Configuring Integrated Routing and Bridging

This module describes the configuration of Integrated Routing and Bridging (IRB). IRB provides the ability to exchange traffic between bridging services and a routed interface using a Bridge-Group Virtual Interface (BVI).

Feature History for IRB

Release	Modification
Release 6.1.1	This feature was introduced.

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.

Figure 1: IRB Functional View and Configuration Elements

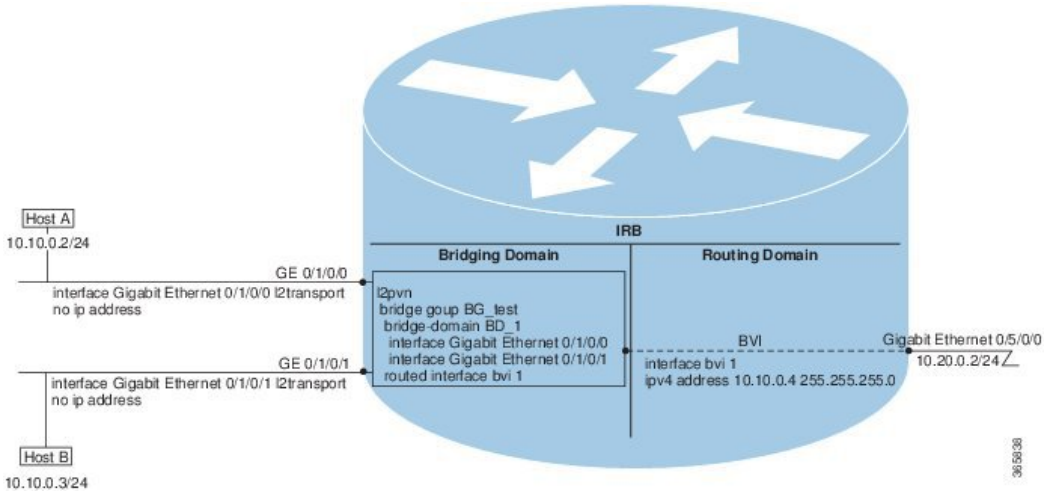

Bridge-Group Virtual Interface

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.

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.
- BVI interfaces support a number range of 1 to 4294967295.

Supported Features on a BVI

- Two-pass packet forwarding model
- These 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**
- **ipv6**
- **mac-address**
- **shutdown**

- The BVI supports IP helper addressing and secondary IP addressing.
- BVI does not support MTU configuration using **mtu** command, which is for physical interfaces. However, **ip mtu** and **ipv6 mtu** commands, which are logical interface commands, are supported.

Two-Pass Forwarding over BVI

Table 1: Feature History Table

Feature Name	Release	Description
Two-pass Forwarding over BVI	Release 7.9.1	With this release, Integrated Routing and Bridging/Bridge-group Virtual Interface (IRB/BVI) supports Layer 2 ACL, QoS, and statistics on BVI-routed packets, using a two-pass forwarding model for packets over BVI. This feature introduces the following changes: • CLI: **hw-module irb** • YANG Data Model: New XPaths for modules Cisco-IOS-XR-fia-hw-profile-cfg.yang and Cisco-IOS-XR-um-hw-module-profile-cfg.yang

The IRB/BVI implementation was originally based on single-pass or collapsed forwarding model in which each packet is processed only once. This forwarding model has some restrictions in supporting Layer 2 accounting and QoS over BVI. With this release, you have the flexibility to choose either the default single-pass model or two-pass forwarding model for packets over BVI. In the two-pass forwarding model L2 and L3 forwarding is split across two paths and packet processing happens in two cycles. This model supports Layer 2 ACL, QoS, and statistics accounting on BVI-routed packets. The earlier implementation with single-pass forwarding did not support these. You can enable the two-pass forwarding using the CLI command **hw-module irb**.

