Bridging Configuration Guide for vEdge Routers, Cisco SD-WAN Release 20

First Published: 2020-04-30

Last Modified: 2022-04-22

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

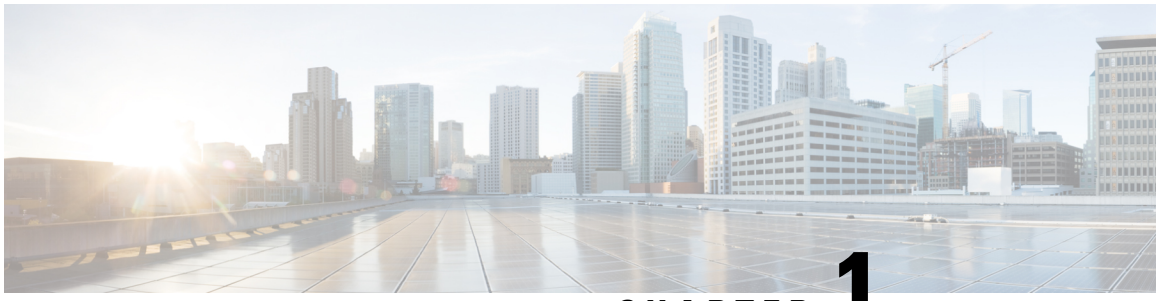

CHAPTER 1

Read Me First

Related References

- Release Notes
- Cisco SD-WAN Controller Compatibility Matrix and Server Recommendations

User Documentation

- Cisco SD-WAN (Cisco vEdge Devices)
- User Documentation for Cisco vEdge Devices

Communications, Services, and Additional Information

- Sign up for Cisco email newsletters and other communications at: Cisco Profile Manager.
- For information on the latest technical, advanced, and remote services to increase the operational reliability of your network visit Cisco Services.
- To browse and discover secure, validated enterprise-class apps, products, solutions, and services, visit Cisco Devnet.
- To obtain general networking, training, and certification titles from Cisco Press Publishers, visit Cisco Press.
- To find warranty information for a specific product or product family, visit Cisco Warranty Finder.
- To view open and resolved bugs for a release, access the Cisco Bug Search Tool.
- To submit a service request, visit Cisco Support.

Documentation Feedback

To provide feedback about Cisco technical documentation use the feedback form available in the right pane of every online document.

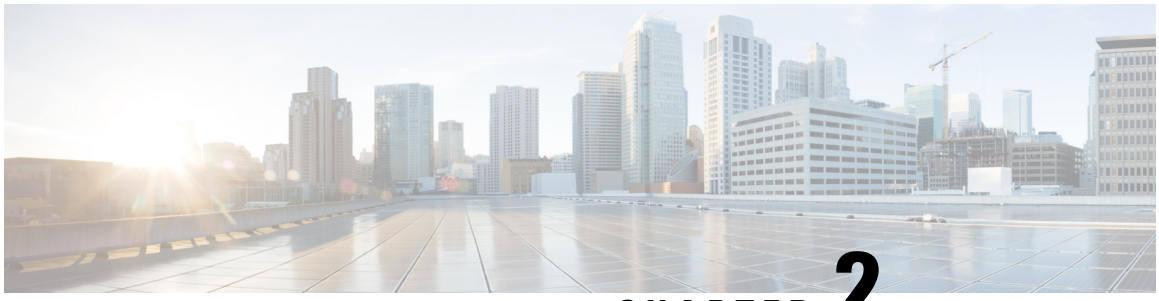

CHAPTER 2

What's New in Cisco SD-WAN

Note The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco is constantly enhancing the SD-WAN solution with every release and we try and keep the content in line with the latest enhancements. The following table lists new and modified features we documented in the Configuration, Command Reference, and Hardware Installation guides. For information on additional features and fixes that were committed to the Cisco SD-WAN solution, see the *Resolved and Open Bugs* section in the Release Notes.

