



Configuring Integrated Routing and Bridging on the Cisco ASR 9000 Series Router

This module describes the configuration of Integrated Routing and Bridging (IRB) on the Cisco ASR 9000 Series Aggregation Services Routers. IRB provides the ability to exchange traffic between bridging services on the Cisco ASR 9000 Series Router and a routed interface using a Bridge-Group Virtual Interface (BVI).

Feature History for IRB

Release	Modification
Release 4.0.1	<p>This feature was introduced on the Cisco ASR 9000 Series Router for the following line cards:</p> <ul style="list-style-type: none">• 2-Port 10-Gigabit Ethernet, 20-Port Gigabit Ethernet Combination Line Cards (A9K-2T20GE-B and A9K-2T20GE-L)• 4-Port 10-Gigabit Ethernet Line Cards (A9K-4T-B, -E, -L)• 8-Port 10-Gigabit Ethernet DX Line Cards (A9K-8T/4-B, -E, -L)• 8-Port 10-Gigabit Ethernet Line Cards (A9K-8T-B, -E, -L)• 16-Port 10-Gigabit Ethernet Line Cards (A9K-16T/8-B, -E, -L)• 40-Port Gigabit Ethernet Line Cards (A9K-40GE-B, -E, -L)

Contents

- [Prerequisites for Configuring IRB, page 140](#)
- [Information About Configuring IRB, page 141](#)
- [How to Configure IRB, page 146](#)
- [Configuration Examples for IRB, page 154](#)
- [Additional References, page 157](#)

Prerequisites for Configuring IRB

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Before configuring IRB, be sure that the following tasks and conditions are met:

- Confirm that you are configuring only the following types of Gigabit Ethernet line cards where you plan to support IRB:
 - 2-Port 10-Gigabit Ethernet, 20-Port Gigabit Ethernet Combination Line Cards (A9K-2T20GE-B and A9K-2T20GE-L)
 - 4-Port 10-Gigabit Ethernet Line Cards (A9K-4T-B, -E, -L)
 - 8-Port 10-Gigabit Ethernet DX Line Cards (A9K-8T/4-B, -E, -L)
 - 8-Port 10-Gigabit Ethernet Line Cards (A9K-8T-B, -E, -L)
 - 16-Port 10-Gigabit Ethernet Line Cards (A9K-16T/8-B, -E, -L)
 - 40-Port Gigabit Ethernet Line Cards (A9K-40GE-B, -E, -L)
- Know the IP addressing and other Layer 3 information to be configured on the bridge virtual interface (BVI).
- Complete MAC address planning if you decide to override the common global MAC address for all BVIs.
- Be sure that the BVI network address is being advertised by running static or dynamic routing on the BVI interface.

Restrictions for Configuring IRB

Before configuring IRB, consider the following restrictions:

- Only one BVI can be configured in any bridge domain.
- The same BVI can not be configured in multiple bridge domains.

**Caution**

If you want to support IRB on a Cisco ASR 9000 Series Router that also has a Cisco ASR 9000 SIP-700 installed, you must be sure to set up your routing configuration to prevent loss of traffic between the SIP-700 and a BVI interface. See the restrictions below for more information.

- IRB is not supported on any SPAs with the Cisco ASR 9000 SIP-700. However, IRB can be implemented on supported line cards in a system where a Cisco ASR 9000 SIP-700 is also installed, with the following restrictions:

- The Cisco ASR 9000 SIP-700 drops any ingress traffic it receives that is destined to a BVI interface (this can only be traffic from the Layer 3 domain to the Layer 2 domain).



Note This drop is counted with other adjacency drops in the system in the “Ipv4AdjDrop” statistic seen in the output of the **show controller pse qfp statistics drop** cisco-support command.