The two-pass forwarding model supports the following features:

- Layer 2 Access Control List (ACL)
- Layer 2 QoS on the BVI-routed packets
- Ingress statistics support for BVI in L2 to L3 packet flow
- Egress statistics support for BVI in L3 to L2 packet flow

The following table shows the configuration commands for different forwarding flows in BVI interfaces.

Flow	Forwarding Model	CLI Command
IRB L2 to L3	Two-Pass	**hw-module irb l2-l3 2-pass**
IRB L3 to L2	Two-Pass	**hw-module irb l3-l2 2-pass**

By default, single-pass forwarding is enabled in both IRB L2 to L3 and IRB L3 to L2 flows.

L2 to L3 Packet Flow

When a packet arrives at the ingress port, the forwarding lookup on ingress line card (LC) points to the egress BVI interface. Based on this egress BVI interface, the packet is queued to the receiving LC. The egress interface is mapped to a physical port.

When the egress BVI bandwidth is available, the receiving LC ports that are ready to receive the packets (based on the packet marking and distribution model) send grants to the ingress ports via the connectors. The ingress ports respond to this permission by transmitting the packets to the receiving LC ports. Then, according to the policy maps (PMs) the packet is queued to the appropriate egress interface. If there is no PM configured, the packet is queued to the main egress interface.

The following support is available:

- Ingress policy map (PM) is supported on both L2 access control (AC) and BVI simultaneously.
- Ingress PM on L2 AC applies to traffic on L2 to L2 direction.
- Ingress PM on BVI interface applies to the traffic on L2 to L3 direction.
- Ingress policer applied on L2 AC can check both L2 to L2 and L2 to L3 flows.
- Ingress policer applied on the BVI interface polices only L2 to L3 flow.
- Setting QoS-group, traffic-class, and discard-class are supported at ingress policy-map.

L3 to L2 Packet Flow

When a Layer 3 packet arrives at the ingress port, the destination IP address is resolved to find the corresponding Layer 2 MAC address of the destination device. Once the MAC address is obtained, a new Layer 2 Ethernet header for the packet is created with the source as the MAC address of the BVI, and the destination as the MAC address of the destination device. The packet is then transmitted over the local network and delivered to the destination device.

The following support is available:

- Egress marking and egress queuing PM are supported on L2 AC.
- No egress policy map is supported on BVI interface.
- Match on QoS-group is supported at egress marking policy-map.
- Match on discard-class is supported only for value 0 at egress marking policy-map.
- Egress queuing policy-map traffic class-based match is supported only for class default.

The two-pass model is supported on routers that have the following Cisco NCS 5700 line cards in native mode:

- NC57-18DD-SE

- NC57-36H-SE

IRB Recycle Performance

The throughput of the BVI IRB recycle port is increased from 400 to 600 Gbps in native mode. The 600 Gbps throughput mode is activated using the **hw-module profile qos irb-recycle-bandwidth** *600*.

```
Router(config)#hw-module profile qos irb-recycle-bandwidth 600
  Router(config)#
```

Note

- You should enable the two-pass forwarding capability before you configure the IRB recycle bandwidth.
- Before disabling the two-pass forwarding capability, you should remove the IRB recycle bandwidth configuration using the **no hw-module profile qos irb-recycle-bandwidth** *600*.

Configuration

To enable the two-pass forwarding capability, use the following sample configuration.

The example shows how to enable the two-pass forwarding of packets from layer 2 to layer 3:

```
Router#configure terminal
Mon Mar 27 05:17:23.887 UTC
Router(config)#hw-module irb L2-L3 2-pass
Mon Mar 27 05:17:31.421 UTC
In order to activate this new IRB model, you must manually reload the chassis/all line cards
```

The example shows how to enable the two-pass forwarding of packets from layer 3 to layer 2:

```
Router#configure terminal
Mon Mar 27 05:17:43.887 UTC
Router(config)#hw-module irb L3-L2 2-pass
Mon Mar 27 05:17:41.751 UTC
In order to activate this new IRB model, you must manually reload the chassis/all line cards
Router(config)#
```

