18 VXLAN Commands

18.1 VXLAN Configuration Commands

18.1 VXLAN Configuration Commands

18.1.1 Command Support

Only the S6730-H, S6730S-H, S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S5731-H, S6720S-EI, and S6720-EI support VXLAN.

18.1.2 advertise l2vpn evpn

Function

The **advertise l2vpn evpn** command enables a device to advertise IP routes imported to a VPN instance to the BGP-EVPN address family.

The **undo advertise l2vpn evpn** command restores the default configuration.

By default, a device is disabled from advertising IP routes imported to a VPN instance to the BGP-EVPN address family.

Format

advertise l2vpn evpn

undo advertise l2vpn evpn

Parameters

None

Views

BGP-VPN instance IPv4 address family view or BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, when you need to establish VTEP tunnels through IP prefix routes and advertise host routes, run the **advertise l2vpn evpn** command to enable a device to advertise IP routes imported to a VPN instance to the BGP-EVPN address family. In this way, the routes imported to the VPN instance can be sent to the EVPN address family and then sent to the remote VTEP through the EVPN peer relationship.

Example

# Enable a device to advertise IP routes imported to VPN instance **vpna** to the BGP-EVPN address family.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family vpn-instance vpna
[HUAWEI-bgp-vpna] advertise l2vpn evpn
```

18.1.3 arp broadcast disable (VBDIF interface view)

Function

The **arp broadcast disable** command disables the ARP broadcast function on a VBDIF interface.

The **undo arp broadcast disable** command enables the ARP broadcast function on a VBDIF interface.

By default, ARP broadcast is enabled on a VBDIF interface.

Format

arp broadcast disable

undo arp broadcast disable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, a VBDIF interface broadcasts ARP packets in a BD. However, in the large Layer 2 aggregation scenario shown in **Figure 18-1**, hosts on the user side are connected to aggregation switch 1 through switches 2-100, and the VBDIF interface is created on switch 1 as the user gateway, allowing hosts on the user side to communicate with external devices. The default ARP broadcast packet processing of the VBDIF interface can cause ARP packet flooding on the user side, which occupies large amounts of resources on the user side and affects normal user services. In addition, when a large number of ARP packets are broadcasted, the performance of switch 1 is affected.

To ensure normal user services and aggregation gateway performance in the large Layer 2 aggregation scenario, you can run this command to disable ARP broadcast on the VBDIF interface of an aggregation gateway.

Figure 18-1 Networking in the Layer 2 aggregation scenario

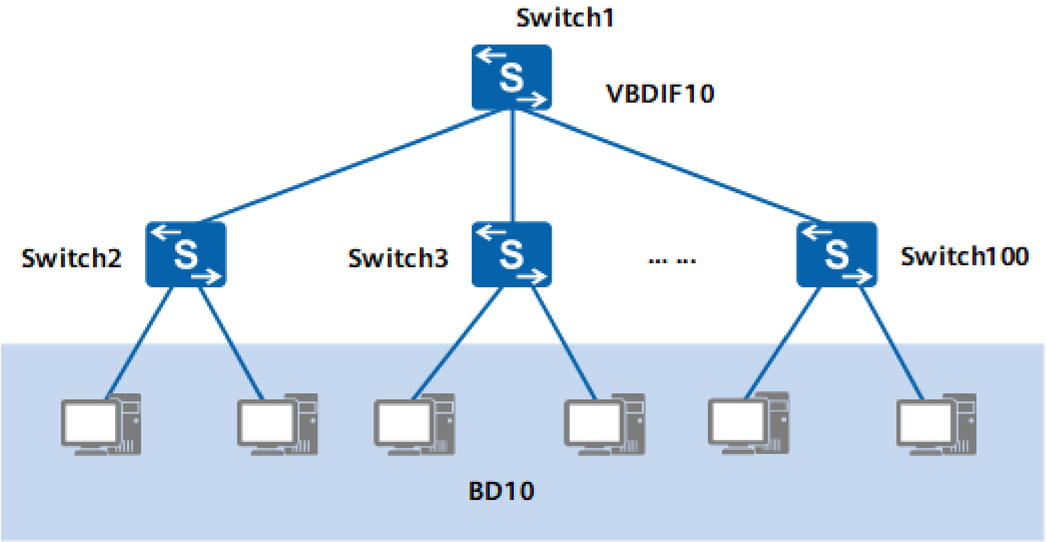

Precautions

Disabling a VBDIF interface from broadcasting ARP packets has impacts on the following scenarios: (Exercise caution before you run this command.)

- ARP proxy scenario. After a VBDIF interface is disabled from broadcasting ARP packets, the proxy does not forward ARP request packets from a host to their destinations even if all proxy conditions are met. As a result, ARP proxy fails.
- Scenarios in which hosts send unicast packets. For example, in **ping** operations, ICMP request packets must be encapsulated with MAC addresses mapped to the destination IP addresses. If a host does not have ARP entries, it has to send ARP request packets to learn the MAC address mapped to the destination IP address. However, the VBDIF interface is disabled from broadcasting ARP packets, and therefore cannot send ARP request packets. Consequently, the host cannot obtain the MAC address mapped to the destination IP address, causing a **ping** operation failure. This problem also occurs in other scenarios in which hosts send unicast packets.
- Strict ARP learning scenarios. In a strict ARP learning scenario, a device learns MAC addresses of only ARP reply packets in response to ARP request packets that it sends. If the VBDIF interface is disabled from broadcasting ARP

packets, it cannot actively send ARP request packets. As a result, strict ARP learning fails.

Example

# Disable the ARP broadcast function on VBDIF interface 10.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] arp broadcast disable
```

18.1.4 arp broadcast-suppress enable

Function

The **arp broadcast-suppress enable** command enables ARP broadcast suppression in a BD.

The **undo arp broadcast-suppress enable** command disables ARP broadcast suppression in a BD.

By default, ARP broadcast suppression is disabled in a BD.

Format

arp broadcast-suppress [**mismatch-discard**] **enable**

undo arp broadcast-suppress [**mismatch-discard**] **enable**

Parameters

Parameter	Description	Value
mismatch-discard	Indicates that the device drops ARP request packets that do not match any entries in the ARP broadcast suppression table.	-

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If a gateway receives a large number of ARP request packets within a short period and broadcasts the ARP request packets in a BD, excessive ARP request packets are forwarded. As a result, excessive network resources are used, traffic congestion may occur, and network performance may decline. ARP broadcast suppression can effectively ease the pressure on the gateway in handling ARP packets. After receiving an ARP request packet, the gateway searches the ARP broadcast suppression table that contains the mapping between the IP address and MAC address of each destination device.

- If a matching entry is found, the gateway replaces the broadcast MAC address in the received ARP request packet with the MAC address of the destination device, and then sends the request packet out through the interface matching the destination MAC address.
- If no matching entry is found:
  - The gateway broadcasts the ARP request packet in the BD if the **mismatch-discard** parameter is not set in the **arp broadcast-suppress enable** command.
  - The gateway drops the ARP request packet if the **mismatch-discard** parameter is set in the **arp broadcast-suppress enable** command.

Precautions

- When ARP broadcast suppression is configured in a VXLAN scenario and a large number of users exist, the gateway receives too many ARP packets and the packet rate exceeds the default CAR value. In this case, you can run the **car packet-type** { *vpls-arp* | *arp-request* } **cir** *cir-value* [**cbs** *cbs-value*] command to adjust the CAR value for ARP request packets. In this scenario, the CAR values specified using the arp-request and vpls-arp parameters share the same CAR resources. If the corresponding VBDIF interface is created on the device and the interface uses an IPv4 address and is Up, specify the arp-request parameter to adjust the CAR value. In other cases, specify the vpls-arp parameter to adjust the CAR value. You can run the **display cpu-defend configuration** command to view the CAR value for packets. If the CAR value is adjusted improperly, network services are affected. To adjust the CAR value for packets, contact technical support personnel.
- If the **arp broadcast-suppress enable** command is configured to enable ARP broadcast suppression in a BD and the **broadcast-suppression** command is configured to enable broadcast suppression in the BD, ARP broadcast suppression in the BD has a higher priority.

Example

# Enable ARP broadcast suppression in BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] arp broadcast-suppress enable
```

# Enable ARP broadcast suppression in BD 20 and configure the switch to discard ARP request packets if no matching entry is found in the ARP broadcast suppression table.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 20
[HUAWEI-bd20] arp broadcast-suppress mismatch-discard enable
```

18.1.5 arp collect host enable

Function

The **arp collect host enable** enables BGP EVPN to collect host information.

The **undo arp collect host enable** command disables BGP EVPN from collecting host information.

By default, BGP EVPN is disabled from collecting host information.

Format

arp collect host enable

undo arp collect host enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

When tenants communicate for the first time, they broadcast ARP request packets to each other. The ARP request packets are broadcasted in the Layer 2 network. To prevent broadcast storms caused by broadcasted ARP request packets, you can enable ARP broadcast suppression on the VXLAN Layer 2 gateway. However, ARP broadcast suppression relies on the host information table (containing the host IP address, MAC address, VTEP address, and VNI ID) on a Layer 3 gateway.

To allow a Layer 2 gateway to obtain the host information table, run this command in the VBDIF interface view to enable BGP EVPN to collect host information.

In VXLAN (BGP EVPN) scenarios, when VXLAN gateways advertise IRB routes to each other, run the **arp collect host enable** command for host information collection.

Example

# Enable BGP EVPN on VBDIF interface 10 to collect host information.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] arp collect host enable
```

18.1.6 arp distribute-gateway enable

Function

The **arp distribute-gateway enable** command enables the distributed gateway function.

The **undo arp distribute-gateway enable** command disables the distributed gateway function.

By default, the distributed gateway function is disabled.

Format

arp distribute-gateway enable

undo arp distribute-gateway enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To configure a gateway to function as a distributed gateway and learn only the ARP packets sent by a user-side host, run the **arp distribute-gateway enable** command to enable the distributed gateway function. After this function is enabled:

- The gateway processes only ARP packets sent by user-side hosts.
- The gateway deletes the network-side ARP entries that have been learned.

Precautions

After the distributed gateway function is enabled:

- Static ARP entries on the tunnel side fail to be configured on the gateway.
- If multiple gateways have the same IP address in the distributed scenario, this gateway does not report an ARP conflict.

Example

# Enable the distributed gateway function on VBDIF interface 10.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] arp distribute-gateway enable
```

18.1.7 arp static bridge-domain

Function

The **arp static bridge-domain** command configures a static ARP entry on an interface of a VXLAN network.

The **undo arp static bridge-domain** command deletes a static ARP entry configured on an interface of a VXLAN network.

By default, no static ARP entry is configured on an interface of a VXLAN network.

Format

arp static *ip-address mac-address* **bridge-domain** *bd-id* [**vid** *vlan-id1* [**cevid** *vlan-id2*]] **interface** *interface-type interface-number.subnum*

undo arp static *ip-address mac-address* **bridge-domain** *bd-id* [**vid** *vlan-id1* [**cevid** *vlan-id2*]] **interface** *interface-type interface-number.subnum*

arp static *ip-address mac-address* **bridge-domain** *bd-id* [**vid** *vlan-id3*] **interface** *interface-type interface-number*

undo arp static *ip-address mac-address* **bridge-domain** *bd-id* [**vid** *vlan-id3*] **interface** *interface-type interface-number*

Parameters

Parameter	Description	Value
ip-address	Specifies a destination IP address.	The value is in dotted decimal notation.
mac-address	Specifies the destination MAC address mapping the destination IP address.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits.
bd-id	Specifies a BD ID.	The value is an integer that ranges from 1 to 16777215.
vid *vlan-id1*	Specifies the outer VLAN ID in the packet received by a sub-interface.	The value is an integer that ranges from 1 to 4094.
cevid *vlan-id2*	Specifies the inner VLAN ID in the packet received by a sub-interface.	The value is an integer that ranges from 1 to 4094.
interface *interface-type interface-number.subnum*	Specifies a sub-interface.	-
vid *vlan-id3*	Specifies the VLAN ID in the packet received by a interface.	The value is an integer that ranges from 1 to 4094.

Parameter	Description	Value
interface *interface-type interface-number*	Specifies an interface.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Static ARP entries are manually configured and maintained. They will not be aged out or overridden by dynamic ARP entries. Therefore, you can run the **arp static bridge-domain** command on an interface of a VXLAN network to configure static ARP entries to increase communication security. Static ARP entries enable the local device and a specified device to communicate with each other using only specified MAC addresses. Attackers cannot modify mappings between IP addresses and MAC addresses in static ARP entries.

Prerequisites

The outbound interface has been added to a VLAN and bound to a BD.

Precautions

- If a static ARP entry already exists, the new configuration cannot be delivered.
- The specified *ip-address* must be in the same network segment as the outbound interface address in the ARP entry.
- To specify the **vid** *vlan-id* and **cevid** *vlan-id* parameters, set the same encapsulation type as that on the interface first.
- When you configure a static ARP entry on an interface of the S6720-EI, S6735-S, and S6720S-EI, you must configure a static MAC address entry for the MAC address in the ARP entry. Otherwise, the switch will broadcast traffic from this MAC address.

Example

# On the outbound interface GE0/0/1, configure a static ARP entry with the IP address and MAC address 10.1.1.2 and aaaa-fccc-1212, respectively.

```
<HUAWEI> system-view
[HUAWEI] vlan 10
[HUAWEI-vlan10] quit
[HUAWEI] interface GigabitEthernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] port link-type trunk
[HUAWEI-GigabitEthernet0/0/1] port trunk allow-pass vlan 10
[HUAWEI-GigabitEthernet0/0/1] quit
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] l2 binding vlan 10
```

```
[HUAWEI-bd10] quit
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ip address 10.1.1.1 255.255.255.0
[HUAWEI-Vbdif10] quit
[HUAWEI] arp static 10.1.1.2 aaaa-fccc-1212 bridge-domain 10 vid 10 interface GigabitEthernet 0/0/1
```

18.1.8 arp static vni

Function

The **arp static vni** command configures a static ARP entry for a VXLAN tunnel.

The **undo arp static vni** command deletes a static ARP entry of a VXLAN tunnel.

By default, no static ARP entry is configured for a VXLAN tunnel.

Format

arp static *ip-address mac-address* **vni** *vni-id* { **source-ip** *ip-address1* **peer-ip** *ip-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* }

undo arp static *ip-address mac-address* **vni** *vni-id* { **source-ip** *ip-address1* **peer-ip** *ip-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* }

NOTE

Only the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, and S5731-H support the **source-ipv6** and **peer-ipv6** parameters.

Parameters

Parameter	Description	Value
ip-address	Specifies a destination IP address.	The value is in dotted decimal notation.
mac-address	Specifies the destination MAC address mapping the destination IP address.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits.
vni-id	Specifies the VNI ID.	The value is an integer that ranges from 1 to 16777215.
source-ip *ip-address1*	Specifies the IP address of the source VTEP.	The value is in dotted decimal notation.
peer-ip *ip-address2*	Specifies the IP address of the destination VTEP.	The value is in dotted decimal notation.

Parameter	Description	Value
source-ipv6 *ipv6-address1*	Specifies the IPv6 address of the source VTEP.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.
peer-ipv6 *ipv6-address2*	Specifies the IPv6 address of the destination VTEP.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Static ARP entries are manually configured and maintained. They will not be aged out or overridden by dynamic ARP entries. Running the **arp static vni** command on a device to configure static ARP entries for a VXLAN tunnel increases communication security. Static ARP entries enable the local device and a specified device to communicate with each other using only specified MAC addresses. Attackers cannot modify mappings between IP addresses and MAC addresses in static ARP entries.

Prerequisites

A VXLAN tunnel and a Layer 3 gateway have been configured.

Precautions

- If a static ARP entry already exists, the new configuration cannot be delivered.
- The specified IP address must be in the same network segment as the outbound interface address in the ARP entry.
- When the VXLAN tunnel is created dynamically, the device does not support to configure a static ARP entry on a VXLAN tunnel-side interface.

Example

# Configure a static ARP entry for a VXLAN tunnel that maps the IP address 10.0.0.2 to the MAC address aaaa-fccc-1212.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] vxlan vni 5000
```

```
[HUAWEI-bd10] quit
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ip address 10.0.0.10 255.255.255.0
[HUAWEI-Vbdif10] quit
[HUAWEI] interface nve 1
[HUAWEI-Nve1] source 10.1.1.1
[HUAWEI-Nve1] vni 5000 head-end peer-list 10.2.2.2
[HUAWEI-Nve1] quit
[HUAWEI] arp static 10.0.0.2 aaaa-fccc-1212 vni 5000 source-ip 10.1.1.1 peer-ip 10.2.2.2
```

18.1.9 arp-proxy local enable

Function

The **arp-proxy local enable** command enables local ARP proxy, realizing interconnection between isolated hosts in a BD.

The **undo arp-proxy local enable** command disables local ARP proxy.

By default, local ARP proxy is disabled.

Format

arp-proxy local enable

undo arp-proxy local enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

On a VXLAN, a BD is a broadcast domain. The member interfaces in a BD broadcast the broadcast, unknown unicast, and multicast (BUM) packets within the BD. To minimize broadcast traffic, network administrators usually run the **isolate enable** command or configure port isolation on the access side to enable access-side isolation. In this way, the access users of the BD are isolated and do not support Layer 2 interconnection. However, with the increase of user services, users have a growing demand for interconnection. To satisfy the demand, network administrators can enable local ARP proxy on a VBDIF interface so that isolated users in a BD can communicate with each other.

Example

# Enable local ARP proxy on VBDIF interface 10.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] arp-proxy local enable
```

18.1.10 bridge-domain (Layer 2 sub-interface view)

Function

The **bridge-domain** command associates a Layer 2 sub-interface with a BD.

The **undo bridge-domain** command restores the default settings.

By default, no Layer 2 sub-interface is associated with a BD.

Format

bridge-domain *bd-id*

undo bridge-domain [*bd-id*]

Parameters

Parameter	Description	Value
bd-id	Specifies the ID of the BD that is associated with a Layer 2 sub-interface.	The value is an integer that ranges from 1 to 16777215.

Views

Layer 2 sub-interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

VXLAN needs to be deployed on a downlink interface to provide access services and an uplink interface to establish a VXLAN tunnel.

On the access side, two methods are available for creating a large Layer 2 BD.

- Based on VLAN: You can associate one or multiple VLANs with a BD to add users in these VLANs to the BD. This VLAN-based mode implements larger-granularity control, but is easy to configure. It applies to VXLAN deployment on a live network.
- Based on encapsulation mode: The device sends packets of different encapsulation modes to different Layer 2 sub-interfaces based on the VLAN tags contained in the packets. You can bind a Layer 2 sub-interface to a BD to add specified users to the BD. This mode implements refined and flexible control but requires more complex configuration. It applies to VXLAN deployment on a new network.

To create a BD based on encapsulation mode, create a Layer 2 sub-interface first. Then run the **encapsulation (Layer 2 sub-interface view)** command to configure a supported encapsulation mode on the sub-interface. After you run the **bridge-domain (Layer 2 sub-interface view)** command to associate a Layer 2 sub-interface with a BD, packets containing the same VLAN tag from different LANs can communicate at Layer 2.

Prerequisites

- Run the command **bridge-domain** to create the BD.
- Run the command **interface** to create the Layer 2 VXLAN sub-interface.

Precautions

One Layer 2 sub-interface can be associated with only one BD.

For the BD that is associated with the Layer 2 sub-interface using **default** encapsulation, the VBDIF interface of this BD cannot be created on the device.

Example

# Associate Layer 2 sub-interface GE0/0/1.1 with BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] quit
[HUAWEI] interface GigabitEthernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] port link-type hybrid
[HUAWEI-GigabitEthernet0/0/1] quit
[HUAWEI] interface gigabitethernet 0/0/1.1 mode l2
[HUAWEI-GigabitEthernet0/0/1.1] bridge-domain 10
```

18.1.11 bridge-domain (system view)

Function

The **bridge-domain** command creates a bridge domain (BD) and displays the BD view, or directly displays the view of an existing BD view.

The **undo bridge-domain** command deletes a BD.

By default, no BD is created.

Format

bridge-domain *bd-id*

undo bridge-domain *bd-id*

Parameters

Parameter	Description	Value
bd-id	Specifies the ID of a BD.	The value is an integer that ranges from 1 to 16777215.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

BDs are Layer 2 bridge domains on a large Layer 2 network constructed using VXLAN. VXLAN packets can be forwarded at Layer 2 within a BD through a VXLAN tunnel.

After you run the **bridge-domain** command to create a BD, you can complete other VXLAN configurations in the BD.

Example

# Create BD 10 and enter the view of BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10]
```

18.1.12 broadcast-suppression (BD view)

Function

The **broadcast-suppression** command enables broadcast traffic suppression in a bridge domain (BD).

The **undo broadcast-suppression** command disables broadcast traffic suppression in a BD.

By default, broadcast traffic suppression is disabled in a BD.

Format

broadcast-suppression cir *cir-value* [**cbs** *cbs-value*]

undo broadcast-suppression

Parameters

Parameter	Description	Value
cir *cir-value*	Specifies the committed information rate (CIR), which is the allowed rate at which traffic can pass through.	The value is an integer that ranges from 0 to 10000000, in kbit/s.

Parameter	Description	Value
cbs *cbs-value*	Specifies the committed burst size (CBS), which is the maximum size of traffic that can pass through.	The value is an integer that ranges from 10000 to 4294967295, in bytes. If the **cbs** is not set, the default *cbs-value* is 125 times the *cir-value*.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

When a large number of broadcast packets are transmitted on a network, a lot of network resources are occupied, and services on the network are affected. You can run the **broadcast-suppression** command to enable broadcast traffic suppression in a BD and configure the maximum number of broadcast packets that can pass through a BD. When the broadcast traffic volume exceeds the specified threshold, the system discards excess broadcast packets.

Example

# Set the CIR value for broadcast traffic in BD 10 to 100 kbit/s.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] broadcast-suppression cir 100
```

18.1.13 description (EVPN instance view)

Function

The **description** command specifies the description of the current EVPN instance.

The **undo description** command deletes the description of the current EVPN instance.

By default, no description is specified for an EVPN instance.

Format

description *description-information*

undo description

Parameters

Parameter	Description	Value
description-information	Specifies the description of an EVPN instance.	The value is a string of 1 to 242 case-sensitive characters with spaces supported.

Views

EVPN instance view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To record the purpose of creating an EVPN instance and the CEs with which the EVPN instance is associated, you can run the **description** command to specify the description of the EVPN instance.

To check the description of an EVPN instance, run the **display evpn vpn-instance** command.

Precautions

If you run the **description** command several times, the latest configuration overrides the previous configurations.

Example

# Specify the description of an EVPN instance named evn10.

```
<HUAWEI> system-view
[HUAWEI] evpn vpn-instance evn10 bd-mode
[HUAWEI-evpn-instance-evn10] description OnlyForAB
```

18.1.14 description (BD view)

Function

The **description** command configures the description of a bridge domain (BD).

The **undo description** command deletes the description of a BD.

By default, no description is configured for a BD.

Format

description *description*

undo description

Parameters

Parameter	Description	Value
description	Specifies the BD description.	The value is a string of 1 to 80 case-sensitive characters without question marks.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

If you have configured multiple BDs using the **bridge-domain (system view)** command, run the **description** command in the corresponding BD view to configure the description for each BD. BD description helps you quickly understand the function of each BD, facilitating service management.

Example

# Configure description **vxlan** for BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] description vxlan
```

18.1.15 display arp broadcast-suppress user bridge-domain

Function

The **display arp broadcast-suppress user bridge-domain** command displays the ARP broadcast suppression table of a BD.

Format

display arp broadcast-suppress user bridge-domain { *bd-id* [*ip-address*] | **all** }

display arp broadcast-suppress user statistics { **bridge-domain** *bd-id* | **all** }

Parameters

Parameter	Description	Value
bd-id	Specifies the ID of the BD of the ARP broadcast suppression table to be queried.	The value is an integer that ranges from 1 to 16777215.