- The Cisco ASR 9000 SIP-700 routes any egress traffic it receives from a BVI interface (this can only be traffic from the Layer 2 domain to the Layer 3 domain).
- The following areas are *not* supported on the BVI:
 - Access Control Lists (ACLs). However, Layer 2 ACLs can be configured on each Layer 2 port of the bridge domain.
 - IP fast reroute (FRR) is not supported.
 - IPv6 addressing.
 - NetFlow.
 - MoFRR.
 - MPLS label switching.
 - mVPNv4.
 - Quality of Service (QoS).
 - Traffic mirroring.
 - Unnumbered interface for BVI.
 - Video monitoring (Vidmon).
- IRB with 802.1ah (BVI and Provider Backbone Bridge (PBB) should not be configured in the same bridge domain).
- PIM snooping. (Need to use selective flood.)
- VRF-aware DHCP relay is not supported.
- BVIs are supported only on bridge domains with the following characteristics:
 - The bridge domain supports single and double-tagged dot1q- and dot1ad-encapsulated EFPs with non-ambiguous or “exact match” EFP encapsulations. Single and double-tagged encapsulation can be specified as long as the **rewrite ingress tag pop symmetric** command is configured.
 - All Layer 2 tags must be removed. VLAN ranges are not supported.
 - Untagged EFPs are supported.

Information About Configuring IRB

This section includes the following topics:

- [IRB Introduction, page 142](#)
- [Bridge-Group Virtual Interface, page 142](#)

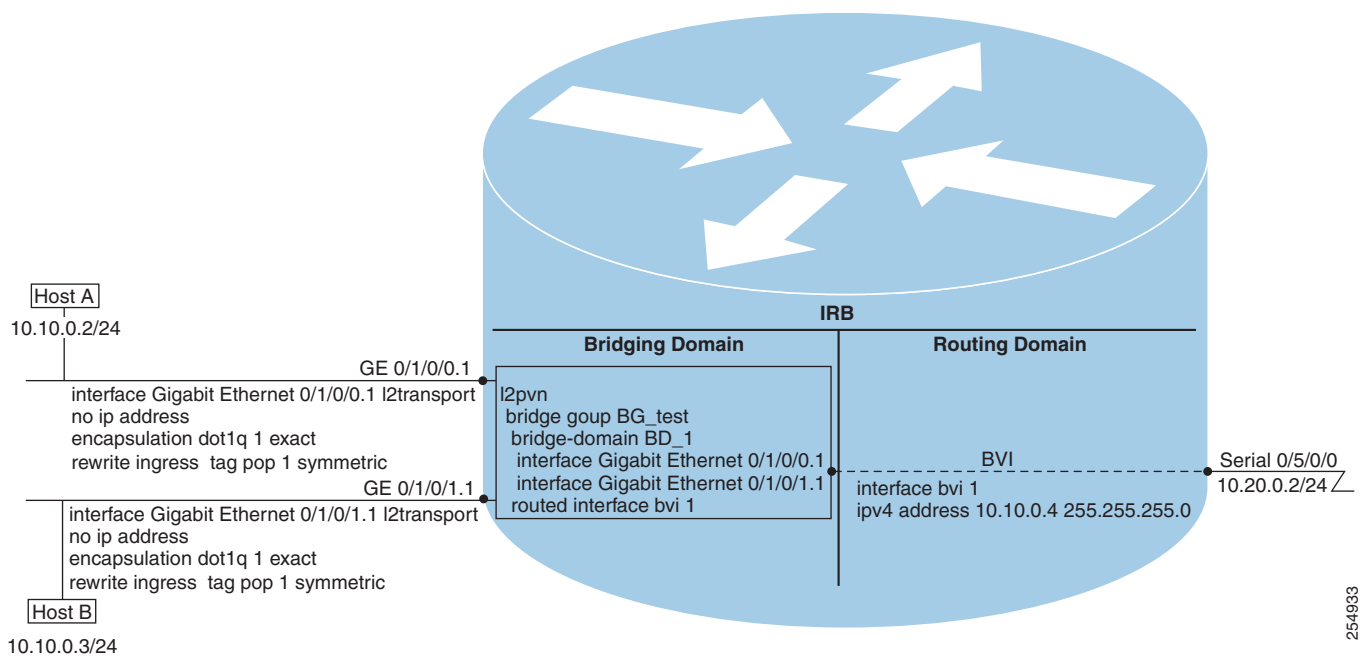
- [Packet Flows Using IRB, page 144](#)
- [Supported Environments for IRB, page 145](#)

IRB Introduction

IRB provides the ability to route between a bridge group and a routed interface using a BVI. The BVI is a virtual interface within the router that acts like a normal routed interface. A BVI is associated with a single bridge domain and represents the link between the bridging and the routing domains on the router. To support receipt of packets from a bridged interface that are destined to a routed interface, the BVI must be configured with the appropriate IP addresses and relevant Layer 3 attributes.

