAN 20 Symmetric property : FALSE Source Address : 2.2.2.2 Multicast Group : - Extend VLAN : 200 VTEP Adjacency Count: 1 VTEP Adjacency List : Interface Source IP Destinaton IP Type ---------- ---------- ---------- ------- OverlayTunnel 4097 2.2.2.2 3.3.3.3 dynamic
TOR3	TOR3# show vxlan 10 VXLAN 10 Symmetric property : FALSE Source Address : 3.3.3.3 Multicast Group : - Router Interface : OverlayRouter 1(non-anycast) VTEP Adjacency Count: 1 VTEP Adjacency List : Interface Source IP Destinaton IP Type ---------- ---------- ---------- ------- OverlayTunnel 4097 3.3.3.3 1.1.1.1 dynamic TOR3# show vxlan 20 VXLAN 20 Symmetric property : FALSE Source Address : 3.3.3.3 Multicast Group : - Router Interface : OverlayRouter 2(non-anycast) VTEP Adjacency Count: 1 VTEP Adjacency List :

```
   Interface            Source IP      Destinaton IP   Type
   -------------------- -------------- -------------- -------
   OverlayTunnel 4097    3.3.3.3        2.2.2.2        dynamic

TOR3#show arp
Protocol  Address         Age(min)  Hardware       Type   Interface
Internet  192.168.1.1      --       001f.ce10.4589  arpa   GigabitEthernet 0/1
Internet  192.168.1.100    11       001f.ce22.33cb  arpa   GigabitEthernet 0/1
Internet  192.168.2.200    12       001f.ce40.3997  arpa   GigabitEthernet 0/2
Internet  192.168.2.1      --       001f.ce0.458a  arpa   GigabitEthernet 0/2
Internet  10.1.1.1         --       001f.ce10.4589  arpa   OverlayRouter 1
Internet  10.1.2.1         --       001f.ce10.4589  arpa   OverlayRouter 2
Internet  10.1.1.2         1        001f.ce10.aaaa  arpa   OverlayRouter 1
Internet  10.1.1.3         1        001f.ce10.bbbb  arpa   OverlayRouter 1
Internet  10.1.2.2         1        001f.ce10.dddd  arpa   OverlayRouter 2
```

Common Errors

Make sure that the device is in VXLAN router mode. You can run the **show vxlan mode** command to display the current mode.

1.4.3 Configuring VXLAN EVPN

Configuration Effect

- Enable the control plane learning function to implement VXLAN tunnel learning, MAC address learning, and route learning via control plane protocols, thereby finally implementing VXLAN bridging, VXLAN routing, and data communication between VXLANs and that between a VXLAN and an external network.
- Support functions such as anycast gateways, symmetric VXLAN instances, and ARP suppression in EVPN control plane mode.

Notes

- The VXLAN instances require support from existing unicast routes on the network. Therefore, an IPv4 unicast routing protocol, for example, the OSPF protocol, must be configured on the network devices.
- The MP-BGP-EVPN protocol is required for VXLANs to implement VXLAN tunnel learning, MAC address learning, and route learning. Therefore, the devices on the network must complete BGP-related configurations.

Configuration Steps

↘ **Configuring Control Plane Mode**

↘ Mandatory. The default mode is EVPN mode.

↘ **Configuring Loopback Interface Associated with Local End**

- Mandatory.

- Configure the loopback interface IP address as the VTEP IP address of the local end. One VTEP device can be associated with only one loopback interface and use the IP address of the loopback interface as the VXLAN VTEP IP address.

Configuring Virtual MAC Address for Anycast Gateways

- Optional.
- Configure a unified virtual MAC address for all anycast gateways on the network. The anycast function can be enabled on the VXLAN overlay router interface of the local device only after the virtual MAC address is configured.

Configuring ARP Suppression

- Optional.
- After ARP suppression is enabled, the switch responds to the ARP request from the host as a proxy, reducing the flooded ARP data.
- ARP suppression is generally enabled on the TOR bridge devices in a centralized deployment scenario, or on the distributed gateways in a distributed deployment scenario.

Configuring the Route Import Function

- Optional.
- You can run the **member add** *vni* command for a VXLAN instance of a device only after the route import function is globally enabled, so that the VXLAN route after VNI inter-import can correctly replace the VNI information of the next hop. This function is required only when VXLAN routes need to be imported in multiple-tenant environments.

Configuring Remote ARP Packet Learning

- Mandatory for centralized gateways and not recommended for other devices.
- After the remote ARP packet learning function is enabled, the VXLAN gateways can learn the VXLAN route entries from the encapsulated VXLAN ARP packets received from the VXLAN tunnels.

Creating VXLAN Instances

- Mandatory.

Associating VXLAN Instance with Overlay Router Interface

- Mandatory for VXLAN gateways.
- Only after the VXLAN is associated with the overlay router interface, the device can provide the VXLAN routing function and serve as a VXLAN IP gateway.

Associating VXLAN Instances with VLANs

- Mandatory for VXLAN devices directly connected to the host.
- Only after a VLAN is associated with a VXLAN instance, packets of the VLAN can be encapsulated into VXLAN packets and then forwarded.

- After a VLAN is associated with a VXLAN, all packets of the VLAN will be encapsulated into VXLAN packets. Therefore, an SVI cannot be used as the VLAN IP gateway on the device.

Configuring the VNI Mapped by a Symmetric VXLAN Instance Route

- Optional.
- You need to run the **import-route enable** command to globally enable the route import function before you can configure the VNI.
- In EVPN mode, if you import a VXLAN route across VRF networks through RD and RT of BGP, you need to run the **import-route enable** and **member add** *vni* commands to ensure that the imported VXLAN route can correctly replace the VNI required for forwarding.

Configuring VXLAN UDP Destination Port

- Optional. As the VXLAN UDP destination port used by earlier devices may not be Port 4789, you can run this command to achieve compatibility. In addition, you can also run this command to customize the VXLAN UDP destination port.
- The VXLAN UDP destination port 4789 designated by IANA is used by default.

Configuring Symmetric Instances

- Optional.
- Symmetric instances need to be configured only in symmetric scenarios. Only one symmetric instance can be configured for each VRF network. After a symmetric instance is configured in a VRF network, L3 forwarding of other asymmetric instances is all taken over by the symmetric instance for implementation.

Configuring Rate Limit for Tunnel Interfaces

- Optional.
- Configure the input/output rate limit on the tunnel interface if you need to limit the rate of the tunnel.

Configuring Static VXLAN Routes

- Optional.
- Configure the static VXLAN routes based on VXLAN instances if required.

Verification

Based on EVPN control plane learning, VXLAN tunnels, VXLAN MAC entries, and VXLAN route entries can be generated. Run the following commands for verification.

- Run the **show vxlan** *vni-number* command to check whether the local and remote VXLAN devices can learn mutual VTEP neighbor relationships.
- Run the **show vxlan mac** command to check whether the VXLAN MAC address is learned.
- Run the **show vxlan arp** command to check whether the ARP entry of the VXLAN IP gateway is learned.
- Run the **show vxlan route** command to check whether route entries are learned by the VXLAN gateway.

- Run the **show vxlan prefix-route** command to display VXLAN route entries.
- Run the **show vxlan udp-port** command to display the VXLAN UDP destination port.

Related Commands

Configuring Loopback Interface Associated with Local End

Command	**source** loopback *loopback-port-id*
Parameter Description	*loopback-port-id*: Indicates the loopback port ID.
Command Mode	VTEP configuration mode
Usage Guide	The local VTEP IP address is the configured loopback interface IP address.

Configuring Virtual MAC Address for Anycast Gateways

Command	**fabric anycast-gateway-mac** *mac-addr*
Parameter Description	*mac-addr*: Indicates the MAC address. The format is xxxx.xxxx.xxxx.
Command Mode	VTEP configuration mode
Usage Guide	All gateways on which the anycast function is enabled use this MAC address as the gateway MAC address. The virtual MAC address for an anycast gateway must not be the same as the local MAC address or the same as the MAC address of any device on the overlay network.

Configuring Remote ARP Packet Learning

Command	**remote arp learn enable**
Parameter Description	N/A
Command Mode	VTEP configuration mode
Usage Guide	Enable or disable the remote ARP packet learning function globally. After this function is enabled, the VXLAN gateways can learn the VXLAN route entries from the encapsulated VXLAN ARP packets received from the VXLAN tunnels.

Configuring ARP Suppression

Command	**arp suppress enable**
Parameter Description	N/A
Command Mode	VTEP configuration mode
Usage Guide	Enable or disable ARP suppression globally. After ARP suppression is enabled, the switch responds to the ARP requests from the host as a proxy.

Configuring the Route Import Function

Command	**import-route enable**
Parameter Description	N/A
Command Mode	VTEP configuration mode
Usage Guide	You can use this command to enable and disable the route import function globally. You can run the **member add** *vni* command for a VXLAN instance of a device only after the route import function is globally enabled on the device, so that the VXLAN route after VNI inter-import can correctly replace the VNI information of the next hop. This function is required only when VXLAN routes need to be imported in multiple-tenant environments.

Creating Overlay Router Interfaces

Command	**interface overlayrouter** *port-id*
Parameter Description	*port-id*: Indicates the ID of the overlay router interface.
Command Mode	Global configuration mode
Usage Guide	Similar to the SVI in a VLAN, this interface serves as the VXLAN IP gateway in the VXLAN routing environment.