Parameter	Description	Value
ip-address	Specifies the IP address of the ARP broadcast suppression table to be queried.	The value is in dotted decimal notation.
statistics	Displays ARP entry statistics.	-
all	Displays ARP entry in all BDs.	-

Views

All views

Default Level

2: Configuration level

Usage Guidelines

ARP broadcast suppression can effectively ease the pressure on the gateway in handling ARP packets. When receiving an ARP request packet, the gateway searches for the host information about the device corresponding to the destination IP address (the IP address and MAC address mapping table of the destination device, also known as the ARP broadcast suppression table).

To display the ARP broadcast suppression table of a BD, run this command.

Example

# Display the ARP broadcast suppression table of BD 10.

```
<HUAWEI> display arp broadcast-suppress user bridge-domain 10
-------------------------------------------------------------------------------
Flag     IP Address       MAC Address       VNI        VTEP Address
-------------------------------------------------------------------------------
local    10.1.1.11        00e0-fc00-0022    10         10.4.1.1
remote   10.1.1.12        00e0-fc98-15db    10         10.3.3.3
-------------------------------------------------------------------------------
Total: 2
```

Table 18-1 Description of the **display arp broadcast-suppress user bridge-domain** command output

Item	Description
Flag	ARP broadcast suppression entry type ● local: ARP entry of the local access device ● remote: ARP entry of the remote access device

Item	Description
IP Address	IP address of an ARP broadcast suppression entry
MAC Address	MAC address of an ARP broadcast suppression entry
VNI	Layer 2 VNI that an ARP broadcast suppression entry belongs
VTEP Address	IP address of the source VTEP of an ARP broadcast suppression entry
Total	Total number of ARP broadcast suppression entries in a BD

# Display statistics on the ARP broadcast suppression table of BD 10.

```
<HUAWEI> display arp broadcast-suppress user statistics bridge-domain 10
Total: 2     Local: 1         Remote: 1
```

Table 18-2 Description of the **display arp broadcast-suppress user bridge-domain** command output

Item	Description
Total	Total number of ARP broadcast suppression entries
Local	Total number of ARP broadcast suppression entries learned from the local end
Remote	Total number of ARP broadcast suppression entries learned from the remote end

18.1.16 display bgp evpn group

Function

The **display bgp evpn group** command displays information about BGP EVPN peer groups.

Format

display bgp evpn group [*group-name*]

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Configuring BGP peer groups simplifies BGP network configuration and improves the route advertisement efficiency.

The **display bgp group** command displays the information about the peer group, including information about the peers in the peer group and configuration information about the peer group. This command is used in the following scenarios:

- Verify the configuration after running the **group** command to configure a peer group.
- Verify the configuration after running the **peer group** command to add a peer to the peer group.
- Verify the configuration after running the **undo peer group** command to delete a peer from a peer group.
- Verify the configuration after modifying the configuration of a peer group.

Example

# Display the information about EVPN peer group **eg1**.

```
<HUAWEI> display bgp evpn group eg1

 BGP peer-group: eg1
 Remote AS: 100
 Authentication type configured: None
 Type : internal
 Configured hold timer value: 180
 Keepalive timer value: 60
 Connect-retry timer value: 32
 Minimum route advertisement interval is 0 seconds
 Connect-interface has been configured
 Tracking has been enabled, and the delay is 9s
 PeerSession Members:
   4.4.4.4          5.5.5.5

 Peer Preferred Value: 0
 No routing policy is configured
```

```
Peer Members:
 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

 4.4.4.4         4         100        4        7     0 00:00:39 Established       2
 5.5.5.5         4         100        4        7     0 00:00:31 Established       2
```

Table 18-3 Description of the **display bgp evpn group** command output

Item	Description
BGP peer-group	Name of a BGP peer-group
Remote AS	Number of the AS where a peer group resides
Authentication type configured	BGP authentication type configured. The value can be the following: ● MD5 ● Keychain (kk), in which kk indicates the name of the configured keychain authentication ● None: no BGP authentication
Type	Type of a peer group: ● internal: The peer group is an IBGP peer group. ● external: The peer group is an EBGP peer group.
Configured hold timer value	Value of the Hold timer, in the unit of seconds
Keepalive timer value	Value of the Keepalive timer, in the unit of seconds
Connect-retry timer value	Value of the Connect-retry timer, in the unit of seconds
Minimum route advertisement interval	Minimum interval between route advertisements
PeerSession Members	Peers that set up sessions
Peer Preferred Value	Preferred value of a peer
Peer Members	Indicating that the following information is about peers
Peer	IP address of a peer
V	BGP version
As	Number of the AS where a member of a peer group resides
MsgRcvd	Number of received messages

Item	Description
MsgSent	Number of sent messages
OutQ	Number of messages to be sent to peers
Up/Down	Period of time during which a BGP session keeps the current state

Item	Description
State	BGP state mechanism: - Idle: indicates that BGP denies any request of entering. This is the initiatory status of BGP. Upon receiving a Start event, BGP initiates a TCP connection to the remote BGP peer, starts the ConnectRetry Timer with the initial value, detects a TCP connection initiated by the remote BGP peer, and changes its state to Connect. - Idle(Admin): indicates that the peer relationship is shut down inactively and no attempt is made to establish the neighbor relationship. If the **peer ignore** command is configured or the peer is set to the Down state through the MIB, the neighbor is in the Idle (Admin) state. - Idle(Ovlmt): indicates that the peer relationship is interrupted because the number of routes exceeds the upper threshold. After a BGP peer relationship is interrupted due to the running of the **peer route-limit** command, the status of the BGP peer relationship is displayed as Idle(Ovlmt). If the **reset bgp** command is not run, the BGP peer relationship will not be reestablished. - Connect: indicates that BGP waits for the TCP connection to be set up before it determines whether to perform other operations. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer

Item	Description
	with the initial value, initiates a TCP connection to the remote BGP peer, and stays in the Connect state. ● Active: indicates that BGP tries to set up a TCP connection. This is the intermediate status of BGP. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value and changes its state to Connect. – If BGP initiates a TCP connection with an unknown IP address, the TCP connection fails. When this occurs, BGP restarts the ConnectRetry Timer with the initial value and stays in the Active state. ● OpenSent: indicates that BGP has sent one Open message to its peer and waits for the other Open message from the peer. – If there are no errors in the Open message received, BGP changes its state to OpenConfirm. – If there are errors in the Open message received, BGP sends a Notification message to the remote peer and changes its state to Idle. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. ● OpenConfirm: indicates that BGP waits for a Notification message or a Keepalive message. – If BGP receives a Notification message, or the TCP connection fails, BGP changes its state to Idle. – If BGP receives a Keepalive message, BGP changes its state to Established.

Item	Description
	● Established: indicates that BGP peers can exchange Update, Notification and Keepalive packets. – If BGP receives an Update or a Keepalive message, its state stays in Established. – If BGP receives a Notification message, BGP changes its state to Idle.
PrefRcv	Indicates the number of route prefixes received by the local peer from the remote peer.

18.1.17 display bgp evpn peer

Function

The **display bgp evpn peer** command displays information about BGP EVPN peers.

Format

display bgp evpn peer [*ipv4-address* **verbose** | **verbose**]

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a peer to be displayed.	It is in dotted decimal notation.
verbose	Indicates to display detailed peer information.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

To query the following information about BGP EVPN peers, run the **display bgp evpn peer** command:

- Status of BGP EVPN connections
- Information about a BGP EVPN peer
- Whether BGP peers are successfully configured using the **peer enable** command
- Whether BGP peers are successfully deleted using the **undo peer enable** command

Precautions

By specifying **verbose**, you can query more BGP peer information, such as the BGP timer information, number of received and transmitted routes, capabilities supported by the peer, number of received and transmitted BGP information messages, and enabled configurations.

Example

# Display the information about BGP EVPN peers.

```
<HUAWEI> display bgp evpn peer

 BGP local router ID : 6.6.6.6
 Local AS number : 100
 Total number of peers : 2              Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  4.4.4.4         4         100        6        9     0 00:02:19 Established       2
  5.5.5.5         4         100        6        9     0 00:02:11 Established       2
```

Table 18-4 Description of the **display bgp evpn peer** command output

Item	Description
BGP Local router ID	Indicates the ID of the BGP local router. **NOTE** If two ends have the same BGP local router ID, no BGP peer relationship can be established between them. In this situation, run the **router id** command in the BGP view on either end to change the BGP local router ID. Changing it to the IP address of a loopback interface is recommended.
local AS number	Indicates the local AS number.
Total number of peers	Indicates the number of the peer.
Peers in established state	Indicates the number of the peer in established state.
Peer	Indicates the IP address of the peer.
V	Indicates the BGP version.
AS	Indicates the number of the AS where a member of a peer group resides.

Item	Description
MsgRcvd	Indicates the number of received messages.
MsgSent	Indicates the number of sent messages.
OutQ	Indicates the number of messages to be sent to peers.
Up/Down	Indicates the period of time during which a BGP session keeps the current state.

Item	Description
State	Indicates the BGP state mechanism: ● Idle: indicates that BGP denies any request of entering. This is the initiatory status of BGP. Upon receiving a Start event, BGP initiates a TCP connection to the remote BGP peer, starts the ConnectRetry Timer with the initial value, detects a TCP connection initiated by the remote BGP peer, and changes its state to Connect. ● Idle(Admin): indicates that the neighbor relationship is shut down initiatively and no attempt is made to establish the neighbor relationship. If the **peer ignore** command is configured or the peer is set to the Down state through the MIB, the neighbor is in the Idle (Admin) state. ● No neg: The negotiation of the capabilities of the BGP peer's address family is not successful. ● Idle(Ovlmt): indicates that the neighbor relationship is interrupted because the number of routes exceeds the upper threshold. After a BGP neighbor relationship is interrupted due to the running of the **peer route-limit** command, the status of the BGP neighbor relationship is displayed as Idle(Ovlmt). If the **reset bgp** command is not run, the BGP neighbor relationship will not be reestablished. ● Connect: indicates that BGP waits for the TCP connection to be set up before it determines whether to perform other operations. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value, initiates a TCP connection to the remote BGP peer, and stays in the Connect state. ● Active: indicates that BGP tries to set up a TCP connection. This is the intermediate status of BGP. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer

Item	Description
	with the initial value and changes its state to Connect. – If BGP initiates a TCP connection with an unknown IP address, the TCP connection fails. When this occurs, BGP restarts the ConnectRetry Timer with the initial value and stays in the Active state. ● OpenSent: indicates that BGP has sent one Open message to its peer and waits for the other Open message from the peer. – If there are no errors in the Open message received, BGP changes its state to OpenConfirm. – If there are errors in the Open message received, BGP sends a Notification message to the remote peer and changes its state to Idle. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. ● OpenConfirm: indicates that BGP waits for a Notification message or a Keepalive message. – If BGP receives a Notification message, or the TCP connection fails, BGP changes its state to Idle. – If BGP receives a Keepalive message, BGP changes its state to Established. ● Established: indicates that BGP peers can exchange Update, Notification and Keepalive packets. – If BGP receives an Update or a Keepalive message, its state stays in Established. – If BGP receives a Notification message, BGP changes its state to Idle.
PrefRcv	Indicates the number of route prefixes sent from the peer.

# Display the detailed information about a BGP EVPN peer.

```
<HUAWEI> display bgp evpn peer 4.4.4.4 verbose

        BGP Peer is 4.4.4.4,  remote AS 100
        Type: IBGP link
        BGP version 4, Remote router ID 4.4.4.4
        Update-group ID: 0
        BGP current state: Established, Up for 00h04m51s
        BGP current event: KATimerExpired
        BGP last state: OpenConfirm
        BGP Peer Up count: 2
        Received total routes: 2
        Received active routes total: 2
        Received mac routes: 0
        Advertised total routes: 2
        Port:  Local - 65472    Remote - 179
```

```
       Configured: Connect-retry Time: 32 sec
       Configured: Min Hold Time: 0 sec
       Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
       Received  : Active Hold Time: 180 sec
       Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
       Peer optional capabilities:
       Peer supports bgp multi-protocol extension
       Peer supports bgp route refresh capability
       Peer supports bgp 4-byte-as capability
       Address family IPv4 Unicast: received
       Address family L2VPN EVPN: advertised and received
Received: Total 8 messages
               Update messages               2
               Open messages                 1
               KeepAlive messages            5
               Notification messages         0
               Refresh messages              0
Sent: Total 11 messages
               Update messages               4
               Open messages                 1
               KeepAlive messages            6
               Notification messages         0
               Refresh messages              0
Authentication type configured: None
Last keepalive received: 2018-03-29 20:52:33+00:00
Last keepalive sent    : 2018-03-29 20:52:33+00:00
Last update    received: 2018-03-29 20:48:33+00:00
Last update    sent    : 2018-03-29 20:48:41+00:00
Minimum route advertisement interval is 0 seconds
Optional capabilities:
Route refresh capability has been enabled
4-byte-as capability has been enabled
Connect-interface has been configured
Peer Preferred Value: 0
Routing policy configured:
No routing policy is configured
Tracking has been enabled, and the delay is 9s
```

Table 18-5 Description of the **display bgp evpn peer** command output

Item	Description
BGP Peer is 4.4.4.4	IP addresses of the BGP peer
remote AS 100	AS number of the BGP peer
Type	BGP link type, which can be **IBGP Link** or **EBGP Link**.
BGP version	BGP version (current version: BGP4)
Remote router ID	Router ID of the peer
Update-group ID	Update the peer group ID

Item	Description
BGP current state	Current BGP status: ● Idle: BGP denies any connection request. This is the initial status of BGP. After BGP receives a start event, BGP initiates a TCP connection to a peer, starts the ConnectRetry timer, and listens to the TCP messages from the peer. BGP then enters the Connect state. ● Idle(Admin): The neighbor relationship is shut down initiatively and no attempt is made to establish the neighbor relationship. If the **peer ignore** command is configured or the peer is set to the down state through the MIB, the neighbor is in this state. ● No neg: The negotiation of the capabilities of the BGP peer's address family is not successful. ● Idle(Ovlmt): The BGP neighbor relationship is interrupted because the number of routes exceeds the upper threshold (configured by running the **peer route-limit** command). When this happens, the neighbor is in this state. If the **reset bgp** command is not run, the BGP neighbor relationship is not reestablished. ● Connect: BGP waits for the TCP connection establishment to complete before performing further operations. – If the TCP connection is successfully established, BGP stops the ConnectRetry timer and sends an Open message to the peer. BGP then enters the Opensent state. – If the TCP connection fails to be established, BGP resets the ConnectRetry timer and listens to the TCP connection initiated by the peer. BGP then enters the Active state. – If the ConnectRetry timer expires, BGP restarts the ConnectRetry timer and attempts to establish a TCP connection with the peer again. At this time, BGP remains in the Connect state. ● Active: BGP attempts to establish a TCP connection. This is the intermediate status of BGP. – If the TCP connection is successfully established, BGP resets the ConnectRetry timer and sends an Open message to the peer. BGP then enters the Opensent state.

Item	Description
	- If the ConnectRetry timer expires, BGP restarts the ConnectRetry timer and enters the Connect state. - If BGP attempts to establish a TCP session with an unknown IP address but fails, BGP resets the ConnectRetry timer and remains in the Active state. ● OpenSent: BGP has sent an Open message to the peer and waits for an Open message from the peer. - If BGP receives a correct Open message, BGP enters the OpenConfirm state. - If BGP receives an incorrect Open message, BGP sends a Notification message to the peer and enters the Idle state. - If BGP receives a TCP connection teardown message, BGP resets the ConnectRetry timer and listens to the TCP connection initiated by the peer. BGP then enters the Active state. ● OpenConfirm: BGP waits for a Notification or Keepalive message. - If BGP receives a Notification or TCP connection teardown message, BGP enters the Idle state. - If BGP receives a Keepalive message, BGP enters the Established state. ● Established: BGP peers can exchange Update, Notification, and Keepalive messages. - If BGP receives an Update or Keepalive message, BGP remains in the Established state. - If BGP receives a Notification message, BGP enters the Idle state.
BGP current event	Current BGP event
BGP last state	Status of the last BGP stage, which can be Idle, Idle(Admin), Idle(Ovlmt), Connect, Active, OpenSent, OpenConfirm, or Established.
BGP Peer Up count	Number of times the BGP peer alternates between Up and Down
Received total routes	Number of route prefixes received
Received active routes total	Number of active route prefixes received
Received mac routes	Number of MAC routes received

Item	Description
Advertised total routes	Number of route prefixes sent
Port	Port number ● Local: local port number, which is fixed to 179 because BGP uses TCP as the transport layer protocol ● Remote: peer port number
Configured	Timers that are locally configured: ● Connect-retry Time: ConnectRetry interval for a peer or peer group, in the unit of seconds. When BGP initiates a TCP connection, the ConnectRetry timer is stopped if the TCP connection is established successfully. If the first attempt to establish a TCP connection fails, BGP tries again to establish the TCP connection after the ConnectRetry timer expires. ● Min Hold Time: minimum hold time configured on the local device, in seconds. ● Active Hold Time: hold time, in seconds. If BGP does not receive any KeepAlive message from the peer in the hold time, BGP considers that the peer is Down and then instructs other peers to remove the routes that are sent from the device. ● Keepalive Time: interval for sending KeepAlive messages to the peer, in seconds. BGP peers send KeepAlive messages at intervals to show that they are working normally.
Received: Active Hold Time	Hold time of the peer
Negotiated	Timer negotiated by peers: ● Active Hold Time: hold time agreed between the BGP peers after capability negotiation ● Keepalive Time: Keepalive message transmission interval agreed between the BGP peers after capability negotiation
Peer optional capabilities	Capability supported by peers (optional) ● Peer supports bgp multi-protocol extension: Indicates the BGP peer supports BGP multi-protocol extension ● Peer supports bgp route refresh capability: Indicates the BGP peer supports BGP route refresh ● Peer supports bgp 4-byte-as capability: Indicates the BGP peer supports BGP 4-byte AS numbers

Item	Description
Address family IPv4 Unicast	IPv4 unicast address family
Address family L2VPN EVPN	BGP EVPN address family
Received	Number of messages received from a peer: ● Total: total number of messages received from the peer ● Update messages: number of Update messages ● Open messages: number of Open messages ● KeepAlive messages: number of Keepalive messages ● Notification messages: number of Notification messages ● Refresh messages: number of route-refresh messages
Sent	Number of messages sent to the peer: ● Total: total number of messages ● Update messages: number of Update messages ● Open messages: number of Open messages ● KeepAlive messages: number of Keepalive messages ● Notification messages: number of Notification messages ● Refresh messages: number of route-refresh messages
Authentication type configured	BGP authentication type configured. The value can be the following: ● Message digest 5 (MD5) ● Keychain (kk): kk indicates the name of the configured keychain authentication. ● None: No BGP authentication is configured.

Item	Description
Last keepalive received	Indicates the time when the Keepalive message is received last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is configured through the **clock timezone** command; DST indicates that the daylight saving time is configured through the **clock daylight-saving-time** command.
Last keepalive sent	Time when the Keepalive message is sent last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is configured through the **clock timezone** command; DST indicates that the daylight saving time is configured through the **clock daylight-saving-time** command.
Last update received	Time when the Update message is received last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is configured through the **clock timezone** command; DST indicates that the daylight saving time is configured through the **clock daylight-saving-time** command.

Item	Description
Last update sent	Time when the Update message is sent last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is configured through the **clock timezone** command; DST indicates that the daylight saving time is configured through the **clock daylight-saving-time** command.
Minimum route advertisement interval is 30 seconds	Minimum route advertisement intervals. The default minimum route advertisement intervals are as follows: ● The minimum interval for advertising EBGP routes is 30 seconds. ● The minimum interval for advertising IBGP routes is 15 seconds.
Optional capabilities	(Optional) Capabilities of the peer
Peer Preferred Value	Preferred value of the peer
Routing policy configured	Whether a routing policy has been configured

18.1.18 display bgp evpn routing-table

Function

The **display bgp evpn routing-table** command displays information about BGP EVPN routes.

Format

display bgp evpn all routing-table statistics

display bgp evpn all routing-table [**inclusive-route** [*inclusive-route*] | **mac-route** [*mac-route*] | **prefix-route** [*prefix-route*]]

display bgp evpn route-distinguisher *route-distinguisher* **routing-table** { **inclusive-route** [*inclusive-route*] | **mac-route** [*mac-route*] | **prefix-route** [*prefix-route*] }

Parameters

Parameter	Description	Value
all	Displays all route information.	-
statistics	Displays route statistics.	-
route-distinguisher *route-distinguisher*	Displays information about EVPN routes with a specified RD.	RD has the following four formats: ● 16-bit AS number:32-bit user-defined number. For example, 101:3. The AS number is an integer ranging from 0 to 65535, and the user-defined number is an integer ranging from 0 to 4294967295. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. ● Integral 4-byte AS number:2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295, and a user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. ● 4-byte AS number in dotted notation: 2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that both range from 0 to 65535. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. ● 32-bit IP address:16-bit user-defined number. For example, 192.168.122.15:1. The IP address ranges from 0.0.0.0 to 255.255.255.255, and the user-defined number ranges from 0 to 65535.
inclusive-route	Displays the information about inclusive multicast routes.	-

Parameter	Description	Value
inclusive-route	Specifies an inclusive multicast route.	The format is M:L:X.X.X.X: ● M: 0 (fixed) ● L: mask length of the source address of the device originating the route ● X.X.X.X: source address of the device originating the route
mac-route	Displays the information about MAC advertisement routes.	-
mac-route	Specifies a MAC advertisement route.	The format is E:M:H-H-H:L:X.X.X.X or E:M:H-H-H:L:[X:X::X:X]: ● E: ID of the VLAN to which the MAC address belongs ● M: 48 (MAC address length) ● H-H-H: MAC address. H is a 4-digit hexadecimal number, such as 00e0 and fc01. If you enter fewer than four alphanumeric characters, 0s are added before the input digits. For example, if e0 is entered, 00e0 is specified. ● L: mask length of the IP address corresponding to the MAC address ● X.X.X.X: IP address corresponding to the MAC address ● X:X::X:X indicates the IPv6 address corresponding to the MAC address.
prefix-route	Displays the information about an IP prefix route.	-
prefix-route	Specifies an IP prefix route.	The format is L:X.X.X.X:M or L:[X:X::X:X]:M: ● L: 0 (fixed) ● X.X.X.X: host IP address ● M: host IP address and mask ● X:X::X:X indicates the IPv6 address of host routes.
statistics	Displays route statistics.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, you can run this command to view EVPN route information and statistics on the device.

Example

# Display all BGP EVPN routing information.

```
<HUAWEI> display bgp evpn all routing-table
Local AS number :100

 BGP Local router ID is 10.35.99.225
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 EVPN address family:
 Number of Mac Routes : 2
 Route Distinguisher: 1:1


        Network(EthTagId/MacAddrLen/MacAddr/IpAddrLen/IpAddr)   NextHop
 *>     0:48:a4dc-be0f-99a6:0:0.0.0.0                           9.9.9.9
 *>     0:48:a4dc-be0f-99a6:32:9.9.9.9                          9.9.9.9

 VPN-Instance vpn1, Router ID 10.35.99.225:

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  9.9.9.9/32         9.9.9.9                   100        0      ?

 EVPN address family:
 Number of Inclusive Multicast Routes : 1
 Route Distinguisher: 1:1


        Network(EthTagId/IpAddrLen/OriginalIp)                  NextHop
 *>     0:32:9.9.9.9                                      9.9.9.9

 EVPN address family:
 Number of Ip Prefix Routes : 2
 Route Distinguisher: 1:1


        Network(EthTagId/IpPrefix/IpPrefixLen)                NextHop
 *>     0:9.9.9.0:24                                      9.9.9.9
 *>     0:[2001:db8:1::]:64                                  9.9.9.9

 VPN-Instance vpn1, Router ID 10.35.99.225:

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  9.9.9.0/24         9.9.9.9         0          100        0         ?
```

```
VPN-Instance vpn1, Router ID 10.35.99.225:

Total Number of Routes: 1

*>  Network   : 2001:db8:1::                  PrefixLen  : 64
    NextHop   : ::FFFF:9.9.9.9                LocPrf     :
    MED       :                               PrefVal    : 0
    Label     : 5010
    Path/Ogn  :200 100 ?
```

Table 18-6 Description of the **display bgp evpn all routing-table** command output

Item	Description
Local AS number	Local AS number
BGP Local router ID	Router ID of the local device
Number of Mac Routes	Number of MAC advertisement routes
Route Distinguisher	Route distinguisher
Network	Reachable address
NextHop	Next hop
VPN-Instance	VPN instance name
Router ID	Router ID
Total Number of Routes	Total number of routes
EVPN address family	BGP EVPN address family
MED	Multi-exit discriminator of route
LocPrf	Local priority
PrefVal	Value preferred by the protocol
Path/Ogn	AS path attribute and origin attribute of the route
Number of Inclusive Multicast Routes	Number of inclusive multicast routes
Number of Ip Prefix Routes	Number of IP prefix routes
PrefixLen	Mask length
Label	Label value

# Display statistics information about all BGP EVPN routes.

```
<HUAWEI> display bgp evpn all routing-table statistics
Total number of routes from all PE: 5
Number of Mac Routes: 2
Number of Inclusive Multicast Routes: 1
Number of Ip Prefix Routes: 2
```

Table 18-7 Description of the **display bgp evpn all routing-table statistics** command output

Item	Description
Total number of routes from all PE	Total number of routes
Number of Mac Routes	Number of MAC routes
Number of Inclusive Multicast Routes	Number of inclusive multicast routes
Number of Ip Prefix Routes	Number of IP prefix routes

18.1.19 display bgp evpn routing-table peer statistics

Function

The **display bgp evpn routing-table peer statistics** command displays statistics about received and advertised BGP EVPN routes.