In software releases before Cisco IOS XR 4.0.1 where IRB is not supported, you would need to implement a physical cabling solution to connect the egress Layer 2 bridge domain interface to a Layer 3 routing domain interface on the same Cisco ASR 9000 Series Router. In Cisco IOS XR Release 4.0.1, IRB accomplishes the same functionality using a BVI and its supporting interface and bridge group configuration shown in [Figure 1](#).

Figure 1 IRB Functional View and Configuration Elements



Bridge-Group Virtual Interface

This section includes the following information:

- [BVI Introduction, page 143](#)
- [Supported Features on a BVI, page 143](#)
- [BVI MAC Address, page 143](#)
- [BVI Interface and Line Protocol States, page 143](#)

BVI Introduction

The BVI is a virtual interface within the router that acts like a normal routed interface. The BVI does not support bridging itself, but acts as a gateway for the corresponding bridge-domain to a routed interface within the router.

Aside from supporting a configurable MAC address, a BVI supports only Layer 3 attributes, and has the following characteristics:

- Uses a MAC address taken from the local chassis MAC address pool, unless overridden at the BVI interface.
- Is configured as an interface type using the **interface bvi** command and uses an IPv4 address that is in the same subnet as the hosts on the segments of the bridged domain. The BVI also supports secondary addresses.
- The BVI identifier is independent of the bridge-domain identifier. These identifiers do not need to correlate like they do in Cisco IOS software.
- Is associated to a bridge group using the **routed interface bvi** command.

Supported Features on a BVI

- The following interface commands are supported on a BVI:
 - **arp purge-delay**
 - **arp timeout**
 - **bandwidth** (The default is 10 Gbps and is used as the cost metric for routing protocols for the BVI)
 - **ipv4**
 - **mac-address**
 - **mtu** (The default is 1500 bytes)
 - **shutdown**
- The BVI supports IP helper addressing and secondary IP addressing.

BVI MAC Address

By default, the Cisco ASR 9000 Series Router uses one MAC address for all BVI interfaces on the router. However, this means that the MAC address is not unique globally. If you want to override the default and specify a unique MAC address at the BVI, then you can configure it at the BVI interface.

BVI Interface and Line Protocol States

Like typical interface states on the router, a BVI has both an Interface and Line Protocol state.

- The BVI interface state is Up when the following occurs:
 - The BVI interface is created.
 - The bridge-domain that is configured with the **routed interface bvi** command has at least one available active bridge port (Attachment circuit [AC] or pseudowire [PW]).

**Note**

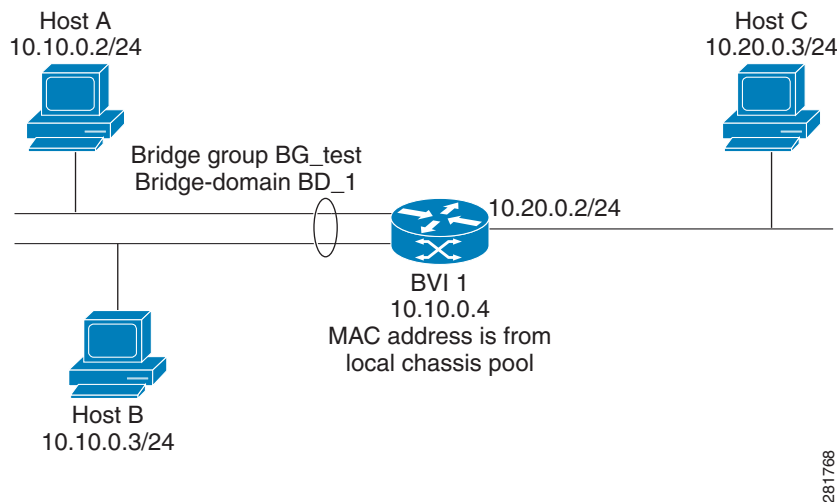
A BVI will be moved to the Down state if all of the bridge ports (Ethernet flow points [EFPs]) associated with the bridge domain for that BVI are down. However, the BVI will remain up if at least one pseudowire is up, even if all EFPs are down.