Configuring IP Address for Overlay Router Interface

Command	**ip address** *ip-address mask*
Parameter Description	*ip-address*: Indicates the IP address of the overlay router interface. *mask*: Indicates the subnet mask.
Command Mode	Interface configuration mode
Usage Guide	Similar to the IP address of the SVI in a VLAN, this IP address serves as the address of the VXLAN IP gateway in the VXLAN routing environment.

Associating Overlay Router Interface with VRF Network

Command	**vrf forwarding** *table name*
Parameter Description	*table name*: Indicates the VRF network associated with the overlay router interface.
Command Mode	Interface configuration mode
Usage Guide	Use this command to associate an overlay router interface with a VRF network in the VXLAN routing environment, to implement VXLAN L3 route isolation.

Creating or Entering VXLAN Instances

Command	**vxlan** *vni-number*
Parameter Description	*vni-number*: Indicates the VNI. The value ranges from **1** to **16777215**.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuring Symmetric Instances

Command	**symmetric**
Parameter Description	N/A
Command Mode	VXLAN configuration mode
Usage Guide	No symmetric instance is configured by default. Symmetric instances are used to manage the L3 forwarding entries of all asymmetric instances of the VRF networks associated with the symmetric instances.

Associating VXLAN Instance with Overlay Router Interface

Command	**router-interface** *interface-name*
Parameter Description	*interface-name*: Indicates the name of the overlay router interface.
Command Mode	VXLAN configuration mode
Usage Guide	Different VXLANs cannot be associated with the same overlay router interface.

Configuring VXLAN UDP Destination Port

Command	**vxlan udp-port** *port-number*
Parameter Description	*port-number*: Indicates the UDP destination port ID. The value ranges from **0** to **65535** and the default value is **4789**.
Command Mode	Global configuration mode
Usage Guide	Note that the UDP destination port cannot be same as commonly used UDP ports.

Configuring the VNI Mapped by a Symmetric VXLAN Instance Route

Command	**member add** *vni-number*
Parameter Description	*vni-number*: Indicates the VNI. The value ranges from **1** to **16777215**.
Command Mode	VXLAN configuration mode
Usage Guide	In EVPN mode, if you import a VXLAN route across VRF networks through RD and RT of BGP, you need to run the **import-route enable** and **member add** *vni* commands to ensure that the imported VXLAN route can

	correctly replace the VNI required for forwarding.

Configuring Rate Limit for Tunnel Interfaces

Command	**vxlan overlaytunnel dip** *ip-address* **rate-limit { output** *rate-num* **| input** *rate-num* **}**
Parameter Description	*Ip-address*: Indicates the VTEP IP address of the peer end of the tunnel interface. *rate-num*: Indicates the rate limit value.
Command Mode	VTEP configuration mode
Usage Guide	N/A

Configuring Static VXLAN Routes

Command	**vxlan ip route** *network net-mask ip-address* **vni** *vni-number*
Parameter Description	*network*: Specifies the address of the target network. *net-mask*: Specifies the mask of the target network. *ip-address*: Specifies the next hop address of the static route. *vni-number*: Indicates the VNI. The value ranges from **1** to **16777215**.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuration Example

1.4.3.1 Configuring EVPN-based Multi-tenant Centralized Scenario

Scenario **Figure 1-24**	

	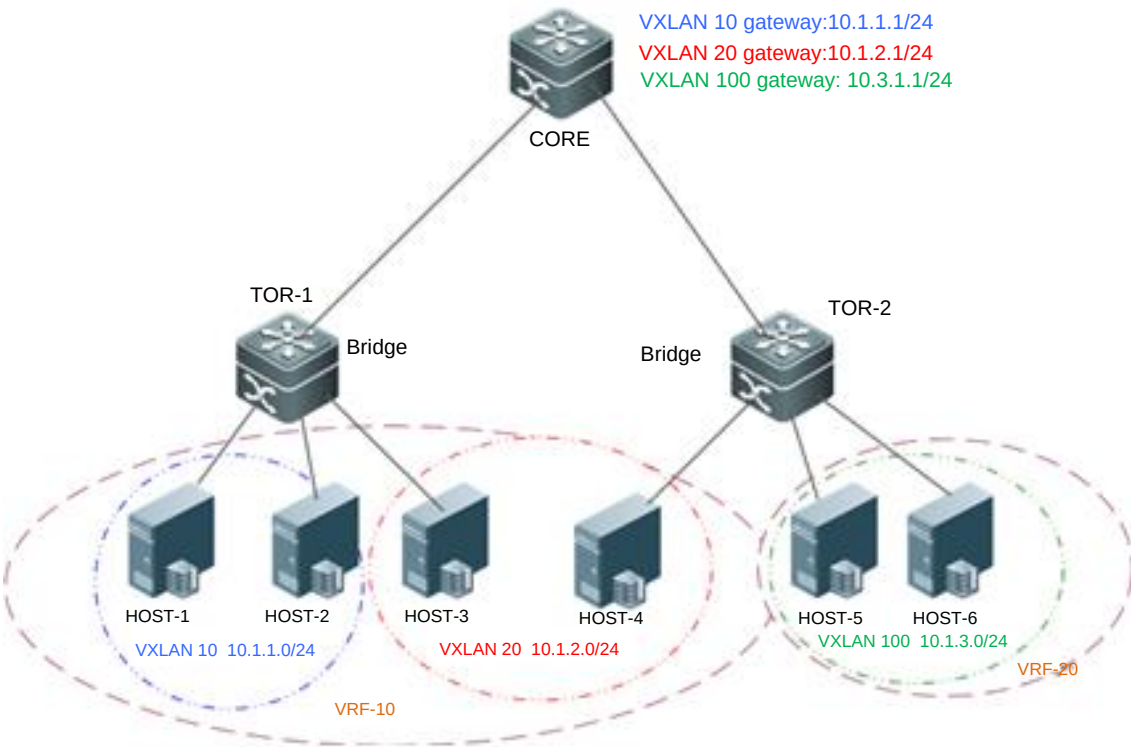
Configuration Steps	● Configure an IPv4 unicast routing protocol such as the OSPF protocol on CORE, TOR-1, and TOR-2 to ensure that unicast routes are reachable. ● Configure the BGP-EVPN routing protocol on CORE, TOR-1, and TOR-2 to establish BGP neighbor relationships between the three devices and to support the EVPN protocol family. ● Configure the EVI for BGP-EVPN on CORE, TOR-1, and TOR-2. For details, see *BGP-EVPN Configuration Guide*. ● Configure a VXLAN on the virtual server and designate the gateway address of the virtual machine. ● Associate the VTEP with the loopback interface on TOR-1, TOR-2, and CORE to establish tunnels. ● Create VXLAN instances on TOR-1, TOR-2, and CORE and associate the VXLAN instances with VLANs. ● Create overlay router interfaces and configure the VXLAN gateway IP address on CORE. Configure different VRF networks for different overlay router interfaces to determine their respective tenants. ● Associate VXLAN instances with overlay router interfaces on CORE to realize VXLAN routing. ● Enable the remote ARP packet learning function on CORE to generate VXLAN routing entries dynamically. ● (Optional) Configure ARP suppression on TOR-1 and TOR-2 to reduce the ARP packets entering the VXLAN.
HOST	Configuring the IP address and gateway according to Figure 1-24 (the detailed configuration on the server is omitted herein).
CORE	The configuration of the OSPF, BGP, and Ethernet interface is omitted herein. The following describe only the VXLAN configuration.

	CORE# configure terminal Enter configuration commands, one per line. End with CNTL/Z. CORE (config)# interface Loopback 1 CORE (config-if- Loopback 1)# ip address 1.1.1.1/32 CORE (config-if- Loopback 1)#exit CORE (config)# vtep CORE (config-vtep)# source loopback 1 CORE (config-vtep)#exit CORE (config)# int overlayrouter 10 CORE (config-if-OverlayRouter 10)# vrf forwarding vrf-10 CORE (config-if-OverlayRouter 10)# ip address 10.1.1.1/24 CORE (config-if-OverlayRouter 10)# exit CORE (config)# vxlan 10 CORE (config-vxlan)# extend-vlan 10 CORE (config-vxlan)# router-interface OverlayRouter 10 CORE (config-vxlan)#exit CORE (config)# int overlayrouter 20 CORE (config-if-OverlayRouter 20)# vrf forwarding vrf-10 CORE (config-if-OverlayRouter 20)# ip address 10.1.2.1/24 CORE (config-if-OverlayRouter 20)# exit CORE (config)# vxlan 20 CORE (config-vxlan)# extend-vlan 20 CORE (config-vxlan)# router-interface OverlayRouter 20 CORE (config-vxlan)#exit CORE (config)# int overlayrouter 100 CORE (config-if-OverlayRouter 100)# vrf forwarding vrf-20 CORE (config-if-OverlayRouter 100)# ip address 10.1.3.1/24 CORE (config-if-OverlayRouter 100)# exit CORE (config)# vxlan 100 CORE (config-vxlan)# extend-vlan 100 CORE (config-vxlan)# router-interface OverlayRouter 100 CORE (config-vxlan)#exit
TOR1	TOR1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR1 (config)# interface Loopback 1 TOR1 (config-if- Loopback 1)# ip address 2.2.2.2/32 TOR1 (config-if- Loopback 1)#exit TOR1 (config)# vtep TOR1 (config-vtep)# source loopback 1 TOR1 (config-vtep)# arp suppress enable

```
TOR1 (config-vtep)#exit
TOR1 (config)# vxlan 10
TOR1 (config-vxlan)# extend-vlan 10
TOR1 (config-vxlan)#exit
TOR1 (config)# vxlan 20
TOR1 (config-vxlan)# extend-vlan 20
TOR1 (config-vxlan)#exit
```

