

H3C S6860 Switch Series

VXLAN Command Reference

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

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

Copyright © 2020, New H3C Technologies Co., Ltd. and its licensors

All rights reserved

No part of this manual may be reproduced or transmitted in any form or by any means without prior written consent of New H3C Technologies Co., Ltd.

Trademarks

Except for the trademarks of New H3C Technologies Co., Ltd., any trademarks that may be mentioned in this document are the property of their respective owners.

Notice

The information in this document is subject to change without notice. All contents in this document, including statements, information, and recommendations, are believed to be accurate, but they are presented without warranty of any kind, express or implied. H3C shall not be liable for technical or editorial errors or omissions contained herein.

Preface

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

GUI conventions

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

Symbols

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

Network topology icons

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

Examples provided in this document

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

Documentation feedback

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

We appreciate your comments.

Contents

VXLAN commands	1
Basic VXLAN commands	1
ac statistics enable	1
arp suppression enable	1
bandwidth (Ethernet service instance view)	2
bandwidth (VSI view)	3
description	4
display arp suppression vsi	4
display hardware-resource vxlan	5
display igmp host group	6
display ipv6 nd suppression vsi	8
display l2vpn mac-address	9
display l2vpn service-instance	11
display l2vpn vsi	13
display vxlan tunnel	16
emulate-ping vxlan	17
encapsulation	18
flooding disable	20
group	21
hardware-resource vxlan	22
igmp host enable	23
ipv6 nd suppression enable	24
ipv6 nd suppression notify-ipsg	24
l2vpn enable	25
l2vpn mac-address software-learning enable	26
learning mode disable	26
mac-address mac-learning priority	27
mac-address static	28
mac-learning enable	29
mac-based ac	30
reserved vxlan	31
reset arp suppression vsi	31
reset ipv6 nd suppression vsi	32
reset l2vpn mac-address	32
reset l2vpn statistics ac	33
reset l2vpn statistics vsi	33
restrain	34
rewrite inbound tag	35
rewrite outbound tag	36
selective-flooding mac-address	37
service-instance	37
shutdown	38
statistics enable (Ethernet service instance view)	39
statistics enable (tunnel interface view)	39
statistics enable (VSI view)	40
tunnel	41
tunnel bfd enable	42
tunnel global source-address	43
tunnel statistics vxlan auto	43
vsi	44
vxlan	45
vxlan default-decapsulation	45
vxlan invalid-udp-checksum discard	46
vxlan invalid-vlan-tag discard	47
vxlan local-mac report	47
vxlan tunnel mac-learning disable	48
vxlan udp-port	48

vxlan vlan-based	49
vxlan vni	50
vxlan-over-vxlan enable	50
xconnect vsi	51
VXLAN IP gateway commands	53
arp distributed-gateway dynamic-entry synchronize	53
arp send-rate	53
bandwidth (VSI interface view)	54
default	55
description	55
display interface vsi-interface	56
distributed-gateway local	58
gateway subnet	59
gateway vsi-interface	60
interface vsi-interface	60
ipv6 nd distributed-gateway dynamic-entry synchronize	61
mac-address	62
mtu	62
reset counters interface vsi-interface	63
shutdown	63
vtep group member local	64
vtep group member remote	65
vxlan tunnel arp-learning disable	65
vxlan tunnel nd-learning disable	66
OVSDB commands	66
ovsdb server bootstrap ca-certificate	66
ovsdb server enable	67
ovsdb server pki domain	68
ovsdb server pssl	69
ovsdb server ptcp	69
ovsdb server ssl	70
ovsdb server tcp	71
vtep access port	72
vtep enable	73
vxlan tunnel flooding-proxy	73

VXLAN commands

Basic VXLAN commands

ac statistics enable

Use **ac statistics enable** to enable packet statistics for Ethernet service instances of a VLAN.

Use **undo ac statistics enable** to disable packet statistics for Ethernet service instances of a VLAN.

Syntax

ac statistics enable

undo ac statistics enable

Default

The packet statistics feature is disabled for Ethernet service instances of a VLAN.

Views

VLAN view

Predefined user roles

network-admin

Usage guidelines

This command enables packet statistics for the Ethernet service instances automatically created for VLAN-based VXLAN assignment.

Before you enable this feature, you must use the **vxlan vlan-based** command to enable VLAN-based VXLAN assignment.

Examples

Map VLAN 10 to VXLAN 100, and enable packet statistics for Ethernet service instances of VLAN 10.

```
<Sysname> system-view
[Sysname] vlan 10
[Sysname-vlan10] vxlan vni 100
[Sysname-vlan10] ac statistics enable
```

Related commands

display l2vpn service-instance

reset l2vpn statistics ac

vxlan vlan-based

arp suppression enable

Use **arp suppression enable** to enable ARP flood suppression.

Use **undo arp suppression enable** to disable ARP flood suppression.

Syntax

arp suppression enable

undo arp suppression enable

Default

ARP flood suppression is disabled.

Views

VSI view

Predefined user roles

network-admin

Usage guidelines

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

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

Examples

```
# Enable ARP flood suppression for VSI vsi1.  
<Sysname> system-view  
[Sysname] vsi vsi1  
[Sysname-vsi-vsi1] arp suppression enable
```

Related commands

display arp suppression vsi

reset arp suppression vsi

bandwidth (Ethernet service instance view)

Use **bandwidth** to set the bandwidth limit for an Ethernet service instance.

Use **undo bandwidth** to restore the default.

Syntax

bandwidth *bandwidth*

undo bandwidth

Default

No bandwidth limit is set for an Ethernet service instance.

Views

Ethernet service instance view

Predefined user roles

network-admin

Parameters

bandwidth: Specifies a bandwidth limit in the range of 64 to 4194303 kbps.

Usage guidelines

The bandwidth limit takes effect only on incoming traffic. An Ethernet service instance drops excess incoming traffic when its bandwidth limit is reached.

If you use the **bandwidth** command for both an Ethernet service instance and its mapped VSI, the Ethernet service instance bandwidth is limited exclusively from the VSI bandwidth. For example, VSI

vs1 has AC 1, AC 2, AC 3, and Tunnel 1. The bandwidth limit of VSI **vs1** is 500 kbps, and the bandwidth limit of AC 1 is 120 kbps. The total bandwidth of AC 2, AC 3, and Tunnel 1 will be limited to 500 kbps.

Examples

```
# Set the bandwidth limit to 10240 kbps for Ethernet service instance 200.
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 200
[Sysname-Ten-GigabitEthernet1/0/1-srv200] encapsulation s-vid 200
[Sysname-Ten-GigabitEthernet1/0/1-srv200] bandwidth 10240
```

Related commands

bandwidth (VSI view)

bandwidth (VSI view)

Use **bandwidth** to set the bandwidth limit for a VSI.

Use **undo bandwidth** to restore the default.

Syntax

bandwidth *bandwidth*

undo bandwidth

Default

No bandwidth limit is set for a VSI.

Views

VSI view

Predefined user roles

network-admin

Parameters

bandwidth: Specifies a bandwidth limit in kbps. For software versions earlier than Release 2612P02, the value range for this argument is 64 to 4194303. For Release 2612P02 and later, the value range for this argument is 64 to 167772159.

Usage guidelines

This command limits the total bandwidth of the incoming traffic on all ACs and VXLAN tunnels mapped to a VSI. A VSI drops excess incoming traffic when its bandwidth limit is reached.

If you use the **bandwidth** command for both an Ethernet service instance and its mapped VSI, the Ethernet service instance bandwidth is limited exclusively from the VSI bandwidth. For example, VSI **vs1** has AC 1, AC 2, AC 3, and Tunnel 1. The bandwidth limit of VSI **vs1** is 500 kbps, and the bandwidth limit of AC 1 is 120 kbps. The total bandwidth of AC 2, AC 3, and Tunnel 1 will be limited to 500 kbps.

If you use both the **restrain** and **bandwidth** commands on a VSI, the **bandwidth** command limits only the bandwidth of the traffic not restrained by the **restrain** command.

As a best practice, do not execute both the **bandwidth** and **selective-flooding mac-address** commands on a VSI. Traffic cannot be forwarded correctly if you use these commands together.

Examples

```
# Set the bandwidth limit to 10240 kbps for VSI vpn1.
```

```
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] bandwidth 10240
```

Related commands

bandwidth (Ethernet service instance view)
display l2vpn vsi

description

Use **description** to configure a description for a VSI.
Use **undo description** to restore the default.

Syntax

```
description text
undo description
```

Default

A VSI does not have a description.

Views

VSI view

Predefined user roles

network-admin

Parameters

text: Specifies a description, a case-sensitive string of 1 to 80 characters.

Examples

```
# Configure a description for VSI vpn1.
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] description vsi for vpn1
```

Related commands

display l2vpn vsi

display arp suppression vsi

Use **display arp suppression vsi** to display ARP flood suppression entries.

Syntax

```
display arp suppression vsi [ name vsi-name ] [ slot slot-number ] [ count ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

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

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays entries on the master device.

count: Displays the number of ARP flood suppression entries that match the command.

Examples

Display ARP flood suppression entries.

```
<Sysname> display arp suppression vsi
IP address      MAC address    Vsi Name      Link ID      Aging
1.1.1.2         000f-e201-0101 vsi1          0x70000     14
1.1.1.3         000f-e201-0202 vsi1          0x80000     18
1.1.1.4         000f-e201-0203 vsi2          0x90000     10
```

Display the number of ARP flood suppression entries.

```
<Sysname> display arp suppression vsi count
Total entries: 3
```

Table 1 Command output

Field	Description
Link ID	Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI.
Aging	Remaining lifetime (in minutes) of the ARP flood suppression entry. When the timer expires, the entry is deleted.

Related commands

arp suppression enable

reset arp suppression vsi

display hardware-resource vxlan

Use **display hardware-resource vxlan** to display the VXLAN hardware resource mode.

Syntax

```
display hardware-resource [ vxlan ]
```

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

vxlan: Specifies the VXLAN hardware resource mode. If you do not specify this keyword, the command displays all hardware resource modes.

Examples

Display the VXLAN hardware resource mode.

```
<Sysname> display hardware-resource vxlan
Vxlan resource(vxlan), all supported modes:
```

```

l2gw          L2 gateway--underlay/overlay 32K/0K
l3gw8k       L3 gateway--underlay/overlay 24K/8K
l3gw16k      L3 gateway--underlay/overlay 16K/16K
l3gw24k      L3 gateway--underlay/overlay 8K/24K
border24k    Border--underlay/overlay 8K/24K
border28k    Border--underlay/overlay 4K/28K

```

```

-----
Default      Current      Next
l2gw         border24k   border24k

```

Table 2 Command output

Field	Description
Vxlan resource(vxlan), all supported modes	VXLAN hardware resource modes supported by the device. <ul style="list-style-type: none"> • l2gw—Layer 2 gateway mode. • l3gw8k—Layer 3 gateway mode that supports 8 K of overlay adjacency table. • l3gw16k—Layer 3 gateway mode that supports 16 K of overlay adjacency table. • l3gw24k—Layer 3 gateway mode that supports 24 K of overlay adjacency table. • border24k—Border mode that supports 24 K of overlay adjacency table. • border28k—Border mode that supports 28 K of overlay adjacency table.
Default	The default VXLAN hardware resource mode.
Current	The current VXLAN hardware resource mode.
Next	The VXLAN hardware resource mode for the next startup.

Related commands

hardware-resource vxlan

display igmp host group

Use **display igmp host group** to display information about the multicast groups that contain IGMP host-enabled interfaces.

Syntax

display igmp host group [*group-address* | **interface** *interface-type interface-number*] [**verbose**]

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

group-address: Specifies a multicast group address in the range of 224.0.1.0 to 239.255.255.255. If you do not specify a multicast group, this command displays information about all multicast groups.

interface interface-type interface-number: Specifies an interface by its type and number. If you do not specify an interface, this command displays multicast group information for all interfaces.

verbose: Displays detailed multicast group information. If you do not specify this keyword, the command displays brief multicast group information.

Usage guidelines

For the VXLAN multicast source interface of a multicast-mode VXLAN to join its VXLAN multicast group, enable the IGMP host feature on the interface. The VXLAN multicast source interface provides the source IP address for multicast VXLAN packets.

Use this command to verify the following information:

- Multicast group information for VXLANs.
- Group membership status of VXLAN multicast source interfaces.

Examples

Display brief information about all multicast groups that contain IGMP host-enabled interfaces.

```
<Sysname> display igmp host group
IGMP host groups in total: 2
Vlan-interface10(1.1.1.20):
  IGMP host groups in total: 2
  Group address      Member state      Expires
  225.1.1.1         Idle              Off
  225.1.1.2         Idle              Off
```

Display detailed information about all multicast groups that contain IGMP host-enabled interfaces.

```
<Sysname> display igmp host group verbose
Vlan-interface10(1.1.1.20):
  IGMP host groups in total: 2
  Group: 225.1.1.1
    Group mode: Exclude
    Member state: Idle
    Expires: Off
    Source list (sources in total: 0):
  Group: 225.1.1.2
    Group mode: Exclude
    Member state: Idle
    Expires: Off
    Source list (sources in total: 0):
```

Table 3 Command output

Field	Description
IGMP host groups in total	Total number of multicast groups that contain IGMP host-enabled interfaces.
Vlan-interface10(1.1.1.20)	Name and IP address of the IGMP host-enabled interface.
IGMP host groups in total	Total number of multicast groups on the interface.
Group address/Group	Address of the multicast group.
Member state	Member state: <ul style="list-style-type: none"> • Delay—The interface has joined the multicast group, and it has started the delay timer for sending IGMP reports. • Idle—The interface has joined the multicast group, but it has not started the delay timer for sending IGMP reports. The delay timer is not user configurable.
Expires	Remaining delay time for the interface to send an IGMP report.

Field	Description
	This field displays Off if the delay timer is disabled.
Group mode	Multicast source filtering mode: <ul style="list-style-type: none"> • Include. • Exclude.
Source list	Multicast sources of the multicast group.
sources in total	Total number of multicast sources.

NOTE:

For more information about the command output, see IGMP in *IP Multicast Configuration Guide*.

Related commands

`igmp host enable`

display ipv6 nd suppression vsi

Use `display ipv6 nd suppression vsi` to display ND flood suppression entries.

Syntax

`display ipv6 nd suppression vsi [name vsi-name] [slot slot-number] [count]`

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

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

slot slot-number: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays entries on the master device.

count: Displays the number of ND flood suppression entries that match the command.

Usage guidelines

This command is available in Release 2612P06 and later.

Examples

Display ND flood suppression entries.

```
<Sysname> display ipv6 nd suppression vsi
IPv6 address      MAC address      VSI Name      Link ID      Aging (min)
1000::2          000f-e201-0101  vsi1          0x70000     5
1000::3          000f-e201-0202  vsi1          0x80000     5
1000::4          000f-e201-0203  vsi2          0x90000     5
```

Display the number of ND flood suppression entries.

