



Configuring Bridge Domain Interfaces

The Cisco 8500 Series Catalyst Edge Platform support the bridge domain interface (BDI) feature for packaging Layer 2 Ethernet segments into Layer 3 IP.

- [Restrictions for Bridge Domain Interfaces, on page 1](#)
- [Information About Bridge Domain Interface, on page 2](#)
- [Configuring Bridge-Domain Virtual IP Interface, on page 10](#)

Restrictions for Bridge Domain Interfaces

The following are the restrictions pertaining to bridge domain interfaces:

- Only 4096 bridge domain interfaces are supported per system.
- For a bridge domain interface, the maximum transmission unit (MTU) size can be configured between 1500 and 9216 bytes.
- Bridge domain interfaces support only the following features:
 - IPv4 Multicast
 - QoS marking and policing. Shaping and queuing are not supported
 - IPv4 VRF
 - IPv6 unicast forwarding
 - Dynamic routing such as BGP, OSPF, EIGRP, RIP, IS-IS, and STATIC
 - Hot Standby Router Protocol (HSRP)
 - Virtual Router Redundancy Protocol (VRRP) from IOS XE 3.8.0 onwards.
- Bridge domain interfaces do not support the following features:
 - PPP over Ethernet (PPPoE)
 - Bidirectional Forwarding Detection (BFD) protocol
 - QoS
 - Network-Based Application Recognition (NBAR) or Advanced Video Coding (AVC)

Information About Bridge Domain Interface

Bridge domain interface is a logical interface that allows bidirectional flow of traffic between a Layer 2 bridged network and a Layer 3 routed network traffic. Bridge domain interfaces are identified by the same index as the bridge domain. Each bridge domain represents a Layer 2 broadcast domain. Only one bridge domain interface can be associated with a bridge domain.

Bridge domain interface supports the following features:

- IP termination
- Layer 3 VPN termination
- Address Resolution Protocol (ARP), G-ARP, and P-ARP handling
- MAC address assignment

Prior to configuring a bridge domain interface, you must understand the following concepts:

- Ethernet Virtual Circuit Overview
- Bridge Domain Interface Encapsulation
- Assigning a MAC Address
- Support for IP Protocols
- Support for IP Forwarding
- Packet Forwarding
- Bridge Domain Interface Statistics

Ethernet Virtual Circuit Overview

An Ethernet Virtual Circuit (EVC) is an end-to-end representation of a single instance of a Layer 2 service that is offered by a provider. It embodies the different parameters on which the service is being offered. In the Cisco EVC Framework, the bridge domains are made up of one or more Layer 2 interfaces known as service instances. A service instance is the instantiation of an EVC on a given port on a given router. Service instance is associated with a bridge domain based on the configuration.

An incoming frame can be classified as service instance based on the following criteria:

- Single 802.1Q VLAN tag, priority-tagged, or 802.1ad VLAN tag
- Both QinQ (inner and outer) VLAN tags, or both 802.1ad S-VLAN and C-VLAN tags
- Outer 802.1p CoS bits, inner 802.1p CoS bits, or both
- Payload Ethernet type (five choices are supported: IPv4, IPv6, PPPoE-all, PPOE-discovery, and PPPoE-session)

Service instance also supports alternative mapping criteria:

- Untagged—Mapping to all the frames lacking a 802.1Q or 802.1ad header
- Default—Mapping to all the frames

Bridge Domain Interface Encapsulation

Security Group classification includes both Source and Destination Group, which is specified by source SGT and DGT. SGT Based PBR feature provides the PBR route-map match clause for SGT/DGT based packet classification. SGT Based PBR feature supports configuration of unlimited number of tags, but it is recommended to configure the tags based on memory available in the platform.

An EVC provides the ability to employ different encapsulations on each Ethernet flow point (EFP) present in a bridge domain. A BDI egress point may not be aware of the encapsulation of an egress packet because the packet may have egressed from one or more EFPs with different encapsulations.

In a bridge domain, if all the EFPs have different encapsulations, the BDI must be untagged (using the `no 802.1Q` tag). Encapsulate all the traffic in the bridge domain (popped or pushed) at the EFPs. Configure rewrite at each EFP to enable encapsulation of the traffic on the bridge domain.