TOR2

```
TOR2# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
TOR2 (config)# interface Loopback 1
TOR2 (config-if- Loopback 1)# ip address 3.3.3.3/32
TOR2 (config-if- Loopback 1)#exit
TOR2 (config)# vtep
TOR2 (config-vtep)# source loopback 1
TOR2 (config-vtep)# arp suppress enable
TOR2 (config-vtep)#exit
TOR2 (config)# vxlan 100
TOR2 (config-vxlan)# extend-vlan 100
TOR2 (config-vxlan)#exit
TOR2 (config)# vxlan 20
TOR2 (config-vxlan)# extend-vlan 20
TOR2 (config-vxlan)#exit
```

Verification

- Verify that HOST-1, HOST-2, HOST-3, and HOST-4 can ping each other.
- Verify that HOST-5 and HOST-6 can ping each other.
- Verify that HOST-1, HOST-2, HOST-3, and HOST-4 cannot ping HOST-5 and HOST-6.
- Verify that the virtual machines can be migrated between the hosts on the same VXLAN and can access the network normally after migration without modifying the configuration.

```
TOR1#sho vxlan
VXLAN Total Count: 2
VXLAN Capacity   : 8000

VXLAN 10
   Symmetric property  : FALSE
   Source Address      : 2.2.2.2
   Multicast Group     : -
   Router Interface    : -
   Extend VLAN         : 10
   VTEP Adjacency Count: 1
   VTEP Adjacency List :
```

```
VXLAN 20
   Symmetric property  : FALSE
   Source Address     : 2.2.2.2
   Multicast Group    : -
   Router Interface   : -
   Extend VLAN        : 20
   VTEP Adjacency Count: 2
   VTEP Adjacency List :

CORE#sho vxlan
VXLAN Total Count: 3
VXLAN Capacity   : 8000

VXLAN 10
   Symmetric property  : FALSE
   Source Address     : 1.1.1.1
   Multicast Group    : -
   Router Interface   : OverlayRouter 10 (non-anycast)
   Extend VLAN        : 10
   VTEP Adjacency Count: 1
   VTEP Adjacency List :

VXLAN 20
   Symmetric property  : FALSE
   Source Address     : 1.1.1.1
   Multicast Group    : -
   Router Interface   : OverlayRouter 20 (non-anycast)
   Extend VLAN        : 20
   VTEP Adjacency Count: 2
   VTEP Adjacency List :

VXLAN 100
   Symmetric property  : FALSE
   Source Address     : 1.1.1.1
   Multicast Group    : -
   Router Interface   : OverlayRouter 100 (non-anycast)
   Extend VLAN        : 100
   VTEP Adjacency Count: 1
   VTEP Adjacency List :
```

Common Errors

Make sure that the devices are in VXLAN device mode, that is router (EVPN) mode. You can run the **show vxlan mode** command to display the current mode.

1.4.3.2 Configuring EVPN-based Multi-tenant Centralized All-active Anycast Gateway Scenario

Scenario **Figure 1-25**	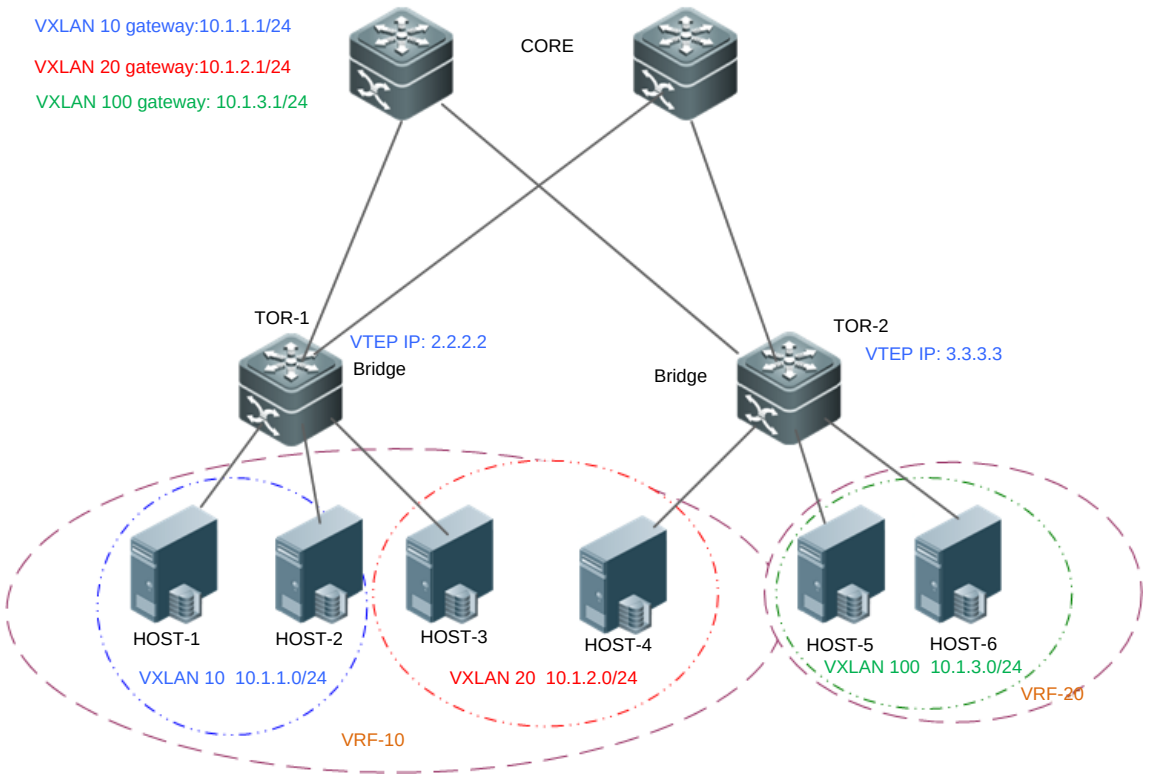
Configuration Steps	● Configure an IPv4 unicast routing protocol such as the OSPF protocol on the COREs and TORs to ensure that unicast routes are reachable. ● Configure the BGP-EVPN routing protocol on the COREs and TORs to establish BGP neighbor relationships between the three devices and to support the EVPN protocol family. ● Configure the EVI for BGP-EVPN on the TORs. For details, see *BGP-EVPN Configuration Guide*. ● Configure a VXLAN on the virtual server and designate the gateway address of the virtual machine.(Omitted) ● Associate the VTEP with the loopback interface on TOR-1 and TOR-2 to establish tunnels. ● Create anycast gateway MAC addresses on TOR-1 and TOR-2 so that all VXLAN anycast gateways on the network use the same virtual MAC address. ● Create VXLAN instances on TOR-1 and TOR-2, and associate the VXLAN instances with VLANs. ● Create overlay router interfaces and configure the VXLAN gateway IP address on TOR-1 and TOR-2.

	Configure different VRF networks for different overlay router interfaces to determine their respective tenants. Configure anycast gateways so that all VXLAN gateways on the network use the same IP/MAC address. As anycast gateway is enabled, the VXLAN gateway IP addresses configured for the overlay router interfaces with which the same VXLAN associates must be consistent on the two TORs. ● Associate VXLAN instances with overlay router interfaces on TOR-1 and TOR-2 to realize VXLAN routing. ● (Optional) Configure ARP suppression on TOR-1 and TOR-2 to reduce the ARP packets entering the VXLAN.
HOST	Configure the IP address and gateway according to Figure 1-25 (the detailed configuration on the server is omitted herein).
CORE	VXLANS may be not configured on the core switches. The configuration of the OSPF and BGP is omitted herein.
TOR1	TOR1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR1 (config)# interface Loopback 1 TOR1 (config-if- Loopback 1)# ip address 2.2.2.2/32 TOR1 (config-if- Loopback 1)#exit TOR1 (config)# vtep TOR1 (config-vtep)# source loopback 1 TOR1 (config-vtep)# arp suppress enable TOR1 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 TOR1 (config-vtep)#exit TOR1 (config)# int overlayrouter 10 TOR1 (config-if-OverlayRouter 10)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24 TOR1 (config-if-OverlayRouter 10)# anycast-gateway TOR1 (config-if-OverlayRouter 10)# exit TOR1 (config)# vxlan 10 TOR1 (config-vxlan)# extend-vlan 10 TOR1 (config-vxlan)# router-interface OverlayRouter 10 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 20 TOR1 (config-if-OverlayRouter 20)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 20)# ip address 10.1.2.1/24 TOR1 (config-if-OverlayRouter 20)# anycast-gateway TOR1 (config-if-OverlayRouter 20)# exit TOR1 (config)# vxlan 20 TOR1 (config-vxlan)# extend-vlan 20 TOR1 (config-vxlan)# router-interface OverlayRouter 20 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 100