Format

display bgp evpn routing-table peer statistics

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To display the statistics of the received and advertised BGP EVPN routes, run this command.

Example

# Display the statistics about the received and advertised BGP EVPN routes.

```
<HUAWEI> display bgp evpn routing-table peer statistics

 BGP local router ID : 6.6.6.6
 Local AS number : 100
 Total number of peers : 2              Number of Peers in established state : 2

  Peer                  Received routes           Advertised routes
 4.4.4.4              2                    2
 5.5.5.5              2                    2
```

Table 18-8 Description of the **display bgp evpn routing-table peer statistics** command output

Item	Description
BGP local router ID	Local router ID
Local AS number	AS number
Total number of peers	Total number of peers
Number of Peers in established state	Number of peers in the established state
Peer	IP address of a peer
Received routes	Total number of routes received from the peer
Advertised routes	Total number of routes advertised to the peer

18.1.20 display bgp evpn update-peer-group

Function

The **display bgp evpn update-peer-group** command displays the BGP update group information about an EVPN address family.

Format

display bgp evpn update-peer-group [**index** *update-peer-group-index*]

Parameters

Parameter	Description	Value
index *update-peer-group-index*	Specifies the index of a BGP update peer-group.	The value is an integer that ranges from 0 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, you can run this command to view the update group information about an EVPN address family.

Example

# Display the update group information about an EVPN address family.

```
<HUAWEI> display bgp evpn update-peer-group

 The EVPN instance's update peer group number : 1
 Keep buffer update peer group number : 0
 BGP Version : 4

 Group ID : 0
 Group Type : internal
 Addr Family : L2-EVPN
 AdvMinTimeVal : 0
 Total Peers : 2
 Leader Peer : 5.5.5.5
 Peers List : 5.5.5.5          4.4.4.4
```

# Display information about the BGP update peer-group with the index 0 in the EVPN address family.

```
<HUAWEI> display bgp evpn update-peer-group index 0

 Group ID : 0
 BGP Version : 4
 Group Type : internal
 Addr Family : L2-EVPN
 AdvMinTimeVal : 0
 Total Peers : 1
 Leader Peer : 1.1.1.1

 Total format packet number : 85
 Total send packet number : 85
 Total replicate packet number : 0
 The replication percentages(%) : 0

 Peers List : 1.1.1.1
```

Table 18-9 Description of the **display bgp evpn update-peer-group** command output

Item	Description
The EVPN instance's update peer group number	Number of update groups in an EVPN instance
Keep buffer update peer group number	Number of packets in update groups saved in the batch buffer
BGP Version	BGP version
Group ID	ID of the update group

Item	Description
Group Type	Type of the update group, which can be one of the following: ● external: EBGP peer group. ● internal: IBGP peer group. ● external-confed: EBGP update group in the confederation ● internal-confed: IBGP update group in the confederation ● unknown: unknown type
Addr Family	Address family
AdvMinTimeVal	Minimum interval for sending Update packets with the same route prefix
Total Peers	Total number of peers in an update group
Leader Peer	Representative of an update-group
Total format packet number	Number of Update messages of a single peer sends in the update peer-group
Total send packet number	Number of Update messages sent to all BGP peers in the update peer-group
Total replicate packet number	Absolute value of the difference between Total send packet number and Total format packet number, that is, absolute value of the difference between the number of Update messages that are sent to all peers in the update peer-group and the number of Update messages sent to a single peer in the update peer-group
The replication percentages(%)	Percentage of the number of replicated Update messages to the total number of sent Update messages, that is, (Total send packet number – Total format packet number)/Total send packet number
Peers List	List of peers

18.1.21 display bridge-domain

Function

The **display bridge-domain** command displays the BD configuration.

Format

display bridge-domain [*bd-id* [**brief** | **verbose**]]

Parameters

Parameter	Description	Value
bd-id	Displays the configuration of a specified BD. If this parameter is not specified, the device displays the configuration of all BDs.	The value is an integer that ranges from 1 to 16777215.
brief	Displays brief BD configuration.	-
verbose	Displays detailed BD configuration.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After BDs are created on a device, you can run the **display bridge-domain** command to view the configuration of a specified BD or all BDs.

Example

# Display the configuration of all BDs.

```
<HUAWEI> display bridge-domain
STAT: Statistics;
--------------------------------------------------------------------------------

BDID     State    STAT       Description
--------------------------------------------------------------------------------
10       down    disable     vxlan
20       up      disable     vxlan
--------------------------------------------------------------------------------
The total number of bridge-domains is : 2
```

# Display the detailed configuration of BD 10.

```
<HUAWEI> display bridge-domain 10 verbose
Bridge-domain ID   :10
Description        :vxlan
State              :Down
Statistics         :Disable
Broadcast-suppression       CIR(kbit/s) :-          CBS(byte) :-
Multicast-suppression       CIR(kbit/s) :10000000   CBS(byte) :4294967295
Unknown-unicast-suppression CIR(kbit/s) :0          CBS(byte) :655355
--------------------------------------------------------------------------------
```

```
Interface                          State
--------------------------------------------------------------------------------
GigabitEthernet1/0/1.1                 down
--------------------------------------------------------------------------------
VLAN                               State
--------------------------------------------------------------------------------
2                                down
3                                down
10                               down
--------------------------------------------------------------------------------
```

Table 18-10 Description of the **display bridge-domain** command output

Item	Description
BDID/Bridge-domain ID	ID of a BD. To set the BD ID, run the **bridge-domain** command in the system view.
State	BD status: ● up: The BD is bound to a Layer 2 sub-interface, VLAN, or VNI, and at least one of the bound Layer 2 sub-interface, VLAN, and VNI is Up. ● down: The BD is not bound to a Layer 2 sub-interface, VLAN, or VNI; alternatively, the BD is bound to a Layer 2 sub-interface, VLAN, or VNI and the bound Layer 2 sub-interface, VLAN, and VNI are Down.
STAT/Statistics	Whether traffic statistics collection is enabled for the BD: ● disable ● enable
Description	BD description. To configure the description of a BD, run the **description (BD view)** command.
The total number of bridge-domains is	Total number of BDs on the device.
Broadcast-suppression CIR(kbit/s) CBS(byte)	CIR and CBS in a BD specified by the broadcast suppression function. The unit is kbit/s and byte, respectively.
Multicast-suppression CIR(kbit/s) CBS(byte)	CIR and CBS in a BD specified by the multicast suppression function. The unit is kbit/s and byte, respectively.
Unknown-unicast-suppression CIR(kbit/s) CBS(byte)	CIR and CBS in a BD specified by the unknown unicast suppression function. The unit is kbit/s and byte, respectively.

Item	Description
Interface State	Member interface in a BD and its status. ● up: The link layer protocol of the interface is in the Up state. ● down: The link layer protocol of the interface is Down.
VLAN State	Status of the VLAN associated with the BD. ● up ● down The status of a VLAN is determined by the status of member interfaces in the VLAN. A VLAN is Up only when at least one member interface in the VLAN is Up.

18.1.22 display bridge-domain statistics

Function

The **display bridge-domain statistics** command displays packet statistics in a BD.

Format

display bridge-domain *bd-id* **statistics**

Parameters

Parameter	Description	Value
bd-id	Displays packet statistics in a specified BD.	The value is an integer that ranges from 1 to 16777215.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bridge-domain statistics** command to view packet statistics in a BD. The information helps you locate faults and simplifies VXLAN network maintenance.

Before using this command to view packet statistics in a BD, run the **statistics enable** command in the BD view to enable packet statistics collection in the BD.

Example

# Display packet statistics in BD 10.

```
<HUAWEI> display bridge-domain 10 statistics
Total:
------------------------------------------------------------------------
Item             Packets                   Bytes
------------------------------------------------------------------------
Inbound            10                      1520
Outbound             10                     1520
------------------------------------------------------------------------
```

Table 18-11 Description of the **display bridge-domain statistics** command output

Item	Description
Slot	Slot ID.
Item	Statistical item.
Packets	Number of packets.
Bytes	Number of bytes.
Inbound	Statistics on packets going in to a BD.
Outbound	Statistics on packets leaving a BD.

18.1.23 display evpn mac routing-table

Function

The **display evpn mac routing-table** command displays the MAC routing entries of the BGP EVPN function.

Format

display evpn mac routing-table { **all-vpn-instance** | **vpn-instance** *evpn-instance-name* [*mac-address*] }

Parameters

Parameter	Description	Value
all-vpn-instance	Displays the MAC routing entries of all the EVPN instances.	-
vpn-instance *evpn-instance-name*	Displays the MAC routing entry of a specified EVPN instance.	The EVPN instance must already exist on the device.

Parameter	Description	Value
mac-address	Displays the MAC routing entry of a specified MAC address.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits. The MAC address cannot be FFFF-FFFF-FFFF, 0000-0000-0000, or a multicast MAC address.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To check MAC routing entries of the BGP EVPN function, run this command.

Example

# Display MAC routing entries of all the EVPN instances.

```
<HUAWEI> display evpn mac routing-table all-vpn-instance
EVPN-Instance Name : evpn10
-------------------------------------------------------------------------------
Flag     MAC Address      Serial number    VNI         VTEP Address
-------------------------------------------------------------------------------
local    0000-5e00-5311  0               1002        10.1.1.1
remote   0000-5e00-5322  0               1002        10.2.2.2
-------------------------------------------------------------------------------
Total Num: 2 Local Num: 1 Remote Num : 1
```

Table 18-12 Description of the **display evpn mac routing-table** command output

Item	Description
EVPN-Instance Name	EVPN instance where a MAC routing entry resides.
Flag	Flag of the MAC routing entry type: ● remote: The MAC routing entry is remotely obtained. ● local: The MAC routing entry is locally obtained.
MAC Address	MAC address of a MAC routing entry.
Serial number	Serial number of a MAC routing entry.
VNI	Layer 2 VNI to which a MAC routing entry belongs.

Item	Description
VTEP Address	VTEP IP address of a MAC routing entry.
Total Num	Total number of MAC routing entries.
Local Num	Total number of MAC routing entries locally obtained.
Remote Num	Total number of MAC routing entries remotely obtained.

18.1.24 display evpn vpn-instance

Function

The **display evpn vpn-instance** command displays the configurations of EVPN instances.

Format

display evpn vpn-instance [**verbose**] [*evpn-instance-name*]

Parameters

Parameter	Description	Value
verbose	Displays detailed information.	-
evpn-instance-name	Specifies the name of an EVPN instance.	The value is the name of an existing VPN instance.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, EVPN instances need to be configured. After EVPN instances are configured, you can run this command to query the configurations of EVPN instances.

If *evpn-instance-name* is not specified, this command displays the information about all configured EVPN instances on the device.

Example

# Display the brief information about all EVPN instances configured on the device.

```
<HUAWEI> display evpn vpn-instance
 Total EVPN-Instances configured      : 1
  VPN-Instance Name                RD
---------------------------------------------------------
  vn10                             10:1
```

Table 18-13 Description of the **display evpn vpn-instance** command output

Item	Description
Total EVPN-Instances configured	Total number of EVPN instances configured on the local end.
VPN-Instance Name	Name of the EVPN instance.
RD	RD of the EVPN instance IPv4 address family.

# Display detailed information about all EVPN instances.

```
<HUAWEI> display evpn vpn-instance verbose
 Total EVPN-Instances configured      : 1

 EVPN-Instance Name and ID : vn10, 1
 Bridge-domain Number : 1
  Bridge-domain List : bridge-domain 10 source 1.1.1.1
  Route Distinguisher : 10:1
  Export VPN Targets : 1:1 100:1
  Import VPN Targets : 100:1
```

Table 18-14 Description of the **display evpn vpn-instance** command output

Item	Description
Total EVPN-Instances configured	Total number of EVPN instances configured on the local end.
EVPN-Instance Name and ID	Name and ID of the EVPN instance. The ID is assigned by the system, which facilitates indexing.
Bridge-domain Number	Number of BDs bound to an EVPN instance
Bridge-domain List	BDs bound to an EVPN and IP address of the source VTEP
Route Distinguisher	RD of the EVPN instance IPv4 address family
Export VPN Targets	Route Target list in the outbound direction

Item	Description
Import VPN Targets	Route Target list in the inbound direction

18.1.25 display evpn vpn-instance bridge-domain

Function

The **display evpn vpn-instance bridge-domain** command displays the information about the BD bound to an EVPN instance.

Format

display evpn vpn-instance bridge-domain

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, EVPN instances need to be configured and bound to BDs. After EVPN instances are configured and bound to BDs, you can run this command to display the information about the BDs bound to EVPN instances.

Example

# Display the information about the BD bound to an EVPN.

```
<HUAWEI> display evpn vpn-instance bridge-domain
 Total EVPN-Instances configured      : 1

 EVPN-Instance Name and ID : vn10, 1
 Bridge-domain Number : 1
  Bridge-domain List : bridge-domain 10 source 1.1.1.1
```

Table 18-15 Description of the **display evpn vpn-instance bridge-domain** command output

Item	Description
Total EVPN-Instances configured	Total number of EVPN instances configured on the device
EVPN-Instance Name and ID	Name and ID of the EVPN instance. The ID is assigned by the system, which facilitates indexing.
Bridge-domain Number	Total number of BDs to which this EVPN instance is bound
Bridge-domain List	IP address of the BD to which the EVPN instance is bound and the local VTEP

18.1.26 display evpn vpn-instance import-vt

Function

The **display evpn vpn-instance import-vt** command displays the information about the EVPN instance that matches an import route target (IRT).

Format

display evpn vpn-instance import-vt *ivt-value*

Parameters

Parameter	Description	Value
ivt-value	Specifies the value of the import VPN-target attribute. The forms of VPN targets are as follows: ● 2-byte AS number: 4-byte user-defined number, for example, 1:3. The AS number ranges from 0 to 65535. The user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot both be 0. That is, a VPN target cannot be 0:0. ● IPv4-address: 2-byte user-defined number, for example, 192.168.122.15:1. The IP address ranges from 0.0.0.0 to 255.255.255.255. The user-defined number ranges from 0 to 65535. ● Integral 4-byte AS number: 2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0. ● 4-byte AS number in dotted notation: 2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of x.y, where x and y are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0.0:0.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

When a device is configured with multiple EVPN instances, you can run this command to display the EVPN instances to which the EVPN route with a specified VPN target can be imported.

Example

# Display the information about the EVPN instance that matches an IRT.

```
<HUAWEI> display evpn vpn-instance import-vt 100:1
 The number of EVPN-Instances matched the import-vt 100:1 : 1
  EVPN-Instance Name and ID : bd10,1
```

Table 18-16 Description of the **display evpn vpn-instance import-vt** command output

Item	Description
The number of EVPN-Instances matched the import-vt	Total number of EVPN instances that matches an IRT
EVPN-Instance Name and ID	Name and ID of the EVPN instance. The ID is assigned by the system, which facilitates indexing.

18.1.27 display evpn vpn-instance peer-list

Function

The **display evpn vpn-instance peer-list** command displays the information about the remote peer of an EVPN instance.

Format

display evpn vpn-instance *evpn-instance-name* **peer-list**

Parameters

Parameter	Description	Value
evpn-instance-name	Specifies the name of an EVPN instance.	The value must be the name of an existing EVPN instance.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, you can run this command to query the information about the remote peer of an EVPN instance.

Example

# Query the information about the remote peer of EVPN instance **vn10**.

```
<HUAWEI> display evpn vpn-instance vn10 peer-list
 Total EVPN-Instance Name and Peer-list Number : vn10, 1
  Remote VNI         Peer-Address       State
-------------------------------------------------
  10                1.1.1.1           Active
```

Table 18-17 Description of the **display evpn vpn-instance peer-list** command output

Item	Description
Total EVPN-Instance Name and Peer-list Number	EVPN instance name and total number of the peers of the EVPN instance
Remote VNI	Layer 2 VNI carried by the peer
Peer-Address	Peer IP address
State	Connection status ● Active ● Inactive

18.1.28 display evpn vpn-instance tunnel-info

Function

The **display evpn vpn-instance tunnel-info** command displays the information about the tunnels constructed through EVPN instances.

Format

.

display evpn vpn-instance [**verbose**] **tunnel-info**

Parameters

Parameter	Description	Value
verbose	Displays detail information	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, the device constructs dynamic Layer 2 VXLAN tunnels through the routes between EVPN instances. You can run this command to display the information about the tunnels constructed through EVPN instances.

Example

# Display the information about the tunnels constructed through EVPN instances.

```
<HUAWEI> display evpn vpn-instance tunnel-info
 Total EVPN-Instances Tunnel          : 2
  Remote VNI         Peer-Address
-------------------------------------
  10                 1.1.1.1
  20                 2.2.2.2
```

Table 18-18 Description of the **display evpn vpn-instance tunnel-info** command output

Item	Description
Total EVPN-Instances Tunnel	Number of tunnels constructed through EVPN instances
Remote VNI	Layer 2 VNI carried by the peer
Peer-Address	Peer IP address

# Display the detailed information about the tunnels constructed through EVPN instances.

```
<HUAWEI> display evpn vpn-instance verbose tunnel-info
 Total EVPN-Instances Tunnel          : 2

 Remote VNI 10         Peer-Address 1.1.1.1
  Export VPN Targets : 1:1 100:1

 Remote VNI 20         Peer-Address 2.2.2.2
  Export VPN Targets : 1:1 200:1
```

Table 18-19 Description of the **display evpn vpn-instance tunnel-info** command output

Item	Description
Total EVPN-Instances Tunnel	Number of tunnels constructed through EVPN instances
Remote VNI	Layer 2 VNI carried by the peer
Peer-Address	Peer IP address
Export VPN Targets	Export RT list

18.1.29 display interface nve

Function

The **display interface nve** command displays Network Virtualization Edge (NVE) interface information.

Format

display interface nve [*nve-number* | **main**]

Parameters

Parameter	Description	Value
nve-number	Specifies the number of an NVE interface. If *nve-number* is not specified, information about all NVE interfaces is displayed.	The value is 1.
main	Displays the running status and statistics of the main interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To monitor the NVE interface status or locate an NVE interface fault on the VXLAN network, run the **display interface nve** command to check the running status and statistics of the NVE interface.

Example

# Display the running status of the NVE interface.

```
<HUAWEI> display interface nve 1
Nve1 current state : UP
Line protocol current state : UP
Description:
Route Port
Internet protocol processing : disabled
Current system time: 2017-03-28 19:50:24
```

Table 18-20 Description of the **display interface nve** command output

Item	Description
Nve1 current state	Physical status of the NVE interface. The physical status of the successfully created NVE interface is always Up.
Line protocol current state	Link layer protocol status of the NVE interface. The link layer protocol status of the successfully created NVE interface is always Up.
Description	Description of the NVE interface.

Item	Description
Route Port	The interface is a Layer 3 interface.
Internet protocol processing	This field displays only **disabled**, and the interface cannot be configured with an IP address.
Current system time	System time.

18.1.30 display interface vbdif

Function

The **display interface vbdif** command displays the status, configuration, and statistics of a VBDIF interface.

Format

display interface vbdif [*bd-id* | **main**]

Parameters

Parameter	Description	Value
bd-id	Displays the status, configuration, and statistics of a VBDIF interface with a specified BD ID. If no BD ID is specified, the status, configuration, and statistics of all VBDIF interfaces are displayed.	The BD ID of a VBDIF interface must already exist.
main	Displays the running status and statistics of the main interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

To monitor an interface or locate an interface fault, you can use the **display interface vbdif** command to view the interface status, interface configuration, and traffic statistics on the interface. The information helps you locate faults in the system or on an interface.

Prerequisites

The specified VBDIF interface has been created.

Example

# Display information of VBDIF interface with BD ID 20.

```
<HUAWEI> display interface vbdif 20
Vbdif20 current state : UP
Line protocol current state : UP
Last line protocol up time : 2015-07-08 11:25:34
Description:
Route Port,The Maximum Transmit Unit is 1500
Internet Address is 192.168.20.1/24
IP Sending Frames' Format is PKTFMT_ETHNT_2, Hardware address is
0000-5e00-0101
Current system time: 2015-07-08 14:09:59
    Input bandwidth utilization  : --
    Output bandwidth utilization : --
```

Table 18-21 Description of the **display interface vbdif** command output

Item	Description
Vbdif20 current state	Physical status of a VBDIF interface. ● UP: The physical status of the interface is Up. ● DOWN: The physical status of the interface is Down. ● Administratively down: The administrator has run the **shutdown** command on the VBDIF interface.
Line protocol current state	Link-layer protocol status of a VBDIF interface. ● UP: The link-layer protocol of the interface is Up. ● DOWN: The link-layer protocol of the interface is Down, or no IP address is assigned to the interface.
Last line protocol up time	Last time the link-layer protocol of an interface goes Up. **NOTE** This field is displayed only when the link-layer protocol status is Up.
Description	Description of a VBDIF interface. The description helps you learn the functions of the interface.
Route Port	Layer 3 interface.
The Maximum Transmit Unit is	Maximum transmit unit (MTU) of an interface. The default MTU is 1500 bytes. Packets whose size is greater than the MTU are fragmented before being transmitted. If non-fragmentation is configured, these packets are discarded.
Internet Address is	IP address of an interface. If no IP address is configured for the current interface, the command output is displayed as "Internet protocol processing: disabled."

Item	Description
IP Sending Frames' Format is	Format of the Ethernet frames sent by a VBDIF interface. The default value is **PKTFMT_ETHNT_2**. A VBDIF interface can identify the received Ethernet frames of the following formats: ● PKTFMT_ETHNT_2 ● Ethernet_SNAP ● 802.2 ● 802.3
Hardware address is	Physical address of an interface.
Current system time	System time.
Input bandwidth utilization	Number of packets received by the interface.
Output bandwidth utilization	Number of packets sent by the interface

18.1.31 display mac-address bridge-domain

Function

The **display mac-address bridge-domain** command displays MAC address entries of a BD.

Format

display mac-address [*mac-address*] **bridge-domain** *bd-id* [**verbose**]

Parameters

Parameter	Description	Value
mac-address	Displays an entry with a specified MAC address.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits. The MAC address cannot be FFFF-FFFF-FFFF, 0000-0000-0000, or a multicast MAC address.
bd-id	Displays MAC address entries of a specified BD.	The value is an integer that ranges from 1 to 16777215.
verbose	Displays detailed information about MAC address entries.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The MAC address table of a switch stores MAC addresses of other devices. The switch queries the MAC address table to quickly locate the outbound interface for data forwarding. You can run the **display mac-address bridge-domain** command to view MAC address entries of a specified BD.

