H3C S6860 Switch Series

EVPN Command Reference

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

Software version: Release 2612 and later
Document version: 6W102-20200419

Preface

This command reference describes EVPN configuration commands.

This preface includes the following topics about the documentation:

- Audience.
- Conventions.
- Documentation feedback.

Audience

This documentation is intended for:

- Network planners.
- Field technical support and servicing engineers.
- Network administrators working with the S6860 switch series.

Conventions

The following information describes the conventions used in the documentation.

Command conventions

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

GUI conventions

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

Symbols

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

Network topology icons

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

Examples provided in this document

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

Documentation feedback

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

We appreciate your comments.

Contents

EVPN commands

address-family evpn (public instance view)

Use **address-family evpn** to enter EVPN view of the public instance.

Use **undo address-family evpn** to delete all settings in EVPN view of the public instance.

Syntax

address-family evpn

undo address-family evpn

Views

Public instance view

Predefined user roles

network-admin

Usage guidelines

You can configure EVPN settings such as route targets in EVPN view of the public instance.

Examples

Enter EVPN view of the public instance.

```
<Sysname> system-view
[Sysname] ip public-instance
[Sysname-public-instance] address-family evpn
[Sysname-public-instance-evpn]
```

address-family evpn (VPN instance view)

Use **address-family evpn** to enter EVPN view of a VPN instance.

Use **undo address-family evpn** to delete all settings in EVPN view of a VPN instance.

Syntax

address-family evpn

undo address-family evpn

Views

VPN instance view

Predefined user roles

network-admin

Usage guidelines

You can configure EVPN settings such as route targets and routing policies in EVPN view of a VPN instance.

Examples

Enter EVPN view of VPN instance **tenant**.

```
<Sysname> system-view
[Sysname] ip vpn-instance tenant
```

```
[Sysname-vpn-instance-tenant] address-family evpn
[Sysname-vpn-evpn-tenant]
```

address-family ipv4 (public instance view)

Use **address-family ipv4** to enter IPv4 VPN view of the public instance.

Use **undo address-family ipv4** to delete all settings in IPv4 VPN view of the public instance.

Syntax

address-family ipv4

undo address-family ipv4

Views

Public instance view

Predefined user roles

network-admin

Examples

Enter IPv4 VPN view of the public instance.

```
<Sysname> system-view
[Sysname] ip public-instance
[Sysname-public-instance] address-family ipv4
[Sysname-public-instance-ipv4]
```

address-family ipv6 (public instance view)

Use **address-family ipv6** to enter IPv6 VPN view of the public instance.

Use **undo address-family ipv6** to delete all settings in IPv6 VPN view of the public instance.

Syntax

address-family ipv6

undo address-family ipv6

Views

Public instance view

Predefined user roles

network-admin

Examples

Enter IPv6 VPN view of the public instance.

```
<Sysname> system-view
[Sysname] ip public-instance
[Sysname-public-instance] address-family ipv6
[Sysname-public-instance-ipv6]
```

address-family l2vpn evpn

Use **address-family l2vpn evpn** to create the BGP EVPN address family and enter its view, or enter the view of the existing BGP EVPN address family.

Use **undo address-family l2vpn evpn** to delete the BGP EVPN address family and all settings in BGP EVPN address family view.

Syntax

address-family l2vpn evpn

undo address-family l2vpn evpn

Default

The BGP EVPN address family does not exist.

Views

BGP instance view

Predefined user roles

network-admin

Usage guidelines

Configuration made in BGP EVPN address family view takes effect only on routes and peers of the BGP EVPN address family that are on the public network.

Examples

Create the BGP EVPN address family and enter its view.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn]
```

advertise l2vpn evpn

Use **advertise l2vpn evpn** to enable BGP EVPN route advertisement to the local site.

Use **undo advertise l2vpn evpn** to disable BGP EVPN route advertisement to the local site.

Syntax

advertise l2vpn evpn

undo advertise l2vpn evpn

Default

BGP EVPN route advertisement to the local site is enabled.

Views

BGP-VPN IPv4 unicast address family view

BGP-VPN IPv6 unicast address family view

Predefined user roles

network-admin

Usage guidelines

This command enables the device to advertise private BGP EVPN routes to the local site after the device adds the routes to the routing table of a VPN instance.

Examples

Enable BGP EVPN route advertisement to the local site for VPN instance **vpn1**.

```
<Sysname> system-view
[Sysname] bgp 100
```

```
[Sysname-bgp-default] ip vpn-instance vpn1
[Sysname-bgp-default-vpn1] address-family ipv4
[Sysname-bgp-default-ipv4-vpn1] advertise l2vpn evpn
```

arp mac-learning disable

Use **arp mac-learning disable** to disable an EVPN instance from learning MAC addresses from ARP information.

Use **undo arp mac-learning disable** to restore the default.

Syntax

arp mac-learning disable

undo arp mac-learning disable

Default

An EVPN instance learns MAC addresses from ARP information.

Views

EVPN instance view

Predefined user roles

network-admin

Usage guidelines

The MAC information and ARP information advertised by a remote VTEP overlap. To avoid duplication, use this command to disable the learning of MAC addresses from ARP information. EVPN will learn remote MAC addresses only from the MAC information advertised from remote sites.

Examples

Disable an EVPN instance from learning MAC addresses from ARP information.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] arp mac-learning disable
```

arp-advertising disable

Use **arp-advertising disable** to disable ARP information advertisement for an EVPN instance.

Use **undo arp-advertising disable** to restore the default.

Syntax

arp-advertising disable

undo arp-advertising disable

Default

ARP information advertisement is enabled for an EVPN instance.

Views

EVPN instance view

Predefined user roles

network-admin

Usage guidelines

In an EVPN network with distributed gateways, you can disable ARP information advertisement for a VXLAN to save resources if all its user terminals use the same EVPN gateway device. The EVPN instance of the VXLAN will stop advertising ARP information through MAC/IP advertisement routes and withdraw advertised ARP information. When ARP information advertisement is disabled, user terminals in other VXLANs still can communicate with that VXLAN through IP prefix advertisement routes.

Examples

Disable ARP information advertisement for an EVPN instance.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] arp-advertising disable
```

dci enable

Use **dci enable** to enable DCI on an interface.

Use **undo dci enable** to disable DCI on an interface.

Syntax

dci enable

undo dci enable

Default

DCI is disabled on an interface.

Views

Interface view

Predefined user roles

network-admin

Usage guidelines

For EDs to automatically establish VXLAN-DCI tunnels, you must enable DCI on the Layer 3 interfaces that interconnect the EDs. You can enable DCI only on Layer 3 Ethernet interfaces and Layer 3 aggregate interfaces.

Subinterfaces of a DCI-enabled interface inherit configuration of the interface.

Examples

Enable DCI on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] dci enable
```

display bgp l2vpn evpn

Use **display bgp l2vpn evpn** to display BGP EVPN routes.