```
<Sysname> display ipv6 nd suppression vsi count
Total entries: 3
```

Table 4 Command output

Field	Description
Link ID	Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI.
Aging (min)	Remaining lifetime (in minutes) of the ND flood suppression entry. When the timer expires, the entry is deleted.

Related commands

ipv6 nd suppression enable

reset ipv6 nd suppression vsi

display l2vpn mac-address

Use **display l2vpn mac-address** to display MAC address entries for VSIs.

Syntax

display l2vpn mac-address [vsi *vsi-name*] [**dynamic**] [**count** | **verbose**]

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.

dynamic: Specifies dynamic MAC address entries learned in the data plane. If you do not specify this keyword, the command displays all MAC address entries, including:

- Dynamic remote- and local-MAC entries.
- Manually added static remote- and local-MAC entries.
- Remote-MAC entries issued through OpenFlow.
- Remote-MAC entries issued through OVSDB.
- Remote-MAC entries advertised through EVPN.

count: Displays the number of MAC address entries.

verbose: Displays detailed information about MAC address entries.

Usage guidelines

If you do not specify the **count** or **verbose** keyword, this command displays brief information about MAC address entries.

Examples

Display brief information about MAC address entries for all VSIs.

```
<Sysname> display l2vpn mac-address
```

```
MAC Address      State    VSI Name      Link ID/Name  Aging
0000-0000-000b  Static  vpn1          Tunnel10     NotAging
0000-0000-000c  Dynamic vpn1          Tunnel160    Aging
0000-0000-000d  Dynamic vpn1          Tunnel199    Aging
--- 3 mac address(es) found ---
```

Display the total number of MAC address entries in all VSIs.

```
<Sysname> display l2vpn mac-address count
3 mac address(es) found
```

Table 5 Command output

Field	Description
State	Entry state: <ul style="list-style-type: none"> • Dynamic—Local- or remote-MAC entry dynamically learned in the data plane. • Static—Static local- or remote-MAC entry. • OpenFlow—Remote-MAC entry issued by a remote controller through OpenFlow. • OVSDB—Remote-MAC entry issued by a remote controller through OVSDB. • EVPN—Remote-MAC entry advertised through EVPN.
Link ID/Name	For a local MAC address, this field displays the name of the interface that hosts the Ethernet service instance for the MAC address. For a remote MAC address, this field displays the tunnel interface name.
Aging	Entry aging state: <ul style="list-style-type: none"> • Aging. • NotAging.

Display detailed information about MAC address entries for all VSIs.

```
<Sysname> display l2vpn mac-address verbose
MAC Address : 0000-0000-000b
VSI Name    : vpn1
VXLAN ID   : 123
Interface   : XGE1/0/1
Link ID     : 1
State       : Dynamic
Aging       : Aging
```

Table 6 Command output

Field	Description
Interface	For a local MAC address, this field displays the name of the interface that hosts the Ethernet service instance for the MAC address. For a remote MAC address, this field displays the tunnel interface name.
Link ID	Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI.
State	Entry state: <ul style="list-style-type: none"> • Dynamic—Local- or remote-MAC entry dynamically learned in the data plane. • Static—Static local- or remote-MAC entry. • OpenFlow—Remote-MAC entry issued by a remote controller through OpenFlow. • OVSDB—Remote-MAC entry issued by a remote controller through OVSDB. • EVPN—Remote-MAC entry advertised through EVPN.
Aging	Entry aging state: <ul style="list-style-type: none"> • Aging. • NotAging.

Related commands

reset l2vpn mac-address

display l2vpn service-instance

Use **display l2vpn service-instance** to display information about Ethernet service instances.

Syntax

```
display l2vpn service-instance [ interface interface-type interface-number [ service-instance instance-id ] ] [ verbose ]
```

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

interface *interface-type interface-number*: Specifies a Layer 2 Ethernet interface or Layer 2 aggregate interface by its interface type and number. If you do not specify an interface, this command displays Ethernet service instance information for all Layer 2 Ethernet interfaces and Layer 2 aggregate interfaces.

service-instance *instance-id*: Specifies an Ethernet service instance by its ID in the range of 1 to 4096. If you do not specify an Ethernet service instance, this command displays information about all Ethernet service instances on the specified interface.

verbose: Displays detailed information about Ethernet service instances. If you do not specify this keyword, the command displays brief information about Ethernet service instances.

Examples

```
# Display brief information about all Ethernet service instances.
```

```
<Sysname> display l2vpn service-instance
```

```
Total number of service-instances: 4, 4 up, 0 down
```

```
Total number of ACs: 2, 2 up, 0 down
```

Interface	SrvID	Owner	LinkID	State	Type
XGE1/0/1	3	vsi12	1	Up	VSI
XGE1/0/1	4	vsi13	1	Up	VSI

Table 7 Command output

Field	Description
Total number of ACs	Total number of attachment circuits (ACs) and the number of ACs in each state (up or down).
Interface	Name of a Layer 2 Ethernet interface or Layer 2 aggregate interface.
SrvID	Ethernet service instance ID.
Owner	VSI name. This field is empty if an Ethernet service instance is not mapped to any VSI.
LinkID	Ethernet service instance's link ID on the VSI.
State	Ethernet service instance state: <ul style="list-style-type: none">• Up.• Down.
Type	L2VPN type of the Ethernet service instance: <ul style="list-style-type: none">• VSI.

Field	Description
	<ul style="list-style-type: none"> • VPWS.

Display detailed information about all Ethernet service instances on Ten-GigabitEthernet 1/0/1.

```
<Sysname> display l2vpn service-instance interface ten-gigabitethernet 1/0/1 verbose
```

```
Interface: XGE1/0/1
```

```
Service Instance: 1
  Type           : Manual
  Encapsulation  : s-vid 16
  Bandwidth      : Unlimited
  VSI Name       : vsi10
  Link ID        : 1
  State          : Up
  Statistics     : Enabled
  Input Statistics:
    Octets       :0
    Packets      :0
  Output Statistics:
    Octets       :0
    Packets      :0
```

Table 8 Command output

Field	Description
Interface	Name of a Layer 2 Ethernet interface or Layer 2 aggregate interface.
Service Instance	Ethernet service instance ID.
Type	Type and traffic match mode of the Ethernet service instance: <ul style="list-style-type: none"> • Dynamic (MAC-based)—Dynamic Ethernet service instance in MAC-based traffic match mode. • Manual—Static Ethernet service instance in VLAN-based traffic match mode.
Encapsulation	Frame match criterion of the Ethernet service instance. If the Ethernet service instance does not contain a match criterion, the command does not display this field.
Bandwidth	Bandwidth limit in kbps. If no bandwidth limit is set for the Ethernet service instance, Unlimited is displayed.
Link ID	Ethernet service instance's link ID on the VSI.
State	Ethernet service instance state: <ul style="list-style-type: none"> • Up. • Down.
Statistics	Packet statistics state: <ul style="list-style-type: none"> • Enabled—The packet statistics feature is enabled for the Ethernet service instance. • Disabled—The packet statistics feature is disabled for the Ethernet service instance.
Input Statistics	Incoming traffic statistics: <ul style="list-style-type: none"> • Octets—Number of incoming bytes. • Packets—Number of incoming packets.
Output Statistics	Outgoing traffic statistics: <ul style="list-style-type: none"> • Octets—Number of outgoing bytes.

Field	Description
	<ul style="list-style-type: none"> Packets—Number of outgoing packets.

Related commands

`service-instance`

display l2vpn vsi

Use `display l2vpn vsi` to display information about VSIs.

Syntax

```
display l2vpn vsi [ name vsi-name ] [ verbose ]
```

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

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

verbose: Displays detailed information about VSIs. If you do not specify this keyword, the command displays brief information about VSIs.

Examples

Display brief information about all VSIs.

```
<Sysname> display l2vpn vsi
Total number of VSIs: 1, 1 up, 0 down, 0 admin down
```

VSI Name	VSI Index	MTU	State
vpna	0	1500	Up

Table 9 Command output

Field	Description
MTU	MTU on the VSI.
State	VSI state: <ul style="list-style-type: none"> Up—The VSI is up. Down—The VSI is down. Admin down—The VSI has been manually shut down by using the shutdown command.

Display detailed information about all VSIs.

```
<Sysname> display l2vpn vsi verbose
VSI Name: vpna
  VSI Index      : 0
  VSI State      : Up
  MTU            : 1500
  Bandwidth      : Unlimited
```

```

Broadcast Restrain      : Unlimited
Multicast Restrain     : Unlimited
Unknown Unicast Restrain: Unlimited
MAC Learning           : Enabled
MAC Table Limit        : -
MAC Learning rate      : -
Drop Unknown           : -
Flooding                : Enabled
Statistics              : Enabled
Input statistics:
  Octets   : 0
  Packets  : 0
  Errors   : 0
  Discards : 0
Output statistics:
  Octets   : 0
  Packets  : 0
  Errors   : 0
  Discards : 0
Gateway Interface      : VSI-interface 100
VXLAN ID               : 10
Tunnels:
  Tunnel Name      Link ID   State  Type      Flood proxy
  Tunnel1          0x5000001 Up     Manual    Disabled
  Tunnel2          0x5000002 Up     Manual    Disabled
  MTunnel0         0x6002710 Up     Auto      Disabled
ACs:
  AC                               Link ID   State  Type
  XGE1/0/1 srv1000                0        Up     Manual
  XGE100/0/1 srv10                 1        Up     Manual

```

Table 10 Command output

Field	Description
VSI Description	Description of the VSI. If the VSI does not have a description, the command does not display this field.
VSI State	VSI state: <ul style="list-style-type: none"> • Up—The VSI is up. • Down—The VSI is down. • Administratively down—The VSI has been manually shut down by using the shutdown command.
MTU	MTU on the VSI.
Bandwidth	Bandwidth limit in kbps. If no bandwidth limit is set for the VSI, Unlimited is displayed.
Broadcast Restrain	Broadcast restraint bandwidth in kbps. If no broadcast restraint bandwidth is set for the VSI, this field displays Unlimited .
Multicast Restrain	Multicast restraint bandwidth in kbps. If no multicast restraint bandwidth is set for the VSI, this field displays Unlimited .

Field	Description
Unknown Unicast Restrain	Unknown unicast restraint bandwidth in kbps. If no unknown unicast restraint bandwidth is set for the VSI, this field displays Unlimited .
MAC Learning	State of the MAC learning feature.
MAC Table Limit	This field is not supported in the current software version. Maximum number of MAC address entries on the VSI.
MAC Learning rate	This field is not supported in the current software version. MAC address entry learning rate of the VSI.
Drop Unknown	This field is not supported in the current software version. Action on source MAC-unknown frames received after the maximum number of MAC entries is reached.
Flooding	State of the VSI's flooding feature: <ul style="list-style-type: none"> • Enabled—Flooding is enabled on the VSI. • Disabled—Flooding is disabled on the VSI.
Statistics	Packet statistics state: <ul style="list-style-type: none"> • Enabled—The packet statistics feature is enabled for the VSI. • Disabled—The packet statistics feature is disabled for the VSI.
Input statistics	Incoming traffic statistics: <ul style="list-style-type: none"> • Octets—Number of incoming bytes. • Packets—Number of incoming packets. • Errors—Number of error packets. • Discards—Number of discarded packets.
Output statistics	Outgoing traffic statistics: <ul style="list-style-type: none"> • Octets—Number of outgoing bytes. • Packets—Number of outgoing packets. • Errors—Number of error packets. • Discards—Number of discarded packets.
Gateway Interface	VSI interface name.
State	Tunnel state: <ul style="list-style-type: none"> • Up—The tunnel is operating correctly. • Blocked—The tunnel is a backup proxy tunnel. Its tunnel interface is up, but the tunnel is blocked because the primary proxy tunnel is operating correctly. • Defect—The tunnel interface is up, but BFD cannot detect the remote VTEP. • Down—The tunnel interface is down.
Type	Tunnel assignment method: <ul style="list-style-type: none"> • Auto—The tunnel was automatically assigned to the VXLAN. For a multicast-mode VXLAN, the tunnel (MTunnel) was automatically created and assigned to the VXLAN to transmit flood traffic. For an EVPN network, VXLAN tunnels are automatically assigned to VXLANs. • Manual—The tunnel was manually assigned to the VXLAN.
Flood proxy	Flood proxy state: <ul style="list-style-type: none"> • Enabled—Flood proxy is enabled. The VTEP sends broadcast, multicast, and unknown unicast traffic to a flood proxy server through the tunnel. The flood proxy server replicates and forwards flood traffic to remote VTEPs. • Disabled—Flood proxy is disabled.
ACs	ACs that are bound to the VSI.

Field	Description
Link ID	AC's link ID on the VSI.
State	AC state: <ul style="list-style-type: none"> • Up. • Down.
Type	Type and traffic match mode of the Ethernet service instance: <ul style="list-style-type: none"> • Dynamic (MAC-based)—Dynamic Ethernet service instance in MAC-based traffic match mode. • Manual—Static Ethernet service instance in VLAN-based traffic match mode.

display vxlan tunnel

Use **display vxlan tunnel** to display VXLAN tunnel information for VXLANs.

Syntax

```
display vxlan tunnel [ vxlan-id vxlan-id ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

vxlan-id vxlan-id: Specifies a VXLAN ID in the range of 0 to 16777215. If you do not specify a VXLAN, this command displays VXLAN tunnel information for all VXLANs.

Examples

Display VXLAN tunnel information for all VXLANs.

```
<Sysname> display vxlan tunnel
```

```
Total number of VXLANs: 1
```

```
VXLAN ID: 10, VSI name: vpna, Total tunnels: 3 (3 up, 0 down, 0 defect, 0 blocked)
```

Tunnel name	Link ID	State	Type	Flood proxy
Tunnel1	0x5000001	Up	Manual	Disabled
Tunnel2	0x5000002	Up	Manual	Disabled
MTunnel0	0x6002710	Up	Auto	Disabled

Display VXLAN tunnel information for VXLAN 10.

```
<Sysname> display vxlan tunnel vxlan-id 10
```

```
VXLAN ID: 10, VSI name: vpna, Total tunnels: 3 (3 up, 0 down, 0 defect, 0 blocked)
```

Tunnel name	Link ID	State	Type	Flood proxy
Tunnel1	0x5000001	Up	Manual	Disabled
Tunnel2	0x5000002	Up	Manual	Disabled
MTunnel0	0x6002710	Up	Auto	Disabled

Table 11 Command output

Field	Description
Link ID	Tunnel's link ID in the VXLAN.
State	Tunnel state: <ul style="list-style-type: none"> • Up—The tunnel is operating correctly. • Blocked—The tunnel is a backup proxy tunnel. Its tunnel interface is up, but the tunnel is blocked because the primary proxy tunnel is operating correctly. • Defect—The tunnel interface is up, but BFD cannot detect the remote VTEP. • Down—The tunnel interface is down.
Type	Tunnel assignment method: <ul style="list-style-type: none"> • Auto—The tunnel was automatically assigned to the VXLAN. For a multicast-mode VXLAN, the tunnel (MTunnel) was automatically created and assigned to the VXLAN to transmit flood traffic. For an EVPN network, VXLAN tunnels are automatically assigned to VXLANs. • Manual—The tunnel was manually assigned to the VXLAN.
Flood proxy	Flood proxy state: <ul style="list-style-type: none"> • Enabled—Flood proxy is enabled. The VTEP sends broadcast, multicast, and unknown unicast traffic to a flood proxy server through the tunnel. The flood proxy server replicates and forwards flood traffic to remote VTEPs. • Disabled—Flood proxy is disabled.