What's New in Cisco SD-WAN (vEdge) Release 20.x

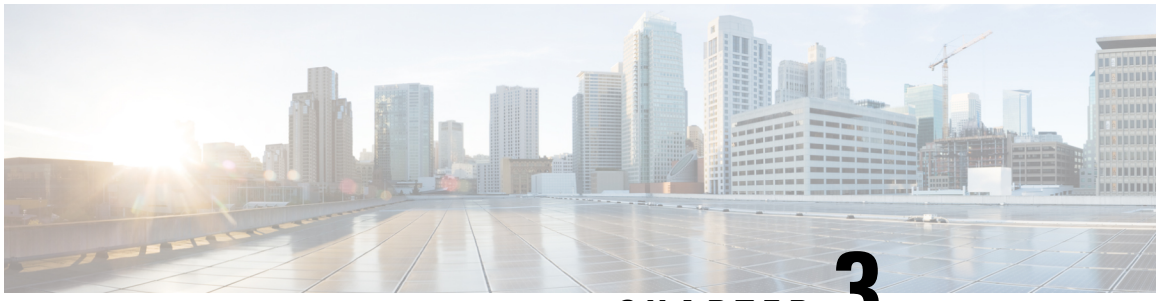

CHAPTER 3

Bridging in Cisco SD-WAN

This chapter contains these topics:

Bridging Overview

A Cisco vEdge device can act as a transparent bridge, switching traffic between LANs that are part of a VLAN at the local device's site. To implement bridging, the Cisco SD-WAN architecture defines the concept of a *bridge domain*. Each bridge domain corresponds to a single VLAN. From a switching point of view, each bridge domain is a separate broadcast domain, and each has its own Ethernet switching table (or MAC table) to use for switching traffic within the broadcast domain. Multiple bridge domains, and hence multiple VLANs, can coexist on a single Cisco vEdge device.

To allow hosts in different bridge domains to communicate with each other, Cisco vEdge devices support *integrated routing and bridging* (IRB). IRB is implemented using *logical IRB interfaces*, which connect a bridge domain to a VPN, or what might better be called a *VPN domain*. The VPN domain provides the Layer 3 routing services necessary so that traffic can be exchanged between different VLANs. Each bridge domain can have a single IRB interface and can connect to a single VPN domain, and a single VPN domain can connect to multiple bridge domains on a vEdge router. The route table in the VPN domain provides reachability between all bridge domains which participate in that VPN domain, whether the bridge domain is located on the local router or on a remote router.

Configure Bridging Using Cisco vManage

To have a Cisco vEdge device act as a transparent bridge, configure bridging domains on the router. A router can have up to 16 bridging domains.

Configure Bridging and Bridge Domains

1. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates** to view your existing feature templates or create a new one.

3. Click **Add Template**. Choose a device from the list of devices. The templates available for the selected device display in the right pane.
4. Choose the **Bridge** template.
5. Enter a name and description for the template.
6. Configure bridging domains under **Basic Configuration**.

Parameter Name	Description
Bridge Name	Enter a text description of the bridging domain. It can be up to 32 characters.
VLAN ID	Enter the VLAN identifier to associate with the bridging domain. *Range:* 0 through 4095
Maximum MAC Addresses	Specify the maximum number of MAC addresses that the bridging domain can learn. *Range:* 0 through 4096 *Default:* 1024
Age-Out Time (Seconds)	Specify how long to store an entry in the MAC table before it ages out. *Range:* 10 through 4096 seconds *Default:* 300 seconds (5 minutes)

7. Click **Save**.

Associate Interfaces with the Bridge Domain

To associate an interface with the bridge domain, click Interface and click **New Interface**.

Parameter Name	Description
Interface Name	Enter the name of the interface to associate with the bridging domain, in the format **ge***slot*/*port*.
Description	Enter a text description of the interface.
Native VLAN Support	Click Enabled to configure the interface to carry untagged traffic. By default, native VLAN is disabled.
Shutdown	Click No to enable the interface. By default, an interface in a bridge domain is disabled.

Parameter Name	Description
Static MAC Address	Click **Add Static MAC Address**, and in the MAC Static Address field that appears, enter a static MAC address entry for the interface in the bridge domain. Click Add MAC Address to add another static MAC address entry for the interface. Click Save to save the MAC address or addresses.

Configure Interface Bridge

Integrated routing and bridging (IRB) allows Cisco vEdge devices in different bridge domains to communicate with each other. To enable IRB, create logical IRB interfaces to connect a bridge domain to a VPN. The VPN provides the Layer 3 routing services necessary so that traffic can be exchanged between different VLANs. Each bridge domain can have a single IRB interface and can connect to a single VPN, and a single VPN can connect to multiple bridge domains on a Cisco vEdge device.

1. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates** to view your existing feature templates or create a new one.
3. Click **Add Template**. Choose a device from the list of devices. The templates available for the selected device display in the right pane.
4. Choose the **VPN Interface Bridge** template.
5. Enter a name and description for the template and enter the parameter. Enter other parameters described in the subsequent sections.

Create a Bridging Interface

To configure an interface to use for bridging servers, select **Basic Configuration** and enter the following details.

Parameter Name	Description
Shutdown	Click No to enable the interface.
Interface Name	Enter the name of the interface, in the format **irb***number*. The IRB interface number can be from 1 through 63, and must be the same as the VPN identifier configured in the Bridge feature template for the bridging domain that the IRB is connected to.
Description	Enter a description for the interface.
IPv4 Address	Enter the IPv4 address of the router.
DHCP Helper	Enter up to eight IP addresses for DHCP servers in the network, separated by commas, to make the interface a DHCP helper. A DHCP helper interface forwards BOOTP (Broadcast) DHCP requests that it receives from the specified DHCP servers.

Parameter Name	Description
Block Non Source IP	Click Yes to have the interface forward traffic only if the source IP address of the traffic matches the interface's IP prefix range.
Secondary IP Address (Maximum: 4)	Click Add to configure up to four secondary IPv4 addresses for a service-side interface.

To save the template, click **Save**.

Apply Access Lists

To apply access lists to IRB interfaces, select ACL and configure the following parameters:

Parameter Name	Description
Ingress ACL - IPv4	Click On, and specify the name of an IPv4 access list to packets being received on the interface.
Egress ACL - IPv4	Click On, and specify the name of an IPv4 access list to packets being transmitted on the interface.

To save the template, click **Save**.

Configure VRRP

To have an interface run the Virtual Router Redundancy Protocol (VRRP), which allows multiple devices to share a common virtual IP address for default gateway redundancy, select VRRP. Then click **New VRRP** and configure the following parameters:

Parameter Name	Description
Group ID	Enter the virtual router ID, which is a numeric identifier of the virtual router. You can configure a maximum of 24 groups.*Range:* 1 through 255
Priority	Enter the priority level of the router. There router with the highest priority is elected as the primary device. If two Cisco vEdge devices have the same priority, the one with the higher IP address is elected as the primary one.*Range:* 1 through 254*Default:* 100
Timer	Specify how often the primary VRRP router sends VRRP advertisement messages. If the subordinate routers miss three consecutive VRRP advertisements, they elect a new primary router.*Range:* 1 through 3600 seconds *Default:* 1 second

Parameter Name	Description
Track OMP Track Prefix List	By default, VRRP uses of the state of the service (LAN) interface on which it is running to determine which Cisco vEdge device is the primary virtual router. If a Cisco vEdge device loses all its WAN control connections, the LAN interface still indicates that it is up even though the router is functionally unable to participate in VRRP. To take WAN side connectivity into account for VRRP, configure one of the following: Track OMP—Click **On** for VRRP to track the Overlay Management Protocol (OMP) session running on the WAN connection. If the primary VRRP router loses all its OMP sessions, VRRP elects a new default gateway from those that have at least one active OMP session. Track Prefix List—Track both the OMP session and a list of remote prefixes, which is defined in a prefix list configured on the local router. If the primary VRRP router loses all its OMP sessions, VRRP failover occurs as described for the Track OMP option. In addition, if reachability to one of the prefixes in the list is lost, VRRP failover occurs immediately, without waiting for the OMP hold timer to expire, thus minimizing the amount of overlay traffic is dropped while the Cisco IOS XE SD-WAN device determines the primary VRRP router.
IP Address	Enter the IP address of the virtual router. This address must be different from the configured interface IP addresses of both the local Cisco IOS XE SD-WAN device and the peer running VRRP.

Add ARP Table Entries

Parameter Name	Description
IP Address	Enter the IP address for the ARP entry in dotted decimal notation or as a fully qualified host name.
MAC Address	Enter the MAC address in colon-separated hexadecimal notation.

To save the ARP configuration, click **Add**.

To save the template, click **Save**.