	TOR1 (config-if-OverlayRouter 100)# vrf forwarding vrf-20 TOR1 (config-if-OverlayRouter 100)# ip address 10.1.3.1/24 TOR1 (config-if-OverlayRouter 100)# anycast-gateway TOR1 (config-if-OverlayRouter 100)# exit TOR1 (config)# vxlan 100 TOR1 (config-vxlan)# extend-vlan 100 TOR1 (config-vxlan)# router-interface OverlayRouter 100 TOR1 (config-vxlan)#exit
TOR2	TOR2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR2 (config)# interface Loopback 1 TOR2 (config-if- Loopback 1)# ip address 3.3.3.3/32 TOR2 (config-if- Loopback 1)#exit TOR2 (config)# vtep TOR2 (config-vtep)# source loopback 1 TOR2 (config-vtep)# arp suppress enable TOR2 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 TOR2 (config-vtep)#exit TOR2 (config)# int overlayrouter 10 TOR2 (config-if-OverlayRouter 10)# vrf forwarding vrf-10 TOR2 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24 TOR2 (config-if-OverlayRouter 10)# anycast-gateway TOR2 (config-if-OverlayRouter 10)# exit TOR2 (config)# vxlan 10 TOR2 (config-vxlan)# extend-vlan 10 TOR2 (config-vxlan)# router-interface OverlayRouter 10 TOR2 (config-vxlan)#exit TOR2 (config)# int overlayrouter 20 TOR2 (config-if-OverlayRouter 20)# vrf forwarding vrf-10 TOR2 (config-if-OverlayRouter 20)# ip address 10.1.2.1/24 TOR2 (config-if-OverlayRouter 20)# anycast-gateway TOR2 (config-if-OverlayRouter 20)# exit TOR2 (config)# vxlan 20 TOR2 (config-vxlan)# extend-vlan 20 TOR2 (config-vxlan)# router-interface OverlayRouter 20 TOR2 (config-vxlan)#exit TOR2 (config)# int overlayrouter 100 TOR2 (config-if-OverlayRouter 100)# vrf forwarding vrf-20 TOR2 (config-if-OverlayRouter 100)# ip address 10.1.3.1/24

```
TOR2 (config-if-OverlayRouter 100)# anycast-gateway
TOR2 (config-if-OverlayRouter 100)# exit
TOR2 (config)# vxlan 100
TOR2 (config-vxlan)# extend-vlan 100
TOR2 (config-vxlan)# router-interface OverlayRouter 100
TOR2 (config-vxlan)#exit
```

Verification

- Verify that HOST-1, HOST-2, HOST-3, and HOST-4 can ping each other.
- Verify that HOST-5 and HOST-6 can ping each other.
- Verify that HOST-1, HOST-2, HOST-3, and HOST-4 cannot ping HOST-5 and HOST-6.
- Verify that virtual machines can be migrated between the hosts on the same VXLAN and can access the network normally after migration without modifying the configuration.

```
TOR1#sho vxlan
VXLAN Total Count: 3
VXLAN Capacity   : 8000

VXLAN 10
   Symmetric property  : FALSE
   Source Address      : 2.2.2.2
   Multicast Group     : -
   Router Interface    : OverlayRouter 10 (anycast)
   Extend VLAN        : 10
   VTEP Adjacency Count: 1
   VTEP Adjacency List :

VXLAN 20
   Symmetric property  : FALSE
   Source Address      : 2.2.2.2
   Multicast Group     : -
   Router Interface    : OverlayRouter 20 (anycast)
   Extend VLAN        : 20
   VTEP Adjacency Count: 1
   VTEP Adjacency List :

VXLAN 100
   Symmetric property  : FALSE
   Source Address      : 2.2.2.2
   Multicast Group     : -
   Router Interface    : OverlayRouter 100(anycast)
   Extend VLAN        : 100
   VTEP Adjacency Count: 1
```

```
VTEP Adjacency List :
```

Common Errors

- Ensure that the device is in VXLAN Device Mode: Router (EVPN) mode. You can run the **show vxlan mode** command to display the current mode.
- When symmetric deployment is disabled, all TOR switches of the same VRF network must have all the VXLAN gateways of the VRF network configured on the TOR switches. For example, if VRF-10 includes VXLAN 10 and VXLAN 20, and therefore all gateways of VXLAN 10 and VXLAN 20 must be configured on TOR-1 and TOR-2. Otherwise, VXLAN 10 and VXLAN 20 cannot communicate with each other. If you expect to deploy only required gateways instead of deploying all gateways on all TOR switches, apply symmetric deployment. For details, see section 1.4.3.3 "Configuring EVPN-based Multi-tenant Distributed Scenario (Symmetric Deployment)."
- Make sure that the global anycast MAC address is not the same as that of any device on the VXLAN.

1.4.3.3 Configuring EVPN-based Multi-tenant Distributed Scenario (Symmetric Deployment)

Scenario

Figure 1-26

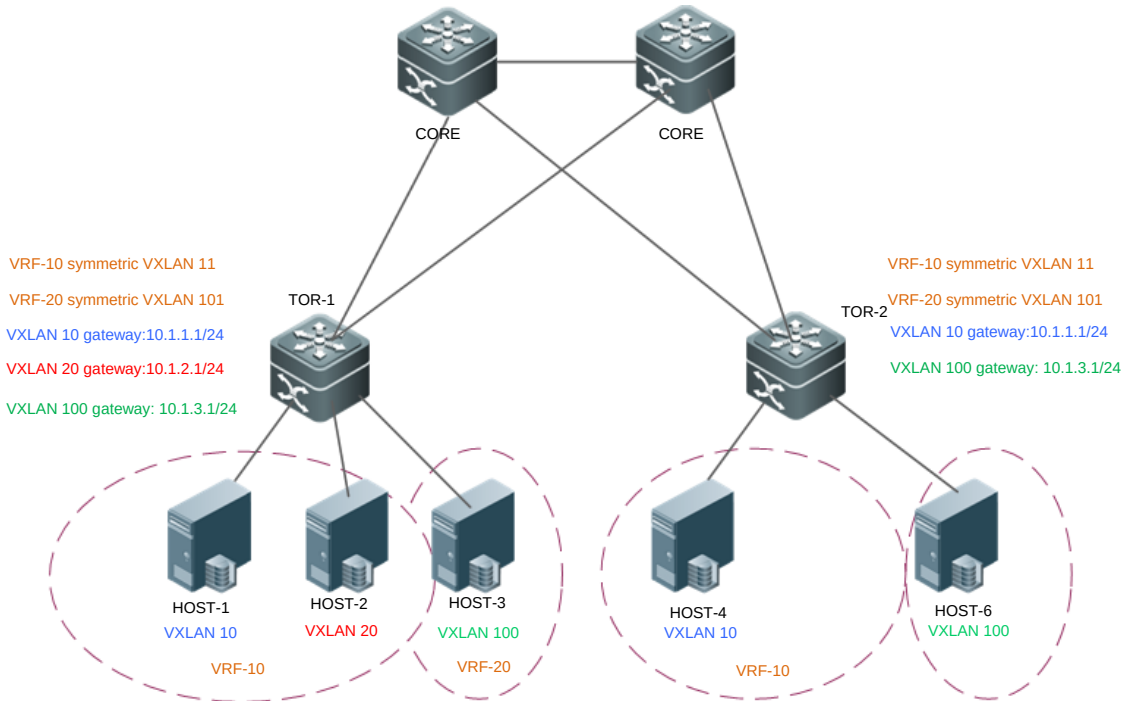

Configuration Steps

- Configure an IPv4 unicast routing protocol such as the OSPF protocol on CORE, TOR-1, and TOR-2 to ensure that unicast routes are reachable.
- Configure the BGP-EVPN routing protocol on CORE, TOR-1, and TOR-2 to establish BGP neighbor relationships between the four devices and to support the EVPN protocol family.
- Configure the EVI for BGP-EVPN on TOR-1 and TOR-2. For details, see *BGP-EVPN Configuration*

	Guide. ● Configure a VXLAN on the virtual server and designate the gateway address of the virtual machine. ● Associate the VTEP with loopback interface on TOR-1 and TOR-2 to establish tunnels. ● Configure the anycast gateway MAC address on TOR-1 and TOR-2 to ensure that all VXLAN anycast gateways on the network use the same MAC address. ● Create VXLAN 10, VXLAN20, and VXLAN 100 on TOR-1 and associate them with VLANs. ● Create VXLAN 10 and VXLAN 100 on TOR-2 and associate them with VLANs. ● Create overlay router interfaces for VXLAN 10, VXLAN 20, and VXLAN 100 on TOR-1 and TOR-2 (TOR-2 do not have VXLAN 20), and configure the VXLAN gateway IP address for them. Configure different VRF networks for different overlay router interfaces to determine their respective tenants. Configure the anycast gateway to ensure that all VXLAN gateways on the network use the same IP address and MAC address. As the anycast gateway function is enabled, the overlay router interfaces associated with the same VXLAN on TOR-1 and TOR-2 must be configured with the same VXLAN gateway IP address. ● Create VXLAN 11 and VXLAN 101 on TOR-1 and TOR-2 and configure them as symmetric VXLANs to serve as L3 routing VXLAN of the corresponding VRF networks. L3 routes between all VXLANs of the same VRF network are advertised via the symmetric VXLANs. In addition, the symmetric VXLANs are also used for L3 routing and forwarding. ● Create overlay router interfaces for VXLAN 11 and VXLAN 101 on TOR-1 and TOR-2. Configure different VRF networks for the overlay router interfaces. VXLAN 11 and VXLAN 101 serve as the symmetric VXLANs of the corresponding VRF networks. ● Associate VXLAN instances with overlay router interfaces on TOR-1 and TOR-2 to realize VXLAN routing. ● (Optional) Configure ARP suppression on TOR-1 and TOR-2 to reduce the ARP packets entering the VXLAN.
HOST	Configuring the IP address and gateway according to Figure 1-26 (the detailed configuration on the server is omitted herein).
CORE	VXLAN may be not configured on the core switches. The configuration of the OSPF and BGP is omitted herein.