Syntax

display bgp [**instance** *instance-name*] **l2vpn evpn** [**peer** *ipv4-address* { **advertised-routes** | **received-routes** } [**statistics**] | **route-distinguisher** *route-distinguisher* [**route-type**

{ **auto-discovery** | **es** | **imet** | **ip-prefix** | **mac-ip** }] [*evpn-route route-length* [**advertise-info**]] | **route-type** { **auto-discovery** | **es** | **imet** | **ip-prefix** | **mac-ip** } | **statistics**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

instance *instance-name*: Specifies a BGP instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays BGP EVPN routes for the default BGP instance.

peer *ipv4-address*: Specifies a peer by its IPv4 address.

advertised-routes: Specifies the routes advertised to the specified peer.

received-routes: Specifies the routes received from the specified peer.

statistics: Displays BGP EVPN route statistics.

route-distinguisher *route-distinguisher*: Specifies a route distinguisher (RD), a string of 3 to 21 characters. The RD can use one of the following formats:

- *16-bit AS number*:*32-bit user-defined number*. For example, 101:3.
- *32-bit IP address*:*16-bit user-defined number*. For example, 192.168.122.15:1.
- *32-bit AS number*:*16-bit user-defined number*. For example, 65536:1. The AS number must be equal to or greater than 65536.

route-type: Specifies a route type.

auto-discovery: Specifies Ethernet auto-discovery routes.

es: Specifies Ethernet segment (ES) routes.

imet: Specifies inclusive multicast Ethernet tag (IMET) routes.

ip-prefix: Specifies IP prefix advertisement routes.

mac-ip: Specifies MAC/IP advertisement routes.

evpn-route: Specifies a BGP EVPN route, a case-insensitive string of 1 to 512 characters.

route-length: Specifies the route length in bits, in the range of 0 to 65535.

advertise-info: Displays advertisement information for BGP EVPN routes.

Usage guidelines

If you do not specify any parameter, this command displays brief information about all BGP EVPN routes.

Examples

Display brief information about all BGP EVPN routes.

```
<Sysname> display bgp l2vpn evpn

 BGP local router ID is 8.8.8.8
 Status codes: * - valid, > - best, d - dampened, h - history,
               s - suppressed, S - stale, i - internal, e - external
               a - additional-path
               Origin: i - IGP, e - EGP, ? - incomplete
```

```
 Total number of routes from all PEs: 3

 Route distinguisher: 1:1
 Total number of routes: 2

     Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

* >i [2][0][48][7010-0000-0001][0][0.0.0.0]/104
                        7.7.7.7         0          100        0       i
*  i                    7.7.7.7         0          100        0       i

Route distinguisher of public instance: 1:15
Total number of routes: 1

     Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

* >e [2][0][48][0011-0022-0033][32][11.22.33.55]/136
                        30.30.1.2       0                     0       100i
```

Table 1 Command output

Field	Description
Status codes	Route status codes: • *** - valid**—Valid route. • **> - best**—Optimal route. • **d - dampened**—Dampened route. • **h - history**—History route. • **i - internal**—Internal route. • **e - external**—External route. • **s - suppressed**—Suppressed route. • **S - Stale**—Stale route. • **a - additional-path**—Add-Path optimal route.
Origin	Origin of the route: • **i – IGP**—Originated in the AS. The origin of routes advertised by using the **network** command is **IGP**. • **e – EGP**—Learned through EGP. • **? – incomplete**—Unknown origin. The origin of routes redistributed from IGP protocols is **incomplete**.

Network	BGP EVPN route/route length. BGP EVPN routes are as follows:
	• [1][ESI][EthernetTagID]
	  ◦ **1**—Ethernet auto-discovery route.
	  ◦ **ESI**—Ethernet segment identifier (ESI).
	  ◦ **EthernetTagID**—Ethernet tag ID.
	• [2][EthernetTagID][MACLength][MAC][IPAddressLength][IPAddress]
	  ◦ **2**—MAC/IP advertisement route.
	  ◦ **EthernetTagID**—Ethernet tag ID.
	  ◦ **MACLength**—MAC address length.
	  ◦ **MAC**—MAC address.
	  ◦ **IPAddressLength**—IP address length.
	  ◦ **IPAddress**—IP address.
	• [3][EthernetTagID][IPAddressLength][IPAddress]

Field	Description
	○ **3**—IMET route. ○ **IPAddressLength**—IP address length. ○ **IPAddress**—IP address of the originating router. • [4][ESI][IPAddressLength][IPAddress] ○ **4**—ES route. ○ **ESI**—ESI. ○ **IPAddressLength**—IP address length. ○ **IPAddress**—IP address of the originating router. • [5][EthernetTagID][IPAddressLength][IPAddress] ○ **5**—IP prefix advertisement route. ○ **EthernetTagID**—Ethernet tag ID. ○ **IPAddressLength**—IP address length. ○ **IPAddress**—IP address of the originating router.
NextHop	Next hop IP address.
MED	Multi-Exit Discriminator (MED) attribute.
LocPrf	Local precedence.
PrefVal	Preferred value.
Path/Ogn	AS_PATH and ORIGIN attributes of the route.

Display detailed information about BGP EVPN route [1][00:00:00:00:00:00:00:00:00:00][5]/120 with RD 1.1.1.1:100.

```
<Sysname> display bgp l2vpn evpn route-distinguisher 1.1.1.1:100
[1][00:00:00:00:00:00:00:00:00:00][5] 120

 BGP local router ID: 172.16.250.133
 Local AS number: 100


 Route distinguisher: 1.1.1.1:100
 Total number of routes: 1
 Paths:   1 available, 1 best

 BGP routing table information of [1][00:00:00:00:00:00:00:00:00:00][5]/120:
 From            : 10.1.1.2 (192.168.56.17)
 Rely nexthop    : 10.1.1.2
 Original nexthop: 10.1.1.2
 OutLabel        : NULL
 Ext-Community   : <RT: 1:2>, <Encapsulation Type: VXLAN >, <ESI Label: Flag 0,
                   Label 1>
 RxPathID        : 0x0
 TxPathID        : 0x0
 AS-path         : 200
 Origin          : igp
 Attribute value : MED 0, pref-val 0
 State           : valid, external, best
 IP precedence   : N/A
 QoS local ID    : N/A
```

```
Traffic index   : N/A
EVPN route type : Ethernet auto-discovery route
ESI             : 00:00:00:00:00:00:00:00:00:00
Ethernet tag ID : 5
MPLS label      : 10
```

Table 2 Command output

Field	Description
Paths	Number of routes: • **available**—Number of valid routes. • **best**—Number of optimal routes.
From	IP address of the BGP peer that advertised the route.
Rely Nexthop	Next hop after route recursion. If no next hop is found, this field displays **not resolved**.
Original nexthop	Original next hop of the route. If the route was obtained from a BGP update message, the original next hop is the next hop IP address in the message.
OutLabel	Outgoing label of the route.
Ext-Community	Extended community attributes: • **RT**. • **Encapsulation Type**. • **ESI Label**.
RxPathID	Add-Path ID in the received route. This field is not supported by the BGP EVPN address family in the current software version.
TxPathID	Add-Path ID in the sent route. This field is not supported by the BGP EVPN address family in the current software version.
AS-path	AS_PATH attribute of the route. This attribute records the ASs the route has passed and avoids routing loops.
Origin	Origin of the route: • **igp**—Originated in the AS. The origin of routes advertised by using the **network** command is **IGP**. • **egp**—Learned through EGP. • **incomplete**—Unknown origin. The origin of routes redistributed from IGP protocols is **incomplete**.
Attribute value	Attributes of the route: • **MED**—MED value for the destination network. • **localpref**—Local preference value. • **pref-val**—Preferred value. • **pre**—Route preference value.
State	Current state of the route: • **valid**. • **internal**. • **external**. • **local**. • **synchronize**. • **best**.
IP precedence	IP precedence in the range of 0 to 7. **N/A** indicates that the IP precedence is invalid.

Field	Description
QoS local ID	QoS local ID in the range of 1 to 4095. **N/A** indicates that the QoS local ID is invalid.
Traffic index	Traffic index in the range of 1 to 64. **N/A** indicates that the traffic index is invalid.
MPLS label	MPLS label. This field is not supported in the current software version.