- The following characteristics determine when the the BVI line protocol state is up:
 - The bridge-domain is in Up state.
 - The BVI IP address is not in conflict with any other IP address on another active interface in the router.

Packet Flows Using IRB

Figure 2 shows a simplified functional diagram of an IRB implementation to describe different packet flows between Host A, B, and C. In this example, Host C is on a network with a connection to the same router. In reality, another router could be between Host C and the router shown.

Figure 2 *IRB Packet Flows Between Hosts*



When IRB is configured on a router, the following processing happens:

- ARP requests are resolved between the hosts and BVI that are part of the bridge domain.
- All packets from a host on a bridged interface go to the BVI if the destination MAC address matches the BVI MAC address. Otherwise, the packets are bridged.
- For packets destined for a host on a routed network, the BVI forwards the packets to the routing engine before sending them out a routed interface.
- All packets either from or destined to a host on a bridged interface go to the BVI first (unless the packet is destined for a host on the bridge domain).
- For packets that are destined for a host on a segment in the bridge domain that come in to the router on a routed interface, the BVI forwards the packet to the bridging engine, which forwards it through the appropriate bridged interface.

Packet Flows When Host A Sends to Host B on the Bridge Domain

When Host A sends data to Host B in the bridge domain on the 10.10.0.0 network, no routing occurs. The hosts are on the same subnet and the packets are bridged between their segment interfaces on the router.

Packet Flows When Host A Sends to Host C From the Bridge Domain to a Routed Interface

Using host information from [Figure 2](#), the following occurs when Host A sends data to Host C from the IRB bridging domain to the routing domain:

- Host A sends the packet to the BVI (as long as any ARP request is resolved between the host and the BVI). The packet has the following information:
 - Source MAC address of host A.
 - Destination MAC address of the BVI.
- Since Host C is on another network and needs to be routed, the BVI forwards the packet to the routed interface with the following information:
 - The Source MAC address of Host A is changed to the MAC address of the BVI.
 - IP destination address is the IP address of Host C (10.20.0.3).
- Interface 10.20.0.2 sees receipt of a packet from the routed BVI 10.10.0.4. The packet is then routed through interface 10.20.0.2 to Host C.

Packet Flows When Host C Sends to Host B From a Routed Interface to the Bridge Domain

Using host information from [Figure 2](#), the following occurs when Host C sends data to Host B from the IRB routing domain to the bridging domain:

- The packet comes into the routing domain with the following information:
 - MAC source address—MAC of Host C.
 - MAC destination address—MAC of the 10.20.0.2 ingress interface.
 - IP source address—IP address of Host C (10.20.0.3).
 - IP destination address—IP address of Host B (10.10.0.3).
- When interface 10.20.0.2 receives the packet, it looks in the routing table and determines that the packet needs to be forwarded to the BVI at 10.10.0.4.
- The routing engine captures the packet that is destined for the BVI and forwards it to the BVI's corresponding bridge domain. The packet is then bridged through the appropriate interface if the destination MAC address for Host B appears in the bridging table, or is flooded on all interfaces in the bridge group if the address is not in the bridging table.

Supported Environments for IRB

The following environments and configuration elements are supported with IRB on the Cisco ASR 9000 Series Router:

- Configuration of one BVI per bridge domain.
- Configuration of up to a maximum of 2000 BVIs.