TOR1	`TOR1# configure terminal`
	`Enter configuration commands, one per line.  End with CNTL/Z.`
	`TOR1 (config)# interface Loopback 1`
	`TOR1 (config-if- Loopback 1)# ip address 2.2.2.2/32`
	`TOR1 (config-if- Loopback 1)#exit`
	`TOR1 (config)# fabric anycast-gateway-mac 0011.2233.2016`
	`TOR1 (config)# vtep`
	`TOR1 (config-vtep)# source loopback 1`
	`TOR1 (config-vtep)# arp suppress enable`
	`TOR1 (config-vtep)#exit`
	`TOR1 (config)# int overlayrouter 10`
	`TOR1 (config-if-OverlayRouter 10)# vrf forwarding vrf-10`

```
TOR1 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24
TOR1 (config-if-OverlayRouter 10)# anycast-gateway
TOR1 (config-if-OverlayRouter 10)# exit
TOR1 (config)# vxlan 10
TOR1 (config-vxlan)# extend-vlan 10
TOR1 (config-vxlan)# router-interface OverlayRouter 10
TOR1 (config-vxlan)#exit
TOR1 (config)# int overlayrouter 20
TOR1 (config-if-OverlayRouter 20)# vrf forwarding vrf-10
TOR1 (config-if-OverlayRouter 20)# ip address 10.1.2.1/24
TOR1 (config-if-OverlayRouter 20)# anycast-gateway
TOR1 (config-if-OverlayRouter 20)# exit
TOR1 (config)# vxlan 20
TOR1 (config-vxlan)# extend-vlan 20
TOR1 (config-vxlan)# router-interface OverlayRouter 20
TOR1 (config-vxlan)#exit
TOR1 (config)# int overlayrouter 100
TOR1 (config-if-OverlayRouter 100)# vrf forwarding vrf-20
TOR1 (config-if-OverlayRouter 100)# ip address 10.1.3.1/24
TOR1 (config-if-OverlayRouter 100)# anycast-gateway
TOR1 (config-if-OverlayRouter 100)# exit
TOR1 (config)# vxlan 100
TOR1 (config-vxlan)# extend-vlan 100
TOR1 (config-vxlan)# router-interface OverlayRouter 100
TOR1 (config-vxlan)#exit
TOR1 (config)# int overlayrouter 11
TOR1 (config-if-OverlayRouter 11)# vrf forwarding vrf-10
TOR1 (config-if-OverlayRouter 11)# exit
TOR1 (config)# vxlan 11
TOR1 (config-vxlan)# symmetric
TOR1 (config-vxlan)# router-interface OverlayRouter 11
TOR1 (config-vxlan)#exit
TOR1 (config)# int overlayrouter 101
TOR1 (config-if-OverlayRouter 101)# vrf forwarding vrf-20
TOR1 (config-if-OverlayRouter 101)# exit
TOR1 (config)# vxlan 101
TOR1 (config-vxlan)# symmetric
TOR1 (config-vxlan)# router-interface OverlayRouter 101
TOR1 (config-vxlan)#exit
```

TOR2

```
TOR2# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
TOR2 (config)# interface Loopback 1
TOR2 (config-if- Loopback 1)# ip address 2.2.2.2/32
TOR2 (config-if- Loopback 1)#exit
TOR2 (config)# fabric anycast-gateway-mac 0011.2233.2016
TOR2 (config)# vtep
TOR2 (config-vtep)# source loopback 1
TOR2 (config-vtep)# arp suppress enable
TOR2 (config-vtep)#exit
TOR2 (config)# int overlayrouter 10
TOR2 (config-if-OverlayRouter 10)# vrf forwarding vrf-10
TOR2 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24
TOR2 (config-if-OverlayRouter 10)# anycast-gateway
TOR2 (config-if-OverlayRouter 10)# exit
TOR2 (config)# vxlan 10
TOR2 (config-vxlan)# extend-vlan 10
TOR2 (config-vxlan)# router-interface OverlayRouter 10
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 100
TOR2 (config-if-OverlayRouter 100)# vrf forwarding vrf-20
TOR2 (config-if-OverlayRouter 100)# ip address 10.1.3.1/24
TOR2 (config-if-OverlayRouter 100)# anycast-gateway
TOR2 (config-if-OverlayRouter 100)# exit
TOR2 (config)# vxlan 100
TOR2 (config-vxlan)# extend-vlan 100
TOR2 (config-vxlan)# router-interface OverlayRouter 100
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 11
TOR2 (config-if-OverlayRouter 11)# vrf forwarding vrf-10
TOR2 (config-if-OverlayRouter 11)# exit
TOR2 (config)# vxlan 11
TOR2 (config-vxlan)# symmetric
TOR2 (config-vxlan)# router-interface OverlayRouter 11
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 101
TOR2 (config-if-OverlayRouter 101)# vrf forwarding vrf-20
TOR2 (config-if-OverlayRouter 101)# exit
TOR2 (config)# vxlan 101
TOR2 (config-vxlan)# symmetric
TOR2 (config-vxlan)# router-interface OverlayRouter 101
```

```
TOR2 (config-vxlan)#exit
```

Verification

- Verify that HOST-1, HOST-2, and HOST-4 can ping each other.
- Verify that HOST-3 and HOST-6 can ping each other.
- Verify that HOST-1, HOST-2, and HOST-4 cannot ping HOST-3 and HOST-6.
- Verify that the virtual machines can be migrated between the hosts on the same VXLAN and can access the network normally after migration without modifying the configuration.

```
TOR1#sho vxlan
VXLAN Total Count: 5
VXLAN Capacity   : 8000

VXLAN 10
   Symmetric property  : FALSE
   Source Address     : 2.2.2.2
   Multicast Group    : -
   Router Interface   : OverlayRouter 10 (anycast)
   Extend VLAN        : 10
   VTEP Adjacency Count: 1

VXLAN 11
   Symmetric property  : TRUE
   Source Address     : 2.2.2.2
   Multicast Group    : -
   Router Interface   : -
   Extend VLAN        : -
   VTEP Adjacency Count: 1

VXLAN 20
   Symmetric property  : FALSE
   Source Address     : 2.2.2.2
   Multicast Group    : -
   Router Interface   : OverlayRouter 20 (anycast)
   Extend VLAN        : 20
   VTEP Adjacency Count: 1

VXLAN 100
   Symmetric property  : FALSE
   Source Address     : 2.2.2.2
   Multicast Group    : -
   Router Interface   : OverlayRouter 100 (anycast)
   Extend VLAN        : 100
```

```
    VTEP Adjacency Count: 1

VXLAN 101
    Symmetric property  : TRUE
    Source Address      : 2.2.2.2
    Multicast Group     : -
    Router Interface    : -
    Extend VLAN         : -
    VTEP Adjacency Count: 1
```

Common Errors

- Make sure that the devices are in VXLAN device mode, that is, router (EVPN) mode. You can run the **show vxlan mode** command to display the current mode.
- Make sure that the global anycast MAC address is not the same as that of any device on the VXLAN.

1.4.3.4 Configuring EVPN-based Single-tenant VXLAN Routing Scenario

Scenario

Figure 1-27

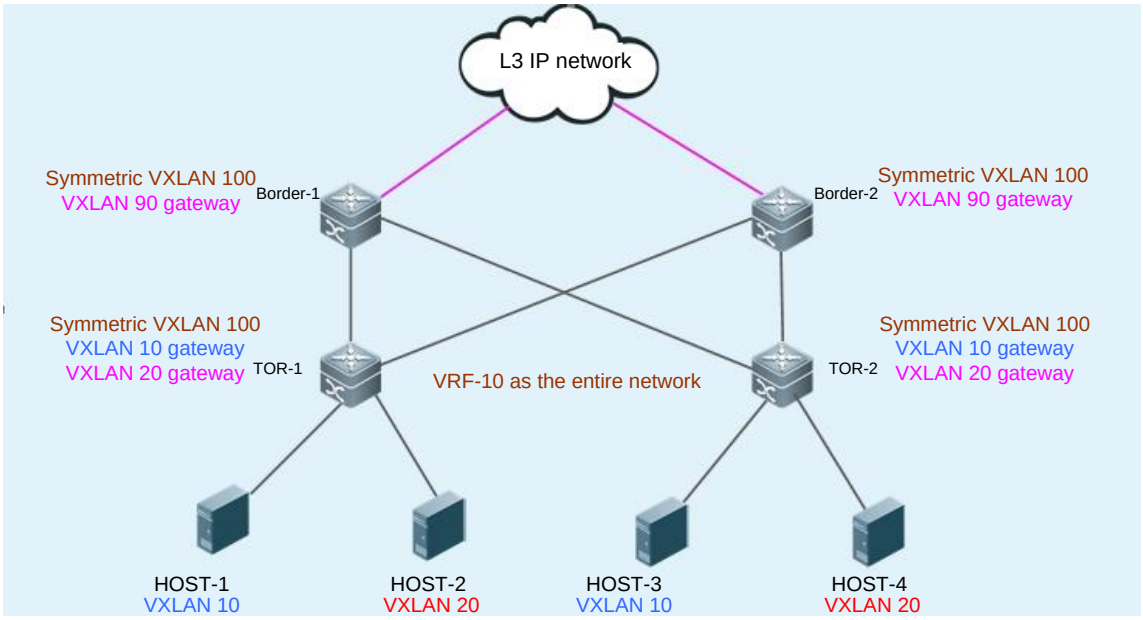

Configuration Steps

- Configure an IPv4 unicast routing protocol such as the OSPF protocol on Border-1, Border-2, TOR-1, and TOR-2 to ensure that unicast routes are reachable.
- Configure the BGP-EVPN routing protocol on Border-1, Border-2, TOR-1, and TOR-2 to establish BGP neighbor relationships between the devices (except between Border-1 and Border-2) and to support