Follow-up Procedure

If any MAC address entry in the command output is incorrect, run the **undo mac-address** command to delete the entry or run the **mac-address static** command to add a correct one.

Example

# Display MAC address entries of BD 20.

```
<HUAWEI> system-view
[HUAWEI] display mac-address bridge-domain 20
-------------------------------------------------------------------------------
MAC Address    VLAN/VSI/BD                       Learned-From        Type
-------------------------------------------------------------------------------
00e0-fc12-0006 -/-/20                           GE0/0/23.5          static

-------------------------------------------------------------------------------
Total items displayed = 1
```

# Display detailed information about MAC address entries of BD 20.

```
<HUAWEI> system-view
[HUAWEI] display mac-address bridge-domain 20 verbose
-------------------------------------------------------------------------------
MAC Address : 00e0-fc12-0006         BD   : 20
Learned-From: GE0/0/23.5             Type : static

-------------------------------------------------------------------------------
Total items displayed = 1
```

Table 18-22 Description of the **display mac-address bridge-domain** command output

Item	Description
MAC Address	MAC address.
VLAN/VSI/BD	ID of the VLAN, name of the virtual switch instance (VSI), or ID of the BD to which the MAC address belongs.
Learned-From	Interface on which a MAC address is learned.

Item	Description
Type	Type of a MAC address entry: ● static: a static MAC address entry, which is manually configured and will not be aged out. ● blackhole: a blackhole MAC address entry, which is manually configured and will not be aged out.

18.1.32 display mac-address evpn

Function

The **display mac-address evpn** command displays information about MAC address entries of the EVPN type.

Format

display mac-address evpn bridge-domain *bd-id* [**verbose**]

Parameters

Parameter	Description	Value
bridge-domain *bd-id*	Displays information about MAC address entries of the EVPN type in a specified BD.	The value is an integer that ranges from 1 to 16777215.
verbose	Displays detailed information about MAC address entries of the EVPN type.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

A device stores other devices' MAC addresses learned by itself into the MAC address table. The device queries the MAC address table to quickly locate the outbound interface for data forwarding.

Follow-up Task

You can run this command to check information about MAC address entries of the EVPN type. If any learned MAC address entry in the command output is incorrect, run the **undo mac-address** command to delete the entry or run the **mac-address static** command to add a correct one.

Example

# Display information about MAC address entries in BD 10.

```
<HUAWEI> display mac-address evpn bridge-domain 10
------------------------------------------------------------------------------
MAC Address    VLAN/VSI/BD                       Learned-From        Type
-------------------------------------------------------------------------------
00e0-fc12-3456 -/-/10                          192.0.2.10          evpn

-------------------------------------------------------------------------------
Total items displayed = 1
```

# Display detailed information about MAC address entries in BD 10.

```
<HUAWEI> display mac-address evpn bridge-domain 10 verbose
-------------------------------------------------------------------------------
MAC Address : 00e0-fc12-3456         BD   : 10
Learned-From: 192.0.2.10             Type : evpn


-------------------------------------------------------------------------------
Total items displayed = 1
```

Table 18-23 Description of the **display mac-address evpn** command output

Item	Description
MAC Address	MAC address information.
VLAN/VSI/BD	ID of the VLAN, name of the virtual switch instance (VSI), or ID of the BD to which the MAC address belongs.
BD	Number of the BD to which an entry belongs.
Learned-From	Remote VTEP IP address of a VXLAN tunnel that learns the MAC address entry of the EVPN type.
Type	MAC address entry type.

18.1.33 display nd multicast-suppress user bridge-domain

Function

The **display nd multicast-suppress user bridge-domain** command displays the ND multicast suppression table of a BD.

Format

display nd multicast-suppress user bridge-domain *bd-id* [*ipv6-address*]

display nd multicast-suppress user statistics { **bridge-domain** *bd-id* | **all** }

Parameters

Parameter	Description	Value
bd-id	Specifies the ID of the BD of the ND multicast suppression table to be queried.	The value is an integer that ranges from 1 to 16777215.
ipv6-address	Specifies the IP address of the ND multicast suppression table to be queried.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.
statistics	Displays ND multicast suppression entry statistics.	-
all	Displays ND multicast suppression entry statistics in all BDs.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

ND multicast suppression can effectively ease the pressure on the gateway in handling NS packets. When receiving an NS request packet, the gateway searches for the host information about the device corresponding to the destination IP address (the IPv6 address and MAC address mapping table of the destination device, also known as the ARP broadcast suppression table).

To display the ND multicast suppression table of a BD, run this command.

Example

# Display the ND multicast suppression table of BD 10.

```
<HUAWEI> display nd multicast-suppress user bridge-domain 10
-------------------------------------------------------------------------------
IPv6 Address : FC00:1::10
MAC Address  : 0487-ea01-0506   VNI : 10         Serial number : 0
VTEP Address : 10.2.2.2         Flag : local
VPN Instance :vpn1
-------------------------------------------------------------------------------
Total Num: 1 Local Num: 1 Remote Num : 0
```

Table 18-24 Description of the **display nd multicast-suppress user bridge-domain** command output

Item	Description
IPv6 Address	IPv6 address of an NS multicast suppression entry
MAC Address	MAC address of an NS multicast suppression entry
VNI	Layer 2 VNI that an NS multicast suppression entry belongs
Serial number	Serial number of an NS multicast suppression entry
VTEP Address	IP address of the source VTEP of an NS multicast suppression entry
Flag	NS multicast suppression entry type ● local: NS multicast suppression entry of the local access device ● remote: NS multicast suppression entry of the remote access device
VPN Instance	VPN instance bound to the gateway to which the NS multicast suppression entries belong
Total Num	Total number of NS multicast suppression entries in a BD
Local Num	Total number of NS multicast suppression entries learned from the local end in a BD
Remote Num	Total number of NS multicast suppression entries learned from the remote end in a BD

# Display statistics on the NS multicast suppression table of BD 10.

```
<HUAWEI> display nd multicast-suppress user statistics bridge-domain 10
Total: 1     Local: 1        Remote: 0
```

Table 18-25 Description of the **display nd multicast-suppress user bridge-domain** command output

Item	Description
Total	Total number of NS multicast suppression entries

Item	Description
Local	Total number of NS multicast suppression entries learned from the local end
Remote	Total number of NS multicast suppression entries learned from the remote end

18.1.34 display vxlan encapsulation

Function

The **display vxlan encapsulation** command displays VXLAN encapsulation information about Layer 2 sub-interfaces of a main interface.

Format

display vxlan encapsulation interface *interface-type interface-number* [**bridge-domain** *bd-id* | **default** | **dot1q** [**vid** *pe-vid*] | **qinq** [**vid** *vlan-vid* [**ce-vid** *ce-vid*]] | **untag**]

Parameters

Parameter	Description	Value
interface *interface-type interface-number*	Displays VXLAN encapsulation information about a specified Layer 2 sub-interface. ● *interface-type* specifies the interface type. ● *interface-number* specifies the interface number.	-
bridge-domain *bd-id*	Displays VXLAN encapsulation information about Layer 2 sub-interfaces by BD ID.	The BD ID must exist.
default	Displays VXLAN encapsulation information about Layer 2 sub-interfaces for the flow encapsulation mode **default**.	-
dot1q	Displays VXLAN encapsulation information about Layer 2 sub-interfaces for the flow encapsulation mode **Dot1q**.	-

Parameter	Description	Value
vid *pe-vid*	Displays VXLAN encapsulation information about Layer 2 sub-interfaces by VLAN ID in packets with the flow encapsulation mode **Dot1q**.	The VLAN ID must exist.
qinq	Displays VXLAN encapsulation information about Layer 2 sub-interfaces for the flow encapsulation mode **QinQ**.	-
vid *vlan-vid* [**ce-vid** *ce-vid*]	Displays VXLAN encapsulation information about Layer 2 sub-interfaces by VLAN ID in packets with the flow encapsulation mode **QinQ**. ● *vlan-vid* specifies the outer VLAN ID. ● *ce-vid* specifies the inner VLAN ID.	The VLAN ID must exist.
untag	Displays VXLAN encapsulation information about Layer 2 sub-interfaces for the flow encapsulation mode **untag**.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display vxlan encapsulation** command to view VXLAN encapsulation information about Layer 2 sub-interfaces of a main interface as well as the bindings between the sub-interface and BD.

Example

# Display VXLAN encapsulation information about Layer 2 sub-interfaces of a main interface.

```
<HUAWEI> display vxlan encapsulation interface gigabitethernet 0/0/1
----------------------------------------------------------------------------------------------------
Interface                EncapType        PeVid        CeVid        BD-ID      RewriteType
----------------------------------------------------------------------------------------------------
GigabitEthernet0/0/1.1      qinq             4093         4093         16000      pop default
----------------------------------------------------------------------------------------------------
Total number to display is 1.
```

Table 18-26 Description of the **display vxlan encapsulation** command output

Item	Description
Interface	Name of the Layer 2 sub-interface.
EncapType	Flow encapsulation type of the Layer 2 sub-interface. ● qinq ● dot1q ● default ● untag To configure the flow encapsulation type, run the **encapsulation** command in the Layer 2 sub-interface view.
PeVid	VLAN ID of packets with the flow encapsulation type **Dot1q** or outer VLAN ID of packets with the flow encapsulation type **QinQ**. To configure the VLAN ID, run the **encapsulation** command in the Layer 2 sub-interface view.
CeVid	Inner VLAN ID of packets with the flow encapsulation type **QinQ**. To configure the VLAN ID, run the **encapsulation** command in the Layer 2 sub-interface view.
BD-ID	BD. To configure the BD, run the **bridge-domain** command in the Layer 2 sub-interface view.
RewriteType	Type of the operation of removing VLAN tags in packets: ● pop default: Indicate the default operation of removing VLAN tags in packets. ● pop single: Remove one VLAN tag. ● pop double: Remove double VLAN tags. ● pop none: Do not remove any VLAN tag. To configure the operation type, run the **rewrite pop** command.
Total number to display is	Total number of encapsulation data records about Layer 2 sub-interfaces of a main interface.

18.1.35 display vxlan peer

Function

The **display vxlan peer** command displays the IP address of the destination VXLAN tunnel endpoint (VTEP) of a Virtual Network Identifier (VNI).

Format

display vxlan peer [**vni** *vni-id*]

Parameters

Parameter	Description	Value
vni *vni-id*	Displays the IP address of the destination VTEP of a specified VNI.	The value is an integer that ranges from 1 to 16777215.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

After completing VXLAN configuration, you can run the **display vxlan peer** command to view information about the source and destination IP address bound to the VNI.

Precautions

Before running the **display vxlan peer** command, ensure that the device has been configured with VNIs. Otherwise, the command output is meaningless.

Example

# Display the IP address of the destination VTEP of a specified VNI.

```
<HUAWEI> display vxlan peer
 Vni ID     Source           Destination        Type
 ---------------------------------------------------------------
 10         10.1.1.2         10.1.1.3           static
 10         10.1.1.2         10.1.1.4           static
 11         10.1.1.2         10.1.1.5           static&l2 dynamic
 ---------------------------------------------------------------
 Number of peers :
 Total: 3   Static: 3   L2 dynamic: 1   L3 dynamic: 0
```

Table 18-27 Description of the **display vxlan peer** command output

Item	Description
Vni ID	VNI ID. To configure or modify a VNI ID, run the **vxlan vni** command in the BD view. For a Layer 3 VXLAN tunnel, the **Vni ID** field displays **-**.
Source	IP address of the source VTEP. To configure or modify the IP address of the source VTEP, run the **source** command in the NVE interface view.

Item	Description
Destination	IP address of the destination VTEP. To configure or modify the IP address of the destination VTEP, run the **vni head-end peer-list** command.
Type	IP address configuration mode of the destination VTEP: ● static: The IP address is manually configured. ● l2 dynamic: Layer 2 VXLAN tunnel IP addresses are dynamically generated through the BGP protocol. ● l3 dynamic: Layer 3 VXLAN tunnel IP addresses are dynamically generated through the BGP protocol. To configure the IP address configuration mode, run the **vni head-end peer-list** command.
Number of peers	Number of IP addresses of destination VTEPs on the device: ● Total: indicates the number of IP addresses of destination VTEPs (If the Type field displays a combination of multiple types, such as static&l2 dynamic, the number of IP addresses of destination VTEPs is counted as 1). ● Static: indicates the number of IP addresses of static destination VTEPs. ● L2 dynamic: indicates the number of IP addresses of Layer 2 dynamic destination VTEPs. ● L3 dynamic: indicates the number of IP addresses of Layer 3 dynamic destination VTEPs.

18.1.36 display vxlan statistics

Function

The **display vxlan statistics** command displays VXLAN tunnel packet statistics.

Format

display vxlan statistics source *source-ip-address* **peer** *peer-ip-address* [**vni** *vni-id*]

display vxlan statistics source *source-ipv6-address* **peer** *peer-ipv6-address* [**vni** *vni-id*]

NOTE

Only the S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, S5731-H, and S6730-S support the **source** and **peer** parameters configured as IPv6 address.

Parameters

Parameter	Description	Value
source *source-ip-address*	Specifies the IPv4 address of the source VTEP.	The value is in dotted decimal notation.
peer *peer-ip-address*	Specifies the IPv4 address of the destination VTEP.	The value is in dotted decimal notation.
source *source-ipv6-address*	Specifies the IPv6 address of the source VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
peer *peer-ipv6-address*	Specifies the IPv6 address of the destination VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
vni *vni-id*	Specifies a VNI ID.	The value is an integer that ranges from 1 to 16777215.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display vxlan statistics** command to view VXLAN tunnel packet statistics. The information helps you locate faults and simplifies VXLAN network maintenance.

Before using this command to view VXLAN tunnel packet statistics, run the **vxlan statistics enable** command on an NVE interface to enable statistics collection on VXLAN tunnel packets.

Example

# Display statistics on VXLAN tunnel packets, with 10.10.1.1 and 10.1.1.1 as the source and destination VTEP IP addresses.

```
<HUAWEI> display vxlan statistics source 10.10.1.1 peer 10.1.1.1
Total:
------------------------------------------------------------------------
Item          Packets              Bytes
------------------------------------------------------------------------
Inbound        5                  760
Outbound        5                  760
------------------------------------------------------------------------
```

Table 18-28 Description of the **display vxlan statistics** command output

Item	Description
Item	Statistical item.
Packets	Number of packets.
Bytes	Number of bytes.
Inbound	Packet statistics in the inbound direction of the VXLAN tunnel.
Outbound	Packet statistics in the outbound direction of the VXLAN tunnel.

18.1.37 display vxlan tunnel

Function

The **display vxlan tunnel** command displays information about VXLAN tunnels.

Format

display vxlan tunnel [*tunnel-id*] [**verbose**]

Parameters

Parameter	Description	Value
tunnel-id	Displays information about the VXLAN tunnel with a specified ID. If this parameter is not specified, the device displays information about all VXLAN tunnels.	The value is an integer that ranges from 1 to 4294967295.
verbose	Displays detailed VXLAN tunnel information.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After VXLAN tunnels are established, you can run the **display vxlan tunnel** command to view VXLAN tunnel information.

Example

# Display VXLAN tunnel information.

```
<HUAWEI> display vxlan tunnel
 Tunnel ID       Source                Destination          State     Type
 ---------------------------------------------------------------------------
 4026531841      10.1.1.2              10.1.1.4             up        static
 ---------------------------------------------------------------------------
 Number of vxlan tunnel :
 Total: 1    Static: 1    L2 dynamic: 0    L3 dynamic: 0
```

# Display detailed VXLAN tunnel information.

```
<HUAWEI> display vxlan tunnel verbose
 Tunnel ID       : 4026531841
 Source          : 10.1.1.2
 Destination     : 10.1.1.4
 State           : up
 Reason          : -
 Type            : static
 ---------------------------------------------------------------------------
 Number of vxlan tunnel :
 Total: 1    Static: 1    L2 dynamic: 0    L3 dynamic: 0
```

Table 18-29 Description of the **display vxlan tunnel** command output

Item	Description
Tunnel ID	ID of a VXLAN tunnel. After a VXLAN tunnel is established, the ID is automatically generated by the device.
Source	Source IP address of the VXLAN tunnel. To configure the source IP address, run the **source** command in the NVE interface view.
Destination	Destination IP address of the VXLAN tunnel.
State	Status of a VXLAN tunnel: ● up: The VXLAN tunnel is reachable. ● down: The VXLAN tunnel is unreachable.
Reason	Reason why the VXLAN tunnel is Down: ● no license: No valid license is available. ● no route: Routes are unreachable.

Item	Description
Type	IP address configuration mode of the destination VTEP: ● static: The IP address is manually configured. ● l2 dynamic: Layer 2 VXLAN tunnel IP addresses are dynamically generated through the BGP protocol. ● l3 dynamic: Layer 3 VXLAN tunnel IP addresses are dynamically generated through the BGP protocol. To configure the IP address configuration mode, run the **vni head-end peer-list** command.
Number of vxlan tunnel	Total number of VXLAN tunnels on the device. ● Total: indicates the number of VXLAN tunnels. ● Static: indicates the number of static VXLAN tunnels. ● L2 dynamic: indicates the number of Layer 2 dynamic VXLAN tunnels. ● L3 dynamic: indicates the number of Layer 3 dynamic VXLAN tunnels.

18.1.38 display vxlan vni

Function

The **display vxlan vni** command displays the VXLAN configuration of a specified VNI or all VNIs.

Format

display vxlan vni [*vni-id* [**verbose**]]

Parameters

Parameter	Description	Value
vni-id	Displays VXLAN information about a specified VNI. If this parameter is not specified, the device displays VXLAN configuration of all VNIs.	The value is an integer that ranges from 1 to 16777215.
verbose	Displays detailed VXLAN information about a specified VNI.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Applications

After VXLAN is configured, you can run the **display vxlan vni** command to view information about BDs associated with VNIs and VNI status.

Precautions

Before running the **display vxlan vni** command, ensure that the device has been configured with VNIs. Otherwise, the command output is meaningless.

Example

# Display VXLAN information about all VNIs.

```
<HUAWEI> display vxlan vni
 VNI               BD-ID            State
 -----------------------------------------
 10                10               up
 50                50               up
 -----------------------------------------
 Number of vxlan vni bound to BD is : 2

 VNI               VRF-ID
 -----------------------------------------
 20                5
 -----------------------------------------
 Number of vxlan vni bound to VPN is : 1
```

# Display detailed VXLAN information about VNI 10.

```
<HUAWEI> display vxlan vni 10 verbose
 BD ID                :10
 State                :up
 Source               :10.1.1.2
 Source IPv6 Address :-
 UDP Port             :4789
 Peer List            :10.1.1.1 10.1.1.3
 IPv6 Peer List       :-
```

# Display detailed VXLAN information about VNI 20.

```
<HUAWEI> display vxlan vni 20 verbose
 VRF-ID               :5
```

Table 18-30 Description of the **display vxlan vni** command output

Item	Description
VNI	ID of a VNI. To configure or modify a VNI ID, run the **vxlan vni** command.
BD-ID (BD ID)	ID of the BD associated with a VNI. To configure or modify a BD ID, run the **bridge-domain (system view)** command.

Item	Description
State	VNI status: ● up ● down To ensure that the VNI status is up, the corresponding VXLAN tunnel must exist and be up for the VNI. If the VNI status is down, check whether **Source** and **Peer List** in this command output are the same as **Source** and **Destination** in the output of the **display vxlan tunnel** command. ● If they are different, no VXLAN tunnel exists for the specified VNI. Run the **source (NVE interface view)** or **vni head-end peer-list** command to change the source or destination IP address of the VXLAN tunnel to ensure that the corresponding VXLAN tunnel exists for the VNI. ● If they are the same, collect related configuration and contact technical support personnel.
Number of vxlan vni bound to BD is	Number of existing VNIs bound to BDs.
VRF-ID	VPN instance ID.
Number of vxlan vni bound to VPN is	Number of existing VNIs bound to VPN instances.
Source	IP address of the source VTEP. To configure the IP address of the source VTEP, run the **source** command in the NVE interface view.
Source IPv6 Address	IPv6 address of the source VTEP. To configure the IPv6 address of the source VTEP, run the **source** command in the NVE interface view.
UDP Port	Destination UDP port. The port number is fixed as 4789.
Peer List	IP address of the destination VTEP. To configure the IP address of the destination VTEP, run the **vni head-end peer-list** command.
IPv6 Peer List	IPv6 address of the destination VTEP. To configure the IPv6 address of the destination VTEP, run the **vni head-end peer-list** command.

18.1.39 encapsulation (Layer 2 sub-interface view)

Function

The **encapsulation** command configures the encapsulation mode of packets allowed to pass a Layer 2 sub-interface.

The **undo encapsulation** command deletes the encapsulation mode of packets allowed to pass a Layer 2 sub-interface.

By default, the encapsulation mode of packets allowed to pass a Layer 2 sub-interface is not configured.

Format

encapsulation { **dot1q vid** *low-pe-vid* [**to** *high-pe-vid*] | **default** | **untag** | **qinq vid** *low-vlan-vid* [**to** *high-vlan-vid*] **ce-vid** *low-ce-vid* [**to** *high-ce-vid*] }

undo encapsulation { **dot1q vid** *low-pe-vid* [**to** *high-pe-vid*] | **default** | **untag** | **qinq vid** *low-vlan-vid* [**to** *high-vlan-vid*] **ce-vid** *low-ce-vid* [**to** *high-ce-vid*] }

Parameters

Parameter	Description	Value
dot1q	Sets encapsulation mode of packets allowed to pass a Layer 2 sub-interface to Dot1q. This mode enables a Layer 2 sub-interface to receive packets with a VLAN tag.	-
vid *low-pe-vid* [**to** *high-pe-vid*]	Specifies the outer VLAN ID in packets allowed to pass a Layer 2 sub-interface in Dot1q encapsulation mode. ● *low-pe-vid*: specifies the start VLAN ID. ● *high-pe-vid*: specifies the end VLAN ID. *high-pe-vid* must be greater than or equal to *low-pe-vid*. *high-pe-vid* and *low-pe-vid* define a range of VLAN IDs. ● If you do not specify **to** *high-pe-vid*, *low-pe-vid* specifies the single VLAN ID carried in packets.	The value is an integer that ranges from 2 to 4094.
default	Sets the encapsulation mode of packets allowed to pass a Layer 2 sub-interface to default. This mode enables a Layer 2 sub-interface to receive all packets, regardless of whether they contain VLAN tags.	-

Parameter	Description	Value
untag	Sets the encapsulation mode of packets allowed to pass a Layer 2 sub-interface to untag. This mode enables a Layer 2 sub-interface to receive packets without VLAN tags.	-
qinq	Sets encapsulation mode of packets allowed to pass a Layer 2 sub-interface to QinQ. This mode enables a Layer 2 sub-interface to receive packets with double VLAN tags.	-
vid *low-vlan-vid* [**to** *high-vlan-vid*]	Specifies the outer VLAN ID in double-tagged packets allowed to pass a Layer 2 sub-interface in QinQ encapsulation mode. ● *low-vlan-vid*: specifies the start VLAN ID. ● *high-vlan-vid*: specifies the end VLAN ID. *high-vlan-vid* must be greater than or equal to *low-vlan-vid*. *high-vlan-vid* and *low-vlan-vid* define a range of VLAN IDs. ● If you do not specify **to** *high-vlan-vid*, *low-vlan-vid* specifies the single VLAN ID carried in packets.	The value is an integer that ranges from 2 to 4094.
ce-vid *low-ce-vid* [**to** *high-ce-vid*]	Specifies the inner VLAN ID in double-tagged packets allowed to pass a Layer 2 sub-interface in QinQ encapsulation mode. ● *low-ce-vid*: specifies the start VLAN ID in an inner tag. ● *high-ce-vid*: specifies the end VLAN ID in an inner tag. *high-ce-vid* must be greater than or equal to *low-ce-vid*. *high-ce-vid* and *low-ce-vid* define a range of VLAN IDs in an inner VLAN tag. ● If you do not specify **to** *high-ce-vid*, *low-ce-vid* specifies the single VLAN ID in the inner VLAN tag carried in packets.	The value is an integer that ranges from 1 to 4094.