Advanced Interface Properties

To configure other interface properties, select Advanced and configure the following parameters:

Parameter Name	Description
MAC Address	Specify a MAC address to associate with the interface, in colon-separated hexadecimal notation.
IP MTU	Specify the maximum MTU size of packets on the interface. *Range:* 576 through 1804 *Default:* 1500 bytes
TCP MSS	Specify the maximum segment size (MSS) of TPC SYN packets passing through the vEdge router. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. *Range:* 552 to 1460 bytes *Default:* None
Clear-Dont-Fragment	Click On to clear the Don't Fragment (DF) bit in the IPv4 packet header for packets being transmitted out the interface. When the DF bit is cleared, packets larger than that interface's MTU are fragmented before being sent.
ARP Timeout	Specify how long it takes for a dynamically learned ARP entry to time out. *Range:* 0 through 2678400 seconds (744 hours) *Default:* 1200 seconds (20 minutes)
ICMP Redirect	Click Disable to disable ICMP redirect messages on the interface. By default, an interface allows ICMP redirect messages.

To save the template, click **Save**.

Configure Bridging Using CLI

Configure Bridging and Bridge Domains

Bridge domains can be marked with a VLAN tag, or they can remain untagged.

Create a Bridge Domain That Uses VLAN Tagging

For a bridge domain that uses VLAN tagging, a tag, called a VLAN ID, is inserted into all frame headers sent by the domain This tag identifies which VLAN the frames belong to, and it is used to determine which interfaces the Cisco vEdge device should send broadcast packets to.

To configure a bridge domain that uses VLAN tagging, create a bridging domain, assign a VLAN tag to that domain, and associate an interface with the domain:

1. Create a bridging domain:

   ```
   Device(config)# bridge bridge-id
   ```

   Each domain is identified by a unique integer, in the range 1 through 63. Each Cisco vEdge device can have up to 16 bridging domains.

2. Tag the bridging domain with a VLAN ID:

   ```
   Device(config-bridge)# vlan number
   ```

   The VLAN identifier can be a value from 1 through 4095.

3. Associate an interface with the bridging domain, and enable that interface:

   ```
   Device(config-bridge)# interface ge slot/port
   vEdge(config-interface)# no shutdown
   ```

   The interface must be a physical interface. You cannot use subinterfaces.

After you have added physical interfaces to a VLAN, if you want to change the VLAN identifier, you must first delete all the interfaces from the VLAN. Then configure a new VLAN identifier, and re-add the interfaces to the VLAN.

You can also configure these optional parameters:

1. Configure a description for the VLAN interface, to help identify the interface in operational command output:

   ```
   Device(config-bridge)#interface ge slot/port
   vEdge(config-interface)# description "text-description"
   ```

2. Configure a static MAC address for the VLAN interface:

   ```
   Device(config-interface)# static-mac-address aa:bb:cc:dd:ee:ff
   ```

3. Configure a name for the VLAN, to help identify the VLAN in operational command output:

   ```
   Device(config-bridge)# name "text description"
   ```

4. By default, a bridging domain can learn up to 1024 MAC addresses. You can modify this to a value from 0 through 4096:

   ```
   Device(config-bridge)# max-macs number
   ```

5. By default, MAC table entries age out after 300 seconds (5 minutes). You can modify this to a value from 10 through 4096 seconds:

   ```
   Device(config-bridge)# age-time seconds
   ```

Here is an example configuration:

```
Device# config
Device(config)# bridge 2
Device(bridge-2)# vlan 27
Device(bridge-2)# interface ge0/4
Device(interface-ge0-4)# no shutdown
Device(interface-ge0-4)# description "VLAN tag = 27"
```

```
Device(interface-ge0/4)# commit and-quit
Device# show running-config bridge
bridge 2
 vlan 27
 interface ge0/4
  description "VLAN tag = 27"
  no native-vlan
  no shutdown
 !
!
Device#
```

After your have configured an interface in a bridge domain, you add or change a VLAN identifier for that domain only by first deleting the bridge domain from the configuration (with a **no bridge** *bridge-id* command) and then reconfiguing the domain with the desired interface name and VLAN tag identifier.