After enabling the two-pass model, apply the ingress PM on both L2 AC and BVI interface. Use the following sample command:

```
Router(config)#
/*Apply ingress PM on both L2 AC and BVI interface*/
Router(config)#int fourHundredGigE 0/5/0/23.601
Router(config-subif)#service-policy input L2AC
Router(config-subif)#commit
Router(config-subif)#exit

Router(config)#int bvi 97
Router(config-if)#service-policy input BVI
Router(config-if)#commit
Router(config-if)#end
```

Verification

Verify the ingress policy map on BVI interface using the **show qos interface bvi** command.

To display the BVI **show qos** output, location keyword is mandatory.

```
Router#show qos interface bvi 97 input location 0/5/CPU0
NOTE:- Configured values are displayed within parentheses
Interface BVI97 ifh 0x20008034  -- input policy
NPU Id:                        0
Total number of classes:       2
Interface Bandwidth:           104857600 kbps
Policy Name:                   BVI
SPI Id:                        0x0
Accounting Type:               Layer2 (Include Layer 2 encapsulation and above)
------------------------------------------------------------------------------
Level1 Class                                =   DSCPAF33
New qos group                               =   3
New traffic class                           =   2

Policer Bucket ID                           =   0x21
Policer Stats Handle                        =   0x0
Policer committed rate                      =   150390 kbps (150 mbits/sec)
Policer peak rate                           =   200195 kbps (200 mbits/sec)
Policer conform burst                       =   186624 bytes (default)
Policer exceed burst                        =   436096 bytes (default)

Level1 Class                                =   class-default

Default Policer Bucket ID                   =   0x20
Default Policer Stats Handle                =   0x0
Policer not configured for this class

Interface BVI97 ifh 0x20008034  -- input policy
NPU Id:                        1
Total number of classes:       2
Interface Bandwidth:           104857600 kbps
Policy Name:                   BVI
SPI Id:                        0x0
Accounting Type:               Layer2 (Include Layer 2 encapsulation and above)
------------------------------------------------------------------------------
Level1 Class                                =   DSCPAF33
New qos group                               =   3
New traffic class                           =   2

Policer Bucket ID                           =   0x21
Policer Stats Handle                        =   0x0
Policer committed rate                      =   150390 kbps (150 mbits/sec)
Policer peak rate                           =   200195 kbps (200 mbits/sec)
Policer conform burst                       =   186624 bytes (default)
Policer exceed burst                        =   436096 bytes (default)

Level1 Class                                =   class-default

Default Policer Bucket ID                   =   0x20
Default Policer Stats Handle                =   0x0
Policer not configured for this class
```

To verify the ingress policy map on L2 AC using the **show qos int** *interface name* **input** command.

```
Router#show qos int fourHundredGigE 0/5/0/23.601 input
NOTE:- Configured values are displayed within parentheses
Interface FourHundredGigE0/5/0/23.601 ifh 0xa00883a  -- input policy
NPU Id:                        1
Total number of classes:       2
Interface Bandwidth:           400000000 kbps
Policy Name:                   L2AC
SPI Id:                        0x0
Accounting Type:               Layer2 (Include Layer 2 encapsulation and above)
```

```
------------------------------------------------------------------------------
Level1 Class                          =   DSCPAF43
New qos group                         =   2
New traffic class                     =   1

Policer Bucket ID                     =   0x9
Policer Stats Handle                  =   0x0
Policer committed rate                =   99609 kbps (100 mbits/sec)
Policer conform burst                 =   124672 bytes (default)

Level1 Class                          =   class-default

Default Policer Bucket ID             =   0x8
Default Policer Stats Handle          =   0x0
Policer not configured for this class
Router#
```

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 configured with the **routed interface bvi** command has at least one active bridge port available, either an Attachment Circuit or a Pseudowire.