Related commands

tunnel
vxlan

emulate-ping vxlan

Use **emulate-ping vxlan** to test the reachability of a remote VM.

Syntax

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

Views

Any view

Predefined user roles

network-admin

Parameters

-c *count*: Specifies the number of ICMP echo requests to be sent. The value range is 1 to 4294967295, and the default value is 5.

-m *interval*: Specifies the interval for sending ICMP echo requests. The value range is 1 to 65535 milliseconds, and the default value is 200 milliseconds.

-s *packet-size*: Specifies the length of ICMP echo requests (IP and ICMP headers not included). The value range is 20 to 8100 bytes, and the default value is 56 bytes.

-t *time-out*: Specifies the ICMP echo reply timeout period. The value range is 0 to 65535 milliseconds, and the default value is 2000 milliseconds. If the device does not receive an ICMP echo reply within the timeout period after sending an ICMP echo request, the device determines that the ICMP echo reply times out.

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

source-mac mac-address: Specifies the source MAC address of ICMP echo requests. The MAC address must be a local VM address in the MAC address table of the specified VXLAN.

destination-mac mac-address: Specifies the destination MAC address of ICMP echo requests. The MAC address must be a remote VM address in the MAC address table of the specified VXLAN.

Usage guidelines

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

Examples

```
# Simulate local VM 1ea3-c0be-0206 to test the reachability of remote VM 1ea3-b77b-0106 in VXLAN 100.
```

```
<Sysname> emulate-ping vxlan vxlan-id 100 source-mac 1ea3-c0be-0206 destination-mac 1ea3-b77b-0106
```

```
Emulate ping in VXLAN 100: source MAC 1ea3-c0be-0206, destination MAC 1ea3-b77b-0106, 56 data bytes.
```

```
Press CTRL_C to break.
```

```
56 bytes from 100.1.2.1: ICMP_seq=0 time=1.114 ms
```

```
56 bytes from 100.1.2.1: ICMP_seq=1 time=1.073 ms
```

```
56 bytes from 100.1.2.1: ICMP_seq=2 time=1.123 ms
```

```
56 bytes from 100.1.2.1: ICMP_seq=3 time=1.781 ms
```

```
56 bytes from 100.1.2.1: ICMP_seq=4 time=0.933 ms
```

```
--- Ping statistics for VXLAN 100 ---
```

```
5 packets transmitted, 5 packets received, 0.0% packet loss
```

```
Round-trip min/avg/max/std-dev=0.933/1.205/1.781/0.296 ms
```

Table 12 Command output

Field	Description
56 bytes from 100.1.2.1: ICMP_seq=0 time=1.114 ms	Statistics about the received ICMP echo reply from the VM at 100.1.2.1. If no reply is received within the timeout period, nothing is displayed. <ul style="list-style-type: none">bytes—Number of bytes in the ICMP echo reply.ICMP_seq—Sequence number, used to determine whether a request is lost, disordered, or repeated.time—Response time.
5 packets transmitted	Number of transmitted ICMP echo requests.
5 packets received	Number of received ICMP echo replies.
0.0% packet loss	Percentage of unacknowledged requests.
Round-trip min/avg/max/std-dev=0.933/1.205/1.781/0.296 ms	Minimum/average/maximum/standard deviation response time in milliseconds.

encapsulation

Use **encapsulation** to configure a frame match criterion for an Ethernet service instance.

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

Syntax

encapsulation s-vid *vlan-id* [**c-vid** { *vlan-id-list* | **all** } | **only-tagged**]

encapsulation s-vid *vlan-id-list* [**c-vid** *vlan-id-list*]

encapsulation { **default** | **tagged** | **untagged** }

undo encapsulation

Default

An Ethernet service instance does not contain a frame match criterion.

Views

Ethernet service instance view

Predefined user roles

network-admin

Parameters

c-vid: Matches frames that are tagged with the specified inner 802.1Q VLAN IDs.

s-vid: Matches frames that are tagged with the specified outer 802.1Q VLAN IDs.

vlan-id: Specifies an 802.1Q VLAN ID in the range of 1 to 4094.

vlan-id-list: Specifies a space-separated list of up to eight VLAN items. Each item specifies a VLAN ID or a range of VLAN IDs in the format of *vlan-id1* **to** *vlan-id2*. The value range for VLAN IDs is 1 to 4094.

only-tagged: Matches tagged frames. If the outer 802.1Q VLAN is not the PVID, the matching result does not differ, whether or not you specify the **only-tagged** keyword. If the outer 802.1Q VLAN is the PVID, the matching result depends on whether or not the **only-tagged** keyword is specified.

- To match only PVID-tagged frames, specify the **only-tagged** keyword.
- To match both untagged frames and PVID-tagged frames, do not specify the **only-tagged** keyword.

all: Specifies all 802.1Q VLAN IDs.

default: Matches frames that do not match any other Ethernet service instance on the interface. On an interface, you can configure this criterion only in one Ethernet service instance. The Ethernet service instance matches any frames if it is the only instance on the interface.

tagged: Matches any frames that have an 802.1Q VLAN tag.

untagged: Matches any frames that do not have an 802.1Q VLAN tag.

Usage guidelines

The **encapsulation s-vid** *vlan-id* [**only-tagged**] and **encapsulation s-vid** *vlan-id* **c-vid** *vlan-id* commands can specify the same outer VLAN ID for two Ethernet service instances on an interface. Frames that match both Ethernet service instances are assigned to the Ethernet service instance configured by using the **encapsulation s-vid** *vlan-id* **c-vid** *vlan-id* command.

If an Ethernet service instance is configured with the **encapsulation untagged** criterion on a Layer 2 Ethernet or aggregate interface, you cannot apply a QoS policy for VLAN tag nesting to that interface. For more information about VLAN tag nesting, see QoS configuration in *ACL and QoS Configuration Guide*.

If you execute the **encapsulation default** command for an Ethernet service instance in Ethernet access mode on a Layer 2 Ethernet interface or Layer 2 aggregate interface, the interface can host multiple Ethernet service instances. Traffic that does not match any other Ethernet service instance matches the Ethernet service instance that uses the default criterion.

An Ethernet service instance can contain only one match criterion. To change the match criterion, first execute the **undo encapsulation** command to remove the original criterion. When you remove the match criterion in an Ethernet service instance, the mapping between the service instance and the VSI is removed automatically.

On an SDN network, after you change the matching outer VLAN ID on the VTEP, you must modify the VLAN ID that the controller issues in the Set-Field action for configuration consistency.

Examples

```
# Configure Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1 to match frames that have an outer 802.1Q VLAN ID of 111.
```

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 1
[Sysname-Ten-GigabitEthernet1/0/1-srv1] encapsulation s-vid 111
```

Related commands

display l2vpn service-instance

flooding disable

Use **flooding disable** to disable flooding for a VSI.

Use **undo flooding disable** to enable flooding for a VSI.

Syntax

flooding disable { all | { broadcast | unknown-multicast | unknown-unicast } * }

undo flooding disable

Default

Flooding is enabled for a VSI.

Views

VSI view

Predefined user roles

network-admin

Parameters

all: Specifies broadcast, unknown unicast, and unknown multicast traffic.

broadcast: Specifies broadcast traffic.

unknown-multicast: Specifies unknown multicast traffic.

unknown-unicast: Specifies unknown unicast traffic.

Usage guidelines

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

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

To confine a kind of flood traffic to the site-facing interfaces, use this command to disable flooding for that kind of flood traffic on the VSI bound to the VXLAN. The VSI will not flood the corresponding frames to VXLAN tunnel interfaces.

Examples

```
# Disable flooding of broadcast traffic for VSI vsi1.  
<Sysname> system-view  
[Sysname] vsi vsi1  
[Sysname-vsi-vsi1] flooding disable broadcast
```

group

Use **group** to assign a VXLAN a multicast group address for flood traffic, and specify a source IP address for multicast VXLAN packets.

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

Syntax

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

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

Default

A VXLAN uses unicast mode (head-end replication) for flood traffic. No multicast group address or source IP address is specified for multicast VXLAN packets.

Views

VXLAN view

Predefined user roles

network-admin

Parameters

group-address: Specifies a multicast address in the range of 224.0.1.0 to 239.255.255.255.

source *source-address*: Specifies a source IP address for multicast VXLAN packets.

Usage guidelines

To reduce traffic sent to the transport network, use multicast mode if the network has dense flood traffic or many VTEPs.

The multicast mode supports the following multicast methods:

- **PIM**—VTEPs and transport network devices run PIM to generate multicast forwarding entries. On a VTEP, you can use the IP address of a loopback interface as the source IP address for multicast VXLAN packets. If the VTEP has multiple transport-facing interfaces, PIM dynamically selects the outgoing interfaces for multicast VXLAN packets.
- **IGMP host**—VTEPs and transport network devices run PIM and IGMP to generate multicast forwarding entries.
 - Transport-facing interfaces of VTEPs act as IGMP hosts.
 - Transport network devices connected to a VTEP run IGMP.
 - All transport network devices run PIM.

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

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

To forward multicast traffic correctly, you must use the source IP address of an up VXLAN tunnel as the source IP address for multicast VXLAN packets.

For multicast-mode VXLANs, transport network devices must maintain multicast group and forwarding information. To reduce the multicast forwarding entries maintained by transport network devices, assign a multicast group address to multiple VXLANs. The VTEP separates traffic between VXLANs by VXLAN IDs.

NOTE:

For VXLANs that use the same multicast group address, you must configure the same source IP address for their multicast VXLAN packets.

If you execute the **group** command multiple times for a VXLAN, the most recent configuration takes effect.

Examples

```
# Set the multicast group address to 233.1.1.1 for flood traffic in VXLAN 100. Set the source IP address to 2.1.1.1 for multicast VXLAN packets.
```

```
<Sysname> system-view
[Sysname] vsi aaa
[Sysname-vsi-aaa] vxlan 100
[Sysname-vsi-aaa-vxlan-100] group 233.1.1.1 source 2.1.1.1
```

Related commands

igmp host enable

pim dm (*IP Multicast Command Reference*)

pim sm (*IP Multicast Command Reference*)

hardware-resource vxlan

Use **hardware-resource vxlan** to set the VXLAN hardware resource mode.

Use **undo hardware-resource vxlan** to restore the default.

Syntax

hardware-resource vxlan { **border24k** | **border28k** | **I2gw** | **I3gw8k** | **I3gw16k** | **I3gw24k** }

undo hardware-resource vxlan

Default

The VXLAN hardware resource mode is Layer 2 gateway.

Views

System view

Predefined user roles

network-admin

Parameters

border24k: Specifies the border mode that supports 24 K of overlay adjacency table.

border28k: Specifies the border mode that supports 28 K of overlay adjacency table.

I2gw: Specifies the Layer 2 gateway mode.

I3gw8k: Specifies the Layer 3 gateway mode that supports 8 K of overlay adjacency table.

I3gw16k: Specifies the Layer 3 gateway mode that supports 16 K of overlay adjacency table.

I3gw24k: Specifies the Layer 3 gateway mode that supports 24 K of overlay adjacency table.

Usage guidelines

Set the hardware resource mode for VXLAN based on the role of the device.

- **I2gw**—Applies to VTEPs that perform only Layer 2 forwarding.
- **I3gw8k**, **I3gw16k**, or **I3gw24k**—Applies to VXLAN IP gateways.
- **border24k** or **border28k**—Applies to Layer 3 border gateways that provide access to external networks.

For the hardware resource mode to take effect, you must reboot the device.

You cannot execute the **arp mode uni** command on interfaces of a Layer 3 border gateway. For more information about this command, see ARP commands *in Layer 3—IP Services Command Reference*.

Examples

```
# Set the VXLAN hardware resource mode to Layer 3 gateway mode that supports 8 K of overlay adjacency table.
```

```
<Sysname> system-view
[Sysname] hardware-resource vxlan l3gw8k
```

Related commands

```
display hardware-resource vxlan
```

igmp host enable

Use **igmp host enable** to enable the IGMP host feature on an interface.

Use **undo igmp host enable** to disable the IGMP host feature on an interface.

Syntax

```
igmp host enable
undo igmp host enable
```

Default

The IGMP host feature is disabled on an interface.

Views

Interface view

Predefined user roles

network-admin

Usage guidelines

You must configure an interface as an IGMP host if its IP address is the source IP address of multicast VXLAN packets. The IGMP host feature enables the interface to send IGMP reports in response to IGMP queries before it can receive traffic from a multicast group.

For this command to take effect, you must use the **multicast routing** command to enable IP multicast routing.

Examples

```
# Enable IP multicast routing, and then enable the IGMP host feature on VLAN-interface 10.
```

```
<Sysname> system-view
[Sysname] multicast routing
[Sysname-mrib] quit
[Sysname] interface vlan-interface 10
```

```
[Sysname-Vlan-interface10] igmp host enable
```

Related commands

display igmp host group

group

multicast routing (*IP Multicast Command Reference*)

ipv6 nd suppression enable

Use **ipv6 nd suppression enable** to enable ND flood suppression.

Use **undo ipv6 nd suppression enable** to disable ND flood suppression.

Syntax

ipv6 nd suppression enable

undo ipv6 nd suppression enable

Default

ND flood suppression is disabled.

Views

VSI view

Predefined user roles

network-admin

Usage guidelines

This command is available in Release 2612P06 and later.

ND flood suppression reduces ND request multicasts by enabling the VTEP to reply to ND requests on behalf of user terminals.

This feature snoops ND packets to populate the ND flood suppression table with local and remote MAC addresses. If an ND request has a matching entry, the VTEP replies to the request on behalf of the user terminal. If no match is found, the VTEP floods the request to both local and remote sites.

Examples