In a bridge domain, if all the EFPs have the same encapsulation, configure the encapsulations on the BDI using the `encapsulation` command. Enabling encapsulation at the BDI ensures effective pushing or popping of tags, thereby eliminating the need for configuring the `rewrite` command at the EFPs. For more information on configuring the encapsulations on the BDI, see the [How to Configure a Bridge Domain Interface](#).

Assigning a MAC Address

All the bridge domain interfaces on the Cisco Catalyst 8500 Series Edge Platforms share a common MAC address. The first bridge domain interface on a bridge domain is allocated a MAC address. Thereafter, the same MAC address is assigned to all the bridge domain interfaces that are created in that bridge domain.



Note You can configure a static MAC address on a bridge domain interface using the `mac-address` command.

Support for IP Protocols

Bridge domain interfaces enable the Cisco 8500 Series Catalyst Edge Platform to act as a Layer 3 endpoint on the Layer 2 bridge domain for the following IP-related protocols:

- ARP
- DHCP
- HTTP
- ICMP
- NTP
- RARP
- SNMP
- TCP
- Telnet
- TFTP
- UDP

Support for IP Forwarding

Bridge domain interface supports the following IP forwarding features:

- IPv4 input and output access control lists (ACL)
- IPv4 input and output QoS policies. The operations supported for the input and output service policies on a bridge domain interface are:
 - Classification
 - Marking

- Policing
- IPv4 L3 VRFs

Packet Forwarding

A bridge domain interface provides bridging and forwarding services between the Layer 2 and Layer 3 network infrastructure.

Layer 2 to Layer 3

During a packet flow from a Layer 2 network to a Layer 3 network, if the destination MAC address of the incoming packet matches the bridge domain interface MAC address, or if the destination MAC address is a multicast address, the packet or a copy of the packet is forwarded to the bridge domain interface.



Note MAC address learning cannot be performed on the bridge domain interface.

Layer 3 to Layer 2

When a packet arrives at a Layer 3 physical interface of a router, a route lookup action is performed. If route lookup points to a bridge domain interface, then the bridge domain interface adds the layer 2 encapsulation and forwards the frame to the corresponding bridge domain. The byte counters are updated.

During a Layer 2 lookup on a bridge domain to which the bridge domain interface belongs, the bridge domain forwards the packets to the correct service instance based on the destination MAC address.

Link States of a Bridge Domain and a Bridge Domain Interface

Bridge domain interface acts as a routable IOS interface on Layer 3 and as a port on a bridge domain. Both bridge domain interfaces and bridge domains operate with individual administrative states.

Shutting down a bridge domain interface stops the Layer 3 data service, but does not override or impact the state of the associated bridge domain.

Shutting down a bridge domain stops Layer 2 forwarding across all the associated members including service instances and bridge domain interfaces. The associated service instances influence operational state of a bridge domain. Bridge domain interface cannot be operational unless one of the associated service instances is up.



Note Because a bridge domain interface is an internal interface, the operational state of bridge domain interface does not affect the bridge domain operational state.

BDI Initial State

The initial administrative state of a BDI depends on how the BDI is created. When you create a BDI at boot time in the startup configuration, the default administrative state for the BDI is up. It will remain in this state unless the startup configuration includes the shutdown command. This behavior is consistent with all the other interfaces. When you create a BDI dynamically at command prompt, the default administrative state is down.

BDI Link State

A BDI maintains a link state that comprises of three states: administratively down, operationally down, and up. The link state of a BDI is derived from two independent inputs: the BDI administrative state set by the corresponding users and the fault indication state from the lower levels of the interface states. It defines a BDI link state based on the state of the two inputs.

Fault Indication State	BDI Admin{start straddle 2 columns}{end straddle 2 columns}	
{start emdash} {end emdash}	Shutdown	No Shutdown
No faults asserted	Admin-down	Up
At least one fault asserted	Admin-down	Operationally-Down

Bridge Domain Interface Statistics

For virtual interfaces, such as the bridge domain interface, protocol counters are periodically queried from the QFP.

When packets flow from a Layer 2 bridge domain network to a Layer 3 routing network through the bridge domain interface, the packets are treated as bridge domain interface input packets and bytes. When packets arrive at a Layer 3 interface and are forwarded through the bridge domain interface to a Layer 2 bridge domain, the packets are treated as output packets and bytes, and the counters are updated accordingly.