- Up to a maximum of 128k IPv4 adjacencies.
- Layer 3 IP multicast, with ability to take ingress IP multicast traffic and bridge it to multiple Layer 2 subinterfaces (Ethernet flow points) on a bridge domain that are part of multicast groups.
- Virtual Private LAN Service (VPLS) virtual forwarding instance (VFI) configuration associated with a bridge domain configured with a BVI.
- BGP PIC edge for BVI-based prefixes.
- Traffic forwarding for the BVI using Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Routing Information Protocol Version 2 (RIPv2), and Border Gateway Protocol (BGP).
- Internet Group Management Protocol (IGMP) static groups.
- Dynamic Host Configuration Protocol (DHCP) relay. When DHCP relay is used from an aggregation node to obtain an IP address, the default gateway will be the IP address configured on the BVI. The BVI IP address should be in a common subnet as the DHCP pool that is being used by the aggregation node to assign IP addresses.
- Virtual Router Redundancy Protocol (VRRP) configuration and priority.
- Hot Standby Router Protocol (HSRP).
- Bridging of non-IP packets on a bridge domain configured with a BVI.
- Parity with stateful protocol support as currently supported on Layer 3 subinterfaces on the Cisco ASR 9000 Series Router.
- IP SLA support as currently supported on Layer 3 subinterfaces on the Cisco ASR 9000 Series Router.
- Load balancing of BVIs as ECMP paths (up to 32 paths).
- Interface-MIB.
- Packet counters for BVI interfaces.
- Multi-chassis link aggregation (LAG) on link bundles that are members of a bridge domain that uses a BVI.
- VRFs for IPv4.

How to Configure IRB

This section includes the following configuration tasks:

- [Configuring the Bridge Group Virtual Interface, page 146](#) (Required)
- [Configuring the Layer 2 AC Interfaces, page 148](#) (Required)
- [Configuring a Bridge Group and Assigning Interfaces to a Bridge Domain, page 150](#) (Required)
- [Associating the BVI as the Routed Interface on a Bridge Domain, page 152](#) (Required)
- [Displaying Information About a BVI, page 154](#) (Optional)

Configuring the Bridge Group Virtual Interface

To configure a BVI, complete the following steps.

Configuration Guidelines

Consider the following guidelines when configuring the BVI:

- The BVI must be assigned an IPv4 address that is in the same subnet as the hosts in the bridged segments.
- If the bridged network has multiple IP networks, then the BVI must be assigned secondary IP addresses for each network.

SUMMARY STEPS

1. **configure**
2. **interface bvi** *identifier*
3. **ipv4 address** *ipv4-address mask* [**secondary**]
4. **arp purge-delay** *seconds*
5. **arp timeout** *seconds*
6. **bandwidth** *rate*
7. **mac-address** *value1.value2.value3*
8. **mtu** *bytes*
9. **end**
or
commit

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface bvi <i>identifier</i> Example: RP/0/RSP0/CPU0:router(config)# interface bvi 1	Specifies or creates a BVI, where <i>identifier</i> is a number from 1 to 65535.
Step 3	ipv4 address <i>ipv4-address mask</i> [secondary] Example: RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.10.0.4 255.255.255.0	Specifies a primary or secondary IPv4 address for an interface.
Step 4	arp purge-delay <i>seconds</i> Example: RP/0/RSP0/CPU0:router(config-if)# arp purge-delay 120	(Optional) Specifies the amount of time (in <i>seconds</i>) to delay purging of Address Resolution Protocol (ARP) table entries when the interface goes down. The range is 1 to 65535. The default is no purge delay is configured.

	Command or Action	Purpose
Step 5	arp timeout <i>seconds</i> Example: RP/0/RSP0/CPU0:router(config-if)# arp timeout 12200	(Optional) Specifies how long dynamic entries learned on the interface remain in the ARP cache. The range is 30 to 2144448000 seconds. The default is 14,400 seconds (4 hours).
Step 6	bandwidth <i>rate</i> Example: RP/0/RSP0/CPU0:router(config-if)# bandwidth 1000000	(Optional) Specifies the amount of bandwidth (in kilobits per second) to be allocated on the interface. This number is used as the cost metric in routing protocols for the BVI. The range is 0 to 4294967295. The default is 10000000 (10 Gbps).
Step 7	mac-address <i>value1.value2.value3</i> Example: RP/0/RSP0/CPU0:router(config-if)# mac-address 1111.2222.3333	(Optional) Specifies the 48-bit MAC address for the BVI as three dotted-hexadecimal values, and overrides use of the default MAC address. The range for each value is 0000 to ffff. A MAC address of all 0s is not supported.
Step 8	mtu <i>bytes</i> Example: RP/0/RSP0/CPU0:router(config-if)# mtu 2000	(Optional) Specifies the maximum transmission unit (MTU) size for packets on the interface. The range is 64 to 65535. The default is 1514.
Step 9	end or commit Example: RP/0/RSP0/CPU0:router(config-if)# end or RP/0/RSP0/CPU0:router(config-if)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?</pre> [cancel]: <ul style="list-style-type: none"> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring the Layer 2 AC Interfaces

To configure the Layer 2 AC interfaces for routing by a BVI, complete the following steps.