```
# Enable ND flood suppression for VSI vsi1.  
<Sysname> system-view  
[Sysname] vsi vsi1  
[Sysname-vsi-vsi1] ipv6 nd suppression enable
```

Related commands

display ipv6 nd suppression vsi

reset ipv6 nd suppression vsi

ipv6 nd suppression notify-ipsg

Use **ipv6 nd suppression notify-ipsg** to enable the device to generate dynamic IPv6SG bindings based on ND flood suppression entries.

Use **undo ipv6 nd suppression notify-ipsg** to disable the device from generating dynamic IPv6SG bindings based on ND flood suppression entries.

Syntax

```
ipv6 nd suppression notify-ipsg
undo ipv6 nd suppression notify-ipsg
```

Default

The device does not generate dynamic IPv6SG bindings based on ND flood suppression entries.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This command is available in Release 2612P06 and later.

After you execute this command, the device notifies the IP source guard module of ND flood suppression entries for it to generate dynamic IPv6SG bindings based on these entries.

For more information about IP source guard, see *Security Configuration Guide*.

Examples

```
# Enable the device to generate dynamic IPv6SG bindings based on ND flood suppression entries.
<Sysname> system-view
[Sysname] ipv6 nd suppression notify-ipsg
```

Related commands

```
display ipv6 source binding (Security Command Reference)
ipv6 nd suppression enable
```

I2vpn enable

Use **I2vpn enable** to enable L2VPN.

Use **undo I2vpn enable** to disable L2VPN.

Syntax

```
I2vpn enable
undo I2vpn enable
```

Default

L2VPN is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

You must enable L2VPN before you can configure L2VPN settings.

Examples

```
# Enable L2VPN.
<Sysname> system-view
```

```
[Sysname] l2vpn enable
```

l2vpn mac-address software-learning enable

Use **l2vpn mac-address software-learning enable** to enable software-based MAC learning on an interface.

Use **undo l2vpn mac-address software-learning enable** to disable software-based MAC learning on an interface.

Syntax

```
l2vpn mac-address software-learning enable
```

```
undo l2vpn mac-address software-learning enable
```

Default

Hardware-based MAC learning is used.

Views

Layer 2 aggregate interface view

Layer 2 Ethernet interface view

Predefined user roles

network-admin

Usage guidelines

This command is applicable to SDN networks.

To reduce broadcast traffic in an SDN network, the controller synchronizes the MAC addresses that each VTEP learns among all VTEPs. On a VTEP, an interface can learn MAC addresses in hardware or software.

- In hardware-based learning mode, the software periodically obtains new MAC addresses from the hardware and advertises the MAC addresses to the controller.
- In software-based learning mode, the software instantly issues new MAC addresses to the hardware and the controller as soon as they are learned.

Software-based MAC learning shortens the interval at which the VTEP advertises MAC address reachability information to the controller. However, this mode is resource intensive. When you use this mode, you must fully understand its impact on the device performance.

Examples

```
# Enable software-based MAC learning on Ten-GigabitEthernet 1/0/1.
```

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] l2vpn mac-address software-learning enable
```

learning mode disable

Use **learning mode disable** to disable MAC address learning for an Ethernet service instance.

Use **undo learning mode** to enable MAC address learning for an Ethernet service instance.

Syntax

```
learning mode disable
```

```
undo learning mode
```

Default

MAC address learning is enabled for Ethernet service instances.

Views

Ethernet service instance view

Predefined user roles

network-admin

Usage guidelines

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

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

Examples

```
# Disable MAC address learning for Ethernet service instance 200.
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] mac-learning enable
[Sysname-vsi-vpn1] quit
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 200
[Sysname-Ten-GigabitEthernet1/0/1-srv200] learning mode disable
```

Related commands

mac-address static

mac-learning enable

mac-address mac-learning priority

Use **mac-address mac-learning priority** to set the MAC learning priority of an Ethernet service instance.

Use **undo mac-address mac-learning priority** to restore the default.

Syntax

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

undo mac-address mac-learning priority

Default

The MAC learning priority of an Ethernet service instance is low.

Views

Ethernet service instance view

Predefined user roles

network-admin

Parameters

high: Specifies high MAC learning priority.

low: Specifies low MAC learning priority.

Usage guidelines

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

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

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

Examples

```
# Set the MAC learning priority to high for Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1.
```

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 1
[Sysname-Ten-GigabitEthernet1/0/1-srv1] mac-address mac-learning priority high
```

mac-address static

Use **mac-address static** to add a static MAC address entry for a VXLAN VSI.

Use **undo mac-address static** to remove a static MAC address entry for a VXLAN VSI.

Syntax

mac-address static *mac-address* { **interface** *interface-type interface-number* **service-instance** *instance-id* | **interface tunnel** *tunnel-number* } **vsi** *vsi-name*

undo mac-address static [*mac-address*] [**interface** *interface-type interface-number* **service-instance** *instance-id* | **interface tunnel** *tunnel-number*] **vsi** *vsi-name*

Default

VXLAN VSIs do not have static MAC address entries.

Views

System view

Predefined user roles

network-admin

Parameters

mac-address: Specifies a MAC address in H-H-H format. Do not specify a multicast MAC address or an all-zeros MAC address. You can omit the consecutive zeros at the beginning of each segment. For example, you can enter **f-e2-1** for **000f-00e2-0001**.

interface *interface-type interface-number* **service-instance** *instance-id*: Specifies an Ethernet service instance on an interface. The *interface-type interface-number* argument specifies the interface by its type and number. The *instance-id* argument specifies the Ethernet service instance by its ID in the range of 1 to 4096. This option applies to local MAC addresses.

interface tunnel *tunnel-number*: Specifies a VXLAN tunnel interface by its tunnel interface number. The specified tunnel interface must already exist. This option applies to remote MAC addresses.

vsi *vsi-name*: Specifies a VSI name, a case-sensitive string of 1 to 31 characters.

Usage guidelines

A local MAC address is the MAC address of a VM in the local site. Local MAC entries include manually added entries and dynamically learned entries.

A remote MAC address is the MAC address of a VM in a remote site. Remote MAC entries include the following types:

- **Static**—Manually added MAC entries.
- **Dynamic**—MAC entries learned in the data plane from incoming traffic on VXLAN tunnels.
- **OpenFlow**—MAC entry issued by a remote controller through OpenFlow.
- **OVSDB**—MAC entry issued by a remote controller through OVSDB.
- **EVPN**—MAC entry advertised through EVPN.

For a MAC address, the manual static entry has higher priority than the dynamic entry.

When you add a local MAC address entry, make sure the specified VSI has been created and the specified Ethernet service instance has been mapped to the VSI. When you add a remote MAC address entry, make sure the VSI's VXLAN has been created and specified on the VXLAN tunnel.

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

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

The **undo mac-address static vsi vsi-name** command removes all static MAC address entries for a VSI.

Examples

Add MAC address **000f-e201-0101** to VSI **vs1**. Specify Tunnel-interface 1 as the outgoing interface.

```
<Sysname> system-view
[Sysname] mac-address static 000f-e201-0101 interface tunnel 1 vsi vs1
```

Add MAC address **000f-e201-0102** of Ethernet service instance 1 to VSI **vs1**. Specify Ten-GigabitEthernet 1/0/1 as the outgoing interface.

```
<Sysname> system-view
[Sysname] mac-address static 000f-e201-0102 interface ten-gigabitethernet 1/0/1
service-instance 1 vsi vs1
```

Related commands

vxlan tunnel mac-learning disable

mac-learning enable

Use **mac-learning enable** to enable MAC address learning for a VSI.

Use **undo mac-learning enable** to disable MAC address learning for a VSI.

Syntax

mac-learning enable

undo mac-learning enable

Default

MAC address learning is enabled for a VSI.

Views

VSI view

Predefined user roles

network-admin

Usage guidelines

If MAC address learning is disabled for a VSI, the VSI does not learn the source MAC addresses of received packets.

Examples

```
# Disable MAC address learning for VSI vpn1.  
<Sysname> system-view  
[Sysname] vsi vpn1  
[Sysname-vsi-vpn1] undo mac-learning enable
```

Related commands

display l2vpn vsi

mac-based ac

Use **mac-based ac** to enable MAC-based traffic match mode for dynamic Ethernet service instances on an interface.

Use **undo mac-based ac** to disable MAC-based traffic match mode for dynamic Ethernet service instances on an interface.

Syntax

mac-based ac

undo mac-based ac

Default

MAC-based traffic match mode is disabled for dynamic Ethernet service instances.

Views

Layer 2 aggregate interface view

Layer 2 Ethernet interface view

Predefined user roles

network-admin

Usage guidelines

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

A dynamic Ethernet service instance matches frames by VLAN ID and source MAC address, which is called MAC-based traffic match mode. To use this mode for dynamic Ethernet service instances, you must enable MAC authentication or 802.1X authentication that uses MAC-based access control.

This command takes effect only on dynamic Ethernet service instances. Static Ethernet service instances created by using the **service-instance** command match traffic only by the VLAN IDs specified by using the **encapsulation** command.

You cannot change the traffic match mode when dynamic Ethernet service instances already exist on an interface.

Examples

```
# Enable MAC-based traffic match mode for dynamic Ethernet service instances on
Ten-GigabitEthernet 1/0/1.
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] mac-based ac
```

Related commands

display l2vpn service-instance

reserved vxlan

Use **reserved vxlan** to specify a reserved VXLAN.

Use **undo reserved vxlan** to restore the default.

Syntax

reserved vxlan *vxlan-id*

undo reserved vxlan

Default

No VXLAN has been reserved.

Views

System view

Predefined user roles

network-admin

Parameters

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

Usage guidelines

You can specify only one reserved VXLAN on the VTEP. The reserved VXLAN cannot be the VXLAN created on any VSI.

Examples

```
# Specify VXLAN 10000 as the reserved VXLAN.
<Sysname> system-view
[Sysname] reserved vxlan 10000
```

reset arp suppression vsi

Use **reset arp suppression vsi** to clear ARP flood suppression entries on VSIs.

Syntax

reset arp suppression vsi [**name** *vsi-name*]

Views

User view

Predefined user roles

network-admin

Parameters

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

Examples

```
# Clear ARP flood suppression entries on all VSIs.
<Sysname> reset arp suppression vsi
This command will delete all entries. Continue? [Y/N]:y
```

Related commands

arp suppression enable
display arp suppression vsi

reset ipv6 nd suppression vsi

Use **reset ipv6 nd suppression vsi** to clear ND flood suppression entries on VSIs.

Syntax

```
reset ipv6 nd suppression vsi [ name vsi-name ]
```

Views

User view

Predefined user roles

network-admin

Parameters

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

Usage guidelines

This command is available in Release 2612P06 and later.

Examples

```
# Clear ND flood suppression entries on all VSIs.
<Sysname> reset ipv6 nd suppression vsi
This command will delete all entries. Continue? [Y/N]:y
```

Related commands

display ipv6 nd suppression vsi
ipv6 nd suppression enable

reset l2vpn mac-address

Use **reset l2vpn mac-address** to clear dynamic MAC address entries on VSIs.

Syntax

```
reset l2vpn mac-address [ vsi vsi-name ]
```

Views

User view

Predefined user roles

network-admin

Parameters

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

Usage guidelines

Use this command when the number of dynamic MAC address entries reaches the limit or the device learns incorrect MAC addresses.

Examples

```
# Clear the dynamic MAC address entries on VSI vpn1.  
<Sysname> reset l2vpn mac-address vsi vpn1
```

Related commands

display l2vpn mac-address vsi

reset l2vpn statistics ac

Use **reset l2vpn statistics ac** to clear packet statistics on ACs.

Syntax

```
reset l2vpn statistics ac [ interface interface-type interface-number service-instance instance-id ]
```

Views

User view

Predefined user roles

network-admin

Parameters

interface *interface-type interface-number*: Specifies an interface by its type and number.

service-instance *instance-id*: Specifies an Ethernet service instance ID in the range of 1 to 4096.

Usage guidelines

If you do not specify any parameters, this command clears packet statistics on all ACs.

Examples

```
# Clear packet statistics for Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1.  
<Sysname> reset l2vpn statistics ac interface ten-gigabitethernet 1/0/1 service-instance  
1
```

Related commands

display l2vpn service-instance verbose

statistics enable (Ethernet service instance view)

reset l2vpn statistics vsi

Use **reset l2vpn statistics vsi** to clear packet statistics on VSIs.

Syntax

```
reset l2vpn statistics vsi [ name vsi-name ]
```

Views

User view

Predefined user roles

network-admin

Parameters

name *vsi-name*: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command clears packet statistics on all VSIs.

Examples

```
# Clear packet statistics on all VSIs.  
<Sysname> reset l2vpn statistics vsi
```

Related commands

statistics enable (VSI view)

restrain

Use **restrain** to set broadcast, multicast, or unknown unicast restraint bandwidth for a VSI.

Use **undo restrain** to restore the default.

Syntax

```
restrain { broadcast | multicast | unknown-unicast } bandwidth  
undo restrain { broadcast | multicast | unknown-unicast }
```

Default

The device does not limit the broadcast restraint bandwidth, multicast restraint bandwidth, and unknown unicast restraint bandwidth.

Views

VSI view

Predefined user roles

network-admin

Parameters

broadcast: Specifies the broadcast restraint bandwidth.

multicast: Specifies the multicast restraint bandwidth.

unknown-unicast: Specifies the unknown unicast restraint bandwidth. Unknown unicast packets refer to the unicast packets whose destination MAC addresses are not in the MAC address table of the VSI.

bandwidth: Specifies the broadcast, multicast, or unknown unicast restraint bandwidth in kbps. For software versions earlier than Release 2612P02, the value can be 0 or in the range of 64 to 4194303. For Release 2612P02 and later, the value can be 0 or in the range of 64 to 167772159. If you set this argument to 0, the command disables a kind of flood traffic.

Usage guidelines

When the bandwidth of incoming broadcast, multicast, or unknown unicast traffic on VXLAN tunnels of a VSI exceeds the configured restraint bandwidth, the VSI drops the exceeding traffic.

If you use both the **restrain** and **bandwidth** commands on a VSI, the **bandwidth** command limits only the bandwidth of the traffic not restrained by the **restrain** command.

If you use both the **restrain** and **selective-flooding mac-address** commands on a VSI, the **restrain** command limits only the bandwidth of the traffic not enabled with selective flood.

Examples

```
# Set the broadcast restraint bandwidth, multicast restraint bandwidth, and unknown unicast restraint bandwidth to 100 kbps for VSI vpn1.
```

```
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] restrain broadcast 100
[Sysname-vsi-vpn1] restrain multicast 100
[Sysname-vsi-vpn1] restrain unknown-unicast 100
```

Related commands

display l2vpn vsi

rewrite inbound tag

Use **rewrite inbound tag** to configure the VLAN tag processing rule for incoming traffic.

Use **undo rewrite inbound tag** to restore the default.

Syntax

```
rewrite inbound tag { remark 1-to-1 s-vid vlan-id | strip s-vid }
```

```
undo rewrite inbound tag
```

Default

VLAN tags of incoming traffic are not processed.

Views

Ethernet service instance view

Predefined user roles

network-admin

Parameters

remark: Maps VLAN tags.

1-to-1: Performs one-to-one mapping to replace one VLAN tag of packets with the specified VLAN tag.

s-vid: Specifies outer VLAN tags.

vlan-id: Specifies a VLAN ID in the range of 1 to 4094.

strip: Removes VLAN tags.

Usage guidelines

Use this command for an Ethernet service instance before you map it to a VSI.

To modify the VLAN tag processing rule for incoming traffic, you must first delete the existing rule by using the **undo rewrite inbound tag** command.

When you use this command for an Ethernet service instance, follow these restrictions:

- The Ethernet service instance must use Ethernet access mode.
- Do not configure the **default**, **tagged**, or **untagged** frame match criterion for the Ethernet service instance.

- Do not specify a VLAN as both the inner and outer VLAN in the frame match criterion of the Ethernet service instance.
- If you use this command for an Ethernet service instance, do not use the **rewrite outbound tag** command for the other Ethernet service instances mapped to the same VSI. Ethernet service instances of a VSI cannot be configured with both incoming and outgoing VLAN tag processing rules.

Examples

Configure Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1 to replace outer VLAN tag 10 with outer VLAN tag 100 for incoming traffic.

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 1
[Sysname-Ten-GigabitEthernet1/0/1-srv1] encapsulation s-vid 10
[Sysname-Ten-GigabitEthernet1/0/1-srv1] rewrite inbound tag remark 1-to-1 s-vid 100
```

rewrite outbound tag

Use **rewrite outbound tag** to configure the VLAN tag processing rule for outgoing traffic.