A BDI maintains a standard set of Layer 3 packet counters as the case with all Cisco IOS interfaces. Use the `show interface` command to view the Layer 3 packet counters.

The convention of the counters is relative to the Layer 3 cloud. For example, input refers to the traffic entry to the Layer 3 cloud from the Layer 2 BD, while output refers to the traffic exit from the Layer 3 cloud to the Layer 2 BD.

Use the `show interfaces accounting` command to display the statistics for the BDI status. Use the `show interface <if-name>` command to display the overall count of the packets and bytes that are transmitted and received.

Creating or Deleting a Bridge Domain Interface

When you define an interface or subinterface for a Cisco IOS router, you name it and specify how it is assigned an IP address. You can create a bridge domain interface before adding a bridge domain to the system. This new bridge domain interface will be activated after the associated bridge domain is configured.



Note When a bridge domain interface is created, a bridge domain is automatically created.

When you create the bridge domain interface and the bridge domain, the system maintains the required associations for mapping the bridge domain-bridge domain interface pair.

The mapping of bridge domain and bridge domain interface is maintained in the system. The bridge domain interface uses the index of the associated bridge domain to show the association.

Bridge Domain Interface Scalability

The following table lists the bridge domain interface scalability numbers, based on the type of Cisco 8500 Series Catalyst Edge Platform Forwarding Processors.

Table 1: Bridge Domain Interface Scalability Numbers Based on the Type of Cisco 8500 Series Catalyst Edge Platform Forwarding Processor

Description
Maximum bridge domain interfaces per router

Bridge-Domain Virtual IP Interface

The Virtual IP Interface (VIF) feature helps to associate multiple BDI interfaces with a BD instance. The BD-VIF interface inherits all the existing L3 features of IOS logical IP interface.



Note You must configure every BD-VIF interface with a unique MAC address and it should belong to a different VRF.

The Virtual IP Interface (VIF) feature has the following limitations:

- BD-VIF interface does not support IP multicast.
- Number of BD-VIF interfaces with automatically generated MAC address varies on the basis of platforms.
- BD-VIF Interface does not support MPLS.
- The maximum number of BD-VIF interfaces per bridge-domain and the total number of BD-VIF interface for per system vary based on the type of platforms.

The maximum number of BD-VIF supported on Cisco Catalyst 8500 Series Edge Platforms are:

- C8500-12X4QC supports maximum 100 BD-VIF for a Bridge Domain
- C8500-12X (support maximum 16 BD-VIF for a Bridge Domain)

From Cisco IOS XE 17.7 release, BD-VIF supports Flexible Netflow (FNF).

How to Configure a Bridge Domain Interface

To configure a bridge domain interface, perform the following steps:

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface BDI** *{interface number}*
4. **encapsulation** *encapsulation dot1q* *<first-tag>* [*second-dot1q* *<second-tag>*]
5. Do one of the following:

6. **match security-group destination tag** *sgt-number*
7. **mac address** *{mac-address}*
8. **no shut**
9. **shut**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3	interface BDI <i>{interface number}</i> Example: <pre>Router(config-if)# interface BDI3</pre>	Specifies a bridge domain interface on a Cisco 8500 Series Catalyst Edge Platform.
Step 4	encapsulation encapsulation dot1q <i><first-tag></i> <i>[second-dot1q <second-tag>]</i> Example: <pre>Router(config-if)# encapsulation dot1Q 1 second-dot1q 2</pre>	Defines the encapsulation type. The example shows how to define dot1q as the encapsulation type.
Step 5	Do one of the following: Example: <pre>ip address ip-address mask</pre> Example: <pre>ipv6 address {X:X:X:X::X link-local X:X:X:X::X/prefix [anycast eui-64] autoconfig [default]}</pre> Example: <pre>Router(config-if)# ip address 2.2.2.1 255.255.255.0</pre> Example: <pre>Example:</pre>	Specifies either the IPv4 or IPv6 address for the bridge domain interface.

Example

	Command or Action	Purpose
	Router(config-if)# ipv6 address AB01:CD1:123:C::/64 eui-64	
Step 6	match security-group destination tag <i>sgt-number</i> Example: Router(config-route-map)# match security-group destination tag 150	Configures the value for security-group destination security tag.
Step 7	mac address {<i>mac-address</i>} Example: Router(config-if)# mac-address 1.1.3	Specifies the MAC address for the bridge domain interface.
Step 8	no shut Example: Router(config-if)# no shut	Enables the bridge domain interface on the Cisco 8500 Series Catalyst Edge Platform.
Step 9	shut Example: Router(config-if)# shut	Disables the bridge domain interface on the Cisco 8500 Series Catalyst Edge Platform.