Prerequisites

The interfaces to be configured as Layer 2 ACs in the bridge domain and routed by a BVI must be located on the following types of cards supporting IRB on the Cisco ASR 9000 Series Router:

- 2-Port 10-Gigabit Ethernet, 20-Port Gigabit Ethernet Combination Line Cards (A9K-2T20GE-B and A9K-2T20GE-L)
- 4-Port 10-Gigabit Ethernet Line Cards (A9K-4T-B, -E, -L)
- 8-Port 10-Gigabit Ethernet DX Line Cards (A9K-8T/4-B, -E, -L)
- 8-Port 10-Gigabit Ethernet Line Cards (A9K-8T-B, -E, -L)
- 40-Port Gigabit Ethernet Line Cards (A9K-40GE-B, -E, -L)

SUMMARY STEPS

1. **configure**
2. **interface { GigabitEthernet | TenGigE } interface-path-id[.subinterface] l2transport**
3. **no ip address**
4. **encapsulation dot1q vlan-id exact**
or
encapsulation dot1ad vlan-id dot1q vlan-id
5. **rewrite ingress tag pop {1 | 2} symmetric**
6. **end**
or
commit

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface [GigabitEthernet TenGigE] interface-path-id[.subinterface] l2transport Example: RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/0.1 l2transport	Enables Layer 2 transport mode on a Gigabit Ethernet or 10-Gigabit Ethernet interface or subinterface and enters interface or subinterface configuration mode, where <i>interface-path-id</i> is specified as the <i>rack/slot/module/port</i> location of the interface and <i>.subinterface</i> is the optional subinterface number.
Step 3	encapsulation dot1q vlan-id [exact] or encapsulation dot1ad vlan-id dot1q vlan-id Example: RP/0/RSP0/CPU0:router(config-if)# encapsulation dot1q 1 exact	(Optional) Specifies IEEE 802.1q encapsulation on the specified VLAN only.

	Command or Action	Purpose
Step 4	rewrite ingress tag pop {1 2} symmetric Example: RP/0/RSP0/CPU0:router(config-if)# rewrite ingress tag pop 1 symmetric	(Required if VLAN tagging configured) Specifies that one or two tags (depending on the network configuration) should be removed from frames arriving at the ingress interface to the bridge domain. Note If configuring double tags using dot1ad and dot1q encapsulation, you need to use the rewrite ingress tag pop 2 symmetric command.
Step 5	end or commit Example: RP/0/RSP0/CPU0:router(config-if)# end or RP/0/RSP0/CPU0:router(config-if)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: <ul style="list-style-type: none"> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring a Bridge Group and Assigning Interfaces to a Bridge Domain

To configure a bridge group and assign interfaces to a bridge domain, complete the following steps.

SUMMARY STEPS

1. **configure**
2. **l2vpn**
3. **bridge group** *name*
4. **bridge-domain** *name*
5. **interface** { **GigabitEthernet** | **TenGigE** } *interface-path-id* [, *subinterface*]
6. **end**
or
commit

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	l2vpn Example: RP/0/RSP0/CPU0:router(config)# l2vpn	Enters L2VPN configuration mode.
Step 3	bridge group <i>bridge-group-name</i> Example: RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group 10	Creates a bridge group and enters L2VPN bridge group configuration mode.
Step 4	bridge-domain <i>bridge-domain-name</i> Example: RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain BD_1	Creates a bridge domain and enters L2VPN bridge group bridge domain configuration mode.

	Command or Action	Purpose
Step 5	interface [GigabitEthernet TenGigE] <i>interface-path-id</i> [<i>.subinterface</i>] Example: RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd) # interface GigabitEthernet 0/1/0/0.1	Associates the Gigabit Ethernet and 10-Gigabit Ethernet interface with the specified bridge domain and enters L2VPN bridge group bridge domain attachment circuit configuration mode, where <i>interface-path-id</i> is specified as the <i>rack/slot/module/port</i> location of the interface and <i>.subinterface</i> is the optional subinterface number. Repeat this step for as many interfaces as you want to associate with the bridge domain.
Step 6	end or commit Example: RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-ac) # end or RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-ac) # commit	Saves configuration changes. <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Associating the BVI as the Routed Interface on a Bridge Domain

To associate the BVI as the routed interface on a bridge domain, complete the following steps.