Use **undo rewrite outbound tag** to restore the default.

Syntax

rewrite outbound tag nest s-vid *vlan-id*

undo rewrite outbound tag

Default

VLAN tags of outgoing traffic are not processed.

Views

Ethernet service instance view

Predefined user roles

network-admin

Parameters

nest: Adds VLAN tags.

s-vid: Specifies an outer VLAN tag.

vlan-id: Specifies a VLAN ID in the range of 1 to 4094.

Usage guidelines

Use this command for an Ethernet service instance before you map it to a VSI.

To modify the VLAN tag processing rule for outgoing traffic, you must first delete the existing rule by using the **undo rewrite outbound tag** command.

When you use this command for an Ethernet service instance, follow these restrictions:

- The Ethernet service instance must use Ethernet access mode.
- This command adds the specified VLAN tag to an untagged packet and replaces the VLAN tag of a single-tagged packet with the specified VLAN tag. This command does not take effect on packets with two layers of VLAN tags.
- If you use this command for an Ethernet service instance, do not use the **rewrite inbound tag** command for the other Ethernet service instances mapped to the same VSI. Ethernet service instances of a VSI cannot be configured with both incoming and outgoing VLAN tag processing rules.

Examples

```
# Configure Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1 to add outer VLAN tag 100 to outgoing traffic.
```

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 1
[Sysname-Ten-GigabitEthernet1/0/1-srv1] encapsulation s-vid 10
[Sysname-Ten-GigabitEthernet1/0/1-srv1] rewrite outbound tag nest s-vid 100
```

selective-flooding mac-address

Use **selective-flooding mac-address** to enable selective flood for a MAC address.

Use **undo selective-flooding mac-address** to disable selective flood for a MAC address.

Syntax

```
selective-flooding mac-address mac-address
```

```
undo selective-flooding mac-address mac-address
```

Default

Selective flood is disabled for all MAC addresses.

Views

VSI view

Predefined user roles

network-admin

Parameters

mac-address: Specifies a MAC address. The MAC address cannot be all Fs.

Usage guidelines

This command excludes a remote unicast or multicast MAC address from the flood suppression done by using the **flooding disable** command. The VTEP will flood the frames destined for the specified MAC address to remote sites when floods are confined to the local site.

If you use both the **restrain** and **selective-flooding mac-address** commands on a VSI, the **restrain** command limits only the bandwidth of the traffic not enabled with selective flood.

As a best practice, do not execute both the **bandwidth** and **selective-flooding mac-address** commands on a VSI. Traffic cannot be forwarded correctly if you use these commands together.

Examples

```
# Enable selective flood for 000f-e201-0101 on VSI vsi1.
```

```
<Sysname> system-view
[Sysname] vsi vsi1
[Sysname-vsi-vsi1] selective-flooding mac-address 000f-e201-0101
```

Related commands

flooding disable

service-instance

Use **service-instance** to create an Ethernet service instance and enter its view, or enter the view of an existing Ethernet service instance.

Use **undo service-instance** to delete an Ethernet service instance.

Syntax

service-instance *instance-id*

undo service-instance *instance-id*

Default

No Ethernet service instances exist.

Views

Layer 2 aggregate interface view

Layer 2 Ethernet interface view

Predefined user roles

network-admin

Parameters

instance-id: Specifies an Ethernet service instance ID in the range of 1 to 4096.

Examples

On Layer 2 Ethernet interface Ten-GigabitEthernet 1/0/1, create Ethernet service instance 1 and enter Ethernet service instance view.

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 1
[Sysname-Ten-GigabitEthernet1/0/1-srv1]
```

Related commands

display l2vpn service-instance

shutdown

Use **shutdown** to shut down a VSI.

Use **undo shutdown** to bring up a VSI.

Syntax

shutdown

undo shutdown

Default

A VSI is not manually shut down.

Views

VSI view

Predefined user roles

network-admin

Usage guidelines

Use this command to temporarily disable a VSI to provide Layer 2 switching services. The shutdown action does not change settings on the VSI. You can continue to configure the VSI. After you bring up the VSI again, the VSI provides services based on the latest settings.

Examples

```
# Shut down VSI vpn1.  
<Sysname> system-view  
[Sysname] vsi vpn1  
[Sysname-vsi-vpn1] shutdown
```

Related commands

display l2vpn vsi

statistics enable (Ethernet service instance view)

Use **statistics enable** to enable packet statistics for an Ethernet service instance.

Use **undo statistics enable** to disable packet statistics for an Ethernet service instance.

Syntax

statistics enable

undo statistics enable

Default

The packet statistics feature is disabled for an Ethernet service instance.

Views

Ethernet service instance view

Predefined user roles

network-admin

Usage guidelines

For this command to take effect, you must configure a frame match criterion for the Ethernet service instance and map it to a VSI. If you modify the frame match criterion or VSI mapping, packet statistics of the instance is cleared.

Examples

```
# Enable packet statistics for Ethernet service instance 200 on Ten-GigabitEthernet 1/0/1.  
<Sysname> system-view  
[Sysname] interface ten-gigabitethernet 1/0/1  
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 200  
[Sysname-Ten-GigabitEthernet1/0/1-srv200] statistics enable
```

Related command

display l2vpn service-instance verbose

reset l2vpn statistics ac

statistics enable (tunnel interface view)

Use **statistics enable** to enable packet statistics for a manually created VXLAN or VXLAN-DCI tunnel.

Use **undo statistics enable** to disable packet statistics for a manually created VXLAN or VXLAN-DCI tunnel.

Syntax

statistics enable

undo statistics enable

Default

The packet statistics feature is disabled for a manually created VXLAN or VXLAN-DCI tunnel.

Views

VXLAN tunnel interface view

VXLAN-DCI tunnel interface view

Predefined user roles

network-admin

Examples

Enable packet statistics for VXLAN tunnel interface **Tunnel 0**.

```
<Sysname> system-view
[Sysname] interface tunnel 0 mode vxlan
[Sysname-Tunnel0] statistics enable
```

Related commands

display interface tunnel (*Layer 3—IP Services Command Reference*)

reset counters interface tunnel (*Layer 3—IP Services Command Reference*)

tunnel statistics vxlan auto

statistics enable (VSI view)

Use **statistics enable** to enable packet statistics for a VSI.

Use **undo statistics enable** to disable packet statistics for a VSI.

Syntax

statistics enable

undo statistics enable

Default

The packet statistics feature is disabled for a VSI.

Views

VSI view

Predefined user roles

network-admin

Examples

Enable packet statistics for VSI **vsi1**.

```
<Sysname> system-view
[Sysname] vsi vsi1
[Sysname-vsi-vsi1] statistics enable
```

Related commands

display l2vpn vsi verbose

reset l2vpn statistics vsi

tunnel

Use **tunnel** to assign VXLAN tunnels to a VXLAN.

Use **undo tunnel** to remove VXLAN tunnels from a VXLAN.

Syntax

```
tunnel { tunnel-number [ backup-tunnel tunnel-number | flooding-proxy ] | all }
```

```
undo tunnel { tunnel-number | all }
```

Default

A VXLAN does not contain VXLAN tunnels.

Views

VXLAN view

Predefined user roles

network-admin

Parameters

tunnel-number: Specifies a tunnel interface number. The value range for this argument is 0 to 15359. The tunnel must be a VXLAN tunnel.

backup-tunnel *tunnel-number*: Specifies a backup tunnel by its tunnel interface number in the range of 0 to 15359. The tunnel must be a VXLAN tunnel.

flooding-proxy: Enables flood proxy on the tunnel for the VTEP to send flood traffic to the flood proxy server. The flood proxy server replicates and forwards flood traffic to remote VTEPs. If you do not specify this keyword, flood proxy is disabled on the tunnel.

all: Specifies all VXLAN tunnels.

Usage guidelines

This command assigns a VXLAN tunnel to a VXLAN to provide Layer 2 connectivity for the VXLAN between two sites.

You can assign multiple VXLAN tunnels to a VXLAN, and configure a VXLAN tunnel to trunk multiple VXLANs. For a unicast-mode VXLAN, the system floods unknown unicast, multicast, and broadcast traffic to each tunnel in the VXLAN.

To assign a pair of primary and backup VXLAN tunnels to the VXLAN, specify the **backup-tunnel** *tunnel-number* option. When the primary VXLAN tunnel is operating correctly, the backup VXLAN tunnel does not forward traffic. When the primary VXLAN tunnel goes down, traffic is switched to the backup VXLAN tunnel.

On a VSI, you can enable flood proxy on multiple VXLAN tunnels. The first tunnel that is enabled with flood proxy works as the primary proxy tunnel to forward broadcast, multicast, and unknown unicast traffic. Other proxy tunnels are backups that do not forward traffic when the primary proxy tunnel is operating correctly.

To change a flood proxy tunnel for a VXLAN, perform the following tasks:

- Use the **undo tunnel** command to remove the flood proxy tunnel.
- Use the **tunnel** command to enable flood proxy on another tunnel and assign the tunnel to the VXLAN.

You cannot use the **tunnel all** command for VXLAN-DCI tunnel assignment. Please assign VXLAN-DCI tunnels to a VXLAN one by one.

Centralized VXLAN IP gateway groups cannot work with the flood proxy feature. Do not use them together in a VXLAN.

If the **tunnel all** command is used for a VXLAN, you cannot remove the VXLAN tunnels one by one. You can only use the **undo tunnel all** command to remove all the VXLAN tunnels.

Examples

```
# Assign VXLAN tunnels 1 and 2 to VXLAN 10000.
<Sysname> system-view
[Sysname] vsi vpna
[Sysname-vsi-vpna] vxlan 10000
[Sysname-vsi-vpna-vxlan-10000] tunnel 1
[Sysname-vsi-vpna-vxlan-10000] tunnel 2
```

Related commands

display vxlan tunnel

tunnel bfd enable

Use **tunnel bfd enable** to enable BFD on a VXLAN tunnel interface.

Use **undo tunnel bfd enable** to disable BFD on a VXLAN tunnel interface.

Syntax

```
tunnel bfd enable destination-mac mac-address
undo tunnel bfd enable
```

Default

BFD is disabled on a VXLAN tunnel interface.

Views

VXLAN tunnel interface view

Predefined user roles

network-admin

Parameters

destination-mac *mac-address*: Specifies a destination MAC address in H-H-H format for BFD control packets. The MAC address can be a remote VTEP address or a multicast address. You can omit the consecutive zeros at the beginning of each segment. For example, you can enter **f-e2-1** for **000f-00e2-0001**.

Usage guidelines

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

For BFD sessions to come up, you must reserve a VXLAN by using the **reserved vxlan** command.

Do not use BFD together with uRPF. When uRPF is enabled, BFD sessions cannot come up. For more information about uRPF, see *Security Configuration Guide*.

Examples

```
# Enable BFD on VXLAN tunnel interface Tunnel 9, and specify 1-1-1 as the destination MAC
address for BFD control packets.
<Sysname> system-view
[Sysname] interface tunnel 9 mode vxlan
```

```
[Sysname-Tunnel9] tunnel bfd enable destination-mac 1-1-1
```

tunnel global source-address

Use **tunnel global source-address** to specify a global source address for VXLAN tunnels.

Use **undo tunnel global source-address** to restore the default.

Syntax

```
tunnel global source-address ip-address
```

```
undo tunnel global source-address
```

Default

No global source address is specified for VXLAN tunnels.

Views

System view

Predefined user roles

network-admin

Parameters

ip-address: Specifies an IP address.

Usage guidelines



IMPORTANT:

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

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

The global source address takes effect only on VXLAN tunnels (VXLAN-DCI tunnels not included).

Examples

```
# Specify 1.1.1.1 as the global source address for VXLAN tunnels.
```

```
<Sysname> system-view
```

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

tunnel statistics vxlan auto

Use **tunnel statistics vxlan auto** to enable packet statistics for automatically created VXLAN tunnels.

Use **undo tunnel statistics vxlan auto** to disable packet statistics for automatically created VXLAN tunnels.

Syntax

```
tunnel statistics vxlan auto
```

```
undo tunnel statistics vxlan auto
```

Default

The packet statistics feature is disabled for automatically created VXLAN tunnels.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This command enables the device to collect packet statistics for all VXLAN tunnels that are automatically created by EVPN or OVSD. This command takes effect on both existing VXLAN tunnels and VXLAN tunnels created after execution of this command.

Examples

```
# Enable packet statistics for automatically created VXLAN tunnels.
```

```
<Sysname> system-view
```

```
[Sysname] tunnel statistics vxlan auto
```

Related commands

display interface tunnel (*Layer 3—IP Services Command Reference*)

reset counters interface tunnel (*Layer 3—IP Services Command Reference*)

statistics enable (tunnel interface view)

VSi

Use **vsi** to create a VSI and enter its view, or enter the view of an existing VSI.

Use **undo vsi** to delete a VSI.

Syntax

```
vsi vsi-name
```

```
undo vsi vsi-name
```

Default

No VSIs exist.

Views

System view

Predefined user roles

network-admin

Parameters

vsi-name: Specifies a VSI name, a case-sensitive string of 1 to 31 characters.