	the EVPN protocol family. ● Configure the EVI for BGP-EVPN on TOR-1 and TOR-2. For details, see *BGP-EVPN Configuration Guide*. ● Configure a VXLAN on the virtual server and designate the gateway address of the virtual machine. ● Associate the VTEP with the loopback interface on TOR-1, TOR-2, Border-1, and Border-2 to establish tunnels. ● Configure the anycast gateway MAC address on TOR-1 and TOR-2 to ensure that all VXLAN anycast gateways on the network use the same MAC address. ● Create VXLAN 10 and VXLAN 20 on TOR-1 and associate them with VLANs. ● Create VXLAN 10 and VXLAN 20 on TOR-2 and associate them with VLANs. ● Create VXLAN 90 on Border-2 and associate VXLAN 90 with a VLAN. ● Create overlay router interfaces for VXLAN 10 and VXLAN 20 on TOR-1 and TOR-2, and configure the VXLAN gateway IP address for them. Configure the same VRF network for the overlay router interfaces to determine their respective tenants. Configure the anycast gateway to ensure that all VXLAN gateways on the network use the same IP address and MAC address. As the anycast gateway function is enabled, the overlay router interfaces associated with the same VXLAN on TOR-1 and TOR-2 must be configured with the same VXLAN gateway IP address. ● Create VXLAN 100 on TOR-1, TOR-2, Border-1, and Border-2. Configure VXLAN 100 as a symmetric VXLAN to serve as the L3 routing VXLAN of the corresponding VRF network. L3 routes between all VXLANs of the same VRF network are advertised via the symmetric VXLAN. In addition, the symmetric VXLAN is also used for L3 routing and forwarding. ● Create overlay router interfaces for VXLAN 100 on TOR-1 and TOR-2 and configure the same VRF network for the overlay router interfaces. VXLAN 100 serves as the symmetric VXLAN of the VRF network. ● Create overlay router interfaces for VXLAN 100 on Border-1 and Border-2, and configure the same VRF network for the overlay router interfaces, so that VXLAN 100 serves as the symmetric VXLAN of the VRF network. Configure VXLAN gateway IP addresses for Border-1 and Border-2 (different IP addresses for different devices). ● Create overlay router interfaces for VXLAN 90 on Border-1 and Border-2. Configure the same VRF network for the overlay router interfaces and configure the VXLAN gateway IP address. ● Associate VXLAN instances with overlay router interfaces on TOR-1, TOR-2, Border-1, and Border-2 to realize VXLAN routing. ● (Optional) Configure ARP suppression on TOR-1 and TOR-2 to reduce the ARP packets entering the VXLAN.
HOST	Configuring the IP address and gateway according to Figure 1-27 (the detailed configuration on the server is omitted herein).

TOR1	`TOR1# configure terminal`
	`Enter configuration commands, one per line.  End with CNTL/Z.`
	`TOR1 (config)# interface Loopback 1`
	`TOR1 (config-if- Loopback 1)# ip address 1.1.1.1/32`
	`TOR1 (config-if- Loopback 1)#exit`

	TOR1 (config)# vtep TOR1 (config-vtep)# source loopback 1 TOR1 (config-vtep)# arp suppress enable TOR1 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 TOR1 (config-vtep)#exit TOR1 (config)# int overlayrouter 10 TOR1 (config-if-OverlayRouter 10)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24 TOR1 (config-if-OverlayRouter 10)# anycast-gateway TOR1 (config-if-OverlayRouter 10)# exit TOR1 (config)# vxlan 10 TOR1 (config-vxlan)# extend-vlan 10 TOR1 (config-vxlan)# router-interface OverlayRouter 10 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 20 TOR1 (config-if-OverlayRouter 20)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 20)# ip address 20.1.1.1/24 TOR1 (config-if-OverlayRouter 20)# anycast-gateway TOR1 (config-if-OverlayRouter 20)# exit TOR1 (config)# vxlan 20 TOR1 (config-vxlan)# extend-vlan 20 TOR1 (config-vxlan)# router-interface OverlayRouter 20 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 100 TOR1 (config-if-OverlayRouter 100)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 100)# exit TOR1 (config)# vxlan 100 TOR1 (config-vxlan)# symmetric TOR1 (config-vxlan)# router-interface OverlayRouter 100 TOR1 (config-vxlan)#exit
TOR2	TOR2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR2 (config)# interface Loopback 1 TOR2 (config-if- Loopback 1)# ip address 2.2.2.2/32 TOR2 (config-if- Loopback 1)#exit TOR2 (config)# vtep TOR2 (config-vtep)# source loopback 1 TOR2 (config-vtep)# arp suppress enable TOR2 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 TOR2 (config-vtep)#exit TOR2 (config)# int overlayrouter 10

	TOR2 (config-if-OverlayRouter 10)# vrf forwarding vrf-10 TOR2 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24 TOR2 (config-if-OverlayRouter 10)# anycast-gateway TOR2 (config-if-OverlayRouter 10)# exit TOR2 (config)# vxlan 10 TOR2 (config-vxlan)# extend-vlan 10 TOR2 (config-vxlan)# router-interface OverlayRouter 10 TOR2 (config-vxlan)#exit TOR2 (config)# int overlayrouter 20 TOR2 (config-if-OverlayRouter 20)# vrf forwarding vrf-10 TOR2 (config-if-OverlayRouter 20)# ip address 20.1.1.1/24 TOR2 (config-if-OverlayRouter 20)# anycast-gateway TOR2 (config-if-OverlayRouter 20)# exit TOR2 (config)# vxlan 20 TOR2 (config-vxlan)# extend-vlan 20 TOR2 (config-vxlan)# router-interface OverlayRouter 20 TOR2 (config-vxlan)#exit TOR2 (config)# int overlayrouter 100 TOR2 (config-if-OverlayRouter 100)# vrf forwarding vrf-10 TOR2 (config-if-OverlayRouter 100)# exit TOR2 (config)# vxlan 100 TOR2 (config-vxlan)# symmetric TOR2 (config-vxlan)# router-interface OverlayRouter 100 TOR2 (config-vxlan)#exit
Border1	Border1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. Border1 (config)# interface Loopback 1 Border1 (config-if- Loopback 1)# ip address 3.3.3.3/32 Border1 (config-if- Loopback 1)#exit Border1 (config)# vtep Border1 (config-vtep)# source loopback 1 Border1 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 Border1 (config-vtep)# arp suppress enable Border1 (config)# int overlayrouter 90 Border1 (config-if-OverlayRouter 90)# vrf forwarding vrf-10 Border1 (config-if-OverlayRouter 90)# ip address 90.1.1.1/24 Border1 (config-if-OverlayRouter 90)# anycast-gateway Border1 (config-if-OverlayRouter 90)# exit Border1 (config)# vxlan 90 Border1 (config-vxlan)# extend-vlan 90 Border1 (config-vxlan)# router-interface OverlayRouter 90

	Border1 (config-vxlan)#exit Border1 (config)# int overlayrouter 100 Border1 (config-if-OverlayRouter 100)# vrf forwarding vrf-10 Border1 (config-if-OverlayRouter 100)# ip address 100.1.1.1/24 Border1 (config-if-OverlayRouter 100)# exit Border1 (config)# vxlan 100 Border1 (config-vxlan)# symmetric Border1 (config-vxlan)# router-interface OverlayRouter 100 Border1 (config-vxlan)#exit
Border2	Border2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. Border2 (config)# interface Loopback 1 Border2 (config-if- Loopback 1)# ip address 4.4.4.4/32 Border2 (config-if- Loopback 1)#exit Border2 (config)# vtep Border2 (config-vtep)# source loopback 1 Border2 (config-vtep)# arp suppress enable Border2 (config)# fabric anycast-gateway-mac 0011.2233.2016 Border2 (config)# int overlayrouter 90 Border2 (config-if-OverlayRouter 90)# vrf forwarding vrf-10 Border2 (config-if-OverlayRouter 90)# ip address 90.1.2.1/24 Border2 (config-if-OverlayRouter 90)# anycast-gateway Border2 (config-if-OverlayRouter 90)# exit Border2 (config)# vxlan 90 Border2 (config-vxlan)# extend-vlan 90 Border2 (config-vxlan)# router-interface OverlayRouter 90 Border2 (config-vxlan)#exit Border2 (config)# int overlayrouter 100 Border2 (config-if-OverlayRouter 100)# vrf forwarding vrf-10 Border2 (config-if-OverlayRouter 100)# ip address 100.1.2.1/24 Border2 (config-if-OverlayRouter 100)# exit Border2 (config)# vxlan 100 Border2 (config-vxlan)# symmetric Border2 (config-vxlan)# router-interface OverlayRouter 100 Border2 (config-vxlan)#exit
Verification	● Verify that HOST-1, HOST-2, HOST-3, and HOST-4 can ping each other. ● Verify that the virtual machines can be migrated between the HOSTs on the same VXLAN and can access the network normally after migration without modifying the configuration.

```
Border1# sh vxlan
VXLAN Total Count: 3
VXLAN Capacity   : 8000

VXLAN 90
   Symmetric property  : FALSE
   Source Address     : 3.3.3.3
   Multicast Group    : -
   Router Interface   : overlayrouter 90 (non-anycast)
   Extend VLAN        : 90
   VTEP Adjacency Count: 1
   Interface            Source IP      Destination IP  Type
   ---------------------- --------------- --------------- -------
   OverlayTunnel 6146    3.3.3.3        2.2.2.2        dynamic

VXLAN 100
   Symmetric property  : TRUE
   Source Address     : 3.3.3.3
   Multicast Group    : -
   Router Interface   : overlayrouter 100 (non-anycast)
   Extend VLAN        : -
   VTEP Adjacency Count: 1
    Interface            Source IP      Destination IP  Type
   ---------------------- --------------- --------------- -------
   OverlayTunnel 6146    3.3.3.3        2.2.2.2        dynamic
```