Display detailed information about BGP EVPN route [2][5][48][0001-0203-0405][32][4.5.5.5]/136 with RD 1.1.1.1:100.

```
<Sysname> display bgp l2vpn evpn route-distinguisher 1.1.1.1:100
[2][5][48][0001-0203-0405][32][5.5.5.5] 136

 BGP local router ID: 172.16.250.133
 Local AS number: 100



 Route distinguisher: 1.1.1.1:100
 Total number of routes: 1
 Paths:   1 available, 1 best

 BGP routing table information of [2][5][48][0001-0203-0405][32][5.5.5.5]/136:
 From            : 10.1.1.2 (192.168.56.17)
 Rely nexthop    : 10.1.1.2
 Original nexthop: 10.1.1.2
 OutLabel        : NULL
 Ext-Community   : <RT: 1:2>, <RT: 1:3>, <RT: 1:4>, <RT: 1:5>, <RT: 1:6>, <RT: 1:7
                   >, <Encapsulation Type: VXLAN>, <Router's Mac: 0006-0708-0910
                   >, <MAC Mobility: Flag 0, SeqNum 2>, <Default GateWay>
 RxPathID        : 0x0
 TxPathID        : 0x0
 AS-path         : 200
 Origin          : igp
 Attribute value : MED 0, pref-val 0
 State           : valid, external, best
 IP precedence   : N/A
 QoS local ID    : N/A
 Traffic index   : N/A
 EVPN route type : MAC/IP advertisement route
 ESI             : 00:00:00:00:00:00:00:00:00:00
 Ethernet tag ID : 5
 MAC address     : 0001-0203-0405
 IP address      : 5.5.5.5/32
 MPLS label1     : 10
 MPLS label2     : 0
```

Table 3 Command output

Field	Description
Ext-Community	Extended community attributes:

Field	Description

	• **RT**.
	• **Encapsulation Type**.
	• **Router's Mac**.
	• **MAC Mobility**—MAC mobility.
	  ○ **Flag**—Indicates whether the MAC address can move. A value of 1 indicates that the MAC address cannot move, and a value of 0 indicates that the MAC address can move.
	  ○ **SeqNum**—Identifies the most recent move of the MAC address.
	• **Default GateWay**—Route for the default gateway.

MPLS label1	VXLAN ID used for Layer 2 forwarding.
MPLS label2	L3 VXLAN ID used for Layer 3 forwarding.

Display detailed information about BGP EVPN route [3][0][32][5.5.5.5]/80 with RD 1.1.1.1:100.

```
<Sysname> display bgp l2vpn evpn route-distinguisher 1.1.1.1:100 [3][0][32][4.5.5.5] 80

BGP local router ID: 172.16.250.133
 Local AS number: 100


 Route distinguisher: 1.1.1.1:100
 Total number of routes: 1
 Paths:   1 available, 1 best

 BGP routing table information of [3][0][32][4.5.5.5]/80:
 From            : 10.1.1.2 (192.168.56.17)
 Rely nexthop    : 10.1.1.2
 Original nexthop: 10.1.1.2
 OutLabel        : NULL
 Ext-Community   : <RT: 1:2>, <Encapsulation Type: VXLAN>
 RxPathID        : 0x0
 TxPathID        : 0x0
 AS-path         : 200
 Origin          : igp
 Attribute value : MED 0,pref-val 0
 State           : valid, external, best
 IP precedence   : N/A
 QoS local ID    : N/A
 Traffic index   : N/A
 EVPN route type : Inclusive multicast Ethernet tag route
 Ethernet tag ID : 0
 Origin address  : 5.5.5.5/32
```

Table 4 Command output

Field	Description
Ext-Community	Extended community attributes: • **RT**. • **Encapsulation Type**.

Field	Description
Origin address	IP address of the originating router.

Display detailed information about BGP EVPN route [4][00:00:00:00:00:00:00:00:00:00][32][4.5.5.5]/128 with RD 1.1.1.1:100.

```
<Sysname> display bgp l2vpn evpn route-distinguisher 1.1.1.1:100
[4][00:00:00:00:00:00:00:00:00:00][32][4.5.5.5] 128

 BGP local router ID: 172.16.250.133
 Local AS number: 100


 Route distinguisher: 1.1.1.1:100
 Total number of routes: 1
 Paths:   1 available, 1 best

 BGP routing table information of [4][00:00:00:00:00:00:00:00:00:00][32][4.5.5.5]/128:
 From            : 10.1.1.2 (192.168.56.17)
 Rely nexthop    : 10.1.1.2
 Original nexthop: 10.1.1.2
 OutLabel        : NULL
 Ext-Community   : <RT: 1:2>, <Encapsulation Type: VXLAN>, <ES-Import RT: 1:1>
 RxPathID        : 0x0
 TxPathID        : 0x0
 AS-path         : 200
 Origin          : igp
 Attribute value : MED 0,pref-val 0
 State           : valid, external, best
 IP precedence   : N/A
 QoS local ID    : N/A
 Traffic index   : N/A
 EVPN route type : Ethernet segment route
 ESI             : 00:00:00:00:00:00:00:00:00:00
 Origin address  : 4.5.5.5/32
```

Table 5 Command output

Field	Description
Ext-Community	Extended community attributes: • **RT**. • **Encapsulation Type**. • **ES-Import RT**.
Origin address	IP address of the originating router.

Display detailed information about BGP EVPN route [5][10][32][4.5.5.5]/80 with RD 1.1.1.1:100.

```
<Sysname> display bgp l2vpn evpn route-distinguisher 1.1.1.1:100 [5][10][32][4.5.5.5] 80

 BGP local router ID: 172.16.250.133
 Local AS number: 100
```

```
Route distinguisher: 1.1.1.1:100
Total number of routes: 1
Paths:   1 available, 1 best

BGP routing table information of [5][10][32][4.5.5.5]/80:
From            : 10.1.1.2 (192.168.56.17)
Rely nexthop    : 10.1.1.2
Original nexthop: 10.1.1.2
OutLabel        : NULL
Ext-Community   : <RT: 1:2>, <Encapsulation Type: VXLAN>, <Router's Mac:
                  0006-0708-0910>
RxPathID        : 0x0
TxPathID        : 0x0
AS-path         : 200
Origin          : igp
Attribute value : MED 0,pref-val 0
State           : valid, external, best
IP precedence   : N/A
QoS local ID    : N/A
Traffic index   : N/A
EVPN route type : IP prefix advertisement route
ESI             : 00:00:00:00:00:00:00:00:00:00
Ethernet tag ID : 10
IP address      : 4.5.5.5/32
Gateway address : 0.0.0.0
MPLS Label      : 1
```

Table 6 Command output

Field	Description
Ext-Community	Extended community attributes: • **RT**. • **Encapsulation Type**. • **Router's Mac**.
IP address	IP address and prefix length.
MPLS Label	L3 VXLAN ID used for Layer 3 forwarding.