Usage guidelines

A VSI acts as a virtual switch to provide Layer 2 switching services for a VXLAN on a VTEP. A VSI has all functions of a physical Ethernet switch, including source MAC address learning, MAC address aging, and flooding.

A VSI can provide services only for one VXLAN.

Examples

```
# Create VSI vxlan10 and enter VSI view.
```

```
<Sysname> system-view
```

```
[Sysname] vsi vxlan10
```

```
[Sysname-vsi-vxlan10]
```

Related commands

display l2vpn vsi

vxlan

Use **vxlan** to create a VXLAN and enter its view, or enter the view of an existing VXLAN.

Use **undo vxlan** to restore the default.

Syntax

vxlan *vxlan-id*

undo vxlan

Default

No VXLANs exist.

Views

VSI view

Predefined user roles

network-admin

Parameters

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

Usage guidelines

You can create only one VXLAN for a VSI. The VXLAN ID for each VSI must be unique.

Examples

Create VXLAN 10000 for VSI **vpna** and enter VXLAN view.

```
<Sysname> system-view
```

```
[Sysname] vsi vpna
```

```
[Sysname-vsi-vpna] vxlan 10000
```

```
[Sysname-vsi-vpna-vxlan-10000]
```

Related commands

vsi

vxlan default-decapsulation

Use **vxlan default-decapsulation** to enable default VXLAN decapsulation.

Use **undo vxlan default-decapsulation** to restore the default.

Syntax

vxlan default-decapsulation source interface *interface-type interface-number*

undo vxlan default-decapsulation source interface

Default

Default VXLAN decapsulation is disabled.

Views

System view

Predefined user roles

network-admin

Parameters

source interface *interface-type interface-number*: Specifies an interface by its type and number.

Usage guidelines

This command is available in Release 2612P02 and later.

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

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

Examples

```
# Enable default VXLAN decapsulation.
```

```
<Sysname> system-view
```

```
[Sysname] vxlan default-decapsulation source interface ten-gigabitethernet 1/0/1
```

vxlan invalid-udp-checksum discard

Use **vxlan invalid-udp-checksum discard** to enable the device to drop the VXLAN packets that fail UDP checksum check.

Use **undo vxlan invalid-udp-checksum discard** to restore the default.

Syntax

```
vxlan invalid-udp-checksum discard
```

```
undo vxlan invalid-udp-checksum discard
```

Default

The device does not check the UDP checksum of VXLAN packets.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This command enables the device to check the UDP checksum of VXLAN packets.

The device always sets the UDP checksum of VXLAN packets to 0. For compatibility with third-party devices, a VXLAN packet can pass the check if its UDP checksum is 0 or correct. If its UDP checksum is incorrect, the VXLAN packet fails the check and is dropped.

Examples

```
# Enable the device to drop the VXLAN packets that fail UDP checksum check.
```

```
<Sysname> system-view
```

```
[Sysname] vxlan invalid-udp-checksum discard
```

Related commands

```
vxlan invalid-vlan-tag discard
```

vxlan invalid-vlan-tag discard

Use **vxlan invalid-vlan-tag discard** to enable the device to drop the VXLAN packets that have 802.1Q VLAN tags in the inner Ethernet header.

Use **undo vxlan invalid-vlan-tag discard** to restore the default.

Syntax

vxlan invalid-vlan-tag discard

undo vxlan invalid-vlan-tag discard

Default

The device does not check whether a VXLAN packet has 802.1Q VLAN tags in the inner Ethernet header.

Views

System view

Predefined user roles

network-admin

Usage guidelines

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

To configure the access mode, use the **xconnect vsi** command.

Examples

```
# Enable the device to drop VXLAN packets that have 802.1Q VLAN tags.
```

```
<Sysname> system-view
```

```
[Sysname] vxlan invalid-vlan-tag discard
```

Related commands

vxlan invalid-udp-checksum discard

xconnect vsi

vxlan local-mac report

Use **vxlan local-mac report** to enable local-MAC change logging.

Use **undo vxlan local-mac report** to restore the default.

Syntax

vxlan local-mac report

undo vxlan local-mac report

Default

Local-MAC change logging is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

Local-MAC change logging enables VXLAN to send a log message to the information center when a local MAC address is added or removed.

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

Examples

```
# Enable local-MAC change logging.
<Sysname> system-view
[Sysname] vxlan local-mac report
```

vxlan tunnel mac-learning disable

Use **vxlan tunnel mac-learning disable** to disable remote-MAC address learning.

Use **undo vxlan tunnel mac-learning disable** to enable remote-MAC address learning.

Syntax

```
vxlan tunnel mac-learning disable
undo vxlan tunnel mac-learning disable
```

Default

Remote-MAC address learning is enabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

When network attacks occur, use this command to prevent the device from learning incorrect remote MAC addresses in the data plane.

Examples

```
# Disable remote-MAC address learning.
<Sysname> system-view
[Sysname] vxlan tunnel mac-learning disable
```

vxlan udp-port

Use **vxlan udp-port** to set the destination UDP port number for VXLAN packets.

Use **undo vxlan udp-port** to restore the default.

Syntax

```
vxlan udp-port port-number
undo vxlan udp-port
```

Default

The destination UDP port number is 4789 for VXLAN packets.

Views

System view

Predefined user roles

network-admin

Parameters

port-number: Specifies a UDP port number in the range of 1 to 65535. As a best practice, specify a port number in the range of 1024 to 65535 to avoid conflict with well-known ports.

Usage guidelines

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

Examples

```
# Set the destination UDP port number to 6666 for VXLAN packets.  
<Sysname> system-view  
[Sysname] vxlan udp-port 6666
```

vxlan vlan-based

Use **vxlan vlan-based** to enable VLAN-based VXLAN assignment.

Use **undo vxlan vlan-based** to disable VLAN-based VXLAN assignment.

Syntax

vxlan vlan-based

undo vxlan vlan-based

Default

VLAN-based VXLAN assignment is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

You can assign customer traffic to a VXLAN by using one of the following methods:

- **Ethernet service instance-to-VSI mapping**—This method uses the frame match criterion of an Ethernet service instance to match a list of VLANs on a site-facing Layer 2 interface. The VTEP assigns customer traffic to a VXLAN by mapping the Ethernet service instance to a VSI.
- **VLAN-based VXLAN assignment**—This method maps a VLAN to a VXLAN. When a VLAN is mapped to a VXLAN and VLAN-based VXLAN assignment is enabled, the device automatically performs the following operations:
 - a. Creates an Ethernet service instance that uses the VLAN ID as its instance ID on each interface in the VLAN.
 - b. Maps the Ethernet service instances to the VSI of the VXLAN.

On an interface, the frame match criterion of the Ethernet service instance is set as follows:

- If the VLAN ID is the PVID of the interface, the Ethernet service instance matches untagged frames.
- If the VLAN ID is not the PVID of the interface, the Ethernet service instance matches frames tagged with an outer VLAN ID that is same as that VLAN ID.

The previous methods are mutually exclusive. You can use only one of them.

Examples

```
# Enable VLAN-based VXLAN assignment.
```

```
<Sysname> system-view
```

```
[Sysname] vxlan vlan-based
```

vxlan vni

Use **vxlan vni** to map a VLAN to a VXLAN.

Use **undo vxlan vni** to remove the VXLAN mapping for a VLAN.

Syntax

```
vxlan vni vxlan-id
```

```
undo vxlan vni
```

Default

A VLAN is not mapped to a VXLAN.

Views

VLAN view

Predefined user roles

network-admin

Parameters

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

Usage guidelines

Before you map VLANs to VXLANs, enable VLAN-based VXLAN assignment by using the **vxlan vlan-based** command.

You cannot map VLAN 1 to any VXLAN. Do not map a VLAN to the L3 VXLAN ID of EVPN.

If you map a VLAN to a nonexistent VXLAN, the configuration takes effect after you create the VXLAN.

Examples

```
# Map VLAN 10 to VXLAN 100.
```

```
<Sysname> system-view
```

```
[Sysname] vlan 10
```

```
[Sysname-vlan10] vxlan vni 100
```

Related commands

```
vxlan vlan-based
```

vxlan-over-vxlan enable

Use **vxlan-over-vxlan enable** to enable VXLAN over VXLAN on an interface.

Use **undo vxlan-over-vxlan enable** to disable VXLAN over VXLAN on an interface.

Syntax

```
vxlan-over-vxlan enable
```

```
undo vxlan-over-vxlan enable
```

Default

VXLAN over VXLAN is disabled on an interface.

Views

Layer 2 aggregate interface view

Layer 2 Ethernet interface view

Predefined user roles

network-admin

Usage guidelines

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

For VXLAN packets received from a non-transport-facing interface on the device to traverse the VXLAN network through VXLAN tunnels, perform the following tasks on the interface:

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

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

Examples

```
# Enable VXLAN over VXLAN on Ten-GigabitEthernet 1/0/1.
```

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet
```

```
[Sysname-Ten-GigabitEthernet1/0/1] vxlan-over-vxlan enable
```

Related commands

vxlan udp-port

xconnect vsi

Use **xconnect vsi** to map an AC to a VSI.

Use **undo xconnect vsi** to restore the default.

Syntax

```
xconnect vsi vsi-name [ access-mode { ethernet | vlan } ] [ track track-entry-number&<1-3> ]
```

```
undo xconnect vsi
```

Default

An AC is not mapped to any VSI.

Views

Ethernet service instance view

Predefined user roles

network-admin

Parameters

vsi-name: Specifies the VSI name, a case-sensitive string of 1 to 31 characters.

access-mode: Specifies an access mode. The default access mode is VLAN.

ethernet: Specifies the Ethernet access mode.

vlan: Specifies the VLAN access mode.

track *track-entry-number*&<1-3>: Specifies a space-separated list of up to three track entry numbers in the range of 1 to 1024. The AC is up only if a minimum of one associated track entry is in positive state.

Usage guidelines

To monitor the status of an AC, associate it with track entries.

To configure this command for an Ethernet service instance, you must first use the **encapsulation** command to add a traffic match criterion to the service instance.

For traffic that matches the Ethernet service instance, the system uses the VSI's MAC address table to make a forwarding decision.

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

- **VLAN access mode**—Ethernet frames received from or sent to the local site must contain 802.1Q VLAN tags.
 - For an Ethernet frame received from the local site, the VTEP removes all its 802.1Q VLAN tags before forwarding the frame.
 - For an Ethernet frame destined for the local site, the VTEP adds 802.1Q VLAN tags to the frame before forwarding the frame.

In VLAN access mode, VXLAN packets sent between VXLAN sites do not contain 802.1Q VLAN tags. VXLAN can provide Layer 2 connectivity for different 802.1Q VLANs between sites. You can use different 802.1Q VLANs to provide the same service in different sites.

- **Ethernet access mode**—The VTEP does not process the 802.1Q VLAN tags of Ethernet frames received from or sent to the local site.
 - For an Ethernet frame received from the local site, the VTEP forwards the frame with the 802.1Q VLAN tags intact.
 - For an Ethernet frame destined for the local site, the VTEP forwards the frame without adding 802.1Q VLAN tags.

In Ethernet access mode, VXLAN packets sent between VXLAN sites contain 802.1Q VLAN tags. VXLAN cannot provide Layer 2 connectivity for different 802.1Q VLANs between sites. You must use the same 802.1Q VLAN to provide the same service between sites.

Examples

```
# On Ten-GigabitEthernet 1/0/1, configure Ethernet service instance 200 to match frames with an outer 802.1Q VLAN tag of 200, and map the instance to VSI vpn1.
```

```
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] quit
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] service-instance 200
[Sysname-Ten-GigabitEthernet1/0/1-srv200] encapsulation s-vid 200
[Sysname-Ten-GigabitEthernet1/0/1-srv200] xconnect vsi vpn1
```

Related commands

display l2vpn service-instance

encapsulation

vsi

VXLAN IP gateway commands

arp distributed-gateway dynamic-entry synchronize

Use **arp distributed-gateway dynamic-entry synchronize** to enable dynamic ARP entry synchronization for distributed VXLAN IP gateways.

Use **undo arp distributed-gateway dynamic-entry synchronize** to disable dynamic ARP entry synchronization for distributed VXLAN IP gateways.

Syntax

arp distributed-gateway dynamic-entry synchronize

undo arp distributed-gateway dynamic-entry synchronize

Default

Dynamic ARP entry synchronization is disabled for distributed VXLAN IP gateways.

Views

System view

Predefined user roles

network-admin

Usage guidelines

When local proxy ARP is enabled on distributed VXLAN IP gateways, each gateway learns ARP information independently. A gateway does not forward ARP packets destined for its local VSI interfaces to other gateways. For distributed VXLAN IP gateways to have the same ARP entries, you must enable dynamic ARP entry synchronization.

A controller or the EVPN feature can also synchronize ARP entries among distributed VXLAN IP gateways. When you use a controller or the EVPN feature, do not enable dynamic ARP entry synchronization.

Examples

```
# Enable dynamic ARP entry synchronization for distributed VXLAN IP gateways.
```

```
<Sysname> system-view
```

```
[Sysname] arp distributed-gateway dynamic-entry synchronize
```

Related commands

distributed-gateway local

local-proxy-arp enable (*Layer 3—IP Services Command Reference*)

arp send-rate

Use **arp send-rate** to set an ARP packet sending rate limit for a VSI interface.

Use **undo arp send-rate** to remove the ARP packet sending rate limit for a VSI interface.

Syntax

arp send-rate *pps*

undo arp send-rate

Default

The ARP packet sending rate is not limited for a VSI interface.

Views

VSI interface view

Predefined user roles

network-admin

Parameters

pps: Specifies a rate limit in the range of 1 to 500 pps.

Usage guidelines

VMs have limited capacity to process packets. To prevent packet processing from degrading VM performance, limit the ARP packet sending rate of the VSI interface for VMs. The VTEP will drop excess ARP packets if the rate limit is exceeded.

Examples

```
# Set the ARP packet sending rate limit to 50 pps for VSI-interface 1.
<Sysname> system
[Sysname] interface vsi-interface 1
[Sysname-Vsi-interface1] arp send-rate 50
```

bandwidth (VSI interface view)

Use **bandwidth** to set the expected bandwidth for a VSI interface.

Use **undo bandwidth** to restore the default.

Syntax

bandwidth *bandwidth-value*

undo bandwidth

Default

The expected bandwidth (in kbps) of a VSI interface equals the interface baud rate divided by 1000.

Views

VSI interface view

Predefined user roles

network-admin

Parameters

bandwidth-value: Specifies the expected bandwidth in the range of 1 to 400000000 kbps.

Usage guidelines

The expected bandwidth is an informational parameter used only by higher-layer protocols for calculation. You cannot adjust the actual bandwidth of an interface by using this command.

Examples