Views

Layer 2 sub-interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a VXLAN network, a Layer 2 sub-interface functions as a VXLAN service access point to forward data packets in a BD.

Packets passing through a physical interface may contain one or two VLAN tags or no VLAN tag. After you run the **encapsulation** command in a Layer 2 sub-interface view to configure the encapsulation mode, the sub-interface can forward only specified types of packets.

Prerequisites

Run the command **interface** *interface-type interface-number.subnum* **mode l2** to create a VXLAN Layer 2 sub-interface

Precautions

When configuring an encapsulation mode on a Layer 2 sub-interface, pay attention to the following points:

- The VLAN ID in **dot1q** mode or outer VLAN ID in **qinq** mode cannot be the same as the allowed VLAN of the corresponding main interface or the global VLAN.
- On the same main interface, the VLAN ID in **dot1q** mode and the outer VLAN ID in **qinq** mode must be different.
- After NAC authentication is configured on the main interface, the traffic encapsulation type on a Layer 2 sub-interface cannot be set to **default**.
- When the encapsulation mode of a Layer 2 sub-interface is **default**, the corresponding main interface cannot be added to any VLAN, including VLAN 1.
- Before the encapsulation mode of a Layer 2 sub-interface is set to **default**, the main interface has only one sub-interface.
- After the encapsulation mode of a Layer 2 sub-interface is set to **default**, no other sub-interface can be created on the main interface.
- When the encapsulation mode of a Layer 2 sub-interface is set to **untag**, the corresponding main interface cannot be added to VLAN 1, and other sub-interfaces of the main interface cannot be set to **untag**.
- You can configure only one encapsulation mode for each Layer 2 sub-interface. If an encapsulation mode has been configured for a Layer 2 sub-interface, run the **undo encapsulation** command to delete the original mode before you configure another mode.
- Before configuring a VLAN segment on a Dot1q or QinQ Layer 2 sub-interface, you must run the **rewrite pop none** command.

Example

# Set the encapsulation mode of packets allowed to pass Layer 2 sub-interface GE0/0/1.1 to Dot1q and the outer VLAN ID in the packets to 10.

```
<HUAWEI> system-view
[HUAWEI] interface GigabitEthernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] port link-type hybrid
[HUAWEI-GigabitEthernet0/0/1] quit
[HUAWEI] interface gigabitethernet 0/0/1.1 mode l2
[HUAWEI-GigabitEthernet0/0/1.1] encapsulation dot1q vid 10
```

18.1.40 evpn binding vpn-instance

Function

The **evpn binding vpn-instance** command binds an EVPN instance to a BD.

The **undo evpn binding vpn-instance** command unbinds an EVPN instance from a BD.

By default, no EVPN instance is bound to a BD.

Format

evpn binding vpn-instance *vpn-instance*

undo evpn binding vpn-instance

Parameters

Parameter	Description	Value
vpn-instance	Specifies an EVPN instance name.	The value must be an existing EVPN instance name.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In the scenario where VXLAN is deployed through BGP EVPN, VXLAN BDs need to be bind to EVPN instances so that VXLAN tunnels can be dynamically created through BGP EVPN. You can run this command to bind an EVPN instance to a BD.

Precautions

Only one EVPN instance can be bound to a BD.

Example

# Bind EVPN instance **evn10** to BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] evpn binding vpn-instance evn10
```

18.1.41 evpn mac-route enable

Function

The **evpn mac-route enable** command enables the MAC route function for BGP EVPN.

The **undo evpn mac-route enable** command disables the MAC route function for BGP EVPN.

By default, the MAC route function is disabled for BGP EVPN.

Format

evpn mac-route enable

undo evpn mac-route enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

On a VXLAN network, if unidirectional isolation from the VXLAN access side to the tunnel side is enabled on a device using the **isolate remote enable** command, the device cannot learn MAC addresses through the received tunnel-side data packets. In this scenario, if communication between the access side and the tunnel side needs to be achieved, run the **evpn mac-route enable** command to enable the MAC route function for BGP EVPN. In this case, MAC address learning can be implemented and communication between the access side and the tunnel side can be achieved.

Example

# Enable the MAC route function for BGP EVPN on the device.

```
<HUAWEI> system-view
[HUAWEI] evpn mac-route enable
```

18.1.42 evpn vpn-instance bd-mode

Function

The **evpn vpn-instance bd-mode** command creates an EVPN instance, and displays the EVPN instance view.

The **undo evpn vpn-instance bd-mode** command deletes an EVPN instance.

By default, no EVPN instance is created on a device.

Format

evpn vpn-instance *vpn-instance* **bd-mode**

undo evpn vpn-instance *vpn-instance* **bd-mode**

Parameters

Parameter	Description	Value
vpn-instance	Specifies an EVPN instance name.	The value is a string of 1 to 31 case-sensitive characters, spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

In the scenario where VXLAN is deployed through BGP EVPN, EVPN instances need to be created so that routes can be advertised between EVPN instances through the BGP protocol. You can run this command to create EVPN instances.

Example

# Create EVPN instance **evn10**.

```
<HUAWEI> system-view
[HUAWEI] evpn vpn-instance evn10 bd-mode
[HUAWEI-evpn-instance-evn10]
```

18.1.43 export route-policy evpn

Function

The **export route-policy evpn** command associates the current VPN instance IPv4 or IPv6 address family with an export routing policy to filter EVPN routes to be advertised by the VPN instance IPv4 or IPv6 address family.

The **undo export route-policy evpn** command cancels the association between the IPv4 or IPv6 address family of a VPN instance and an export routing policy.

By default, the IPv4 or IPv6 address family of a VPN instance is not associated with any export routing policy.

Format

export route-policy *policy-name* **evpn**

undo export route-policy *policy-name* **evpn**

Parameters

Parameter	Description	Value
policy-name	Specifies the name of the export routing policy to be associated with the VPN instance IPv4 or IPv6 address family.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

VPN instance view, VPN instance IPv4 address family view, VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, the IPv4 or IPv6 address family of a VPN instance adds all EVPN-VPN targets in the export EVPN-VPN target list to EVPN routes to be advertised by the IPv4 or IPv6 address family of the VPN instance to the EVPN address family. To control route export more precisely, run the **export route-policy** *policy-name* **evpn**command, specify the export routing policy to filter the routes to be advertised, and set attributes for eligible routes.

Prerequisites

- An RD has been configured for the IPv4 or IPv6 address family of the VPN instance by running the **route-distinguisher** *route-distinguisher* command.
- If no associated routing policy exists, a routing policy needs to be configured by running the **route-policy** command.

Precautions

- The IPv4 or IPv6 address family of a VPN instance can be associated with only one export routing policy. If the **export route-policy evpn** command is run more than once, the latest configuration overrides the previous one.
- The export routing policy configured using the **export route-policy evpn** command does not affect the export routing policy associated with another VPN instance using the **export route-policy** *policy-name* command.

Example

# Associate the IPv4 address family of VPN instance **vrf1** with an export routing policy **policy-2** to filter EVPN routes to be advertised by the IPv4 address family of the VPN instance to the other VPN instances.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vrf1
[HUAWEI-vpn-instance-vrf1] ipv4-family
[HUAWEI-vpn-instance-vrf1-af-ipv4] route-distinguisher 100:1
[HUAWEI-vpn-instance-vrf1-af-ipv4] export route-policy policy-2 evpn
```

18.1.44 hub-mode enable (VXLAN)

Function

The **hub-mode enable** command sets the access side mode to hub.

The **undo hub-mode enable** command cancels the hub mode on the access side.

By default, the access side mode is not set to hub.

NOTE

Only the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, and S5731-H switches support this command.

Format

hub-mode enable

undo hub-mode enable

Parameters

None

Views

VLAN view, Layer 2 sub-interface view

Default Level

2: Configuration level

Usage Guidelines

On a VXLAN, users connected to the same BD can directly communicate with each other. If the **isolate enable** command is run in the BD to isolate users on the access side. To enable users connected to a BD through a VLAN or a Layer 2 sub-interface to communicate with other users in the BD, you can run the **hub-mode enable** command in the VLAN or Layer 2 sub-interface view to set the access side mode to hub.

Example

# Set the access side mode of VLAN 10 to hub.

```
<HUAWEI> system-view
[HUAWEI] vlan 10
[HUAWEI-vlan10] hub-mode enable
```

# Set the access-side mode of Layer 2 sub-interface GE0/0/1.1 to hub.

```
<HUAWEI> system-view
[HUAWEI] interface GigabitEthernet 0/0/1.1 mode l2
[HUAWEI-GigabitEthernet0/0/1.1] hub-mode enable
```

18.1.45 import route-policy evpn

Function

The **import route-policy evpn** command associates the IPv4 or IPv6 address family of a VPN instance with an import routing policy to filter EVPN routes imported.

The **undo import route-policy evpn** command cancels the association between the IPv4 or IPv6 address family of a VPN instance and an import routing policy.

By default, the IPv4 or IPv6 address family of a VPN instance is not associated with any import routing policy.

Format

import route-policy *policy-name* **evpn**

undo import route-policy *policy-name* **evpn**

Parameters

Parameter	Description	Value
policy-name	Specifies the name of the import routing policy to be associated with the IPv4 or IPv6 address family of a VPN instance.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

VPN instance view, VPN instance IPv4 address family view, VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, a VPN instance matches the export EVPN-VPN targets of the routes against the import EVPN-VPN targets of the IPv4 or IPv6 address family of the VPN instance to determine whether to import these IPv4 or IPv6 address family routes. To control EVPN route import more precisely, run the **import route-policy** *policy-name* **evpn** command to associate the VPN instance with an import routing policy and set attributes for eligible routes.

Prerequisites

- An RD has been configured for the IPv4 or IPv6 address family of the VPN instance by running the **route-distinguisher** *route-distinguisher* command.
- If no associated routing policy exists, a routing policy needs to be configured by running the **route-policy** command.

Precautions

- The IPv4 or IPv6 address family of a VPN instance can be associated with only one import routing policy. If the **import route-policy evpn** command is run several times, the latest configuration overrides the previous configurations.
- The import routing policy configured using the **import route-policy evpn** command does not affect the import routing policy applied to the VPN instance using the **import route-policy** *policy-name* command.

Example

# Associate the IPv4 address family of VPN instance **vrf1** with an import routing policy **policy-1** to filter EVPN routes to be advertised by other VPN instances to the IPv4 address family of the VPN instance.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vrf1
[HUAWEI-vpn-instance-vrf1] ipv4-family
[HUAWEI-vpn-instance-vrf1-af-ipv4] route-distinguisher 100:1
[HUAWEI-vpn-instance-vrf1-af-ipv4] import route-policy policy-1 evpn
```

18.1.46 interface nve

Function

The **interface nve** command creates a Network Virtualization Edge (NVE) interface and displays the NVE interface view.

The **undo interface nve** command deletes a specified NVE interface.

By default, no NVE interface is created.

Format

interface nve *nve-number*

undo interface nve *nve-number*

Parameters

Parameter	Description	Value
nve-number	Specifies the number of an NVE interface.	The value is 1.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To fully use advantages of server virtualization, you can deploy VXLAN to connect to multiple tenants. VXLAN tunnel information needs to be configured on an NVE interface, so the **interface nve** command needs to be executed to create the NVE interface.

Precautions

After a VXLAN tunnel is configured, running the **undo interface nve** command will delete the specified NVE interface and all the configuration of the NVE interface.

Example

# Create an NVE interface.

```
<HUAWEI> system-view
[HUAWEI] interface nve 1
```

18.1.47 interface vbdif

Function

The **interface vbdif** command creates a VBDIF interface and displays the VBDIF interface view, or displays the view of an existing VBDIF interface.

The **undo interface vbdif** command deletes a VBDIF interface.

By default, no VBDIF interface is created.

Format

interface vbdif *bd-id*

undo interface vbdif *bd-id*

Parameters

Parameter	Description	Value
bd-id	Specifies the ID of a BD.	The value is an integer that ranges from 1 to 16777215.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IP routes are required for communication between VXLAN networks on different network segments and between VXLAN and non-VXLAN networks.

To enable the communication, run the **interface vbdif** command to create a VBDIF interface for each BD and assign an IP address to the VBDIF interface. A VBDIF interface is a Layer 3 logical interface and can be configured with an IP address.

Prerequisites

- The specified BD has been created.
- If a BD contains user-side interfaces of the **default** subinterface type or the subinterface with **rewrite pop none**, the VBDIF interface cannot be configured in this BD.

Example

# Create a VBDIF interface for BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] quit
[HUAWEI] interface vbdif 10
```

18.1.48 ipv6 nd collect host enable

Function

The **ipv6 nd collect host enable** command enables BGP EVPN to collect host IPv6 information.

The **undo ipv6 nd collect host enable** command disables BGP EVPN from collecting host IPv6 information.

By default, BGP EVPN is disabled from collecting host IPv6 information on a device.

Format

ipv6 nd collect host enable

undo ipv6 nd collect host enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To enable a Layer 3 gateway to obtain the host information table, run the command in the VBDIF interface view so that EVPN can collect host IPv6 information.

In distributed VXLAN gateway deployment (BGP EVPN mode), if the VXLAN gateways advertise IRBv6 or ND routes to each other, run the command to advertise host IPv6 routes.

Prerequisites

IPv6 has been enabled in the VBDIF interface view using the **ipv6 enable** command.

Example

# On VBDIF 10, enable BGP EVPN to collect host IPv6 information.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ipv6 enable
[HUAWEI-Vbdif10] ipv6 nd collect host enable
```

18.1.49 ipv6 nd distribute-gateway enable

Function

The **ipv6 nd distribute-gateway enable** command enables the IPv6 distributed gateway function on an interface.

The **undo ipv6 nd distribute-gateway enable** command disables the IPv6 distributed gateway function on an interface.

By default, IPv6 distributed gateway is disabled on an interface.

Format

ipv6 nd distribute-gateway enable

undo ipv6 nd distribute-gateway enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If you want a distributed gateway to learn ND packets only from hosts on the user side, run the **ipv6 nd distribute-gateway enable** command. After distributed gateway is enabled:

- The gateway processes only ND packets received from user hosts and generates host routes.
- The gateway deletes the ND packets learned from the network side and the hosts routes for the ND packets.

Prerequisites

IPv6 has been enabled in the VBDIF interface view using the **ipv6 enable** command.

Configuration Impact

After distributed gateway is enabled:

- Tunnel-side static ND entries cannot be configured on the gateway.
- In a distribution scenario, if multiple gateways use the same IPv6 address, the gateways do not report the address conflict.

Example

# Enable IPv6 distributed gateway on VBDIF 10.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ipv6 enable
[HUAWEI-Vbdif10] ipv6 nd distribute-gateway enable
```

18.1.50 ipv6 nd multicast-suppress enable

Function

The **ipv6 nd multicast-suppress enable** command enables NS multicast suppression in a BD.

The **undo ipv6 nd multicast-suppress enable** command disables NS multicast suppression in a BD.

By default, NS multicast suppression is disabled in a BD.

Format

ipv6 nd multicast-suppress [**mismatch-discard**] **enable**

undo ipv6 nd multicast-suppress [**mismatch-discard**] **enable**

Parameters

Parameter	Description	Value
mismatch-discard	Indicates that the device drops NS packets that do not match any entries in the NS multicast suppression table.	-

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a network device receives an NS message to implement address resolution, the device forwards the NS message within its own BD. This NS message is received only by the nodes with the last 24 bits the same as the multicast address. If a device receives a large number of NS messages within a specified period, forwarding these NS messages uses excessive network resources and leads to network congestion and deteriorated network performance. As a result, user services are affected.

To resolve this problem, run the ipv6 nd multicast-suppress enable command to enable NS multicast suppression. With NS multicast suppression enabled, upon receipt of an NS message, a device checks whether the NS message contains information about the end user. If so, the device simply implements converts multicast streams to unicast streams. If not, the device forwards the NS message based on the original process. This reduces or suppresses message flooding occurred during address resolution.

Precautions

- When NS multicast suppression is configured in the VXLAN scenario with a large number of users and the NS packets received by the gateway exceeds the committed access rate (CAR), run the **car packet-type** *vpls-arp* **cir** *cir-value* [**cbs** *cbs-value*] command to adjust the CAR of the NS request packets. In the VXLAN scenario, the CAR for NS request packets is the same as that for vpls-arp packets. You can run the **display cpu-defend configuration** command to query the CAR. If the central processor CAR (CPCAR) is adjusted to an improper value, network services are affected. To adjust the CPCAR for packets, contact Huawei technical support.
- If the **ipv6 nd multicast-suppress enable** command is configured to enable NS multicast suppression in a BD and the **multicast-suppress enable** command is configured to enable multicast suppression in the BD, NS multicast suppression in the BD has a higher priority.

Example

# Enable NS multicast suppression in BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] ipv6 nd multicast-suppress enable
```

# Enable NS multicast suppression in BD 20 and configure the switch to discard ARP request packets if no matching entry is found in the NS multicast suppression table.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 20
[HUAWEI-bd20] ipv6 nd multicast-suppress mismatch-discard enable
```

18.1.51 ipv6 nd proxy enable

Function

The **ipv6 nd proxy enable** command enables routed proxy ND on an interface.

The **undo ipv6 nd proxy enable** command disables routed proxy ND on an interface.

By default, routed proxy ND is disabled on an interface.

Format

ipv6 nd proxy enable

undo ipv6 nd proxy enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If two hosts physically belong to different network segments and no gateway is configured, run the **ipv6 nd proxy enable** command on the VBDIF interface of the device that connects the two hosts. This command enables routed proxy ND, which parses IPv6 addresses between the hosts.

Prerequisites

IPv6 has been enabled in the VBDIF interface view using the **ipv6 enable** command.

Precautions

After routed proxy ND is configured on an interface, the device counts the number of received NS unicast packets twice. As a result, the number of NS packets displayed in the **display cpu-defend statistics packet-type nd** command output may be greater than the actual number of NS packets sent to the CPU.

Example

# Enable routed proxy ND on an interface.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ipv6 enable
[HUAWEI-Vbdif10] ipv6 nd proxy enable
```

18.1.52 ipv6 nd proxy local enable

Function

The **ipv6 nd proxy local enable** command enables intra-BD local proxy ND on an interface.

The **undo ipv6 nd proxy local enable** command disables intra-BD local proxy ND on an interface.

By default, intra-BD local proxy ND is disabled on an interface.

Format

ipv6 nd proxy local enable

undo ipv6 nd proxy local enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the overlay network of a VXLAN network is an IPv6 network, if two hosts belong to the same BD but they are isolated, to enable the hosts to communicate with each other, run the **ipv6 nd proxy local enable** command on the VBDIF interface. The command enables intra-BD local proxy ND on the interface.

Prerequisites

IPv6 has been enabled in the VBDIF interface view using the **ipv6 enable** command.

Precautions

After routed proxy ND is configured on an interface, the device counts the number of received NS unicast packets twice. As a result, the number of NS packets displayed in the **display cpu-defend statistics packet-type nd** command output may be greater than the actual number of NS packets sent to the CPU.

Example

# Enable intra-BD local proxy ND on an interface.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ipv6 enable
[HUAWEI-Vbdif10] ipv6 nd proxy local enable
```

18.1.53 ipv6 neighbor (VXLAN)

Function

The **ipv6 neighbor** command configures a static IPv6 neighbor entry on an interface of a VXLAN network.

The **undo ipv6 neighbor** command deletes a static IPv6 neighbor entry configured on an interface of a VXLAN network.

By default, no static IPv6 neighbor entry is configured on an interface of a VXLAN network.

Format

ipv6 neighbor *ipv6-address mac-address* **vni** *vni-id* { **source-ip** *ipv4-address1* **peer-ip** *ipv4-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* }

undo ipv6 neighbor *ipv6-address mac-address* **vni** *vni-id* { **source-ip** *ipv4-address1* **peer-ip** *ipv4-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* }

ipv6 neighbor *ipv6-address mac-address* { **vid** *vlan-id1* [**cevid** *vlan-id2*] } **interface** *interface-type interface-num.subnum*

undo ipv6 neighbor *ipv6-address mac-address* { **vid** *vlan-id1* [**cevid** *vlan-id2*] } **interface** *interface-type interface-num.subnum*

ipv6 neighbor *ipv6-address mac-address* **vid** *vlan-id3* **interface** *interface-type interface-num*

undo ipv6 neighbor *ipv6-address mac-address* **vid** *vlan-id3* **interface** *interface-type interface-num*

NOTE

Only the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, and S5731-H support the **source-ipv6** and **peer-ipv6** parameters.

Parameters

Parameter	Description	Value
ipv6-address	Specifies a destination IPv6 address.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.
mac-address	Specifies the destination MAC address mapping the destination IP address.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits.
vni *vni-id*	Specifies a VNI ID.	The value is an integer that ranges from 1 to 16777215.
source-ip *ipv4-address1*	Specifies the IP address of the source VTEP.	The value is in dotted decimal notation.
peer-ip *ipv4-address2*	Specifies the IP address of the destination VTEP.	The value is in dotted decimal notation.
source-ipv6 *ipv6-address1*	Specifies the IPv6 address of the source VTEP.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
peer-ipv6 *ipv6-address2*	Specifies the IPv6 address of the destination VTEP.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.
vid *vlan-id1*	Specifies the outer VLAN ID in the packet received by a interface.	The value is an integer that ranges from 1 to 4094.
cevid *vlan-id2*	Specifies the inner VLAN ID in the packet received by a interface.	The value is an integer that ranges from 1 to 4094.
interface *interface-type interface-number.subnum*	Specifies an L2 sub-interface.	-
vid *vlan-id3*	Specifies the VLAN ID in the packet received by a interface.	The value is an integer that ranges from 1 to 4094.
interface *interface-type interface-number*	Specifies an interface.	-

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a VXLAN, if a Layer 3 gateway on a VXLAN is not enabled to send ND protocol packets, run the ipv6 neighbor command to configure a static IPv6 neighbor entry. To filter out invalid ND protocol packets, you can also run this command to bind the destination IPv6 addresses of these packets to nonexistent MAC addresses. To configure a static IPv6 neighbor entry for a VXLAN tunnel:

To configure static IPv6 neighbor entries on the VXLAN tunnel side, run the **ipv6 neighbor** *ipv6-address mac-address* **vni** *vni-id* { **source-ip** *ipv4-address1* **peer-ip** *ipv4-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* } command. If the current device has an access-side L2 sub-interface added to a BD, run the **ipv6 neighbor** *ipv6-address mac-address* { **vid** *vlan-id1* [**cevid** *vlan-id2*] } **interface** *interface-type interface-num.subnum* command to configure a static

IPv6 neighbor entry for the sub-interface. If the current device has an access-side interface added to a BD based a VLAN, run the **ipv6 neighbor** *ipv6-address mac-address* **vid** *vlan-id3* **interface** *interface-type interface-num* command to configure a static IPv6 neighbor entry for the interface.

Prerequisites

The IPv6 function has been enabled on a VBDIF interface using the **ipv6 enable** command.

In addition, before configuring a static IPv6 neighbor entry for the VXLAN tunnel side, ensure that a VNI ID has been associated with a BD using the **vxlan vni** command in the BD view on a VBDIF interface.

Precautions

- If the IPv6 address or MAC address specified in the **ipv6 neighbor** command is incorrect, communication with this neighbor fails.
- When the VXLAN tunnel is created dynamically, the device does not support to configure a static IPv6 neighbor entry on a VXLAN tunnel-side interface by the command **ipv6 neighbor** *ipv6-address mac-address* **vni** *vni-id* { **source-ip** *ipv4-address1* **peer-ip** *ipv4-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* }.