Display detailed information about BGP EVPN route [4][00:00:00:00:00:00:00:00:00:00][32][4.5.5.5]/128 with RD 1.1.1.1:100.

```
<Sysname> display bgp l2vpn evpn route-distinguisher 1.1.1.1:100
[4][00:00:00:00:00:00:00:00:00:00] [32][4.5.5.5] 128 advertise-info

BGP local router ID: 172.16.250.133
Local AS number: 100

Route distinguisher: 1.1.1.1:100
Total number of routes: 1
```

```
Paths:   1 best

BGP routing table information of [4][00:00:00:00:00:00:00:00:00:00][32][4.5.5.5]/128:
Advertised to peers (1 in total):
   10.2.1.2
```

Table 7 Command output

Field	Description
Paths	Number of optimal routes.
Advertised to peers (1 in total)	Peers to whom the route has been advertised and the number of the peers.

display evpn auto-discovery

Use **display evpn auto-discovery** to display information about peers that are automatically discovered through BGP.

Syntax

display evpn auto-discovery { **imet** [**peer** *ip-address*] [**vsi** *vsi-name*] | **macip-prefix** [**nexthop** *next-hop*] [**count**] }

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

imet: Specifies peers discovered through IMET routes.

peer *ip-address*: Specifies a peer by its IP address. If you do not specify this option, the command displays information about all automatically discovered peers.

vsi *vsi-name*: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays peer information for all VSIs.

macip-prefix: Specifies peers discovered through MAC/IP advertisement routes and IP prefix advertisement routes.

nexthop *next-hop*: Specifies a next hop. If you do not specify this option, the command displays peer information for all next hops.

count: Displays the number of peers. If you do not specify this keyword, the command displays detailed peer information.

Examples

Display information about peers discovered through IMET routes.

```
<Sysname> display evpn auto-discovery imet
Total number of automatically discovered peers: 2

VSI name: vpna
RD                   PE_address      Tunnel_address  Tunnel mode VXLAN ID
1:10                 2.2.2.2         2.2.2.2         VXLAN       10
2:100                3.3.3.3         3.3.3.3         VXLAN       10
```

Table 8 Command output

Field	Description
PE_address	Identifier of the remote VTEP on the VSI.
Tunnel_address	Tunnel destination IP address.
Tunnel mode	Tunnel mode: • **VXLAN**. • **VXLAN-DCI**.

Display information about peers discovered through MAC/IP advertisement routes and IP prefix advertisement routes.

```
<Sysname> display evpn auto-discovery macip-prefix
Destination IP  Source IP       L3VNI           Tunnel mode Outgoing interface
1.1.1.1         3.3.3.3         200             VXLAN       Vsi-interface3
2.2.2.2         3.3.3.3         200             VXLAN       Vsi-interface3
```

Display the total number of peers discovered through MAC/IP advertisement routes and IP prefix advertisement routes.

```
<Sysname> display evpn auto-discovery macip-prefix count
Total number of entries: 2
```

Table 9 Command output

Field	Description
Destination IP	Tunnel destination IP address.
Source IP	Tunnel source IP address.
L3VNI	L3 VXLAN ID used for Layer 3 forwarding.
Tunnel mode	Tunnel mode: • **VXLAN**. • **VXLAN-DCI**.
Outgoing interface	VSI interface associated with the L3 VXLAN ID.

display evpn drni synchronized-mac

Use **display evpn drni synchronized-mac** to display DR-synchronized MAC address entries.

Syntax

display evpn drni synchronized-mac [**vsi** *vsi-name*] [**count**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

vsi *vsi-name*: Specifies a VSI name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command displays MAC address entries for all VSIs.

count: Displays the number of MAC address entries that match the command. If you do not specify this keyword, the command displays detailed information about MAC address entries.

Usage guidelines

To ensure VM reachability information consistency in a DR system, DR member devices synchronize MAC address entries and ARP packets with each other through an IPL. This command displays the synchronized MAC address entries from a DR peer.

Examples

Display all DR-synchronized MAC address entries.

```
<Sysname> display evpn drni synchronized-mac
VSI name: bbb
MAC address           Link ID      Interface
0000-0000-000a        1            BAGG10
0000-0000-0009        0            Tunnel1
```

Display the total number of DR-synchronized MAC address entries.

```
<Sysname> display evpn drni synchronized-mac count
Total number of entries: 2
```

Table 10 Command output

Field	Description
Link ID	Link ID of an AC or VXLAN tunnel on a VSI.
Interface	Outgoing interface of a MAC address.

display evpn route arp

Use **display evpn route arp** to display EVPN ARP entries.

Syntax

display evpn route arp [**local** | **remote**] [**public-instance** | **vpn-instance** *vpn-instance-name*] [**count**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

local: Specifies local ARP entries.

remote: Specifies remote ARP entries.

public-instance: Specifies the public instance.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

count: Displays the number of ARP entries. If you do not specify this keyword, the command displays detailed information about ARP entries.

Usage guidelines

If you do not specify the **local** or **remote** keyword, this command displays both local and remote EVPN ARP entries.

If you do not specify the **public-instance** keyword or the **vpn-instance** *vpn-instance-name* option, this command displays EVPN ARP entries for the public instance and all VPN instances.

Examples

Display all EVPN ARP entries.

```
<Sysname> display evpn route arp
Flags: D - Dynamic   B - BGP      L - Local active
       G - Gateway   S - Static   M - Mapping

VPN instance: vpn1                              Interface: Vsi-interface1
IP address      MAC address     Router MAC      VSI Index   Flags
10.1.1.1        0003-0003-0003  a0ce-7e40-0400  0           GL
10.1.1.11       0001-0001-0001  a0ce-7e40-0400  0           DL
10.1.1.12       0001-0001-0011  a0ce-7e41-0401  0           B
10.1.1.13       0001-0001-0021  a0ce-7e42-0402  0           B

Public instance                                 Interface: Vsi-interface2
IP address      MAC address     Router MAC      VSI index   Flags
11.1.1.1        0033-0033-0033  a0ce-7e40-0400  0           GL
11.1.1.11       0011-0011-0011  a0ce-7e40-0400  0           DL
```

Display the total number of EVPN ARP entries.

```
<Sysname> display evpn route arp count
Total number of entries: 6
```

Table 11 Command output

Field	Description
Interface	VSI interface.
Flags	ARP entry type: • **D**—The entry is dynamically learned. • **B**—The entry is learned from BGP EVPN routes. • **L**—The local entry is active. If this flag is not set and the B flag is set, the entry learned from BGP EVPN routes is active. • **G**—The entry for the gateway is active. • **S**—The static entry is active. • **M**—The entry from a remote VXLAN mapped to a local VXLAN is active.

display evpn route arp suppression

Use **display evpn route arp suppression** to display EVPN ARP flood suppression entries.

Syntax

display evpn route arp suppression [**local** | **remote**] [**vsi** *vsi-name*] [**count**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

local: Specifies local ARP flood suppression entries.

remote: Specifies remote ARP flood suppression entries.

vsi *vsi-name*: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays ARP flood suppression entries for all VSIs.

count: Displays the number of ARP flood suppression entries. If you do not specify this keyword, the command displays detailed information about ARP flood suppression entries.

Usage guidelines

If you do not specify the **local** or **remote** keyword, this command displays both local and remote EVPN ARP flood suppression entries.

Examples

Display all EVPN ARP flood suppression entries.

```
<Sysname> display evpn route arp suppression
Flags: D - Dynamic   B - BGP      L - Local active
       G - Gateway   S - Static   M - Mapping

VSI name: vpna
IP address      MAC address     Flags
10.1.1.12       0002-0002-0002  B
```

Display the total number of ARP flood suppression entries.

```
<Sysname> display evpn route arp suppression count
Total number of entries: 1
```

Table 12 Command output

Field	Description
Flags	ARP flood suppression entry type: • **D**—The entry is dynamically learned. • **B**—The entry is learned from BGP EVPN routes. • **L**—The local entry is active. If this flag is not set and the B flag is set, the entry learned from BGP EVPN routes is active. • **G**—The entry for the gateway is active. • **S**—The static entry is active. • **M**—The entry from a remote VXLAN mapped to a local VXLAN is active.

display evpn route mac

Use **display evpn route mac** to display EVPN MAC address entries.