```
# Set the expected bandwidth to 10000 kbps for VSI-interface 100.
<Sysname> system-view
[Sysname] interface vsi-interface 100
[Sysname-Vsi-interface100] bandwidth 10000
```

default

Use **default** to restore the default settings for a VSI interface.

Syntax

default

Views

VSI interface view

Predefined user roles

network-admin

Usage guidelines

CAUTION:

The **default** command might interrupt ongoing network services. Make sure you are fully aware of the impact of this command when you use it on a live network.

This command might fail to restore the default settings for some commands for reasons such as command dependencies and system restrictions.

To resolve this problem:

1. Use the **display this** command in interface view to identify these commands.
2. Use their **undo** forms or follow the command reference to restore their default settings.
3. If the restoration attempt still fails, follow the error message instructions to resolve the problem.

Examples

```
# Restore the default settings for VSI-interface 100.
<Sysname> system-view
[Sysname] interface vsi-interface 100
[Sysname-Vsi-interface100] default
This command will restore the default settings. Continue? [Y/N]:y
```

description

Use **description** to configure the description of a VSI interface.

Use **undo description** to restore the default.

Syntax

description *text*

undo description

Default

The description of a VSI interface is *interface-name* plus **Interface** (for example, **Vsi-interface100 Interface**).

Views

VSI interface view

Predefined user roles

network-admin

Parameters

text: Specifies a description, a case-sensitive string of 1 to 255 characters.

Examples

```
# Configure the description as gateway for VXLAN 10 for VSI-interface 100.
<Sysname> system-view
[Sysname] interface vsi-interface 100
[Sysname-Vsi-interface100] description gateway for VXLAN 10
```

display interface vsi-interface

Use **display interface vsi-interface** to display information about VSI interfaces.

Syntax

```
display interface [ vsi-interface [ vsi-interface-id ] ] [ brief [ description | down ] ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

vsi-interface [*vsi-interface-id*]: Specifies VSI interfaces. If you specify a VSI interface, this command displays information about the specified interface. If you specify only the **vsi-interface** keyword, this command displays information about all VSI interfaces. If you do not specify the **vsi-interface** [*vsi-interface-id*] option, this command displays information about all interfaces. Make sure the specified VSI interfaces have been created on the device.

brief: Display brief interface information. If you do not specify this keyword, the command displays detailed interface information.

description: Displays complete interface descriptions. If you do not specify this keyword, the command displays only the first 27 characters of interface descriptions.

down: Displays interfaces that are physically down as well as the down reason. If you do not specify this keyword, the command does not filter output by physical interface state.

Examples

```
# Display information about VSI-interface 100.
<Sysname> display interface vsi-interface 100
Vsi-interface100
Current state: UP
Line protocol state: UP
Description: Vsi-interface100 Interface
Bandwidth: 1000000 kbps
Maximum transmission unit: 1500
Internet address: 10.1.1.1/24 (primary)
IP packet frame type: Ethernet II, hardware address: 0011-2200-0102
IPv6 packet frame type: Ethernet II, hardware address: 0011-2200-0102
Physical: Unknown, baudrate: 1000000 kbps
Last clearing of counters: Never
Input (total): 0 packets, 0 bytes
```

Output (total): 0 packets, 0 bytes

Table 13 Command output

Field	Description
Current state	Physical link state of the interface: <ul style="list-style-type: none"> Administratively DOWN—The interface has been shut down by using the shutdown command. DOWN—The interface is administratively up, but its physical state is down. UP—The interface is both administratively and physically up.
Line protocol state	Data link layer state of the interface: <ul style="list-style-type: none"> UP—The data link layer protocol is up. UP(spoofing)—The data link layer protocol is up, but the link is an on-demand link or does not exist. DOWN—The data link layer protocol is down.
Description	Description of the interface.
Bandwidth	Expected bandwidth of the interface.
Maximum transmission unit	MTU of the interface.
Internet protocol processing: Disabled	The interface is not assigned an IP address and cannot process IP packets.
Internet address	IP address of the interface. The primary attribute indicates that the address is the primary IP address.
IP packet frame type	IPv4 packet framing format.
hardware address	MAC address.
IPv6 packet frame type	This field is not supported in the current software version. IPv6 packet framing format.
Physical	Physical type of the interface, which is fixed at Unknown .
baudrate	Interface baudrate in kbps.
Last clearing of counters	Last time when the reset counters interface vsi-interface command was used to clear interface statistics. This field displays Never if the reset counters interface vsi-interface command has never been used on the interface since the device startup.
Input (total): 0 packets, 0 bytes	Incoming traffic statistics on the interface: <ul style="list-style-type: none"> Number of incoming packets. Number of incoming bytes.
Output (total): 0 packets, 0 bytes	Outgoing traffic statistics on the interface: <ul style="list-style-type: none"> Number of outgoing packets. Number of outgoing bytes.

Display brief information about all VSI interfaces.

```
<Sysname> display interface vsi-interface brief
Brief information on interfaces in route mode:
Link: ADM - administratively down; Stby - standby
Protocol: (s) - spoofing
Interface          Link Protocol Primary IP      Description
Vsi100            DOWN DOWN      --
```

```

# Display brief information and complete description for VSI-interface 100.
<Sysname> display interface vsi-interface 100 brief description
Brief information on interfaces in route mode:
Link: ADM - administratively down; Stby - standby
Protocol: (s) - spoofing
Interface          Link Protocol Primary IP      Description
Vsi100            UP    UP      1.1.1.1      VSI-interface100

# Displays interfaces that are physically down and the down reason.
<Sysname> display interface brief down
Brief information on interfaces in route mode:
Link: ADM - administratively down; Stby - standby
Interface          Link  Cause
Vsi100            DOWN  Administratively
Vsi200            DOWN  Administratively

```

Table 14 Command output

Field	Description
Interface	Abbreviated interface name.
Link	Physical link state of the interface: <ul style="list-style-type: none"> • UP—The interface is physically up. • DOWN—The interface is physically down. • ADM—The interface has been shut down by using the shutdown command. To restore the physical state of the interface, use the undo shutdown command.
Protocol	Data link layer protocol state of the interface: <ul style="list-style-type: none"> • UP—The data link layer protocol of the interface is up. • UP (s)—The data link layer protocol of the interface is up, but the link is an on-demand link or does not exist. The (s) attribute represents the spoofing flag. • DOWN—The data link layer protocol of the interface is down.
Primary IP	Primary IP address of the interface. This field displays two hyphens (--) if the interface does not have an IP address.
Description	Description of the interface.
Cause	Cause for the physical link state of an interface to be DOWN : <ul style="list-style-type: none"> • Administratively—The interface has been manually shut down by using the shutdown command. To restore the physical state of the interface, use the undo shutdown command. • Not connected—The interface is not mapped to any VSI, or the mapped VSI does not have any AC or VXLAN tunnel.

Related commands

reset counters interface vsi-interface

distributed-gateway local

Use **distributed-gateway local** to specify a VSI interface as a distributed gateway to provide services for the local site.

Use **undo distributed-gateway local** to restore the default.

Syntax

```
distributed-gateway local  
undo distributed-gateway local
```

Default

A VSI interface is not a distributed gateway.

Views

VSI interface view

Predefined user roles

network-admin

Usage guidelines

If a VXLAN uses distributed gateway services, you must assign the same IP address to the VXLAN's VSI interfaces on different VTEPs. To avoid IP address conflicts, you must specify the VSI interface on each VTEP as a distributed gateway.

Examples

```
# Specify VSI-interface 100 as a distributed gateway.  
<Sysname> system-view  
[Sysname] interface vsi-interface 100  
[Sysname-Vsi-interface100] distributed-gateway local
```

gateway subnet

Use **gateway subnet** to assign a subnet to a VSI.

Use **undo gateway subnet** to remove a subnet from a VSI.

Syntax

```
gateway subnet { ipv4-address wildcard-mask | ipv6-address prefix-length }  
undo gateway subnet { ipv4-address wildcard-mask | ipv6-address prefix-length }
```

Default

No subnet is assigned to a VSI.

Views

VSI view

Predefined user roles

network-admin

Parameters

ipv4-address: Specifies an IPv4 subnet address in dotted-decimal notation.

wildcard-mask: Specifies a wildcard mask in dotted decimal notation. In contrast to a network mask, the 0 bits in a wildcard mask represent "do care" bits, and the 1 bits represent "don't care" bits. If the "do care" bits in a packet's IP address are identical to the "do care" bits in the specified subnet address, the packet is assigned to the VSI. All "don't care" bits are ignored. The 0s and 1s in a wildcard mask can be noncontiguous. For example, 0.255.0.255 is a valid wildcard mask.

ipv6-address prefix-length: Specifies an IPv6 subnet address and the address prefix length in the range of 1 to 128.

Usage guidelines

You can assign a maximum of eight subnets to a VSI.

You must specify a gateway interface for a VSI before you can assign subnets to the VSI. If you remove the gateway interface from the VSI, the VSI's subnet settings are automatically deleted.

Examples

```
# Assign subnet 100.0.10.0/24 to VSI vxlan.
<Sysname> system-view
[Sysname] vsi vxlan
[Sysname-vsi-vxlan] gateway subnet 100.0.10.0 0.0.0.255
```

gateway vsi-interface

Use **gateway vsi-interface** to specify a gateway interface for a VSI.

Use **undo gateway vsi-interface** to restore the default.

Syntax

```
gateway vsi-interface vsi-interface-id
undo gateway vsi-interface
```

Default

No gateway interface is specified for a VSI.

Views

VSI view

Predefined user roles

network-admin

Parameters

vsi-interface-id: Specifies a VSI interface by its number. The value range for this argument is 1 to 99999. The VSI interface must already exist.

Usage guidelines

When you delete a VSI interface by using the **undo interface vsi-interface** command, the gateway interface setting of the VSI interface is also deleted.

A VSI can have only one gateway interface. Multiple VSIs cannot share a gateway interface.

Examples

```
# Specify VSI-interface 100 as the gateway interface for VSI vpna.
<Sysname> system-view
[Sysname] vsi vpna
[Sysname-vsi-vpna] gateway vsi-interface 100
```

Related commands

```
interface vsi-interface
```

interface vsi-interface

Use **interface vsi-interface** to create a VSI interface and enter its view, or enter the view of an existing VSI interface.

Use **undo interface vsi-interface** to delete a VSI interface.

Syntax

```
interface vsi-interface vsi-interface-id  
undo interface vsi-interface vsi-interface-id
```

Default

No VSI interfaces exist.

Views

System view

Predefined user roles

network-admin

Parameters

vsi-interface-id: Specifies a VSI interface number. The value range for this argument is 1 to 99999.

Usage guidelines

When you delete a VSI interface by using the **undo interface vsi-interface** command, the gateway interface setting of the VSI interface is also deleted.

Examples

```
# Create VSI-interface 100 and enter VSI interface view.  
<Sysname> system-view  
[Sysname] interface vsi-interface 100  
[Sysname-Vsi-interface100]
```

Related commands

gateway vsi-interface

ipv6 nd distributed-gateway dynamic-entry synchronize

Use **ipv6 nd distributed-gateway dynamic-entry synchronize** to enable dynamic ND entry synchronization for distributed VXLAN IP gateways.

Use **undo ipv6 nd distributed-gateway dynamic-entry synchronize** to disable dynamic ND entry synchronization for distributed VXLAN IP gateways.

Syntax

```
ipv6 nd distributed-gateway dynamic-entry synchronize  
undo ipv6 nd distributed-gateway dynamic-entry synchronize
```

Default

Dynamic ND entry synchronization is disabled for distributed VXLAN IP gateways.

Views

System view

Predefined user roles

network-admin

Usage guidelines

When local ND proxy is enabled on distributed VXLAN IP gateways, each gateway learns ND information independently. A gateway does not forward ND packets destined for its local VSI interfaces to other gateways. For distributed VXLAN IP gateways to have the same ND entries, you must enable dynamic ND entry synchronization.

A controller or the EVPN feature can also synchronize ND entries among distributed VXLAN IP gateways. When you use a controller or the EVPN feature, do not enable dynamic ND entry synchronization.

Examples

```
# Enable dynamic ND entry synchronization for distributed VXLAN IP gateways.
```

```
<Sysname> system-view
```

```
[Sysname] ipv6 nd distributed-gateway dynamic-entry synchronize
```

Related commands

distributed-gateway local

local-proxy-nd enable (*Layer 3—IP Services Command Reference*)

mac-address

Use **mac-address** to assign a MAC address to a VSI interface.

Use **undo mac-address** to restore the default.

Syntax

```
mac-address mac-address
```

```
undo mac-address
```

Default

All VSI interfaces on a device use a fixed MAC address. The MAC addresses of VSI interfaces on different devices are different.

Views

VSI interface view

Predefined user roles

network-admin

Parameters

mac-address: Specifies a MAC address in H-H-H format.

Examples

```
# Assign MAC address 0001-0001-0001 to VSI-interface 100.
```

```
<Sysname> system-view
```

```
[Sysname] interface vsi-interface 100
```

```
[Sysname-Vsi-interface100] mac-address 1-1-1
```

mtu

Use **mtu** to set the MTU for a VSI interface.

Use **undo mtu** to restore the default.

Syntax

```
mtu size
```

```
undo mtu
```

Default

The default MTU of a VSI interface is 1444 bytes.

Views

VSI interface view

Predefined user roles

network-admin

Parameters

size: Specifies an MTU value in the range of 46 to 1500 bytes.

Usage guidelines

Make sure the MTU of a VSI interface is a minimum of 36 bytes less than the MTU of the physical outgoing interface.

Examples

```
# Set the MTU to 1430 bytes for VSI-interface 100.
<Sysname> system-view
[Sysname] interface vsi-interface 100
[Sysname-Vsi-interface100] mtu 1430
```

reset counters interface vsi-interface

Use **reset counters interface vsi-interface** to clear packet statistics on VSI interfaces.

Syntax

```
reset counters interface [ vsi-interface [ vsi-interface-id ] ]
```

Views

User view

Predefined user roles

network-admin

Parameters

vsi-interface [*vsi-interface-id*]: Specifies VSI interfaces. If you specify a VSI interface, this command clears packet statistics on the specified interface. If you specify only the **vsi-interface** keyword, this command clears packet statistics on all VSI interfaces. If you do not specify the **vsi-interface** [*vsi-interface-id*] option, this command clears packet statistics on all interfaces. Make sure the specified VSI interfaces have been created on the device.

Usage guidelines

Use this command to clear history statistics before you collect traffic statistics for a time period.

Examples

```
# Clear packet statistics on VSI-interface 100.
<Sysname> reset counters interface vsi-interface 100
```

Related commands

display interface vsi-interface

shutdown

Use **shutdown** to shut down a VSI interface.

Use **undo shutdown** to bring up a VSI interface.

Syntax

shutdown
undo shutdown

Default

A VSI interface is not manually shut down.