Example

The following example shows the configuration of a bridge domain interface at IP address 2.2.2.1 255.255.255.0:

```
Router# configure terminal
Router(config)# interface BDI3
Router(config-if)# encapsulation dot1Q 1 second-dot1q 2
Router(config-if)# ip address 2.2.2.1 255.255.255.0
Router(config-if)# mac-address 1.1.3
Router(config-if)# no shut
Router(config-if)# exit
```

Displaying and Verifying Bridge Domain Interface Configuration

SUMMARY STEPS

1. enable
2. show interfaces bdi
3. show platform software interface fp active name
4. show platform hardware qfp active interface if-name
5. debug platform hardware qfp feature
6. platform trace runtime process forwarding-manager module

7. platform trace boottime process forwarding-manager module interfaces

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	show interfaces bdi Example: Router# show interfaces BDI3	Displays the configuration summary of the corresponding BDI.
Step 3	show platform software interface fp active name Example: Router# show platform software interface fp active name BDI4	Displays the bridge domain interface configuration in a Forwarding Processor.
Step 4	show platform hardware qfp active interface if-name Example: Router# show platform hardware qfp active interface if-name BDI4	Displays the bridge domain interface configuration in a data path.
Step 5	debug platform hardware qfp feature Example: Router# debug platform hardware qfp active feature l2bd client all	The selected CPP L2BD Client debugging is on.
Step 6	platform trace runtime process forwarding-manager module Example: Router(config)# platform trace runtime slot F0 bay 0 process forwarding-manager module interfaces level info	Enables the Forwarding Manager Route Processor and Embedded Service Processor trace messages for the Forwarding Manager process.
Step 7	platform trace boottime process forwarding-manager module interfaces Example: Router(config)# platform trace boottime slot	Enables the Forwarding Manager Route Processor and Embedded Service Processor trace messages for the Route Processor Forwarding Manager process during bootup.

Command or Action	Purpose
R0 bay 1 process forwarding-manager forwarding-manager level max	

What to do next

For additional information on the commands and the options available with each command, see the Cisco IOS Configuration Fundamentals Command Reference Guide located at:

{start hypertext} http://www.cisco.com/en/US/docs/ios/fundamentals/command/reference/cf_book.html {end hypertext}

Configuring Bridge-Domain Virtual IP Interface

```
enable
configure terminal
[no] interface BD-VIF interface-number
  [[no] vrf forwarding vrf-name]
  [[no] mac address mac-address]
  [[no] ip address ip-address mask]
  [[no] ipv6 address {X:X:X:X::X link-local | X:X:X:X::X/prefix [anycast | eui-64] | autoconfig [default]}]

exit
```

To delete BD-VIF interface, use the 'no' form of the command.

Associating VIF Interface with a Bridge Domain

```
enable
configure terminal
bridge-domain bridge-domain number
[no] member BD-VIF interface-number
exit
```

Verifying Bridge-Domain Virtual IP Interface

All existing show commands for interface and IP interface can be used for the BD-VIF interface.

```
show interface bd-vif bd-vif-id
```

```
show ip interface bd-vif bd-vif-id
```

```
show bd-vif interfaces in fman-fp
```

```
show pla sof inter fp ac brief | i BD_VIF
```

Example Configuration Bridge-Domain Virtual IP Interface

Detail sample:

```
interface Port-channell
mtu 9000
no ip address
```

```
!Ethernet service endpoint one per neutron network
service instance 1756 ethernet
  description 4e8e5957-649f-477b-9e5b-f1f75b21c03c
  encapsulation dot1q 1756
  rewrite ingress tag pop 1 symmetric
  bridge-domain 1756
!
interface BD-VIF5001
no shutdown
vrf forwarding vrf5001
ip address 10.0.0.1 255.255.255.0
interface BD-VIF5002
no shutdown
vrf forwarding vrf5002
ip address 10.0.0.2 255.255.255.0

bridge-domain 1756
member Port-channell1 service-instance 1756
member bd-vif5001
member bd-vif5002
```