Syntax

display evpn route mac [**local** | **remote**] [**vsi** *vsi-name*] [**count**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

local: Specifies local MAC address entries.

remote: Specifies remote MAC address entries.

vsi *vsi-name*: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays MAC address entries for all VSIs.

count: Displays the number of MAC address entries. If you do not specify this keyword, the command displays detailed information about MAC address entries.

Usage guidelines

If you do not specify the **local** or **remote** keyword, this command displays both local and remote EVPN MAC address entries.

Examples

Display all EVPN MAC address entries.

```
<Sysname> display evpn route mac
Flags: D - Dynamic   B - BGP      L - Local active
       G - Gateway   S - Static   M - Mapping

VSI name: bbb
MAC address     Link ID/Name    Flags   Next hop
0000-0000-000a  1               DB      -
0000-0000-0009  Tunnel1         B       2.2.2.2
```

Display the total number of EVPN MAC address entries.

```
<Sysname> display evpn route mac count
Total number of entries: 2
```

Table 13 Command output

Field	Description
Link ID/Name	For a local MAC address, this field displays the AC's link ID on the VSI. For a remote MAC address, this field displays the tunnel interface name.
Flags	MAC address entry type: • **D**—The entry is dynamically learned. • **B**—The entry is learned from BGP EVPN routes. • **L**—The local entry is active. If this flag is not set and the B flag is set, the entry learned from BGP EVPN routes is active. • **G**—The entry for the gateway is active. • **S**—The static entry is active. • **M**—The entry from a remote VXLAN mapped to a local VXLAN is active.
Next hop	IP address of the remote VTEP. If the MAC address entry is a local entry, a hyphen (**-**) is displayed.

display evpn route nd

Use **display evpn route nd** to display EVPN ND entries.

Syntax

display evpn route nd [**local** | **remote**] [**public-instance** | **vpn-instance** *vpn-instance-name*] [**count**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

local: Specifies local ND entries.

remote: Specifies remote ND entries.

public-instance: Specifies the public instance.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

count: Displays the number of ND entries. If you do not specify this keyword, the command displays detailed information about ND entries.

Usage guidelines

If you do not specify the **local** or **remote** keyword, this command displays both local and remote EVPN ND entries.

If you do not specify the **public-instance** keyword or the **vpn-instance** *vpn-instance-name* option, this command displays EVPN ND entries for the public instance and all VPN instances.

Examples

Display all EVPN ND entries.

```
<Sysname> display evpn route nd
Flags: D - Dynamic   B - BGP      L - Local active
       G - Gateway   S - Static   M - Mapping

VPN instance: vpn1                          Interface: Vsi-interface1
IPv6 address :   AD80:0300:1000:0050:0200:0300:0100:0012
MAC address  :   0001-0001-0001       Router MAC   :   a0ce-7e40-0400
VSI index    :   0                    Flags        :   GL

IPv6 address :   AD10:0300:1000:0020:0200:0300:0100:0022
MAC address  :   0001-0001-0002       Router MAC   :   a0ce-7e40-0411
VSI index    :   0                    Flags        :   GL

Public instance                             Interface: Vsi-interface1
IPv6 address :   BC80:0300:1000:0050:0200:0300:0100:0033
MAC address  :   0002-0002-0001       Router MAC   :   a0ce-7e40-0422
VSI index    :   0                    Flags        :   GL

IPv6 address :   BC10:0300:1000:0020:0200:0300:0100:0034
MAC address  :   0002-0002-0002       Router MAC   :   a0ce-7e40-0433
VSI index    :   0                    Flags        :   GL
```

Display the total number of EVPN ND entries.

```
<Sysname>display evpn route nd count
Total number of entries: 2
```

Table 14 Command output

Field	Description
Interface	VSI interface.
Flags	ND entry type: • **D**—The entry is dynamically learned. • **B**—The entry is learned from BGP EVPN routes. • **L**—The local entry is active. If this flag is not set and the B flag is set, the entry learned from BGP EVPN routes is active. • **G**—The entry for the gateway is active. • **S**—The static entry is active. This type is not supported in the current software version. • **M**—The entry from a remote VXLAN mapped to a local VXLAN is active.

display evpn routing-table

Use **display evpn routing-table** to display the EVPN routing table for a VPN instance.

Syntax

display evpn routing-table [**ipv6**] { **public-instance** | **vpn-instance** *vpn-instance-name* } [**count**]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

ipv6: Specifies IPv6 information. If you do not specify this keyword, the command displays IPv4 information.

public-instance: Specifies the public instance.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

count: Displays the number of entries in the routing table. If you do not specify this keyword, the command displays detailed information about the routing table.

Examples

Display the EVPN IPv4 routing table for VPN instance **vpn1**.

```
<Sysname> display evpn routing-table vpn-instance vpn1

VPN instance name: vpn1                          Local L3VNI: 7
IP address         Nexthop           Outgoing interface       NibID
10.1.1.11          1.1.1.1           Vsi-interface3           0x18000000
10.1.1.12          2.2.2.2           Vsi-interface3           0x18000001
```

Display the EVPN IPv4 routing table for the public instance.

```
<Sysname> display evpn routing-table public-instance

Public instance                                    Local L3VNI: 3900
IP address         Nexthop           Outgoing interface       NibID
10.1.1.11          1.1.1.1           Vsi-interface3           0x18000000
10.1.1.12          2.2.2.2           Vsi-interface3           0x18000001
```

Display the number of EVPN route entries in the IPv4 routing table for VPN instance **vpn1**.

```
<Sysname> display evpn routing-table vpn-instance vpn1 count
Total number of entries: 2
```

Display the EVPN IPv6 routing table for VPN instance **vpna**.

```
<Sysname> display evpn routing-table ipv6 vpn-instance vpna

VPN instance: vpna                              Local L3VNI: 7
IPv6 address        :      BC10:0300:1000:0020:0200:0300:0100:0034
Next hop            :      1.1.1.1
Outgoing interface  :      Vsi-interface3
NibID               :      0x18000000

IPv6 address        :      BC10:0300:1000:0020:0200:0300:0100:0035
Next hop            :      2.2.2.2
Outgoing interface  :      Vsi-interface3
NibID               :      0x18000001
```

Table 15 Command output

Field	Description
Local L3VNI	L3 VXLAN ID associated with the VPN instance or the public instance.
NibID	Next hop ID.

evpn drni group

Use **evpn drni group** to enable EVPN distributed relay and specify the virtual VTEP address.

Use **undo evpn drni group** to restore the default.

Syntax

evpn drni group *virtual-vtep-ip*

undo evpn drni group

Default

EVPN distributed relay is disabled.

Views

System view

Predefined user roles

network-admin

Parameters

virtual-vtep-ip: Specifies the virtual VTEP address.

Usage guidelines

EVPN distributed relay virtualizes two VTEPs or EVPN gateways into one DR system to avoid single points of failure. The VTEPs or EVPN gateways use a virtual VTEP address to establish VXLAN tunnels to remote devices.

For the device to re-establish VXLAN tunnels, you must execute the **address-family l2vpn evpn** command in BGP instance view after you perform one of the following tasks:

- Modify the virtual VTEP address.
- Enable or disable EVPN distributed relay.