Example

# Configure a static IPv6 neighbor entry for a VXLAN tunnel-side interface, and set the IP address to fc00:1::10 and MAC address to aaaa-fccc-1212.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] vxlan vni 5000
[HUAWEI-bd10] quit
[HUAWEI] interface nve 1
[HUAWEI-Nve1] source 10.1.1.1
[HUAWEI-Nve1] vni 5000 head-end peer-list 10.2.2.2
[HUAWEI-Nve1] quit
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ipv6 enable
[HUAWEI-Vbdif10] ipv6 address fc00:1::1/64
[HUAWEI-Vbdif10] ipv6 neighbor fc00:1::10 aaaa-fccc-1212 vni 5000 source-ip 10.1.1.1 peer-ip 10.2.2.2
[HUAWEI-Vbdif10] quit
```

# Configure a static IPv6 neighbor entry in which the outbound interface is GE0/0/1, and set the IP address to fc00:1::11 and MAC address to aaaa-fccc-1212.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] vlan 10
[HUAWEI-vlan10] quit
[HUAWEI] interface GigabitEthernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] port link-type trunk
[HUAWEI-GigabitEthernet0/0/1] port trunk allow-pass vlan 10
[HUAWEI-GigabitEthernet0/0/1] quit
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] l2 binding vlan 10
[HUAWEI-bd10] quit
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] ipv6 enable
[HUAWEI-Vbdif10] ipv6 address fc00:1::1/64
[HUAWEI-Vbdif10] ipv6 neighbor fc00:1::11 aaaa-fccc-1212 vid 10 interface GigabitEthernet 0/0/1
[HUAWEI-Vbdif10] quit
```

18.1.54 isolate enable (BD View)

Function

The **isolate enable** command enables isolation of users connected to an access-side BD.

The **undo isolate enable** command disables isolation of users connected to an access-side BD.

By default, isolation of users connected to an access-side BD is disabled.

NOTE

Only the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, and S5731-H switches support this command.

Format

isolate enable

undo isolate enable

Parameters

None

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

In a VXLAN, users connected to the same BD can communicate. To isolate users connected to a BD, you can run the **isolate enable** command.

Example

# Enables isolation of users connected to BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] isolate enable
```

18.1.55 isolate remote enable (BD view)

Function

The **isolate remote enable** command configures unidirectional isolation from the access side to the tunnel side in a BD.

The **undo isolate remote enable** command disables unidirectional isolation from the access side to the tunnel side in a BD.

By default, unidirectional isolation from the access side to the tunnel side is disabled in a BD.

NOTE

Only the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, and S5731-H support this command.

Format

isolate remote enable

undo isolate remote enable

Parameters

None

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

On a VXLAN network, users in the same BD can directly communicate with each other. To isolate unidirectional traffic from the access side to the tunnel side in a BD, run this command in the BD view.

Example

# Configure isolation from the access side to the tunnel side in BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] isolate remote enable
```

18.1.56 l2 binding vlan

Function

The **l2 binding vlan** command associates a specified VLAN with a BD.

The **undo l2 binding vlan** command restores the default settings.

By default, a VLAN is not associated with a BD.

Format

l2 binding vlan *vlan-id*

undo l2 binding vlan *vlan-id*

Parameters

Parameter	Description	Value
vlan-id	Specifies a VLAN ID.	The value is an integer that ranges from 1 to 4094. Currently, VLAN 1 cannot be associated with a BD.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

VXLAN needs to be deployed on a downlink interface to provide access services and an uplink interface to establish a VXLAN tunnel.

On the access side, two methods are available for creating a large Layer 2 BD.

- Based on VLAN: You can associate one or multiple VLANs with a BD to add users in these VLANs to the BD. This VLAN-based mode implements larger-granularity control, but is easy to configure. It applies to VXLAN deployment on a live network.
- Based on encapsulation mode: The device sends packets of different encapsulation modes to different Layer 2 sub-interfaces based on the VLAN tags contained in the packets. You can bind a Layer 2 sub-interface to a BD to add specified users to the BD. This mode implements refined and flexible control but requires more complex configuration. It applies to VXLAN deployment on a new network.

After you run this command to associate specified VLANs with a BD, different VLANs associated with the same BD form a large Layer 2 network. Users belong to these VLANs can communicate at Layer 2 through VXLAN tunnels.

Prerequisites

The VLAN to be bound to the BD has been created using the **vlan** command.

Precautions

- One VLAN can be associated with only one BD, but one BD can be associated with multiple VLANs.
- After a global VLAN is associated with a BD, you need to add corresponding interfaces to the VLAN.

- If a VLAN is configured as a voice VLAN on the S6735-S, S6720-EI, S6720S-EI, the VLAN cannot be associated with a BD.
- In NAC authentication scenarios, if there are online users in a VLAN, running the **undo l2 binding vlan** command to unbind the VLAN from a BD makes the users go offline.
- If a VLAN is an ISP VLAN authorized to users and users exist in the VLAN on the device, the VLAN cannot be associated with a BD.
- If a VLAN is used as the management VLAN of a Fit AP, it is not recommended that the VLAN be associated with a BD.

Example

# Associate VLAN 10 with BD 10.

```
<HUAWEI> system-view
[HUAWEI] vlan 10
[HUAWEI-vlan10] quit
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] l2 binding vlan 10
```

18.1.57 l2vpn-family evpn

Function

The **l2vpn-family evpn** command enables and displays the BGP-EVPN address family view.

The **undo l2vpn-family evpn** command disables the BGP-EVPN address family view.

By default, the BGP-EVPN address family view is disabled.

Format

l2vpn-family evpn

undo l2vpn-family evpn

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Before performing configurations in the BGP-EVPN address family view, run the **l2vpn-family evpn** command in the BGP view to enable and display the BGP-EVPN address family view.

Example

# Enable and display the BGP-EVPN address family view.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn]
```

18.1.58 l3-interface virtual-mac compression

Function

The **l3-interface virtual-mac compression disable** command disables the virtual MAC address compression function. After the virtual MAC address compression function is disabled for a VBDIF interface on the S6720-EI, S6735-S and S6720S-EI, the device can still forward packets with destination MAC addresses in the range of 0000-5e00-0100 to 0000-5e00-01ff at Layer 2 when a MAC address is configured for the VBDIF interface.

The **l3-interface virtual-mac compression enable** command enables the virtual MAC address compression function. After a MAC address is configured for a VBDIF interface on the S6720-EI, S6735-S and S6720S-EI, the device cannot forward packets with destination MAC addresses in the range of 0000-5e00-0100 to 0000-5e00-01ff at Layer 2.

The **undo l3-interface virtual-mac compression disable** command restores the default setting of the virtual MAC address compression function.

By default, the virtual MAC address compression function is enabled.

NOTE

This command is supported only by the S6720-EI, S6735-S and S6720S-EI.

Format

l3-interface virtual-mac compression disable

l3-interface virtual-mac compression enable

undo l3-interface virtual-mac compression disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

In a distributed gateway scenario where multiple Layer 3 gateways need to be simulated into one gateway, you need to run the **mac-address** command to configure the same MAC address for VBDIF interfaces on different Layer 3 gateways. In this way, terminals connect to the same gateway, ensuring normal traffic forwarding and VM migration. By default, the virtual MAC address compression function is enabled on the S6720-EI, S6735-S and S6720S-EI. After a MAC address is configured for a VBDIF interface, the device cannot forward packets with destination MAC addresses in the range of 0000-5e00-0100 to 0000-5e00-01ff at Layer 2. To enable the device to forward packets with destination MAC addresses in the range of 0000-5e00-0100 to 0000-5e00-01ff at Layer 2, run the **l3-interface virtual-mac compression disable** command to disable the virtual MAC address compression function.

Example

# Disable the virtual MAC address compression function on the device.

```
<HUAWEI> system-view
[HUAWEI] l3-interface virtual-mac compression disable
```

18.1.59 mac rib-only

Function

The **mac rib-only** command configures a device not to deliver MAC address entries after receiving remote MAC routes.

The **undo mac rib-only** command configures a device to deliver MAC address entries after receiving remote MAC routes.

By default, a device is triggered to deliver MAC address entries after receiving remote MAC routes.

Format

mac rib-only

undo mac rib-only

Parameters

None

Views

EVPN instance view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

For a Layer 3 VXLAN gateway, if Layer 2 unicast traffic forwarding is not involved, you can run the **mac rib-only** command to configure a device not to deliver local MAC address entries after receiving VNI ID-based MAC routes advertised by the EVPN peer, saving forwarding entry resources.

Precautions

When this command is configured in an EVPN instance bound to a BD, the BD does not deliver MAC address entries after receiving remote MAC routes. Therefore, in this scenario, MAC address flapping in this BD does not depend on the MAC routes advertised by the remote BGP peer.

Example

# Configure a device not to deliver MAC address entries after receiving remote MAC routes in evpn10.

```
<HUAWEI> system-view
[HUAWEI] evpn vpn-instance evpn10 bd-mode
[HUAWEI-evpn-instance-evpn10] mac rib-only
```

18.1.60 mac-address update host enable

Function

The **mac-address update host enable** command enables user host information update triggered by a MAC address entry change.

The **undo mac-address update host enable** command disables user host information update triggered by a MAC address entry change.

By default, user host information update triggered by a MAC address entry change is disabled.

Format

mac-address update host enable

undo mac-address update host enable

Parameters

None

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If user hosts migrate between different gateways in a distributed VXLAN gateway scenario and do not send gratuitous ARP or NA packets after the migration, you can run the **mac-address update host enable** command to configure the switch to be triggered to send ARP or NS requests to update user host information by changes in the user hosts' MAC address entries on the switch before and after the migration. This configuration ensures that user hosts can successfully go online after the migration.

Precautions

When the distributed gateway is the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, or S5731-H, the switch cannot learn a user host's MAC address from Layer 3 packets (the destination MAC address of the packets is the MAC address of the switch) sent by the user host. Therefore, the switch cannot be triggered to update user host information by Layer 3 traffic send by user hosts in this scenario. In this scenario, the switch can only be triggered to update user host information by Layer 2 traffic (the destination MAC address of the packets is not the MAC address of the switch) send by user hosts.

Example

# Enable user host information update triggered by a MAC address entry change on VBDIF 10.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] mac-address update host enable
[HUAWEI-Vbdif10] quit
```

18.1.61 mac-address static bridge-domain

Function

The **mac-address static bridge-domain** command configures a static MAC address entry on a VXLAN access-side interface.

The **undo mac-address static bridge-domain** command deletes a static MAC address entry on a VXLAN access-side interface.

By default, no static MAC address entry is configured on a VXLAN access-side interface.

Format

mac-address static *mac-address interface-type interface-number.subnum* **bridge-domain** *bd-id* { **default** | **untag** | **vid** *vlan-id1* [**ce-vid** *vlan-id2*] }

undo mac-address static *mac-address interface-type interface-number.subnum* **bridge-domain** *bd-id* { **default** | **untag** | **vid** *vlan-id1* [**ce-vid** *vlan-id2*] }

mac-address static *mac-address interface-type interface-number* **bridge-domain** *bd-id* **vid** *vlan-id3*

undo mac-address static *mac-address interface-type interface-number* **bridge-domain** *bd-id* **vid** *vlan-id3*

Parameters

Parameter	Description	Value
mac-address	Specifies the MAC address in the static MAC address entry.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits. The MAC address cannot be FFFF-FFFF-FFFF, 0000-0000-0000, or a multicast MAC address.
interface-type interface-number.subnum	Specifies that the outbound interface in the static MAC address entry is a Layer 2 sub-interface.	-
bd-id	Specifies the BD to which the outbound interface belongs.	The value is an integer that ranges from 1 to 16777215.
default	Specifies that the outbound interface allows packets of the **default** type to pass.	-
untag	Specifies that the outbound interface allows packets of the **untag** type to pass.	-
vid *vlan-id1*	Specifies the outer VLAN ID in the packets allowed to pass the outbound interface.	The value is an integer that ranges from 1 to 4094.
ce-vid *vlan-id2*	Specifies the inner VLAN ID in the packets allowed to pass the outbound interface.	The value is an integer that ranges from 1 to 4094.
interface-type interface-number	Specifies that the outbound interface in the static MAC address entry is a specified interface.	-
vid *vlan-id3*	Specifies the ID of the VLAN to which the outbound interface belongs.	The value is an integer that ranges from 1 to 4094.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the device creates a MAC address table by learning source MAC addresses, the device cannot distinguish packets from authorized and unauthorized users. This threatens network security. If an unauthorized user uses the MAC address of an authorized user as the source MAC address of attack packets and connects to another interface of the device, the device learns an incorrect MAC address entry. The device incorrectly forwards the packets to the unauthorized user. Actually, the packets should be forwarded to the authorized user. You can run the **mac-address static bridge-domain** command to add a static MAC address entry to the MAC address table on the VXLAN access side. The static MAC address entry binds the MAC address to a specified interface, which prevents unauthorized users from intercepting data of authorized users. In addition, a manually configured static MAC address entry improves the unicast packet forwarding efficiency and saves bandwidth.

Prerequisites

- The interface has been added to a BD.

Example

# Configure a static MAC address entry on a VXLAN access-side interface. In the entry, the destination MAC address is 00e0-fc12-3456 and the flow encapsulation type of the outbound interface is **dot1q**.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] port link-type hybrid
[HUAWEI-GigabitEthernet0/0/1] quit
[HUAWEI] bridge-domain 20
[HUAWEI-bd20] quit
[HUAWEI] interface gigabitethernet 0/0/1.1 mode l2
[HUAWEI-GigabitEthernet0/0/1.1] encapsulation dot1q vid 6
[HUAWEI-GigabitEthernet0/0/1.1] bridge-domain 20
[HUAWEI-GigabitEthernet0/0/1.1] quit
[HUAWEI] mac-address static 00e0-fc12-3456 GigabitEthernet 0/0/1.1 bridge-domain 20 vid 6
```

# Configure a static MAC address entry on the VXLAN access-side interface. In the entry, the destination MAC address is 00e0-fc12-3457 and the outbound interface is added to a BD by the VLAN.

```
<HUAWEI> system-view
[HUAWEI] vlan 8
[HUAWEI-vlan8] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] port link-type hybrid
[HUAWEI-GigabitEthernet0/0/1] port hybrid tagged vlan 8
[HUAWEI-GigabitEthernet0/0/1] quit
[HUAWEI] bridge-domain 30
```

```
[HUAWEI-bd30] l2 binding vlan 8
[HUAWEI-bd30] quit
[HUAWEI] mac-address static 00e0-fc12-3457 GigabitEthernet 0/0/1 bridge-domain 30 vid 8
```

18.1.62 mac-address static bridge-domain vni

Function

The **mac-address static bridge-domain vni** command configures a static MAC address entry on a VXLAN tunnel-side interface.

The **undo mac-address static bridge-domain vni** command deletes a static MAC address entry on a VXLAN tunnel-side interface.

By default, no static MAC address entry is configured on a VXLAN tunnel-side interface.

Format

mac-address static *mac-address* **bridge-domain** *bd-id* { { **source** *ip-address1* **peer** *ip-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* } } **vni** *vni-id*

undo mac-address static *mac-address* **bridge-domain** *bd-id* { { **source** *ip-address1* **peer** *ip-address2* } | { **source-ipv6** *ipv6-address1* **peer-ipv6** *ipv6-address2* } } **vni** *vni-id*

Parameters

Parameter	Description	Value
mac-address	Specifies the MAC address in the static MAC address entry.	The value is in H-H-H format. An H is a hexadecimal number of 1 to 4 digits. The MAC address cannot be FFFF-FFFF-FFFF, 0000-0000-0000, or a multicast MAC address.
bd-id	Specifies the BD to which the outbound interface belongs.	The value is an integer that ranges from 1 to 16777215.
source *ip-address1*	Specifies the source IP address of the VXLAN tunnel.	The value is in dotted decimal notation.
peer *ip-address2*	Specifies the remote IP address of the VXLAN tunnel.	The value is in dotted decimal notation.

Parameter	Description	Value
source-ipv6 *ipv6-address1*	Specifies the source IPv6 address of the VXLAN tunnel.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.
peer-ipv6 *ipv6-address2*	Specifies the remote IPv6 address of the VXLAN tunnel.	The value consists of 128 bits, which are classified into 8 groups. Each group contains 4 hexadecimal numbers in the format of X:X:X:X:X:X:X:X.
vni-id	Specifies the ID of a VXLAN tunnel.	The value is an integer that ranges from 1 to 16777215.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the device creates a MAC address table by learning source MAC addresses, the device cannot distinguish packets from authorized and unauthorized users. This threatens network security. If an unauthorized user uses the MAC address of an authorized user as the source MAC address of attack packets and connects to another interface of the device, the device learns an incorrect MAC address entry. The device incorrectly forwards the packets to the unauthorized user. Actually, the packets should be forwarded to the authorized user. You can run the **mac-address static bridge-domain vni** command to add a static MAC address entry to the MAC address table on the VXLAN tunnel side. The static MAC address entry binds the MAC address to a specified interface, which prevents unauthorized users from intercepting data of authorized users. In addition, a manually configured static MAC address entry improves the unicast packet forwarding efficiency and saves bandwidth.

Prerequisites

- A VXLAN tunnel has been created.
- When the VXLAN tunnel is created dynamically, the device does not support to configure a static MAC address entry on a VXLAN tunnel-side interface.

Precautions

If a static MAC address entry is configured on a VXLAN tunnel-side interface and the VXLAN tunnel is Down, the static MAC address entry is not displayed in the output of the **display mac-address** command. When the VXLAN tunnel is Up, the static MAC address entry is displayed in the output of the **display mac-address** command.

Example

# On a VXLAN tunnel-side interface, configure a static MAC address entry with the destination MAC address 00e0-fc12-3456.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 20
[HUAWEI-bd20] vxlan vni 2000
[HUAWEI-bd20] quit
[HUAWEI] interface nve 1
[HUAWEI-Nve1] source 10.1.1.2
[HUAWEI-Nve1] vni 2000 head-end peer-list 10.1.2.2
[HUAWEI-Nve1] quit
[HUAWEI] mac-address static 00e0-fc12-3456 bridge-domain 20 source 10.1.1.2 peer 10.1.2.2 vni 2000
```

18.1.63 mac-address (VBDIF interface view)

Function

The **mac-address** command configures the MAC address of a VBDIF interface.

The **undo mac-address** command restores the default MAC address of a VBDIF interface.

By default, the MAC address of a VBDIF interface is the system MAC address.

Format

mac-address *mac-address*

undo mac-address

Parameters

Parameter	Description	Value
mac-address	Specifies the MAC address of a VBDIF interface.	The value is in the format of H-H-H. Each H is a 4-digit hexadecimal number, such as 00e0 or fc01. If you enter fewer than four alphanumeric characters, 0s are added before the input digits. For example, if e0 is entered, 00e0 is specified. A MAC address cannot be set to all 0s or all 1s. The MAC address of a VBDIF interface ranges from 0000-5e00-0100 to 0000-5e00-01ff.

Views

VBDIF interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a distributed gateway scenario, multiple Layer 3 gateways are simulated into one gateway. You need to run the **mac-address** command to configure the same MAC address for the VBDIF interfaces on these Layer 3 gateways to ensure normal traffic forwarding and VM migration. Using the same MAC address, the Layer 3 gateways act like one gateway for terminals.

Precautions

- After the MAC address of a VBDIF interface is changed, the switch proactively sends gratuitous ARP packets to update the mapping relationship between MAC addresses and ports.
- The MAC address configured for a VBDIF interface cannot be the same as the virtual MAC address of a VRRP group on the switch. You can run the **display vrrp** and **display vrrpv6** commands to view virtual MAC addresses of VRRP groups configured on the switch.

Example

# Set the MAC address of VBDIF interface 10 to 0000-5e00-0101.

```
<HUAWEI> system-view
[HUAWEI] interface vbdif 10
[HUAWEI-Vbdif10] mac-address 0000-5e00-0101
```

18.1.64 mac-route no-advertise

Function

The **mac-route no-advertise** command configures a device not to advertise local MAC routes.

The **undo mac-route no-advertise** command configures a device to advertise local MAC routes.

By default, a device advertises local MAC routes.

Format

mac-route no-advertise

undo mac-route no-advertise

Parameters

None

Views

EVPN instance view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

For a Layer 3 VXLAN gateway, if Layer 2 unicast traffic forwarding is not involved, you can run the **mac-route no-advertise** command to configure the gateway not to advertise local MAC routes. In this case, the remote gateway cannot receive MAC routes, saving the memory for learning unnecessary MAC routes.

Precautions

When this command is configured in an EVPN instance bound to a BD, the corresponding MAC routes in the BD are no longer externally advertised. Therefore, in this scenario, MAC address flapping on the device where the corresponding BGP peer resides does not depend on the MAC routes advertised by the local device.

Example

# Configure a device not to advertise local MAC routes in evpn10.

```
<HUAWEI> system-view
[HUAWEI] evpn vpn-instance evpn10 bd-mode
[HUAWEI-evpn-instance-evpn10] mac-route no-advertise
```

18.1.65 multicast-suppression (BD view)

Function

The **multicast-suppression** command enables multicast traffic suppression in a BD.

The **undo multicast-suppression** command disables multicast traffic suppression in a BD.

By default, multicast traffic suppression is disabled in a BD.

Format

multicast-suppression cir *cir-value* [**cbs** *cbs-value*]

undo multicast-suppression

Parameters

Parameter	Description	Value
cir *cir-value*	Specifies the committed information rate (CIR), which is the allowed rate at which traffic can pass through.	The value is an integer that ranges from 0 to 10000000, in kbit/s.

Parameter	Description	Value
cbs *cbs-value*	Specifies the committed burst size (CBS), which is the maximum size of traffic that can pass through.	The value is an integer that ranges from 10000 to 4294967295, in bytes. If the **cbs** is not set, the default *cbs-value* is 125 times the *cir-value*.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

When a large number of multicast packets are transmitted on a network, a lot of network resources are occupied, and services on the network are affected. You can run the **multicast-suppression** command to enable multicast traffic suppression in a BD and configure the maximum number of multicast packets that can pass through a BD. When the multicast traffic volume exceeds the specified threshold, the system discards excess multicast packets.

Example

# Set the CIR value for multicast traffic in BD 10 to 100 kbit/s.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] multicast-suppression cir 100
```

18.1.66 peer advertise

Function

The **peer advertise** command configures a device to advertise ARP, integrated routing and bridging (IRB), ND, or IRBv6 routes to its BGP EVPN peers.

The **undo peer advertise** command restores the default configurations.

By default, a device cannot advertise ARP, IRB, ND, or IRBv6 routes to its BGP EVPN peers.

Format

peer { *ipv4-address* | *group-name* } **advertise** { **arp** | **irb** }

undo peer { *ipv4-address* | *group-name* } **advertise** { **arp** | **irb** }

peer { *ipv4-address* | *group-name* } **advertise** { **nd** | **irbv6** }

undo peer { *ipv4-address* | *group-name* } **advertise** { **nd** | **irbv6** }

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a peer.	The value is in dotted decimal notation.
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
arp	Specifies the ARP routes to be advertised.	-
irb	Specifies the IRB routes to be advertised.	-
nd	Specifies the ND routes to be advertised.	-
irbv6	Specifies the IRBv6 routes to be advertised.	-

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the overlay network is an IPv4 network, to allow a device to advertise ARP or IRB routes to its BGP EVPN peers, run the **peer advertise** command. This command allows VTEPs to implements ARP broadcast suppression on networks. If you specify **irb**, VTEPs can also transmit host routes.

When the overlay network is an IPv6 network, to allow a device to advertise ND or IRBv6 routes to its BGP EVPN peers, run the **peer advertise** command. This command allows VTEPs to implements NS broadcast suppression on networks. If you specify **irbv6**, VTEPs can also transmit host routes.

Precautions

You cannot specify both **arp** and **irb** in the same BGP-EVPN address family view.

You cannot specify both **nd** and **irbv6** in the same BGP-EVPN address family view.

Example

# Configure a device to advertise ARP routes to its BGP EVPN peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] peer 1.1.1.1 advertise arp
```

18.1.67 peer advertise route-reoriginated

Function

The **peer advertise route-reoriginated** command enables a device to advertise IP prefix routes re-encapsulated by the EVPN address family to a BGP EVPN peer.

The **undo peer advertise route-reoriginated** command restores the default configuration.

By default, the local device does not advertise IP prefix routes re-encapsulated by the EVPN address family to a BGP EVPN peer.