1.4.3.5 Configuring EVPN-based Multi-tenant VXLAN Routing Scenario

Scenario	
Figure 1-28	

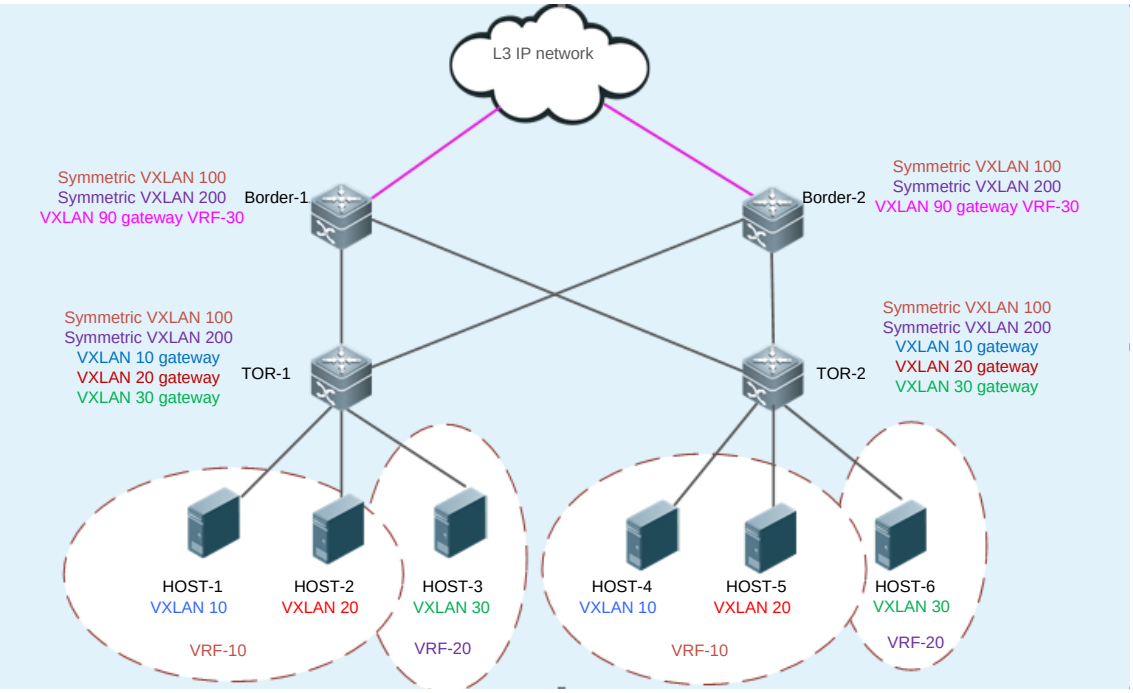

Configuration Steps

- Configure an IPv4 unicast routing protocol such as the OSPF protocol on Border-1, Border-2, TOR-1, and TOR-2 to ensure that unicast routes are reachable.
- Configure the BGP-EVPN routing protocol on Border-1, Border-2, TOR-1, and TOR-2 to establish BGP neighbor relationships between the devices (except between Border-1 and Border-2) and to support the EVPN protocol family.
- Configure the EVI for BGP-EVPN on TOR-1 and TOR-2. For details, see *BGP-EVPN Configuration Guide*.
- Configure a VXLAN on the virtual server and designate the gateway address of the virtual machine.
- Associate the VTEP with the loopback interface on TOR-1, TOR-2, Border-1, and Border-2 to establish tunnels.
- Configure the anycast gateway MAC address on TOR-1 and TOR-2 to ensure that all VXLAN anycast gateways on the network use the same MAC address.
- Create VXLAN 10, VXLAN 20, and VXLAN 30 on TOR-1 and associate them with VLANs.
- Create VXLAN 10, VXLAN 20, and VXLAN 30 on TOR-2 and associate them with VLANs.
- Create VXLAN 90 on Border-2 and associate VXLAN 90 with a VLAN.
- Create overlay router interfaces for VXLAN 10, VXLAN 20, and VXLAN 30 on TOR-1 and TOR-2 and configure the VXLAN gateway IP address for them. Configure different VRF networks for different overlay router interfaces to determine their respective tenants. Configure the anycast gateway to ensure that all VXLAN gateways on the network use the same IP address and MAC address. As the anycast gateway function is enabled, the overlay router interfaces associated with the same VXLAN on TOR-1 and TOR-2 must be configured with the same VXLAN gateway IP address.
- Create VXLAN 100 and VXLAN 200 on TOR-1, TOR-2, Border-1, and Border-2. Configure VXLAN 100 and VXLAN 200 as symmetric VXLANs to serve as L3 routing VXLANs of the corresponding VRF

	networks. L3 routes between all VXLANs of the same VRF network are advertised via the symmetric VXLANs. In addition, the symmetric VXLANs are used for L3 routing and forwarding. ● Create overlay router interfaces for VXLAN 100 and VXLAN 200 on TOR-1 and TOR-2 and configure different VRF networks for the overlay router interfaces. VXLAN 100 and VXLAN 200 serve as the symmetric VXLANs of the corresponding VRF networks. ● Create overlay router interfaces for VXLAN 100 and VXLAN 200 on Border-1 and Border-2. Configure different VRF networks for the overlay router interfaces so that VXLAN 100 and VXLAN 200 serve as the symmetric VXLANs of the corresponding VRF networks. Configure VXLAN gateway IP addresses for Border-1 and Border-2 (different IP addresses for different devices). ● Create overlay router interfaces for VXLAN 90 on Border-1 and Border-2. Configure different VRF networks for the OverlayRouter interfaces and configure the VXLAN gateway IP address. ● Enable the global import-route function on the Border-1 and Border2 devices, and configure the route-mapped VNI in the symmetric instance. ● Associate VXLAN instances with overlay router interfaces on TOR-1, TOR-2, Border-1, and Border-2 to realize VXLAN routing. ● (Optional) Configure ARP suppression on TOR-1 and TOR-2 to reduce the ARP packets entering the VXLAN.
HOST	Configuring the IP address and gateway according to Figure 1-28 (the detailed configuration on the server is omitted herein).
TOR1	TOR1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR1 (config)# interface Loopback 1 TOR1 (config-if- Loopback 1)# ip address 1.1.1.1/32 TOR1 (config-if- Loopback 1)#exit TOR1 (config)# vtep TOR1 (config-vtep)# source loopback 1 TOR1 (config-vtep)# arp suppress enable TOR1 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 TOR1 (config-vtep)#exit TOR1 (config)# int overlayrouter 10 TOR1 (config-if-OverlayRouter 10)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24 TOR1 (config-if-OverlayRouter 10)# anycast-gateway TOR1 (config-if-OverlayRouter 10)# exit TOR1 (config)# vxlan 10 TOR1 (config-vxlan)# extend-vlan 10 TOR1 (config-vxlan)# router-interface OverlayRouter 10 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 20 TOR1 (config-if-OverlayRouter 20)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 20)# ip address 20.1.1.1/24