Examples

Enable EVPN distributed relay and specify the virtual VTEP address as 1.1.1.1.

```
<Sysname> system-view
[Sysname] evpn drni group 1.1.1.1
```

evpn edge group

Use **evpn edge group** to configure a virtual ED address.

Use **undo evpn edge group** to restore the default.

Syntax

evpn edge group *group-ip*

undo evpn edge group

Default

No virtual ED address is configured.

Views

System view

Predefined user roles

network-admin

Parameters

group-ip: Specifies the virtual ED address.

Usage guidelines

For high availability and load sharing, you can deploy two EDs at a data center. To virtualize the redundant EDs into one device, you must configure the same virtual ED address on them. The redundant EDs use the virtual ED address to establish tunnels with VTEPs and remote EDs.

Redundant EDs cannot provide access service for local VMs. They can act only as EDs. For correct communication, do not redistribute external routes on only one of the redundant EDs. However, you can redistribute the same external routes on both EDs.

On a redundant ED, the virtual ED address must be the IP address of a loopback interface, and it cannot be the BGP peer IP address of the ED.

EVPN-DCI dual-homing is mutually exclusive with EVPN distributed relay. Do not use the **evpn edge group** and **evpn drni group** commands together.

Examples

Configure 1.2.3.4 as the virtual ED address.

```
<Sysname> system-view
[Sysname] evpn edge group 1.2.3.4
```

evpn encapsulation

Use **evpn encapsulation** to create an EVPN instance and enter its view, or enter the view of an existing EVPN instance.

Use **undo evpn encapsulation** to restore the default.

Syntax

evpn encapsulation vxlan

undo evpn encapsulation

Default

No EVPN instance exists.

Views

VSI view

Predefined user roles

network-admin

Parameters

vxlan: Specifies VXLAN encapsulation.

Usage guidelines

Before you can configure EVPN settings, you must create an EVPN instance.

Examples

Create an EVPN instance and enter its view.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan]
```

export route-policy

Use **export route-policy** to apply an export routing policy to EVPN on a VPN instance.

Use **undo export route-policy** to restore the default.

Syntax

export route-policy *route-policy*

undo export route-policy

Default

No export routing policy is applied to EVPN on a VPN instance.

Views

EVPN view of a VPN instance

Predefined user roles

network-admin

Parameters

route-policy: Specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters.

Usage guidelines

You can specify an export routing policy to filter advertised routes or modify their route attributes for EVPN.

If you execute this command multiple times, the most recent configuration takes effect.

EVPN can use an export routing policy specified in VPN instance view or in EVPN view of the VPN instance. Export routing policy configuration in EVPN view takes precedence over that in VPN instance view.

Examples

Apply export routing policy **poly-1** to EVPN on VPN instance **vpn1**.

```
<Sysname> system-view
[Sysname] ip vpn-instance vpn1
[Sysname-vpn-instance-vpn1] address-family evpn
[Sysname-vpn-evpn-vpn1] export route-policy poly-1
```

Related commands

route-policy (*Layer 3—IP Routing Command Reference*)

import route-policy

Use **import route-policy** to apply an import routing policy to EVPN on a VPN instance.

Use **undo import route-policy** to restore the default.

Syntax

import route-policy *route-policy*

undo import route-policy

Default

No import routing policy is applied to EVPN on a VPN instance. The VPN instance accepts a route when the export route targets of the route match local import route targets.

Views

EVPN view of a VPN instance

Predefined user roles

network-admin

Parameters

route-policy: Specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters.

Usage guidelines

You can specify an import routing policy to filter received routes or modify their route attributes for EVPN.

If you execute this command multiple times, the most recent configuration takes effect.

EVPN can use an import routing policy specified in VPN instance view or in EVPN view of the VPN instance. Import routing policy configuration in EVPN view takes precedence over that in VPN instance view.

Examples

Apply import routing policy **poly-1** to EVPN on VPN instance **vpn1**.

```
<Sysname> system-view
[Sysname] ip vpn-instance vpn1
```

```
[Sysname-vpn-instance-vpn1] address-family evpn
[Sysname-vpn-evpn-vpn1] import route-policy poly-1
```

Related commands

route-policy (*Layer 3—IP Routing Command Reference*)

ip forwarding-conversational-learning

Use **ip forwarding-conversational-learning** to enable conversational learning for host route FIB entries.

Use **undo ip forwarding-conversational-learning** to disable conversational learning for host route FIB entries.

Syntax

ip forwarding-conversational-learning [**aging** *aging-time*]

undo ip forwarding-conversational-learning

Default

Conversational learning is disabled for host route FIB entries.

Views

System view

Predefined user roles

network-admin

Parameters

aging *aging-time*: Specifies an aging timer in minutes for host route FIB entries, in the range of 60 to 1440. The default value is 60.

Usage guidelines

Use this command only on an EVPN network.

By default, the device issues a host route FIB entry to the hardware after the entry is generated. This feature enables the device to issue a host route FIB entry to the hardware only when the entry is required for packet forwarding. This feature saves hardware resources on the device.

Set an appropriate aging timer for host route FIB entries according to your network. A much longer or shorter aging timer will degrade the device performance.

- If the aging timer is too long, the device will save many outdated host route FIB entries and fail to accommodate the most recent network changes. These entries cannot be used for correct packet forwarding and exhaust FIB resources.
- If the aging timer is too short, the device will delete the valid host route FIB entries that can still be effective for packet forwarding. As a result, FIB entry flapping will occur, and the device performance will be affected.

Examples

Enable conversational learning for host route FIB entries.

```
<Sysname> system-view
[Sysname] ip forwarding-conversational-learning
```

ip public-instance

Use **ip public-instance** to create the public instance and enter its view, or enter the view of the existing public instance.

Use **undo ip public-instance** to delete the public instance.

Syntax

ip public-instance

undo ip public-instance

Default

The public instance does not exist.

Views

System view

Predefined user roles

network-admin

Usage guidelines

A distributed EVPN gateway uses the public instance to perform Layer 3 forwarding for the public network and to enable communication between private and public networks. The public instance is similar to a VPN instance. A distributed EVPN gateway processes traffic of the public instance in the same way it does for a VPN instance.

Examples

Create the public instance and enter its view.

```
<Sysname> system-view
[Sysname] ip public-instance
[Sysname-public-instance]
```

ip-prefix-route generate disable

Use **ip-prefix-route generate disable** to disable generation of IP prefix advertisement routes for the subnets of a VSI interface.

Use **undo ip-prefix-route generate disable** to enable generation of IP prefix advertisement routes for the subnets of a VSI interface.

Syntax

ip-prefix-route generate disable

undo ip-prefix-route generate disable

Default

The device only generates MAC/IP advertisement routes for a VSI interface that provides centralized VXLAN IP gateway service. The device generates IP prefix advertisement routes for the subnets of a VSI interface that provides distributed VXLAN IP gateway service.

Views

VSI interface view

Predefined user roles

network-admin

Usage guidelines

This command takes effect only on a VSI interface that provides distributed VXLAN IP gateway service (configured by using the **distributed-gateway local** command). It does not take effect on VSI interfaces that provide centralized VXLAN IP gateway service.

Examples

Disable generation of IP prefix advertisement routes for the subnets of VSI-interface 1.