Format

peer { *ipv4-address* | *group-name* } **advertise route-reoriginated evpn ip**

undo peer { *ipv4-address* | *group-name* } **advertise route-reoriginated evpn ip**

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a BGP EVPN peer.	The value is in dotted decimal notation.
group-name	Specifies the name of a BGP EVPN peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
evpn	Re-encapsulates the received EVPN routes.	-
ip	Re-encapsulates the IP prefix routes in the received EVPN routes.	-

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In three-segment VXLAN scenarios, to re-encapsulate EVPN routes received from a DC into IP prefix routes and then send them to the BGP EVPN peer in another DC, run the **peer advertise route-reoriginated** command. This configuration allows communication between VMs in different DCs.

During EVPN route re-encapsulation, a border leaf node receives an EVPN route from a DC and changes the route next hop to the VTEP address of the leaf node. Additionally, the leaf node replaces the source MAC address in the gateway MAC attribute of the host route with its own MAC address and the L3VNI with that of an L3VPN instance.

Prerequisites

The device has been enabled to add a regeneration flag to the routes received from BGP EVPN peers using the **peer** { *ipv4-address* | *group-name* } **import reoriginate** command.

Example

# Enable a device to advertise regenerated IP prefix routes to a BGP EVPN peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 2.2.2.2 as-number 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] peer 2.2.2.2 enable
[HUAWEI-bgp-af-evpn] peer 2.2.2.2 import reoriginate
[HUAWEI-bgp-af-evpn] peer 2.2.2.2 advertise route-reoriginated evpn ip
```

18.1.68 peer enable (BGP-EVPN address family view)

Function

The **peer enable** command enables a device to exchange route information with a peer or peer group in the address family view.

The **undo peer enable** command disables a device from exchanging route information with a peer or peer group.

By default, only the peers of the BGP-IPv4 unicast address family are automatically enabled.

Format

peer { *group-name* | *ipv4-address* } **enable**

undo peer { *group-name* | *ipv4-address* } **enable**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	The value is in dotted decimal notation.

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, only the peers of the BGP-IPv4 unicast address family are automatically enabled. After the **peer as-number** command is run in the BGP view, the system automatically runs the **peer enable** command to enable a peer. In the BGP EVPN address family view, the **peer enable** command must be manually run to enable a peer.

Prerequisites

The **peer as-number** command has been run to create a peer or peer group.

Example

# Configure a peer and enable the peer in the BGP-EVPN address family view.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 1.1.1.2 as-number 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] peer 1.1.1.2 enable
```

18.1.69 peer import reoriginate

Function

The **peer import reoriginate** command enables a device to add a regeneration flag to the routes received from BGP EVPN peers.

The **undo peer import reoriginate** command restores the default configuration.

By default, a device does not add a regeneration flag to the routes received from BGP EVPN peers.

Format

peer { *ipv4-address* | *group-name* } **import reoriginate**

undo peer { *ipv4-address* | *group-name* } **import reoriginate**

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a BGP EVPN peer.	The value is in dotted decimal notation.
group-name	Specifies the name of a BGP EVPN peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a three-segment VXLAN scenario, by default, an edge node that connects to a carrier network does not re-encapsulate the routes received from BGP EVPN peers, causing the EVPN routes to be terminated on the edge node. As a result, the EVPN routes from one DC cannot be advertised to the BGP EVPN peers of another DC. To address this problem, run the **peer import reoriginate** command to enable the edge node to add a regeneration flag to the routes received from BGP EVPN peers. The edge node then re-encapsulates the EVPN routes received from one DC before sending them to another DC for inter-DC VM communication.

Prerequisites

A device has been enabled to exchange routes with a specified peer or peer group using the **peer** { *group-name* | *ipv4-address* } **enable** command.

Example

# Enable a device to add a regeneration flag to the routes received from BGP EVPN peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 1.1.1.1 as-number 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] peer 1.1.1.1 enable
[HUAWEI-bgp-af-evpn] peer 1.1.1.1 import reoriginate
```

18.1.70 peer mac-limit

Function

The **peer mac-limit** command limits the number of MAC advertisement routes received from a peer.

The **undo peer mac-limit** command restores the default configuration.

By default, the number of MAC advertisement routes received from a peer is not limited.

Format

peer { *group-name* | *ipv4-address* } **mac-limit** *number* [**idle-forever** | **idle-timeout** *times*]

undo peer { *group-name* | *ipv4-address* } **mac-limit**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	The value is in dotted decimal notation.
number	Specifies the maximum number of MAC advertisement routes received from a peer.	The value is an integer, the value varies according to different devices.
idle-forever	Indicates that a terminated connection is not automatically re-established after the number of routes exceeds the maximum limit.	-
idle-timeout *times*	Specifies a timer for automatically re-establishing a terminated connection after the number of routes exceeds the maximum limit. No connection will be automatically re-established before the timer expires.	The value is an integer that ranges from 1 to 1200, in minutes.

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A BGP-EVPN instance may import many invalid MAC advertisement routes from peers. In this case, To prevent these routes from occupying a large proportion of MAC advertisement routes, run the **peer mac-limit** command to configure the maximum number of MAC advertisement routes allowed to be received from each peer.

Precautions

- After this command is run, excessive routes in the BGP-EVPN address family view may be discarded.
- When the number of MAC advertisement routes exceeds the maximum limit, run the **undo peer mac-limit** command. The device then receives routes from each peer and adds them to the EVPN routing table.

Example

# Configure the maximum number of MAC advertisement routes received by a device from peer 2.2.2.2 as 1000.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] peer 2.2.2.2 mac-limit 1000
```

18.1.71 peer route-policy (BGP-EVPN address family view)

Function

The **peer route-policy** command specifies a routing policy for filtering routes received from an EVPN peer or peer group or routes to be advertised to an EVPN peer or peer group.

The **undo peer route-policy** command deletes a specified routing policy.

By default, no routing policy is used for filtering routes received from an EVPN peer or peer group or routes to be advertised to an EVPN peer or peer group.

Format

peer { *group-name* | *ipv4-address* } **route-policy** *route-policy-name* { **import** | **export** }

undo peer { *group-name* | *ipv4-address* } **route-policy** *route-policy-name* { **import** | **export** }

Parameters

Parameter	Description	Value
group-name	Specifies the name of an EVPN peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of an EVPN peer.	The value is in dotted decimal notation.
route-policy-name	Specifies the name of a routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
import	Applies a routing policy to routes to be imported from an EVPN peer or peer group.	-
export	Applies a routing policy to routes to be advertised to an EVPN peer or peer group.	-

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After a routing policy is created, the **peer route-policy** command is run to apply a routing policy to an EVPN peer or peer group so that the routes advertised to or received from the peer or peer group can be controlled. To be specific, only the necessary routes are received from or advertised to the peer or peer group. In this manner, route management is implemented, the scale of the routing table is reduced, and fewer network resources are consumed.

Prerequisites

- The ability to exchange EVPN routes with the peer or a group has been enabled using the **peer** { *group-name* | *ipv4-address* } **enable**.
- The corresponding routing policy has been established. By default, nonexistent routing policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the

system view and a nonexistent routing policy is referenced using this command, all routes are advertised to neighbors or all routes are received.

Precautions

If a routing policy is specified for an EVPN peer, all the peers in the EVPN peer group inherit the configuration.

Example

# Apply routing policy **test-rp** to the routes received from EVPN peer 1.1.1.9.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] peer 1.1.1.9 enable
[HUAWEI-bgp-af-evpn] peer 1.1.1.9 route-policy test-rp import
```

18.1.72 policy vpn-target (BGP-EVPN address family view)

Function

The **policy vpn-target** command enables a device to filter received EVPN routes by the EVPN-VPN target.

The **undo policy vpn-target** command cancels EVPN-VPN target-based filtering of received EVPN routes.

By default, EVPN-VPN target-based filtering of received EVPN routes is enabled.

Format

policy vpn-target

undo policy vpn-target

Parameters

None

Views

BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

After EVPN-VPN target-based filtering is enabled, only EVPN routes whose export EVPN-VPN target attribute matches the import EVPN-VPN target attribute are received.

Example

# Enable a device to implement EVPN-VPN target-based filtering for received EVPN routes.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] l2vpn-family evpn
[HUAWEI-bgp-af-evpn] policy vpn-target
```

18.1.73 port nvo3 mode access

Function

The **port nvo3 mode access** command configures an interface as a VXLAN access-side interface to allow common IP packets carrying the destination UDP port number 4789 of VXLAN packets to enter a VXLAN network.

The **undo port nvo3 mode access** command restores the default setting.

By default, an interface is not configured as a VXLAN access-side interface and cannot forward common IP packets carrying the destination UDP port number 4789 of VXLAN packets to a VXLAN network.

NOTE

Only the S6720-EI, S6735-S, and S6720S-EI switches support this command.

Format

port nvo3 mode access

undo port nvo3 mode access

Parameters

None

Views

XGE interface view, 40GE interface view, Eth-Trunk interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, if an access-side interface is on an S6720-EI, S6735-S, or S6720S-EI and the VXLAN function is configured, when the interface receives common IP packets carrying the destination UDP port number 4789 of VXLAN packets, the packets are discarded due to VXLAN decapsulation. In scenarios where common IP packets carrying the destination UDP port number 4789 of VXLAN packets need to enter a VXLAN network or other transparent transmission scenarios (such as QinQ and

VLAN mapping scenarios), you can run the **port nvo3 mode access** command on an interface to configure the interface as a VXLAN access-side interface. In this case, the interface can forward the received packets with UDP port number 4789.

Precautions

After the **port nvo3 mode access** command is configured on an interface, the interface cannot perform VXLAN decapsulation for received VXLAN-encapsulated packets. Therefore, this command can only be used on access-side interfaces.

Example

# Specify XGE0/0/1 as a VXLAN access-side interface.

```
<HUAWEI> system-view
[HUAWEI] interface XGigabitEthernet 0/0/1
[HUAWEI-XGigabitEthernet0/0/1] port nvo3 mode access
```

18.1.74 refresh bgp evpn

Function

The **refresh bgp evpn** command soft resets an EVPN connection.

Format

refresh bgp evpn { **all** | *ipv4-address* | **group** *group-name* | **external** | **internal** } { **export** | **import** }

Parameters

Parameter	Description	Value
all	Soft resets all EVPN connections.	-
ipv4-address	Specifies the IPv4 address of an EVPN peer.	The value is in dotted decimal notation.
group *group-name*	Specifies the name of an EVPN peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
external	Specifies the EVPN peer of an EBGP connection.	-
internal	Specifies the EVPN peer of an IBGP connection.	-
export	Triggers a soft reset in the export direction.	-
import	Triggers a soft reset in the import direction.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can run the **refresh bgp evpn** command to manually soft reset an EVPN connection. EVPN soft resets can update the EVPN routing table without interrupting a BGP connection and apply new filtering policies.

If a policy application delay time is configured:

- When the routing policy is set to neighbor import by running the **peer route-policy import** command, you can run the **refresh bgp evpn all** command to trigger BGP to re-apply the routing policy.
- When the routing policy is set to neighbor export by running the **peer route-policy export** command, the routing policy is not re-applied even if the **refresh bgp evpn all** command is run.

Example

# Soft reset all BGP-EVPN connections in the import direction to make new configurations take effect.

```
<HUAWEI> refresh bgp evpn all import
```

18.1.75 reset bgp evpn

Function

The **reset bgp evpn** command resets the BGP connections of a BGP EVPN address family.

Format

reset bgp evpn { **all** | *as-number-plain* | *as-number-dot* | **group** *group-name* | *ipv4-address* }

Parameters

Parameter	Description	Value
all	Resets all BGP connections of a BGP EVPN address family.	-
as-number-plain	Specifies an integral AS number.	The value is an integer that ranges from 1 to 4294967295.

Parameter	Description	Value
as-number-dot	Specifies an AS number in dotted notation.	The value is in the format of x.y, where x and y are integers that range from 1 to 65535 and from 0 to 65535, respectively.
ipv4-address	Resets the connection with a specified BGP peer.	The value is in dotted decimal notation.
group *group-name*	Resets the connection with a specified BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Applicable Environment

When the BGP EVPN configuration is changed, you can run the **reset bgp evpn** command to make the new configuration take effect immediately.

NOTICE

After the command is run, the TCP connection established by the BGP device is reset and the corresponding peer relationship is re-established. Therefore, exercise caution before you run this command.

Example

# Reset all BGP connections of a BGP EVPN address family.

```
<HUAWEI> reset bgp evpn all
```

18.1.76 reset bridge-domain statistics

Function

The **reset bridge-domain statistics** command clears packet statistics in a BD.

Format

reset bridge-domain *bd-id* **statistics**

Parameters

Parameter	Description	Value
bd-id	Specifies the ID of a BD packet statistics of which are to be deleted.	The value is an integer that ranges from 1 to 16777215.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Before you collect packet statistics in a specified BD within a specific period, run the **reset bridge-domain statistics** command to clear the existing statistics in the BD to ensure statistics accuracy.

Precautions

Packet statistics cannot be restored after you clear them; therefore, exercise caution before you run the **reset bridge-domain statistics** command.

Example

# Clear packet statistics in BD 10.

```
<HUAWEI> reset bridge-domain 10 statistics
```

18.1.77 reset vxlan statistics

Function

The **reset vxlan statistics** command clears VXLAN tunnel packet statistics.

Format

reset vxlan statistics source *source-ip-address* **peer** *peer-ip-address* [**vni** *vni-id*]

reset vxlan statistics source *source-ipv6-address* **peer** *peer-ipv6-address* [**vni** *vni-id*]

NOTE

Only the S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, S5731-H, and S6730-S support the **source** and **peer** parameters configured as IPv6 address.

Parameters

Parameter	Description	Value
source *source-ip-address*	Specifies the IPv4 address of the source VTEP.	The value is in dotted decimal notation.
peer *peer-ip-address*	Specifies the IPv4 address of the destination VTEP.	The value is in dotted decimal notation.
source *source-ipv6-address*	Specifies the IPv6 address of the source VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
peer *peer-ipv6-address*	Specifies the IPv6 address of the destination VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
vni *vni-id*	Specifies a VNI ID.	The value is an integer that ranges from 1 to 16777215.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Before you collect VXLAN tunnel packet statistics within a specific period, run the **reset vxlan statistics** command to clear the existing statistics to ensure statistics accuracy.

Precautions

Packet statistics cannot be restored after you clear them; therefore, exercise caution before you run the **reset vxlan statistics** command.

Example

# Clear statistics on VXLAN tunnel packets, with 10.10.1.1 and 10.1.1.1 as the source and destination VTEP IP addresses.

```
<HUAWEI> reset vxlan statistics source 10.10.1.1 peer 10.1.1.1
```

18.1.78 rewrite pop

Function

The **rewrite pop** command enables the device to remove VLAN tags from packets received by Layer 2 sub-interfaces.

The **undo rewrite pop** command disables the device from removing VLAN tags from packets received by Layer 2 sub-interfaces.

By default, the device removes two VLAN tags from packets received by Layer 2 sub-interfaces that use QinQ encapsulation, removes one VLAN tag from packets received by Layer 2 sub-interfaces that use Dot1q encapsulation.

Format

rewrite pop { **single** | **double** | **none** }

undo rewrite pop { **single** | **double** | **none** }

Parameters

Parameter	Description	Value
single	Configures the device to remove one VLAN tag from received packets. This parameter can only be configured on Layer 2 sub-interfaces that use Dot1q encapsulation and no VLAN segment can be configured for Layer 2 sub-interfaces.	-
double	Configures the device to remove one VLAN tag from received packets. This parameter can only be configured on Layer 2 sub-interfaces that use QinQ encapsulation and no VLAN segment can be configured for Layer 2 sub-interfaces.	-
none	Configures the device not to remove VLAN tags from received packets. That is, the device transparently transmits received packets without modifying them. This parameter can only be configured on Layer 2 sub-interfaces that use QinQ or Dot1q encapsulation.	-

Views

Layer 2 sub-interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If a Layer 2 sub-interface with the encapsulation type being Dot1q or QinQ is configured as a VXLAN service access point on a VXLAN, to enable the sub-interface to remove the VLAN tag from received packets, run the **rewrite pop** command.

Precautions

Before running the **encapsulation** command to configure a VLAN segment on a Dot1q or QinQ Layer 2 sub-interface, you must run the **rewrite pop** command to specify the operation of removing VLAN tags as **none**.

Example

# Configure the device to remove two VLAN tags from packets received by a Layer 2 sub-interface.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1.1 mode l2
[HUAWEI-GigabitEthernet0/0/1.1] rewrite pop double
```

18.1.79 route-distinguisher

Function

The **route-distinguisher** command configures a route distinguisher (RD) for an EVPN instance address family.

By default, no RD is configured for an EVPN instance address family.

Format

route-distinguisher *route-distinguisher*

Parameters

Parameter	Description	Value
route-distinguisher	Specifies the value of an RD. The RD can be in the following forms: ● 2-byte AS number: 4-byte user-defined number, for example, 101:3. The AS number ranges from 0 to 65535, and the user-defined number ranges from 0 to 4294967295. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. ● Integral 4-byte AS number: 2-byte user-defined number, for example, 0:3 or 65537:3. The AS number ranges from 65536 to 4294967295, and the user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. ● 4-byte AS number in dotted notation: 2-byte user-defined number, for example, 0.0:3 or 0.1:0. The 4-byte AS number in dotted notation is in the format of x.y, where x and y are integers that range from 0 to 65535, and the user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. ● IPv4 address: 2-byte user-defined number, for example, 192.168.122.15:1. The IP address ranges from 0.0.0.0 to 255.255.255.255, and the user-defined number ranges from 0 to 65535.	-

Views

EVPN instance view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After an EVPN instance is created, run this command to configure an RD in the EVPN instance view.

Different EVPN instances may have the same route prefix. RDs can be configured for EVPN instances so that EVPN routes sent to the peer device are added with RD attributes. In this way, EVPN routes are unique globally and Rx devices can easily distinguish EVPN instances.

Precautions

After you configure the RD for an EVPN instance address family, the RD cannot be modified or deleted. When you need to modify an RD, delete the corresponding EVPN instance and reconfigure the RD.

Example

# Configure an RD for EVPN instance **evn1**.

```
<HUAWEI> system-view
[HUAWEI] evpn vpn-instance evn1 bd-mode
[HUAWEI-evpn-instance-evn1] route-distinguisher 22:1
```

18.1.80 service type vxlan-tunnel

Function

The **service type vxlan-tunnel** command configures an Eth-Trunk as a VXLAN loopback interface.

The **undo service type vxlan-tunnel** command cancels the configuration.

By default, an Eth-Trunk is not a VXLAN loopback interface.

NOTE

Only the S6720-EI, S6735-S, and S6720S-EI switches support this command.

Format

service type vxlan-tunnel

undo service type vxlan-tunnel

Parameters

None

Views

Eth-Trunk interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The S6720-EI, S6735-S, and S6720S-EI switches can decapsulate received VXLAN packets and forward them at Layer 3 only after a VXLAN loopback interface is configured on them. As a result, you need to configure an Eth-Trunk as the VXLAN loopback interface when the S6720-EI, S6735-S, and S6720S-EI switches function as the Layer 3 VXLAN gateway.

Follow-up Procedure

Run the **trunkport** command to add member interfaces to the VXLAN loopback interface.

Precautions

- After an Eth-Trunk is configured as a VXLAN loopback interface, STP is automatically disabled on the Eth-Trunk. The Eth-Trunk then does not support STP configuration commands. After the configuration is canceled, STP is automatically enabled on the Eth-Trunk.
- Only one Eth-Trunk on a switch can be configured as the VXLAN loopback interface. VXLAN packets from all VBDIF interfaces are encapsulated and decapsulated by this loopback interface.
- An Eth-Trunk containing member interfaces cannot be configured as a VXLAN loopback interface.
- The configurations allowed on an Eth-Trunk to be configured as a loopback interface include **description**, **enable snmp trap updown**, **jumboframe enable**, **mixed-rate link enable**, **qos phb marking enable**, **set flow-stat interval**, **shutdown**, **local-preference enable**, **traffic-policy (interface view)**, and **trust**. If other configurations exist on the Eth-Trunk, the Eth-Trunk cannot be configured as a loopback interface.
- After an Eth-Trunk is configured as a loopback interface, the Eth-Trunk supports only the following configurations: **authentication open ucl-policy enable**, **description**, **enable snmp trap updown**, **jumboframe enable**, **mixed-rate link enable**, **qos phb marking enable**, **set flow-stat interval**, **shutdown**, **local-preference enable**, **statistic enable (interface view)**, **traffic-policy (interface view)**, **vcmp disable**, and **trust**.
- Before running the **undo service type vxlan-tunnel** command, delete all the member interfaces of the Eth-Trunk interface and all VBDIF interfaces on the device.

Example

# Configure Eth-Trunk 1 as a VXLAN loopback interface.

```
<HUAWEI> system-view
[HUAWEI] interface Eth-Trunk 1
[HUAWEI-Eth-Trunk1] service type vxlan-tunnel
```

18.1.81 set vxlan resource super-mode

Function

The **set vxlan resource super-mode** command sets the super VXLAN resource mode.

The **undo set vxlan resource super-mode** command restores the default VXLAN resource mode.

By default, the device supports 4095 BDs.

NOTE

Only the S5731S-H, S6730-H, S6730S-H, S5731-H, and S5732-H support this command.

Format

set vxlan resource super-mode

undo set vxlan resource super-mode

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When VXLAN is configured, the device supports 4095 BDs by default. If you want more than 4095 BDs, run the **set vxlan resource super-mode** command to set the super VXLAN resource mode. After this command is configured, the device supports 16000 BDs.

Precautions

- After setting the super VXLAN resource mode, save the configuration and then restart the device to make the configuration take effect.
- When the super VXLAN resource mode is configured, the forwarding performance of some services may degrade, such as the IP multicast, VPLS, VLAN mapping, Layer 3 traffic forwarding of sub-interfaces, and VLAN stacking services.

Example

# Set the super VXLAN resource mode.

```
<HUAWEI> system-view
[HUAWEI] set vxlan resource super-mode
```

18.1.82 source (NVE interface view)

Function

The **source** command configures an IP address for the source VXLAN tunnel endpoint (VTEP) of a VXLAN tunnel.

The **undo source** command deletes the IP address of the source VTEP of a VXLAN tunnel.

By default, no IP address is configured for the source VTEP of a VXLAN tunnel.

Format

source *ip-address*

undo source [*ip-address*]

source *ipv6-address*

undo source [*ipv6-address*]

NOTE

Only the S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, S5731-H, and S6730-S support the *ipv6-address* parameter.

Parameters

Parameter	Description	Value
ip-address	Specifies the IPv4 address of a source VTEP.	The value is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a source VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Views

NVE interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

VXLAN needs to be deployed on a downlink interface to provide access services and an uplink interface to establish a VXLAN tunnel.

To establish a VXLAN tunnel, configure IP addresses for the source and destination VTEPs.

To configure an IP address for a source VTEP, run this command. When access service packets reach a Network Virtualization Edge (NVE), the VTEP encapsulates the packets based on the IP addresses of source and destination VTEPs and forwards them.

Precautions

- You can specify the IP address of a physical interface or a loopback interface as the IP address of the source VTEP. The address of a loopback interface is recommended.

- Generally, the NVE interfaces of different devices must be configured with different VTEP addresses; otherwise, traffic forwarding errors may occur.
- The IP address of a VBDIF interface cannot be configured as the IP address of the source VTEP of a VXLAN tunnel.

Follow-up Procedure

After running the **source** command, you can run the **vni head-end peer-list** command to configure an IP address for the destination VTEP.

Example

# Set the IP address of the source VTEP of a VXLAN tunnel to 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] interface nve 1
[HUAWEI-Nve1] source 10.1.1.2
```

18.1.83 statistics enable (BD view)

Function

The **statistics enable** command enables statistics collection in a BD.

The **undo statistics enable** command disables statistics collection in a BD.

By default, statistics collection is disabled in a BD.

Format

statistics enable

undo statistics enable

Parameters

None

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

Packet statistics in a BD provide information about the number of packets going into and leaving a BD.