	TOR1 (config-if-OverlayRouter 20)# anycast-gateway TOR1 (config-if-OverlayRouter 20)# exit TOR1 (config)# vxlan 20 TOR1 (config-vxlan)# extend-vlan 20 TOR1 (config-vxlan)# router-interface OverlayRouter 20 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 30 TOR1 (config-if-OverlayRouter 30)# vrf forwarding vrf-20 TOR1 (config-if-OverlayRouter 30)# ip address 30.1.1.1/24 TOR1 (config-if-OverlayRouter 30)# anycast-gateway TOR1 (config-if-OverlayRouter 30)# exit TOR1 (config)# vxlan 30 TOR1 (config-vxlan)# extend-vlan 30 TOR1 (config-vxlan)# router-interface OverlayRouter 30 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 100 TOR1 (config-if-OverlayRouter 100)# vrf forwarding vrf-10 TOR1 (config-if-OverlayRouter 100)# exit TOR1 (config)# vxlan 100 TOR1 (config-vxlan)# symmetric TOR1 (config-vxlan)# router-interface OverlayRouter 100 TOR1 (config-vxlan)#exit TOR1 (config)# int overlayrouter 200 TOR1 (config-if-OverlayRouter 200)# vrf forwarding vrf-20 TOR1 (config-if-OverlayRouter 200)# exit TOR1 (config)# vxlan 200 TOR1 (config-vxlan)# symmetric TOR1 (config-vxlan)# router-interface OverlayRouter 200 TOR1 (config-vxlan)#exit
TOR2	TOR2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. TOR2 (config)# interface Loopback 1 TOR2 (config-if- Loopback 1)# ip address 2.2.2.2/32 TOR2 (config-if- Loopback 1)#exit TOR2 (config)# vtep TOR2 (config-vtep)# source loopback 1 TOR2 (config-vtep)# arp suppress enable TOR2 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 TOR2 (config-vtep)#exit TOR2 (config)# int overlayrouter 10

```
TOR2 (config-if-OverlayRouter 10)# vrf forwarding vrf-10
TOR2 (config-if-OverlayRouter 10)# ip address 10.1.1.1/24
TOR2 (config-if-OverlayRouter 10)# anycast-gateway
TOR2 (config-if-OverlayRouter 10)# exit
TOR2 (config)# vxlan 10
TOR2 (config-vxlan)# extend-vlan 10
TOR2 (config-vxlan)# router-interface OverlayRouter 10
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 20
TOR2 (config-if-OverlayRouter 20)# vrf forwarding vrf-10
TOR2 (config-if-OverlayRouter 20)# ip address 20.1.1.1/24
TOR2 (config-if-OverlayRouter 20)# anycast-gateway
TOR2 (config-if-OverlayRouter 20)# exit
TOR2 (config)# vxlan 20
TOR2 (config-vxlan)# extend-vlan 20
TOR2 (config-vxlan)# router-interface OverlayRouter 20
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 30
TOR2 (config-if-OverlayRouter 30)# vrf forwarding vrf-20
TOR2 (config-if-OverlayRouter 30)# ip address 30.1.1.1/24
TOR2 (config-if-OverlayRouter 30)# anycast-gateway
TOR2 (config-if-OverlayRouter 30)# exit
TOR2 (config)# vxlan 30
TOR2 (config-vxlan)# extend-vlan 30
TOR2 (config-vxlan)# router-interface OverlayRouter 30
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 100
TOR2 (config-if-OverlayRouter 100)# vrf forwarding vrf-10
TOR2 (config-if-OverlayRouter 100)# exit
TOR2 (config)# vxlan 100
TOR2 (config-vxlan)# symmetric
TOR2 (config-vxlan)# router-interface OverlayRouter 100
TOR2 (config-vxlan)#exit
TOR2 (config)# int overlayrouter 200
TOR2 (config-if-OverlayRouter 200)# vrf forwarding vrf-20
TOR2 (config-if-OverlayRouter 200)# exit
TOR2 (config)# vxlan 200
TOR2 (config-vxlan)# symmetric
TOR2 (config-vxlan)# router-interface OverlayRouter 200
TOR2 (config-vxlan)#exit
```

Border1

```
Border1# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Border1 (config)# interface Loopback 1
Border1 (config-if- Loopback 1)# ip address 3.3.3.3/32
Border1 (config-if- Loopback 1)#exit
Border1 (config)# vtep
Border1 (config-vtep)# source loopback 1
Border1 (config-vtep)# import-route enable
Border1 (config-vtep)# arp suppress enable
Border1 (config-vtep )# fabric anycast-gateway-mac 0011.2233.2016
Border1 (config-vtep )# exit
Border1 (config)# int overlayrouter 90
Border1 (config-if-OverlayRouter 90)# vrf forwarding vrf-30
Border1 (config-if-OverlayRouter 90)# ip address 90.1.1.1/24
Border1 (config-if-OverlayRouter 90)# anycast-gateway
Border1 (config-if-OverlayRouter 90)# exit
Border1 (config)# vxlan 90
Border1 (config-vxlan)# extend-vlan 90
Border1 (config-vxlan)# router-interface OverlayRouter 90
Border1 (config-vxlan)#exit
Border1 (config)# int overlayrouter 100
Border1 (config-if-OverlayRouter 100)# vrf forwarding vrf-10
Border1 (config-if-OverlayRouter 100)# ip address 100.1.1.1/24
Border1 (config-if-OverlayRouter 100)# exit
Border1 (config)# vxlan 100
Border1 (config-vxlan)# symmetric
Border1 (config-vxlan)# router-interface OverlayRouter 100
Border1 (config-vxlan)# member add 100
Border1 (config-vxlan)# member add 10
Border1 (config-vxlan)# member add 20
Border1 (config-vxlan)#exit
Border1 (config)# int overlayrouter 200
Border1 (config-if-OverlayRouter 200)# vrf forwarding vrf-20
Border1 (config-if-OverlayRouter 200)# ip address 200.1.1.1/24
Border1 (config-if-OverlayRouter 200)# exit
Border1 (config)# vxlan 200
Border1 (config-vxlan)# symmetric
Border1 (config-vxlan)# router-interface OverlayRouter 200
Border1 (config-vxlan)# member add 200
Border1 (config-vxlan)# member add 30
Border1 (config-vxlan)#exit
```

Border2	Border2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. Border2 (config)# interface Loopback 1 Border2 (config-if- Loopback 1)# ip address 4.4.4.4/32 Border2 (config-if- Loopback 1)#exit Border2 (config)# vtep Border2 (config-vtep)# source loopback 1 Border2 (config-vtep)# import-route enable Border2 (config-vtep)# arp suppress enable Border2 (config-vtep)# fabric anycast-gateway-mac 0011.2233.2016 Border2 (config-vtep)# exit Border2 (config)# int overlayrouter 90 Border2 (config-if-OverlayRouter 90)# vrf forwarding vrf-30 Border2 (config-if-OverlayRouter 90)# ip address 90.1.2.1/24 Border2 (config-if-OverlayRouter 90)# anycast-gateway Border2 (config-if-OverlayRouter 90)# exit Border2 (config)# vxlan 90 Border2 (config-vxlan)# extend-vlan 90 Border2 (config-vxlan)# router-interface OverlayRouter 90 Border2 (config-vxlan)#exit Border2 (config)# int overlayrouter 100 Border2 (config-if-OverlayRouter 100)# vrf forwarding vrf-10 Border2 (config-if-OverlayRouter 100)# ip address 100.1.2.1/24 Border2 (config-if-OverlayRouter 100)# exit Border2 (config)# vxlan 100 Border2 (config-vxlan)# symmetric Border2 (config-vxlan)# router-interface OverlayRouter 100 Border2 (config-vxlan)# member add 100 Border2 (config-vxlan)# member add 10 Border2 (config-vxlan)# member add 20 Border2 (config-vxlan)#exit Border2 (config)# int overlayrouter 200 Border2 (config-if-OverlayRouter 200)# vrf forwarding vrf-20 Border2 (config-if-OverlayRouter 200)# ip address 200.1.2.1/24 Border2 (config-if-OverlayRouter 200)# exit Border2 (config)# vxlan 200 Border2 (config-vxlan)# symmetric Border2 (config-vxlan)# router-interface OverlayRouter 200 Border2 (config-vxlan)# member add 200 Border2 (config-vxlan)# member add 30

```
Border2 (config-vxlan)#exit
```

Verification

- Verify that HOST-1, HOST-2, and HOST-4 can ping each other.
- Verify that HOST-3 and HOST-6 can ping each other.
- Verify that HOST-1, HOST-2, and HOST-4 cannot ping HOST-3 and HOST-6.
- Verify that the virtual machines can be migrated between the HOSTs on the same VXLAN and can access the network normally after migration without modifying the configuration.

```
Border1# sh vxlan
VXLAN Total Count: 3
VXLAN Capacity   : 8000

VXLAN 90
   Symmetric property  : FALSE
   Source Address     : 3.3.3.3
   Multicast Group    : -
   Router Interface   : overlayrouter 90 (non-anycast)
   Extend VLAN       : 90
   VTEP Adjacency Count: 1
   Interface           Source IP      Destination IP  Type
   ---------------------- --------------- --------------- -------
   OverlayTunnel 6146    3.3.3.3       2.2.2.2       dynamic

VXLAN 100
   Symmetric property  : TRUE
   Source Address     : 3.3.3.3
   Multicast Group    : -
   Router Interface   : overlayrouter 100 (non-anycast)
   Extend VLAN       : -
   VTEP Adjacency Count: 1
    Interface           Source IP      Destination IP  Type
   ---------------------- --------------- --------------- -------
   OverlayTunnel 6146    3.3.3.3       2.2.2.2       dynamic

VXLAN 200
   Symmetric property  : TRUE
   Source Address     : 3.3.3.3
   Multicast Group    : -
   Router Interface   : overlayrouter 200 (non-anycast)
   Extend VLAN       : -
   VTEP Adjacency Count: 1
   VTEP Adjacency List :
```

```
Interface            Source IP      Destination IP  Type
-------------------- -------------- --------------- -------
OverlayTunnel 6146    3.3.3.3        2.2.2.2        dynamic
```

1.5 Monitoring

Displaying

Description	Command
Displays the VXLAN configuration and status of the device.	**show vxlan** *vni-number*
Displays the MAC addresses learned by the device.	**show vxlan mac [vni** *vni-number*] **[address** *mac-address*]
Displays the routing entries learned by the device.	**show vxlan route [remote | local] [vni** *vni-number*] **[vrf** *vrf-id*]
Displayes the effective routing entries in VXLAN network.	**show vxlan prefix-route [remote | local] [vni** *vni-number*] **[vrf** *vrf-id*]
Displays the VXLAN ARP entries learned by the device.	**show vxlan arp table [vni** *vni-number*]
Displays the global configurations of the device, such as the VTEP IP address and anycast MAC address.	**show vxlan global**
Displays the ARP suppression status of the device.	**show vxlan arp suppress**
Displays the number of MAC addresses learned by the device.	**show vxlan mac count**
Displays the VXLAN UDP destination port of the device.	**show vxlan udp-port**