```
<Sysname> system-view
[Sysname] interface vsi-interface 1
[Sysname-Vsi-interface1] ip-prefix-route generate disable
```

l3-vni

Use **l3-vni** to configure an L3 VXLAN ID for a VSI interface or for the public instance.

Use **undo l3-vni** to remove the L3 VXLAN ID for a VSI interface or for the public instance.

Syntax

l3-vni *vxlan-id*

undo l3-vni

Default

No L3 VXLAN ID is configured for a VSI interface or for the public instance.

Views

VSI interface view

Public instance view

Predefined user roles

network-admin

Parameters

vxlan-id: Specifies a VXLAN ID in the range of 0 to 16777215.

Usage guidelines

On distributed EVPN gateways, you must configure L3 VXLAN IDs for the gateways to differentiate traffic of different VPN instances.

To forward Layer 3 traffic of a VPN instance, you must assign an L3 VXLAN ID to the VSI interface of the VPN instance. To forward Layer 3 traffic of the public network, you must assign the same L3 VXLAN ID to the public instance and the VSI interface of the public instance.

To modify the L3 VXLAN ID for the public instance, you must first delete the original L3 VXLAN ID.

Examples

Configure the L3 VXLAN ID as 1000 for VSI-interface 100.

```
<Sysname> system-view
[Sysname] interface vsi-interface 100
[Sysname-Vsi-interface100] l3-vni 1000
```

mac-address forwarding-conversational-learning

Use **mac-address forwarding-conversational-learning** to enable conversational learning for remote MAC address entries.

Use **undo mac-address forwarding-conversational-learning** to disable conversational learning for remote MAC address entries.

Syntax

mac-address forwarding-conversational-learning

undo mac-address forwarding-conversational-learning

Default

Conversational learning is disabled for remote MAC address entries.

Views

System view

Predefined user roles

network-admin

Usage guidelines

Use this command only on an EVPN network.

By default, the device issues a remote MAC address entry to the hardware after the remote MAC address is advertised to the local site by BGP EVPN routes. This feature enables the device to issue a remote MAC address entry to the hardware only when the entry is required for packet forwarding. This feature saves hardware resources on the device.

With this feature enabled, the device generates a blackhole MAC address entry for an unknown MAC address if receiving 50 frames destined for that MAC address within the MAC aging time. For more information about the MAC aging time and blackhole MAC address entries, see MAC address table configuration in *Layer 2—LAN Switching Configuration Guide*.

Examples

Enable conversational learning for remote MAC address entries.

```
<Sysname> system-view
[Sysname] mac-address forwarding-conversational-learning
```

mac-advertising disable

Use **mac-advertising disable** to disable MAC address advertisement and withdraw advertised MAC addresses.

Use **undo mac-advertising disable** to restore the default.

Syntax

mac-advertising disable

undo mac-advertising disable

Default

MAC address advertisement is enabled.

Views

EVPN instance view

Predefined user roles

network-admin

Usage guidelines

The MAC information and ARP information advertised by the VTEP overlap. To avoid duplication, use this command to disable MAC address advertisement and withdraw the MAC addresses advertised to remote VTEPs.

Examples

Disable MAC address advertisement and withdraw advertised MAC addresses for an EVPN instance.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] mac-advertising disable
```

mapping vni

Use **mapping vni** to map a local VXLAN to a remote VXLAN.

Use **undo mapping vni** to restore the default.

Syntax

mapping vni *vxlan-id*

undo mapping vni

Default

A local VXLAN is not mapped to any remote VXLAN.

Views

EVPN instance view

Predefined user roles

network-admin

Parameters

vxlan-id: Specifies a remote VXLAN ID in the range of 0 to 16777215.

Usage guidelines

The VXLAN mapping feature provides Layer 2 connectivity for a tenant subnet that uses different VXLAN IDs in multiple data centers.

VXLAN mapping includes the following types:

- **Non-intermediate VXLAN mapping**—When two data centers use different VXLAN IDs for a subnet, map the local VXLAN to the remote VXLAN on the ED of one data center. For example, for VXLAN 10 of data center 1 to communicate with VXLAN 20 of data center 2, map VXLAN 10 to VXLAN 20 on the ED of data center 1.
- **Intermediate VXLAN mapping**—When multiple data centers use different VXLAN IDs for a subnet, map the VXLANs to an intermediate VXLAN on all EDs. For example, data center 1 uses VXLAN 10, data center 2 uses VXLAN 20, and data center 3 uses VXLAN 30. To provide connectivity for the VXLANs, map them to intermediate VXLAN 500 on EDs of the data centers. You must use intermediate VXLAN mapping if more than two data centers use different VXLAN IDs.

You must create mapped remote VXLANs on the device, create an EVPN instance for each remote VXLAN, and configure RD and route target settings for the EVPN instances.

Examples

Map local VXLAN 100 to remote VXLAN 200.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] vxlan 100
[Sysname-vsi-aaa-vxlan-100] quit
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] mapping vni 200
```

nd mac-learning disable

Use **nd mac-learning disable** to disable an EVPN instance from learning MAC addresses from ND information.

Use **undo nd mac-learning disable** to restore the default.

Syntax

nd mac-learning disable

undo nd mac-learning disable

Default

An EVPN instance learns MAC addresses from ND information.

Views

EVPN instance view

Predefined user roles

network-admin

Usage guidelines

The MAC information and ND information advertised by a remote VTEP overlap. To avoid duplication, use this command to disable the learning of MAC addresses from ND information. EVPN will learn remote MAC addresses only from the MAC information advertised from remote sites.

Examples

Disable an EVPN instance from learning MAC addresses from ND information.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] nd mac-learning disable
```

peer next-hop-invariable

Use **peer next-hop-invariable** to configure the device to not change the next hop of routes advertised to an EBGP peer or peer group.

Use **undo peer next-hop-invariable** to configure the device to use its address as the next hop of routes advertised to an EBGP peer or peer group.

Syntax

peer { *group-name* | *ipv4-address* [*mask-length*] } **next-hop-invariable**

undo peer { *group-name* | *ipv4-address* [*mask-length*] } **next-hop-invariable**

Default

The device uses its address as the next hop of routes advertised to EBGP peers or peer groups.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters.

ipv4-address: Specifies a peer by its IPv4 address.

mask-length: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments.

Usage guidelines

This command is exclusive with the **peer next-hop-local** command.

The next hop in BGP EVPN routes is the IP address of the originating VTEP. By default, the device replaces the next hop of IBGP routes with its address when advertising the routes to an EBGP peer. If the device is a transport network device, it will modify the next hop of BGP EVPN routes. For VTEPs to learn one another's IP address, you must configure the device to not change the next hop of routes advertised to EBGP peers.

Examples

Configure the device to not change the next hop of routes advertised to EBGP peer 1.1.1.1.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] peer 1.1.1.1 next-hop-invariable
```

Related commands

peer next-hop-local (*Layer 3—IP Routing Command Reference*)

peer router-mac-local

Use **peer router-mac-local** to enable route router MAC replacement for a peer or peer group.

Use **undo peer router-mac-local** to cancel route router MAC replacement configuration for a peer or peer group.

Syntax

peer { *group-name* | *ipv4-address* [*mask-length*] } **router-mac-local**

undo peer { *group-name* | *ipv4-address* [*mask-length*] } **router-mac-local**

Default

The device does not modify the router MAC address of routes before advertising the routes.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must exist.

ipv4-address: Specifies a peer by its IPv4 address. The peer must exist.

mask-length: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments

Usage guidelines

This command enables an ED to use its router MAC address to replace the router MAC address of routes received from and advertised to a peer or peer group in the local data center. The router MAC replacement process is as follows:

- For routes received from the peer or peer group, the ED performs router MAC replacement and advertises the routes to remote EDs.
- For routes received from a remote data center, the ED performs router MAC replacement and advertises the routes to the peer or peer group.

Examples

In BGP EVPN address family view, enable route router MAC replacement for peer 1.1.1.1.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] peer 1.1.1.1 router-mac-local
```

policy vpn-target

Use **policy vpn-target** to enable route target filtering for BGP EVPN routes.