To see which interfaces bridging is running on, use the **show bridge interface** command:

```
Device# show bridge interface
                         ADMIN   OPER    ENCAP                 RX    RX      TX    TX

BRIDGE  INTERFACE  VLAN  STATUS  STATUS  TYPE   IFINDEX  MTU   PKTS  OCTETS  PKTS  OCTETS

-------------------------------------------------------------------------------------------
2       ge0/4      27    Up      Up      vlan   41       1500  4     364     0     0
```

"Up" in the Admin Status column indicates that the interface has been configured, and "Up" in the Oper Status column indicates that bridging is running on the interface.

Create a Bridge Domain with an Untagged VLAN

All frames in an untagged VLAN are sent with no VLAN tag, or VLAN ID, in the frame header. For frames that already contain a tag, the tag is removed before it is sent.

In the minimal configuration for a tagged VLAN, you simply create a bridging domain that contains an interface:

1. Create a bridging domain. This domain is identified by a unique integer.

   ```
   Device(config)# bridge number
   ```

   On each Cisco vEdge router, you can configure up to 16 bridging domains.

2. Associate an interface with the bridging domain, and enable that interface:

   ```
   Device(config-bridge)# interface interface-name
   Device(config-interface)# no shutdown
   ```

You can also configure the optional parameters described in the previous section.

Configure a Native VLAN

In the minimal configuration for a native VLAN, you create a bridging domain that contains an interface, and you mark that interface as a native VLAN interface:

1. Create a bridging domain. This domain is identified by a unique integer.

   ```
   Device(config)# bridge number
   ```

   On each vEdge router, you can configure up to 16 bridging domains.

2. Associate an interface with the bridging domain, and enable that interface:

```
Device(config-bridge)# interface interface-name
Device(config-interface)# no shutdown
```

3. Enabled native VLAN on the interface:

```
Device(config-interface)# native-vlan
```

You can also configure the optional parameters described in the section about creating a tagged VLAN.

Configure IRB

With bridging, all frame traffic remains within its VLAN. To allow frames to be passed among different VLANs, you enable integrated routing and bridging (IRB). To do this, you create a logical IRB interface in a VPN domain that connects to the bridge domain. Frames with destinations in other VLANs travel over the IRB interface to the VPN domain, and the Layer 3 route table is used to forward the frames toward their destination. The route table learns the routes to other IRB interfaces. With IRB, communication can be established between VLANs that are connected to the same VPN. The VLANs can be both on the local vEdge router and on a remote router.

In a minimal configuration to configure IRB, you create an IRB interface and assign it an IP address:

1. In the desired VPN, create an IRB interface:

```
Device(config)# vpn number
Device(config-vpn)# interface irb number
```

   The VPN number can be any number from 1 through 65530, which correspond to service VPNs, except for 512 (which is the management VPN). You cannot place IRB interfaces in either the transport VPN (VPN 0) or the management VPN (VPN 512). The IRB interface type is **irb**. The IRB interface number is a number from 1 through 63, and it must be the same number as the the identifier of the bridging domain that the IRB is connected to. For example, if you configure a bridging domain with an identifier of 2 (with the command **bridge 2**), the IRB interface number must be 2, and so you must configure **interface irb2**.

2. Configure an IP address for the IRB interface. This address is the subnet for the VLAN in the connected bridge domain:

```
Device(config-irb)# ip address prefix/length
```

3. Enable the interface:

```
Device(config-irb)# no shutdown
```

In all respects, the logical IRB interfaces is just another interface. This means, for instance, that you can configure additional interfaces properties as desired. (Note, however, that you cannot configure autonegotiation on IRB interfaces.) It also means that you can ping a logical IRB interface from another device in the same VPN, and you can ping the interface regardless of whether a corresponding bridge exists for that IRB interface. That is, if you configure interface **irb4**, but there is no corresponding **bridge 4**, you are still able to ping **irb4**.

Here is an example IRB configuration:

```
Device# show running-config vpn 1
vpn 1
 interface ge0/4
  ip address 10.20.24.15/24
  no shutdown
 !
 interface irb1
```

```
  ip address 10.1.1.15/24
  no shutdown
  access-list IRB_ICMP in
  access-list IRB_ICMP out
 !
 interface irb50
  ip address 3.3.3.15/24
  no shutdown
 !
!
vEdge# show running-config vpn 2
vpn 2
 interface irb2
  ip address 2.2.2.15/24
  no shutdown
 !
!
```