SUMMARY STEPS

1. **configure**
2. **l2vpn**
3. **bridge group** *bridge-group-name*
4. **bridge-domain** *bridge-domain-name*
5. **routed interface bvi** *identifier*
6. **end**
or
commit

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	l2vpn Example: RP/0/RSP0/CPU0:router(config)# l2vpn	Enters L2VPN configuration mode.
Step 3	bridge group <i>bridge-group-name</i> Example: RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group BG_test	Creates a bridge group and enters L2VPN bridge group configuration mode.
Step 4	bridge-domain <i>bridge-domain-name</i> Example: RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain 1	Creates a bridge domain and enters L2VPN bridge group bridge domain configuration mode.
Step 5	routed interface bvi <i>identifier</i> Example: RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# routed interface bvi 1	Associates the specified BVI as the routed interface for the interfaces assigned to the bridge domain.
Step 6	end or commit Example: RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# end or RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Displaying Information About a BVI

To display information about BVI status and packet counters, use the following commands:

Command	Purpose
show interfaces bvi <i>identifier</i> [accounting brief description detail]	Displays interface status, line protocol state, and packet counters for the specified BVI.
show adjacency bvi <i>identifier</i> [detail remote]	Displays packet and byte transmit counters per adjacency to the specified BVI.
show l2vpn bridge-domain detail	Displays the reason that a BVI is down.

Configuration Examples for IRB

This section provides the following configuration examples:

- [Basic IRB Configuration: Example, page 154](#)
- [IRB Using ACs With VLANs: Example, page 155](#)
- [IPv4 Addressing on a BVI Supporting Multiple IP Networks: Example, page 155](#)
- [Comprehensive IRB Configuration with BVI Bundle Interfaces and Multicast Configuration: Example, page 156](#)
- [IRB With BVI and VRRP Configuration: Example, page 157](#)

Basic IRB Configuration: Example

The following example shows how to perform the most basic IRB configuration:

```
! Configure the BVI and its IPv4 address
!
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface bvi 1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.10.0.4 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# exit
!
! Configure the Layer 2 AC interface
!
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/0 l2transport
RP/0/RSP0/CPU0:router(config-if)# exit
!
! Configure the L2VPN bridge group and bridge domain and assign interfaces
!
RP/0/RSP0/CPU0:router(config)# l2vpn
RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group 10
RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain 1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-if)# exit
!
! Associate a BVI to the bridge domain
!
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# routed interface bvi 1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# commit
```


IRB Using ACs With VLANs: Example

The following example shows how to configure IRB on a bridge domain with Layer 2 ACs using 802.1q-encapsulated VLANs:

```
! Configure the BVI and its IPv4 address
!
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface bvi 1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.10.0.4 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# exit
!
! Configure the Layer 2 AC interfaces using dot1q encapsulation on a VLAN
!
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/0.1 l2transport
RP/0/RSP0/CPU0:router(config-if)# no ip address
RP/0/RSP0/CPU0:router(config-if)# encapsulation dot1q 1 exact
RP/0/RSP0/CPU0:router(config-if)# rewrite ingress tag pop 1 symmetric
RP/0/RSP0/CPU0:router(config-if)# exit
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1.1 l2transport
RP/0/RSP0/CPU0:router(config-if)# no ip address
RP/0/RSP0/CPU0:router(config-if)# encapsulation dot1q 1 exact
RP/0/RSP0/CPU0:router(config-if)# rewrite ingress tag pop 1 symmetric
RP/0/RSP0/CPU0:router(config-if)# exit
!
! Configure the L2VPN bridge group and bridge domain and assign interfaces
!
RP/0/RSP0/CPU0:router(config)# l2vpn
RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group 10
RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain 1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface GigabitEthernet 0/1/0/0.1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface GigabitEthernet 0/1/0/1.1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-if)# exit
!
! Associate a BVI to the bridge domain
!
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# routed interface bvi 1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# commit
```