Use **undo policy vpn-target** to disable route target filtering for BGP EVPN routes.

Syntax

policy vpn-target

undo policy vpn-target

Default

Route target filtering is enabled for BGP EVPN routes.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Usage guidelines

When route target filtering is enabled for BGP EVPN routes, the EVPN routing table accepts only BGP EVPN routes of which the export route targets match the local import route targets. If the device must save all BGP EVPN routes, use the **undo policy vpn-target** command to disable route target filtering for BGP EVPN routes.

Examples

Disable route target filtering for BGP EVPN routes.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] undo policy vpn-target
```

route-distinguisher (EVPN instance view)

Use **route-distinguisher** to configure an RD for an EVPN instance.

Use **undo route-distinguisher** to restore the default.

Syntax

route-distinguisher { *route-distinguisher* | **auto** }

undo route-distinguisher

Default

No RD is configured for an EVPN instance.

Views

EVPN instance view

Predefined user roles

network-admin

Parameters

route-distinguisher: Specifies an RD, a string of 3 to 21 characters. The RD cannot be all zeros and can use one of the following formats:

- *16-bit AS number*:*32-bit user-defined number*. For example, 101:3.
- *32-bit IP address*:*16-bit user-defined number*. For example, 192.168.122.15:1.
- *32-bit AS number*:*16-bit user-defined number*. For example, 65536:1. The AS number must be equal to or greater than 65536.

auto: Automatically generates an RD in the *N*:*VXLAN ID* format. The initial value of *N* is 1. If *N*:*VXLAN ID* is already in use, the system increases the value of *N* by 1 until the RD is available.

Usage guidelines

EVPN uses MP-BGP to advertise BGP EVPN routes for automatic VTEP discovery, MAC reachability information advertisement, and host route advertisement. MP-BGP uses the RD to differentiate BGP EVPN routes of different EVPN instances.

Examples

Configure 22:1 as the RD of an EVPN instance.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] route-distinguisher 22:1
```

route-distinguisher (public instance view)

Use **route-distinguisher** to configure an RD for the public instance.

Use **undo route-distinguisher** to restore the default.

Syntax

route-distinguisher *route-distinguisher*

undo route-distinguisher

Default

No RD is configured for the public instance.

Views

Public instance view

Predefined user roles

network-admin

Parameters

route-distinguisher: Specifies an RD, a string of 3 to 21 characters. The RD can use one of the following formats:

- *16-bit AS number*:*32-bit user-defined number*. For example, 101:3.

- *32-bit IP address*:*16-bit user-defined number*. For example, 192.168.122.15:1.
- *32-bit AS number*:*16-bit user-defined number*. For example, 65536:1. The AS number must be equal to or greater than 65536.

Usage guidelines

To modify the RD of the public instance, first execute the **undo route-distinguisher** command to remove the original RD.

Examples

Configure 22:1 as the RD of the public instance.

```
<Sysname> system-view
[Sysname] ip public-instance
[Sysname-public-instance] route-distinguisher 22:1
```

rr-filter

Use **rr-filter** to create a route reflector (RR) reflection policy.

Use **undo rr-filter** to restore the default.

Syntax

rr-filter *ext-comm-list-number*

undo rr-filter

Default

An RR does not filter reflected BGP EVPN routes.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Parameters

ext-comm-list-number: Specifies an extended community attribute list by its number in the range of 1 to 199.

Usage guidelines

This command enables an RR to reflect only received BGP EVPN routes that match the attributes in the specified extended community attribute list.

If a cluster contains multiple RRs, you can configure different reflection policies on the RRs for load sharing among the RRs.

For more information about the extended community attribute list, see *Layer 3—IP Routing Configuration Guide*.

Examples

Configure a reflection policy for the device to reflect BGP EVPN routes that match extended community attribute list 10.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] rr-filter 10
```

Related commands

ip extcommunity-list (*Layer 3—IP Routing Command Reference*)

vpn-route cross multipath

Use **vpn-route cross multipath** to enable ECMP VPN route redistribution.

Use **undo vpn-route cross multipath** to disable ECMP VPN route redistribution.

Syntax

vpn-route cross multipath

undo vpn-route cross multipath

Default

ECMP VPN route redistribution is disabled. If multiple routes have the same prefix and RD, BGP only imports the optimal route into the EVPN routing table.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Usage guidelines

ECMP VPN route redistribution enables BGP to import all routes that have the same prefix and RD into the EVPN routing table.

Examples

Enable ECMP VPN route redistribution.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] vpn-route cross multipath
```

vpn-target

Use **vpn-target** to configure route targets for EVPN.

Use **undo vpn-target** to delete route targets for EVPN.

Syntax

In EVPN instance view:

vpn-target { *vpn-target*&<1-8> | **auto** } [**both** | **export-extcommunity** | **import-extcommunity**]

undo vpn-target { *vpn-target*&*<1-8>* | **auto** | **all** } [**both** | **export-extcommunity** | **import-extcommunity**]

In EVPN view of a VPN instance, public instance view, IPv4 VPN view of the public instance, IPv6 VPN view of the public instance, or EVPN view of the public instance:

vpn-target *vpn-target*&<1-8> [**both** | **export-extcommunity** | **import-extcommunity**]

undo vpn-target { **all** | *vpn-target*&<1-8> [**both** | **export-extcommunity** | **import-extcommunity**] }

Default

EVPN does not have route targets.

Views

EVPN instance view

EVPN view of a VPN instance

Public instance view

EVPN view of the public instance

IPv4 VPN view of the public instance

IPv6 VPN view of the public instance

Predefined user roles

network-admin

Parameters

vpn-target&<1-8>: Specifies a space-separated list of up to eight route targets. Each route target is a string of 3 to 21 characters in one of the following formats:

- *16-bit AS number*:*32-bit user-defined number*. For example, 101:3.
- *32-bit IP address*:*16-bit user-defined number*. For example, 192.168.122.15:1.
- *32-bit AS number*:*16-bit user-defined number*. For example, 65536:1. The AS number must be equal to or greater than 65536.

auto: Automatically generates a route target in the format of *BGP AS number*:*VXLAN ID*.

both: Uses the specified route targets as both import and export targets. If you do not specify the **both**, **export-extcommunity**, or **import-extcommunity** keyword, the **both** keyword applies.

export-extcommunity: Uses the specified route targets as export targets.

import-extcommunity: Uses the specified route targets as import targets.

all: Specifies all route targets.

Usage guidelines

EVPN uses MP-BGP to advertise BGP EVPN routes for automatic VTEP discovery, MAC reachability information advertisement, and host route advertisement. MP-BGP uses route targets to control the advertisement and acceptance of BGP EVPN routes.

A VTEP sets the export targets for BGP EVPN routes before advertising the routes to remote VTEPs. The VTEP checks the export targets of BGP EVPN routes from remote VTEPs and imports only BGP EVPN routes of which the export targets match the local import targets.

Examples

Configure import route targets 10:1, 100:1, and 1000:1 for an EVPN instance.

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] evpn encapsulation vxlan
[Sysname-vsi-aaa-evpn-vxlan] vpn-target 10:1 100:1 1000:1 import-extcommunity
```