To view packet statistics in a BD, run the **statistics enable** and **display bridge-domain statistics** commands in sequence. The information helps you locate faults.

Example

# Enable packet statistics collection in BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] statistics enable
```

18.1.84 undo mac-address bridge-domain

Function

The **undo mac-address bridge-domain** command deletes MAC address entries of a BD.

Format

undo mac-address bridge-domain *bd-id*

Parameters

Parameter	Description	Value
bd-id	Deletes MAC address entries of a specified BD.	The value is an integer that ranges from 1 to 16777215.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

The MAC address table space is limited on a device. If the number of MAC address entries reaches the upper limit, the device cannot learn new MAC address entries before some entries in the MAC address table are aged out. In this case, the device broadcasts packets from new users, wasting network resources. To solve the problem, you can run the **undo mac-address bridge-domain** command to manually delete unnecessary MAC address entries of a specified BD.

Example

# Delete MAC address entries of BD 20.

```
<HUAWEI> system-view
[HUAWEI] undo mac-address bridge-domain 20
```

18.1.85 unknown-unicast-suppression (BD view)

Function

The **unknown-unicast-suppression** command enables unknown unicast traffic suppression in a BD.

The **undo unknown-unicast-suppression** command disables unknown unicast traffic suppression in a BD.

By default, unknown unicast traffic suppression is disabled in a BD.

Format

unknown-unicast-suppression cir *cir-value* [**cbs** *cbs-value*]

undo unknown-unicast-suppression

Parameters

Parameter	Description	Value
cir *cir-value*	Specifies the committed information rate (CIR), which is the allowed rate at which traffic can pass through.	The value is an integer that ranges from 0 to 10000000, in kbit/s.
cbs *cbs-value*	Specifies the committed burst size (CBS), which is the maximum size of traffic that can pass through.	The value is an integer that ranges from 10000 to 4294967295, in bytes. If the **cbs** is not set, the default *cbs-value* is 125 times the *cir-value*.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

When a large number of unknown unicast packets are transmitted on the network, a lot of network resources are occupied, and services on the network are affected. You can run the **unknown-unicast-suppression** command to enable unknown unicast traffic suppression in a BD and configure the maximum number of unknown unicast packets that can pass through a BD. When the unknown unicast traffic volume exceeds the specified threshold, the system discards excess unknown unicast packets.

Example

# Set the CIR value for unknown unicast traffic in BD 10 to 100 kbit/s.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] unknown-unicast-suppression cir 100
```

18.1.86 vni head-end peer-list

Function

The **vni head-end peer-list** command configures an ingress replication list that contains the IP addresses of those remote VTEPs for a VXLAN network identifier (VNI).

The **undo vni head-end peer-list** command deletes the ingress replication list of a VNI.

By default, no ingress replication list is configured for any VNI.

Format

vni *vni-id* **head-end peer-list** *ip-address* &<1-10>

undo vni *vni-id* [**head-end peer-list** *ip-address* &<1-10>]

vni *vni-id* **head-end peer-list** *ipv6-address* &<1-10>

undo vni *vni-id* [**head-end peer-list** *ipv6-address* &<1-10>]

vni *vni-id* **head-end peer-list protocol bgp**

undo vni *vni-id* **head-end peer-list protocol bgp**

NOTE

Only the S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, S5731-H, and S6730-S support the **head-end peer-list** parameter configured as IPv6 address.

Parameters

Parameter	Description	Value
vni-id	Specifies a VNI ID.	The value is an integer that ranges from 1 to 16777215.
ip-address	Specifies the IPv4 address of a remote VTEP.	The value is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a remote VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
protocol bgp	Specifies BGP for establishing Layer 2 VXLAN tunnels.	-

Views

NVE interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the ingress of a VXLAN tunnel receives broadcast, unknown unicast, and multicast (BUM) packets, it replicates these packets and sends a copy to each VTEP in the ingress replication list. The ingress replication list is a collection of remote VTEP IP addresses to which the ingress of a VXLAN tunnel should send replicated BUM packets.

If a source VTEP on a VXLAN connects to multiple remote VTEPs on the same VXLAN segment, run the **vni head-end peer-list** command to configure an ingress replication list that contains the IP addresses of those remote VTEPs. After the source NVE receives BUM packets, the local VTEP sends a copy of the BUM packets to every VTEP in the list.

If an underlay network is an IPv6 network, run the **vni head-end peer-list** *ipv6-address* &<1-10> command to configure an ingress replication list that contains IPv6 addresses of remote VTEPs, which is used to forward BUM traffic.

To use BGP to dynamically establish Layer 2 VXLAN tunnels, run the **vni** *vni-id* **head-end peer-list protocol bgp** command.

Precautions

- You need to run the **vni head-end peer-list** command to configure the corresponding VTEP address even if the source VTEP matches only one destination VTEP.
- Run the **ping** command to check whether a reachable route exists between two ends of the tunnel. If there is a reachable route, the tunnel can be established and packets can be normally forwarded. If the two devices have a route to each other but the route is unreachable, the tunnel can still go Up but packets cannot be forwarded.
- Currently, the device can forward BUM packets only through ingress replication. To establish a VXLAN tunnel between the device and a non-Huawei device, ensure that ingress replication is also configured on the non-Huawei device. Otherwise, the VXLAN tunnel cannot be established.

Example

# Configure an ingress replication list for VNI 10, with the remote VTEPs' IP addresses being 10.1.1.1 and 10.1.1.3.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] vxlan vni 10
[HUAWEI-bd10] quit
[HUAWEI] interface nve 1
[HUAWEI-Nve1] vni 10 head-end peer-list 10.1.1.1 10.1.1.3
```

18.1.87 vpn-target (EVPN instance view)

Function

The **vpn-target** command configures the export or import VPN target extended community attribute for an EVPN instance address family.

The **undo vpn-target** command deletes the VPN target extended community attribute for the EVPN instance address family.

By default, no export or import VPN target extended community attribute list is configured for the EVPN instance address family.

Format

vpn-target *vpn-target* &<1-8> [**both** | **export-extcommunity** | **import-extcommunity**]

undo vpn-target { **all** | *vpn-target* &<1-8> [**both** | **export-extcommunity** | **import-extcommunity**] }

Parameters

Parameter	Description	Value
vpn-target	Adds a VPN target to the VPN target extended community attribute list of the EVPN instance address family. The forms of VPN targets are as follows: ● 2-byte AS number: 4-byte user-defined number, for example, 1:3. The AS number ranges from 0 to 65535, and the user-defined number ranges from 0 to 4294967295. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0. ● IPv4 address: 2-byte user-defined number, for example, 192.168.122.15:1. The IP address ranges from 0.0.0.0 to 255.255.255.255, and the user-defined number ranges from 0 to 65535. ● Integral 4-byte AS number: 2-byte user-defined number, for example, 65537:3. The AS number ranges from 65536 to 4294967295, and the user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0. ● 4-byte AS number in dotted notation: 2-byte user-defined number, for example, 0.0:3 or 0.1:0. The 4-byte AS number in dotted notation is in the format of x.y, where x and y are integers that range from 0 to 65535, and the user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0.0.	-
both	Adds a VPN target to the import or export VPN target extended community attribute list of the EVPN instance address family. If keywords **both**, **export-extcommunity**, and **import-extcommunity** are not specified, **both** is used by default.	-
export-extcommunity	Adds a VPN target to the export VPN target extended community attribute list of the EVPN instance address family.	-
import-extcommunity	Adds a VPN target to the import VPN target extended community attribute list of the EVPN instance address family.	-
all	Deletes all the VPN targets of the EVPN instance address family.	-

Views

EVPN instance view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a local device advertises EVPN routes to other devices, the EVPN routes carry all VPN target attributes in the export VPN target attribute list of the local EVPN instance. Only when the VPN target attribute carried in an EVPN route is included in the import VPN target attribute list, the route can be added to the EVPN instance routing table.

NOTE

> You can run the **vpn-target** command to add VPN targets to the VPN target extended community attribute list of an EVPN instance through more than one time. A maximum of eight VPN targets can be added at one time. To configure more VPN targets for an EVPN instance, you can run the **vpn-target** command multiple times.

Prerequisites

The RD of the EVPN instance has been configured by running the **route-distinguisher** command.

Precautions

If this command is not run, all received EVPN routes cannot be added to the local EVPN instance routing table.

If all the VPN targets of an EVPN instance are deleted using the **undo vpn-target** command, all routes learned by the EVPN instance are deleted.

Example

# Add 3:3 to the export VPN target extended community attribute list and 4:4 to the import VPN target extended community attribute list of EVPN instance **evn3**.

```
<HUAWEI> system-view
[HUAWEI] evpn vpn-instance evn3 bd-mode
[HUAWEI-evpn-instance-evn3] route-distinguisher 100:1
[HUAWEI-evpn-instance-evn3] vpn-target 3:3 export-extcommunity
[HUAWEI-evpn-instance-evn3] vpn-target 4:4 import-extcommunity
```

18.1.88 vpn-target evpn

Function

The **vpn-target evpn** command configures the export or import VPN target extended community attribute when a VPN instance address family advertises routes to an EVPN instance.

The **undo vpn-target** command deletes the VPN target extended community attribute from a VPN instance address family.

By default, no export or import VPN target extended community attribute list is configured for a VPN instance address family.

Format

vpn-target *vpn-target* &<1-8> [**both** | **export-extcommunity** | **import-extcommunity**] **evpn**

undo vpn-target { **all** | *vpn-target* &<1-8> [**both** | **export-extcommunity** | **import-extcommunity**] [**evpn**] }

Parameters

Parameter	Description	Value
vpn-target	Adds a VPN target to the VPN target extended community attribute list of the VPN instance address family. The forms of VPN targets are as follows: ● 2-byte AS number: 4-byte user-defined number, for example, 1:3. The AS number ranges from 0 to 65535, and the user-defined number ranges from 0 to 4294967295. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0. ● IPv4 address: 2-byte user-defined number, for example, 192.168.122.15:1. The IP address ranges from 0.0.0.0 to 255.255.255.255, and the user-defined number ranges from 0 to 65535. ● Integral 4-byte AS number: 2-byte user-defined number, for example, 65537:3. The AS number ranges from 65536 to 4294967295, and the user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0. ● 4-byte AS number in dotted notation: 2-byte user-defined number, for example, 0.0:3 or 0.1:0. The 4-byte AS number in dotted notation is in the format of x.y, where x and y are integers that range from 0 to 65535, and the user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, a VPN target cannot be 0:0.0.	-
both	Adds a VPN target to the import or export VPN target extended community attribute list of the VPN instance address family. If keywords **both**, **export-extcommunity**,and **import-extcommunity**are not specified, **both** is used.	-
export-extcommunity	Adds a VPN target to the export VPN target extended community attribute list of the VPN instance address family.	-

Parameter	Description	Value
import-extcommunity	Adds a VPN target to the import VPN target extended community attribute list of the VPN instance address family.	-
all	Deletes all the VPN targets of the VPN instance address family.	-

Views

VPN instance view, VPN instance IPv4 address family view, VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When you configure a VPN instance on a device, you must run the **vpn-target evpn** command to configure a VPN target attribute of EVPN routes for the address family of the VPN instance.

The VPN target controls route learning of EVPN routes. A VPN target attribute is either an import one or an export one. An export VPN target is carried in a EVPN route to be advertised to a remote MP-BGP peer. When receiving a EVPN route, a peer compares the received VPN target attribute with the import VPN target attribute to determine whether the EVPN route can be added to the routing table of the local VPN instance address family.

In the VXLAN deployed through BGP EVPN scenario, to add EVPN routes to the routing table of an VPN instance, run the **vpn-target** command to configure the corresponding VPN target.

Prerequisites

The **route-distinguisher** command have been run to configure the VPN instance RD.

Precautions

A VPN target configured using the **vpn-target evpn** command will not overwrite any previously configured VPN target. If the number of configured VPN targets has reached the maximum limit, no VPN target can be added by running the **vpn-target evpn** command.

After a VPN target is configured for the address family of a VPN instance, only the routes that match the VPN target are accepted by the address family of the VPN instance.

If all the VPN targets of the address family of a VPN instance are deleted using the **undo vpn-target evpn** command, all EVPN routes learned by the address family of the VPN instance from other VPN instances are deleted.

Multiple VPN targets can be configured for the address family of a VPN instance. The **vpn-target evpn** command can configure a maximum of eight VPN targets at a time. To configure more VPN targets in the VPN instance view, run the **vpn-target evpn** command multiple times. When EVPN routes are advertised between VPN instances, if one of the VPN targets carried in the EVPN routes matches the import VPN target of the address family of a local VPN instance, the routes are added to the routing table of the local VPN instance.

Example

# Add 3:3 to the export VPN target EVPN list and 4:4 to the import VPN target evpn list of VPN instance **vrf1**.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vrf1
[HUAWEI-vpn-instance-vrf1] ipv4-family
[HUAWEI-vpn-instance-vrf1-af-ipv4] route-distinguisher 100:1
[HUAWEI-vpn-instance-vrf1-af-ipv4] vpn-target 3:3 export-extcommunity evpn
[HUAWEI-vpn-instance-vrf1-af-ipv4] vpn-target 4:4 import-extcommunity evpn
```

18.1.89 vxlan fragment-reassemble enable

Function

The **vxlan fragment-reassemble enable** command enables the function of reassembling fragmented VXLAN packets.

The **undo vxlan fragment-reassemble enable** command disables the function of reassembling fragmented VXLAN packets.

By default, the function of reassembling fragmented VXLAN packets is disabled.

NOTE

Only the S6730-S, S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, and S5731-H support this command.

Format

vxlan fragment-reassemble enable

undo vxlan fragment-reassemble enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a switch receives fragmented VXLAN packets in a VXLAN scenario, it does not reassemble the fragmented packets by default. To reassemble fragmented VXLAN packets, run the **vxlan fragment-reassemble enable** command to enable the function of reassembling fragmented VXLAN packets.

If the function of reassembling fragmented VXLAN packets is disabled on a switch, the switch generates the ADPVXLAN_1.3.6.1.4.1.2011.5.25.227.2.1.40 hwNotsuppDecapVxlanFragPackets alarm when receiving fragmented VXLAN packets.

Precautions

- Enabling the function of reassembling fragmented VXLAN packets may affect processing of other IP packets. Therefore, exercise caution when using this command.
- If the system software of a switch is upgraded from V200R011C10 to V200R012C00 or a later version and the **interface nve** and **vxlan vni (BD view)** commands are configured on the switch before the upgrade, the function of reassembling fragmented VXLAN packets is enabled after the upgrade. If the **interface nve** and **vxlan vni (BD view)** commands are not configured on the switch before the upgrade, the function of reassembling fragmented VXLAN packets is disabled after the upgrade.

Example

# Enable the function of reassembling fragmented VXLAN packets.

```
<HUAWEI> system-view
[HUAWEI] vxlan fragment-reassemble enable
```

18.1.90 vxlan regularly-refresh

Function

The **vxlan regularly-refresh** command sets the interval at which VXLAN-related entries are updated periodically, the interval between two rounds of updates, and the number of entries updated in each round.

The **undo vxlan regularly-refresh** command restores the default configuration.

By default, the interval at which VXLAN-related entries are updated periodically is 1 minute, the interval between two rounds of updates is 1 second, and the number of entries updated in each round is 50.

Format

vxlan regularly-refresh { **interval** *interval* | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* }

undo vxlan regularly-refresh { **interval** | **entry-number** | **cycle-interval** }

Parameters

Parameter	Description	Value
interval *interval*	Specifies the interval between two rounds of VXLAN-related entry updates.	The value is an integer in the range from 1 to 300, in seconds.
entry-number *entry-number*	Specifies the number of VXLAN-related entries updated in each round.	The value is an integer in the range from 1 to 100.
cycle-interval *cycle-interval*	Specifies the interval at which VXLAN-related entries are updated periodically.	The value is an integer in the range from 1 to 1440, in minutes.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The device supports periodic update of VXLAN-related entries to ensure the real-timeliness and validity of the underlying forwarding data. Periodic update of VXLAN-related entries consumes system resources, including CPU and memory resources. A longer update interval and a smaller number of updated entries have less impact on other services in the system. You can flexibly configure the interval at which VXLAN-related entries are updated periodically and the number of updated entries based on system resource usage.

Prerequisites

Before running this command, ensure that periodic update of VXLAN-related entries has been enabled. By default, this function is enabled. If this function is disabled, run the **undo vxlan regularly-refresh disable** command in the system view to enable it.

Example

# Set the interval between two rounds of VXLAN-related entry updates to 5 seconds.

```
<HUAWEI> system-view
[HUAWEI] vxlan regularly-refresh interval 5
```

18.1.91 vxlan regularly-refresh disable

Function

The **vxlan regularly-refresh disable** command disables periodic update of VXLAN-related entries.

The **undo vxlan regularly-refresh disable** command enables periodic update of VXLAN-related entries.

By default, periodic update of VXLAN-related entries is enabled.

Format

vxlan regularly-refresh disable

undo vxlan regularly-refresh disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, the device updates VXLAN-related entries for the first time 30 minutes after it is powered on, and the update is performed periodically at an interval of 1 minute to ensure the correctness of underlying data. Periodic update of VXLAN-related entries consumes system resources. If the CPU usage is high in a short period of time due to periodic update of VXLAN-related entries and the processing of normal services is affected, you can run the **vxlan regularly-refresh** { **interval** *interval* | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* } command to prolong the periodic update interval and the interval between two rounds of updates, and reduce the number of entries updated in each round to reduce system resource consumption. If the CPU usage remains high, run the **vxlan regularly-refresh disable** command to immediately disable periodic update of VXLAN-related entries.

Precautions

Disabling periodic update of VXLAN-related entries may cause a failure to rectify certain hardware faults in a timely manner. Exercise caution when running the **vxlan regularly-refresh disable** command.

Example

# Enables periodic update of VXLAN-related entries.

```
<HUAWEI> system-view
[HUAWEI] undo vxlan regularly-refresh disable
```

18.1.92 vxlan statistics enable

Function

The **vxlan statistics enable** command enables statistics collection on VXLAN tunnel packets.

The **undo vxlan statistics enable** command disables statistics collection on VXLAN tunnel packets.

By default, statistics collection on VXLAN tunnel packets is disabled.

Format

vxlan statistics peer *peer-ip-address* [**vni** *vni-id*] **enable**

undo vxlan statistics peer *peer-ip-address* [**vni** *vni-id*] **enable**

vxlan statistics peer *peer-ipv6-address* [**vni** *vni-id*] **enable**

undo vxlan statistics peer *peer-ipv6-address* [**vni** *vni-id*] **enable**

NOTE

Only the S6730S-S, S5732-H, S5731-S, S5731S-S, S5731S-H, S6730-H, S6730S-H, S5731-H, and S6730-S support the **source** and **peer** parameters configured as IPv6 address.

Parameters

Parameter	Description	Value
peer *peer-ip-address*	Specifies the IPv4 address of the destination VTEP.	The value is in dotted decimal notation.
peer *peer-ipv6-address*	Specifies the IPv6 address of the destination VTEP.	The address is a 32-bit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
vni *vni-id*	Specifies a VNI ID. If **vni** *vni-id* is specified, the device collects packet statistics based on the VXLAN tunnel and VNI. If **vni** *vni-id* is not specified, the device collects packet statistics based on the VXLAN tunnel only.	The value is an integer that ranges from 1 to 16777215.

Views

NVE interface view

Default Level

2: Configuration level

Usage Guidelines

VXLAN tunnel packet statistics provide information about the number of packets going into and leaving a VXLAN tunnel.

To view VXLAN tunnel packet statistics, run the **vxlan statistics enable** and **display vxlan statistics** commands in sequence.

Packet statistics collection based on the VXLAN tunnel and VNI and packet statistics collection based on the VXLAN tunnel are mutually exclusive. For example, if the **vxlan statistics peer 10.1.1.1 vni 10 enable** command is configured, do not configure the **vxlan statistics peer 10.1.1.1 enable** command. If the **vxlan statistics peer 10.1.1.1 enable** command is configured, do not configure the **vxlan statistics peer 10.1.1.1 vni 10 enable** command.

Example

# Enable statistics collection on VXLAN tunnel packets with 10.1.1.1 as the destination VTEP IP address.

```
<HUAWEI> system-view
[HUAWEI] interface nve 1
[HUAWEI-Nve1] vxlan statistics peer 10.1.1.1 enable
```

18.1.93 vxlan tunnel-status track exact-route

Function

The **vxlan tunnel-status track exact-route** command enables subscription to the status of the exact route to a VXLAN tunnel destination.

The **undo vxlan tunnel-status track exact-route** command disables subscription to the status of the exact route to a VXLAN tunnel destination.

By default, subscription to the status of the exact route to a VXLAN tunnel destination is disabled.

Format

vxlan tunnel-status track exact-route

undo vxlan tunnel-status track exact-route

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

By default, a VXLAN tunnel is considered Up if its source IP address and the network segment where its destination IP address resides are reachable. In the real-world situation, when multiple destination IP addresses share the same network segment, this network segment is reachable if only one such destination IP address is reachable. As such, a VXLAN tunnel may be incorrectly considered Up even if its destination IP address is unreachable. Network problems then cannot be immediately identified. To resolve this problem, run the **vxlan tunnel-status track exact-route** command to enable subscription to the status of the exact route to a VXLAN tunnel destination. The VXLAN tunnel is then reported Up only when the exact route to the VTEP destination IP address is reachable.

NOTE

The VXLAN tunnel status can be checked by running the **display vxlan tunnel** command.

Example

# Enable subscription to the status of the exact route to a VXLAN tunnel destination.

```
<HUAWEI> system-view
[HUAWEI] vxlan tunnel-status track exact-route
```

18.1.94 vxlan vni (BD view)

Function

The **vxlan vni** command associated a specified VNI with a BD.

The **undo vxlan vni** command restores the default settings.

By default, no VNI is associated with a BD.

Format

vxlan vni *vni-id*

undo vxlan vni *vni-id*

Parameters

Parameter	Description	Value
vni-id	Specifies a VNI ID.	The value is an integer that ranges from 1 to 16777215.

Views

BD view

Default Level

2: Configuration level

Usage Guidelines

Similar to VLAN ID, VNI is used to distinguish VXLAN segments.

On the VXLAN network, one VXLAN segment is a large Layer 2 BD; therefore, VNI and BD have a one-to-one mapping relationship.

You can run this command to configure the mapping relationship between VNIs and BDs. In this way, the VTEP can forward received packets through a correct VXLAN tunnel based on the mapping between BDs and VNIs.

Example

# Set the mapping between VNI 10 and BD 10.

```
<HUAWEI> system-view
[HUAWEI] bridge-domain 10
[HUAWEI-bd10] vxlan vni 10
```

18.1.95 vxlan vni (VPN instance view)

Function

The **vxlan vni** command configures the VXLAN VNI corresponding to a VPN instance.

The **undo vxlan vni** command restores the default settings.

By default, no VNI is configured for a VPN instance.

Format

vxlan vni *vni-id*

undo vxlan vni *vni-id*

Parameters

Parameter	Description	Value
vni-id	Specifies a VNI ID.	The value is an integer that ranges from 1 to 16777215.

Views

VPN instance view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To isolate tenants at Layer 3, VPN is generally used. In a distributed VXLAN gateway scenario, to implement Layer 3 communication through a Layer 3 gateway, the Layer 3 gateway must be bound to a VPN instance.

The Layer 3 gateway assigns a VNI to each subnet and a Layer 3 VNI to each tenant identified by a VPN instance. To bind a VNI to a VPN instance, run the **vxlan vni** command. During Layer 3 communication through the Layer 3 gateway, the VNI ID bound to the VPN instance is transmitted to the remote Layer 3 gateway through the VXLAN tunnel. The remote Layer 3 gateway identifies VPNs based on tenants' VNI IDs to determine whether tenants belong to the same VPN for communication or isolation purposes.

Precautions

- A VNI can be bound only to one VPN instance.
- The VNI bound to a VPN instance cannot be bound to a BD.

Example

# Configure the ID of the VNI corresponding to VPN instance **vpn1** as **10**.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vpn1
[HUAWEI-vpn-instance-vpn1] vxlan vni 10
```