Views

VSI interface view

Predefined user roles

network-admin

Examples

```
# Shut down VSI-interface 100.  
<Sysname> system-view  
[Sysname] interface vsi-interface 100  
[Sysname-Vsi-interface100] shutdown
```

vtep group member local

Use **vtep group member local** to assign the local VTEP to a VTEP group.

Use **undo vtep group member local** to remove the local VTEP from a VTEP group.

Syntax

vtep group *group-ip* **member local** *member-ip*
undo vtep group *group-ip* **member local**

Default

A VTEP is not assigned to any VTEP group.

Views

System view

Predefined user roles

network-admin

Parameters

group-ip: Specifies a VTEP group by its group IP address. The IP address must already exist on the local VTEP.

member-ip: Specifies the member VTEP IP address for the local VTEP. The IP address must already exist on the local VTEP.

Usage guidelines

Member VTEPs in a VTEP group cannot use the group IP address or share an IP address.

Examples

```
# Assign the local VTEP to VTEP group 1.1.1.1, and specify 2.2.2.2 as the member VTEP IP address  
of the local VTEP.  
<Sysname> system-view  
[Sysname] vtep group 1.1.1.1 member local 2.2.2.2
```

Related commands

vtep group member remote

vtep group member remote

Use **vtep group member remote** to specify a VTEP group and its member VTEPs.

Use **undo vtep group member remote** to remove a VTEP group and its member VTEPs.

Syntax

vtep group *group-ip* **member remote** *member-ip*&<1-8>

undo vtep group *group-ip* **member remote**

Default

No VTEP group is specified.

Views

System view

Predefined user roles

network-admin

Parameters

group-ip: Specifies a VTEP group by its group IP address.

member-ip&<1-8>: Specifies a space-separated list of up to eight member VTEP IP addresses.

Examples

```
# Specify VTEP group 1.1.1.1 and its member VTEPs at 2.2.2.2, 3.3.3.3, and 4.4.4.4.
```

```
<Sysname> system-view
```

```
[Sysname] vtep group 1.1.1.1 member remote 2.2.2.2 3.3.3.3 4.4.4.4
```

Related commands

vtep group member local

vxlan tunnel arp-learning disable

Use **vxlan tunnel arp-learning disable** to disable remote ARP learning for VXLANs.

Use **undo vxlan tunnel arp-learning disable** to enable remote ARP learning for VXLANs.

Syntax

vxlan tunnel arp-learning disable

undo vxlan tunnel arp-learning disable

Default

Remote ARP learning is enabled for VXLANs.

Views

System view

Predefined user roles

network-admin

Usage guidelines

By default, the device learns ARP information of remote VMs from packets received on VXLAN tunnel interfaces. To save resources on VTEPs in an SDN transport network, you can temporarily disable remote ARP learning when the controller and VTEPs are synchronizing entries. After the entry synchronization is completed, use the **undo vxlan tunnel arp-learning disable** command to enable remote ARP learning.

As a best practice, disable remote ARP learning for VXLANs only when the controller and VTEPs are synchronizing entries.

Examples

```
# Disable remote ARP learning for VXLANs.
<Sysname> system
[Sysname] vxlan tunnel arp-learning disable
```

vxlan tunnel nd-learning disable

Use **vxlan tunnel nd-learning disable** to disable remote ND learning for VXLANs.

Use **undo vxlan tunnel nd-learning disable** to enable remote ND learning for VXLANs.

Syntax

```
vxlan tunnel nd-learning disable
undo vxlan tunnel nd-learning disable
```

Default

Remote ND learning is enabled for VXLANs.

Views

System view

Predefined user roles

network-admin

Usage guidelines

By default, the device learns ND information of remote VMs from packets received on VXLAN tunnel interfaces. To save resources on VTEPs in an SDN transport network, you can temporarily disable remote ND learning when the controller and VTEPs are synchronizing entries. After the entry synchronization is completed, use the **undo vxlan tunnel nd-learning disable** command to enable remote ND learning.

As a best practice, disable remote ND learning for VXLANs only when the controller and VTEPs are synchronizing entries.

Examples

```
# Disable remote ND learning for VXLANs.
<Sysname> system
[Sysname] vxlan tunnel nd-learning disable
```

OVSDB commands

ovsdb server bootstrap ca-certificate

Use **ovsdb server bootstrap ca-certificate** to specify a CA certificate file for establishing OVSDB SSL connections.

Use **undo ovssdb server bootstrap ca-certificate** to restore the default.

Syntax

ovssdb server bootstrap ca-certificate *ca-filename*

undo ovssdb server bootstrap ca-certificate

Default

SSL uses the CA certificate file in the PKI domain.

Views

System view

Predefined user roles

network-admin

Parameters

ca-filename: Specifies the CA certificate file name, a case-insensitive string. The file name cannot contain the **slot** string.

Usage guidelines

For the specified certificate to take effect, you must execute the **ovssdb server enable** command to enable the OVSSDB server. You must disable and then re-enable the OVSSDB server if it has been enabled before OVSSDB connections are configured, or OVSSDB connection changes or SSL version changes are made.

If the specified CA certificate file does not exist, the device obtains a self-signed certificate from the controller. The obtained file uses the name specified for the *ca-filename* argument.

Examples

```
# Specify CA certificate file ca-new for establishing OVSSDB SSL connections.
```

```
<Sysname> system-view
```

```
[Sysname] ovssdb server bootstrap ca-certificate ca-new
```

Related commands

ovssdb server enable

ovssdb server pki domain

ovssdb server pssl

ovssdb server ssl

ovssdb server enable

Use **ovssdb server enable** to enable the OVSSDB server.

Use **undo ovssdb server enable** to disable the OVSSDB server.

Syntax

ovssdb server enable

undo ovssdb server enable

Default

The OVSSDB server is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

To obtain configuration data from controllers, you must enable the OVSDB server.

Before you enable the OVSDB server, you must establish an OVSDB SSL or TCP connection with a minimum of one controller.

Examples

```
# Enable the OVSDB server.
<Sysname> system-view
[Sysname] ovssdb server enable
```

ovssdb server pki domain

Use **ovssdb server pki domain** to specify a PKI domain for establishing OVSDB SSL connections.

Use **undo ovssdb bootstrap server pki domain** to restore the default.

Syntax

ovssdb server pki domain *domain-name*

undo ovssdb server pki domain

Default

No PKI domain is specified for establishing OVSDB SSL connections.

Views

System view

Predefined user roles

network-admin

Parameters

domain-name: Specifies a PKI domain name, a case-sensitive string of 1 to 31 characters. The PKI domain must already exist and contain a complete certificate and key.

Usage guidelines

To communicate with controllers through SSL, you must specify a PKI domain.

For the specified PKI domain to take effect, you must execute the **ovssdb server enable** command to enable the OVSDB server. You must disable and then re-enable the OVSDB server if it has been enabled before OVSDB connections are configured, or OVSDB connection changes or SSL version changes are made.

For more information about PKI domains, see PKI in *Security Configuration Guide*.

Examples

```
# Specify PKI domain ovssdb_test for establishing OVSDB SSL connections.
<Sysname> system-view
[Sysname] ovssdb server pki domain ovssdb_test
```

Related commands

ovssdb server bootstrap ca-certificate

ovssdb server enable

ovssdb server pssl

ovsdb server ssl

ovsdb server pssl

Use **ovsdb server pssl** to enable the device to listen for OVSDB SSL connection requests.

Use **undo ovsdb server pssl** to restore the default.

Syntax

ovsdb server pssl [**port** *port-number*]

undo ovsdb server pssl

Default

The device does not listen for OVSDB SSL connection requests.

Views

System view

Predefined user roles

network-admin

Parameters

port *port-number*: Specifies a port to listen for OVSDB SSL connection requests. The value range for the *port-number* argument is 1 to 65535. If you do not specify a port, the device uses the port number 6640.

Usage guidelines

Before you use this command, you must specify a PKI domain for SSL.

You can specify only one port to listen for OVSDB SSL connection requests. If you execute this command multiple times, the most recent configuration takes effect.

For the specified port setting to take effect, you must execute the **ovsdb server enable** command to enable the OVSDB server. You must disable and then re-enable the OVSDB server if it has been enabled before OVSDB connections are configured, or OVSDB connection changes or SSL version changes are made.

Examples

```
# Enable the device to listen for OVSDB SSL connection requests on port 6640.
```

```
<Sysname> system-view
```

```
[Sysname] ovsdb server pssl
```

Related commands

ovsdb server bootstrap ca-certificate

ovsdb server enable

ovsdb server pki domain

ovsdb server ssl

ovsdb server ptcp

Use **ovsdb server ptcp** to enable the device to listen for OVSDB TCP connection requests.

Use **undo ovsdb server ptcp** to restore the default.

Syntax

ovsdb server ptcp [**port** *port-number*] [**acl** *acl-number*]

undo ovssdb server tcp

Default

The device does not listen for OVSSDB TCP connection requests.

Views

System view

Predefined user roles

network-admin

Parameters

port *port-number*. Specifies a port to listen for OVSSDB TCP connection requests. The value range for the *port-number* argument is 1 to 65535. If you do not specify a port, the device uses the port number 6640.

acl *acl-number*. Specifies a basic IPv4 ACL by its ACL number to filter OVSSDB TCP connection requests. The value range for the *acl-number* argument is 2000 to 2999. This option is available in Release 2612P06 and later.

Usage guidelines

You can specify only one port to listen for OVSSDB TCP connection requests. If you execute this command multiple times, the most recent configuration takes effect.

For the specified port setting to take effect, you must execute the **ovssdb server enable** command to enable the OVSSDB server. You must disable and then re-enable the OVSSDB server if it has been enabled before OVSSDB connections are configured, or OVSSDB connection changes or SSL version changes are made.

If you specify a basic IPv4 ACL for this command, the OVSSDB server accepts only the OVSSDB TCP connection requests that are permitted by the ACL. The OVSSDB server rejects any OVSSDB TCP connection requests that do not match the ACL or that are denied by the ACL.

The OVSSDB server does not filter OVSSDB TCP connection requests in any of the following situations:

- No IPv4 ACL is specified for this command.
- The specified IPv4 ACL does not exist.
- The specified IPv4 ACL does not have rules.

Modification to rules in the specified IPv4 ACL takes effect only on OVSSDB TCP connection requests that are received after the modification. If you specify another IPv4 ACL, the IPv4 ACL also takes effect only on OVSSDB TCP connection requests that are received after the modification.

Examples

```
# Enable the device to listen for OVSSDB TCP connection requests on port 6640.  
<Sysname> system-view  
[Sysname] ovssdb server tcp
```

Related commands

ovssdb server enable

ovssdb server tcp

ovssdb server ssl

Use **ovssdb server ssl** to set up an active OVSSDB SSL connection to a controller.

Use **undo ovssdb server ssl** to remove an OVSSDB SSL connection from a controller.

Syntax

```
ovsdb server ssl ip ip-address port port-number  
undo ovsdb server ssl ip ip-address port port-number
```

Default

The device does not have active OVSDB SSL connections to a controller.

Views

System view

Predefined user roles

network-admin

Parameters

ip *ip-address*: Specifies the destination IP address for the SSL connection.

port *port-number*: Specifies the destination port for the SSL connection. The value range for the *port-number* argument is 1 to 65535.

Usage guidelines

Before you use this command, you must specify a PKI domain for SSL.

The device can have a maximum of eight active SSL connections.

To establish the connection, you must execute the **ovsdb server enable** command. You must disable and then re-enable the OVSDB server if it has been enabled before OVSDB connections are configured, or OVSDB connection changes or SSL version changes are made.

Examples

```
# Set up an active SSL connection to port 6632 at 192.168.12.2.  
<Sysname> system-view  
[Sysname] ovsdb server ssl ip 192.168.12.2 port 6632
```

Related commands

ovsdb server bootstrap ca-certificate

ovsdb server enable

ovsdb server pki domain

ovsdb server pssl

ovsdb server tcp

Use **ovsdb server tcp** to set up an active OVSDB TCP connection to a controller.

Use **undo ovsdb server tcp** to remove an OVSDB TCP connection.

Syntax

```
ovsdb server tcp ip ip-address port port-number  
undo ovsdb server tcp ip ip-address port port-number
```

Default

The device does not have active OVSDB TCP connections.

Views

System view

Predefined user roles

network-admin

Parameters

ip *ip-address*: Specifies the destination IP address for the TCP connection.

port *port-number*: Specifies the destination port for the TCP connection. The value range for the *port-number* argument is 1 to 65535.

Usage guidelines

The device can have a maximum of eight active OVSDb TCP connections.

To establish the connection, you must execute the **ovsdb server enable** command. You must disable and then re-enable the OVSDb server if it has been enabled before OVSDb connections are configured, or OVSDb connection changes or SSL version changes are made.

Examples

```
# Set up an active OVSDb TCP connection to port 6632 at 192.168.12.2.
```

```
<Sysname> system-view
```

```
[Sysname] ovsdb server tcp ip 192.168.12.2 port 6632
```

Related commands

ovsdb server enable

ovsdb server ptcp

vtep access port

Use **vtep access port** to specify a site-facing interface as a VTEP access port.

Use **undo vtep access port** to restore the default.

Syntax

vtep access port

undo vtep access port

Default

An interface is not a VTEP access port.

Views

Layer 2 aggregate interface view

Layer 2 Ethernet interface view

Predefined user roles

network-admin

Usage guidelines

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

Examples

```
# Specify Ten-GigabitEthernet 1/0/1 as a VTEP access port.
```

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] vtep access port
```

vtep enable

Use **vtep enable** to enable the OVSDDB VTEP service.

Use **undo vtep enable** to disable the OVSDDB VTEP service.

Syntax

vtep enable

undo vtep enable

Default

The OVSDDB VTEP service is disabled.

Views

System view

Predefined user roles

network-admin

Examples

```
# Enable the OVSDDB VTEP service.
```

```
<Sysname> system-view
```

```
[Sysname] vtep enable
```

vxlan tunnel flooding-proxy

Use **vxlan tunnel flooding-proxy** to enable flood proxy on multicast VXLAN tunnels.

Use **undo vxlan tunnel flooding-proxy** to disable flood proxy on multicast VXLAN tunnels.

Syntax

vxlan tunnel flooding-proxy

undo vxlan tunnel flooding-proxy

Default

Flood proxy is disabled on multicast VXLAN tunnels.

Views

System view

Predefined user roles

network-admin

Usage guidelines

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

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

The **vxlan tunnel flooding-proxy** command and its **undo** form affect only VXLAN tunnels that are issued after the **vxlan tunnel flooding-proxy** command.

Examples

```
# Enable flood proxy on all multicast VXLAN tunnels.
```

```
<Sysname> system  
[Sysname] vxlan tunnel flooding-proxy
```