IPv4 Addressing on a BVI Supporting Multiple IP Networks: Example

The following example shows how to configure secondary IPv4 addresses on a BVI that supports bridge domains for the 10.10.10.0/24, 10.20.20.0/24, and 10.30.30.0/24 networks. In this example, the BVI must have an address on each of the bridge domain networks:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface bvi 1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.10.10.4 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.20.20.4 255.255.255.0 secondary
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.30.30.4 255.255.255.0 secondary
RP/0/RSP0/CPU0:router(config-if)# commit
```

Comprehensive IRB Configuration with BVI Bundle Interfaces and Multicast Configuration: Example

NOTE: Router prompts will be added time permitting or in future update.

The following example shows a more comprehensive router configuration with IRB and BVI multicast support:

```
interface Bundle-Ether25
  ipv4 address 10.21.0.2 255.255.255.0
!
interface Loopback0
  ipv4 address 10.5.5.5 255.255.255.255
!
interface GigabitEthernet0/0/0/1
  negotiation auto
!
interface GigabitEthernet0/0/0/1.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag pop 1 symmetric
!
interface GigabitEthernet0/0/0/1.2 l2transport
  encapsulation dot1q 2
  rewrite ingress tag pop 1 symmetric
!

interface GigabitEthernet0/0/0/9
  bundle id 25 mode active
!
interface GigabitEthernet0/0/0/19
  bundle id 25 mode active
!
interface GigabitEthernet0/0/0/29
  bundle id 25 mode active
!

interface GigabitEthernet0/0/0/39
  bundle id 25 mode active

interface BVI1
  ipv4 address 10.1.1.1 255.255.255.0
!
interface BVI2
  ipv4 address 10.1.2.1 255.255.255.0

router ospf 100
  router-id 10.5.5.5
  area 0
    interface Bundle-Ether25
    interface Loopback0
    interface BVI1
    interface BVI2
  !
!2vpn
bridge group irb
  bridge-domain irb1
  igmp snooping profile irb_snoop
  interface GigabitEthernet0/0/0/1.1
  !
  routed interface BVI1
  !
  bridge-domain irb2
```

```
    igmp snooping profile irb_snoop
    interface GigabitEthernet0/0/0/1.2
    !
    routed interface BVI2

multicast-routing
  address-family ipv4
    interface all enable
  igmp snooping profile irb_snoop
  report-suppression disable
  !
router pim
  address-family ipv4
    rp-address 10.10.10.10
```

IRB With BVI and VRRP Configuration: Example

NOTE: Router prompts will be added time permitting or in future update.

The following example shows a partial router configuration for the relevant configuration areas for IRB support of a BVI and VRRP:

```
l2vpn
  bridge group irb
    bridge-domain irb-edge
      interface GigabitEthernet0/0/0/8
    !
    routed interface BVI 100
  !
  !
  !
interface GigabitEthernet0/0/0/8
  l2transport
  !
interface BVI100
  ipv4 address 10.21.1.1 255.255.255.0
  !
router vrrp
  interface BVI 100
    vrrp 1 ipv4 10.21.1.100
    vrrp 1 priority 100
  !
```

Additional References

The following sections provide references related to configuring IRB on the Cisco ASR 9000 Series Router.

Related Documents

Related Topic	Document Title
Ethernet L2VPN	<i>Cisco ASR 9000 Series Aggregation Services Router L2VPN and Ethernet Services Configuration Guide</i> <i>Cisco ASR 9000 Series Aggregation Services Router L2VPN and Ethernet Services Command Reference</i>
Cisco IOS XR master command reference	<i>Cisco ASR 9000 Series Aggregation Services Router Master Command Listing, Release 4.0</i>
Cisco IOS XR interface configuration commands	<i>Cisco ASR 9000 Series Aggregation Services Router Interface and Hardware Component Command Reference</i>
Cisco IOS XR multicast configuration	<i>Cisco ASR 9000 Series Aggregation Services Router Multicast Configuration Guide</i>

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
IF-MIB	To locate and download MIBs for selected platforms using Cisco IOS XR Software, use the Cisco MIB Locator found at the following URL: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