To display information about the IRB interfaces, use the **show interface** command. The IRB interfaces are listed in the Interface column, and the Encapsulation Type columns marks these interfaces as "vlan".

```
Device# show interface

                                IF      IF                                                              TCP
                                ADMIN   OPER    ENCAP                                           SPEED           MSS                  RX      TX
VPN  INTERFACE  IP ADDRESS      STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR                 MBPS    DUPLEX  ADJUST  UPTIME       PACKETS PACKETS
-------------------------------------------------------------------------------------------------------------------------------------------------
0    ge0/0      10.1.15.15/24   Up      Up      null   transport  1500  00:0c:29:cb:4f:9c  10      full    0       0:02:48:12   1467    1460
0    ge0/1      -               Up      Up      null   service    1500  00:0c:29:cb:4f:a6  10      full    0       0:02:48:12   0       0
0    ge0/2      -               Up      Up      null   service    1500  00:0c:29:cb:4f:b0  10      full    0       0:02:48:03   0       0
0    ge0/3      10.0.20.15/24   Up      Up      null   service    1500  00:0c:29:cb:4f:ba  10      full    0       0:02:48:12   0       0
0    ge0/5      -               Up      Up      null   service    1500  00:0c:29:cb:4f:ce  10      full    0       0:02:48:03   0       0
0    ge0/6      -               Up      Up      null   service    1500  00:0c:29:cb:4f:d8  10      full    0       0:02:48:03   0       0
0    ge0/7      10.0.100.15/24  Up      Up      null   service    1500  00:0c:29:cb:4f:e2  10      full    0       0:02:48:12   0       0
0    system     172.16.255.15/32 Up     Up      null   loopback   1500  00:00:00:00:00:00  10      full    0       0:02:48:12   0       0
1    ge0/4      10.20.24.15/24  Up      Up      null   service    1500  00:0c:29:cb:4f:c4  10      full    0       0:02:48:00   92      14
1    irb1       10.1.1.15/24     Up      Up      vlan   service    1500  00:0c:00:00:aa:00  10      full    0       0:02:48:00   1178    0
1    irb50      3.3.3.15/24     Up      Up      vlan   service    1500  00:0c:00:00:aa:00  10      full    0       0:02:48:00   0       0
2    irb2       2.2.2.15/24     Up      Up      vlan   service    1500  00:0c:00:00:aa:00  10      full    0       0:02:48:01   0       0
512  eth0       10.0.1.15/24    Up      Up      null   service    1500  00:50:56:00:01:05  1000    full    0       0:02:48:01   210     148
```

Configuration and Monitoring Commands

CLI commands for configuring and monitoring Layer 2 bridging and Layer 3 integrated routing and bridging (IRB) on Cisco vEdge routers.

Bridging Configuration Commands

Use the following commands to configure bridging on a vEdge router.

```
bridge bridge-id
  age-time seconds
  interface interface-name
    description "text description"
    native-vlan
    [no] shutdown
    static-mac-address mac-address
  max-macs number
  name text
  vlan number
```

Bridging Monitoring Commands

Use the following commands to monitor Layer 2 bridging on a vEdge router:

- **clear bridge mac** — Clear the MAC addresses that the vEdge router has learned.
- **clear bridge statistics** —Clear the bridging statistics.
- **show bridge interface** —List information about the interfaces on which bridging is configured.
- **show bridge mac** —List the MAC addresses that the vEdge router has learned.
- **show bridge table** —List the information in the bridge forwarding table.

IRB Configuration Commands

Use the following commands to configure IRB within a VPN on a vEdge router:

```
vpn vpn-id
  interface irbnumber
    access-list acl-list
    arp
      ip address ip-address mac mac-address
    arp-timeout seconds
    autonegotiate
    clear-dont-fragment
    description "text description"
    dhcp-server (on vEdge routers only)
      address-pool prefix/length
      exclude ip-address
      lease-time minutes
      max-leases number
      offer-time minutes
      options
        default-gateway ip-address
        dns-servers ip-address
        domain-name domain-name
        interface-mtu mtu
        tftp-servers ip-address
      static-lease mac-address
    ip address address/subnet
    mac-address mac-address
    mtu bytes
    [no] shutdown
    tcp-mss-adjust bytes
```

IRB Monitoring Commands

Use the following commands to monitor IRB:

- **show interface** —List information about the interfaces on which IRB is enabled.