    Attachment Circuit (AC) is a physical or logical interface that connects a customer network to a service provider network. Pseudowire (PW) is a virtual connection that emulates a physical wire, enabling data transport across a packet-switched network.

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 bridgeport is up, even if all EFPs are down.

- These 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.

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 these tasks and conditions are met:

- 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 these restrictions:

- Only one BVI can be configured in any bridge domain.
- The same BVI can not be configured in multiple bridge domains.
- MTU configuration and fragmentation of packets is not supported on BVI interfaces.
- IGP ECMP path list with mix of BVI and non-BVI paths is not supported.
- 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)
  - TI-LFA
  - SR
  - LDP
  - NetFlow
  - MoFRR
  - 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
- The following areas are *not* supported on the Layer2 bridging (with BVI):
  - Static mac entry configuration in Bridge.
  - Mac ageing configuration at global config mode.

- MAC Learning Disable.
- Vlan rewrite.

- QOS configuration on BVI interface is not supported for egress.
- Label allocation mode per-CE with BVI is not supported in an access network along with PE-CE protocols enabled.

How to Configure IRB

This section includes the following configuration tasks:

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 or IPv6 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**] **ipv6 address** *ipv6-prefix/prefix-length* [**eui-64**] [**route-tag** *route-tag value*]
4. **arp purge-delay** *seconds*
5. **arp timeout** *seconds*
6. **bandwidth** *rate*
7. **end** or **commit**

DETAILED STEPS

Procedure

Step 1 **configure**

Example:

```
Router# configure
```

Enters the global configuration mode.

Step 2 **interface bvi** *identifier*

Example:

```
Router(config)# interface bvi 1
```

Specifies or creates a BVI, where *identifier* is a number from 1 to 65535.

Step 3 **ipv4 address** *ipv4-address mask* [**secondary**] **ipv6 address** *ipv6-prefix/prefix-length* [**eui-64**] [**route-tag** *route-tag value*]

Example:

```
Router(config-if)# ipv4 address 10.10.0.4 255.255.255.0
```

Specifies a primary or secondary IPv4 address or an IPv6 address for an interface.

Step 4 **arp purge-delay** *seconds*

Example:

```
Router(config-if)#arp purge-delay 120
```

(Optional) Specifies the amount of time (in *seconds*) to delay purging of Address Resolution Protocol (ARP) table entries when the interface goes down.

The range is 1 to 65535. By default purge delay is not configured.

Step 5 **arp timeout** *seconds*

Example:

```
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** *rate*

Example:

```
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 **end** or **commit**

Example:

```
Router(config-if)# end
```

or

```
Router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

  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 Attachment Circuit (AC) interfaces for routing by a BVI, complete these steps.

SUMMARY STEPS

1. **configure**
2. **interface** [**HundredGigE | TenGigE**] **l2transport**
3. **end** or **commit**

DETAILED STEPS

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface** [**HundredGigE | TenGigE**] **l2transport**

Example:

```
RP/0/RP0/CPU0:router(config)# interface TenGigE 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.

Step 3 **end** or **commit**

Example:

```
RP/0/RP0/CPU0:router(config-if)# end
```

or

```
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- 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** *bridge-group-name*
4. **bridge-domain** *bridge-domain-name*
5. **interface** [**HundredGigE | TenGigE**
6. **end** or **commit**

DETAILED STEPS

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **l2vpn**

Example:

```
RP/0/RP0/CPU0:router(config)# l2vpn
```

Enters L2VPN configuration mode.

Step 3 **bridge group** *bridge-group-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn)# bridge group 10
```

Creates a bridge group and enters L2VPN bridge group configuration mode.

Step 4 **bridge-domain** *bridge-domain-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain BD_1
```

Creates a bridge domain and enters L2VPN bridge group bridge domain configuration mode.

Step 5 **interface** [**HundredGigE | TenGigE**

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# interface HundredGigE 0/1/0/0.1
```

Associates the 100-Gigabit Ethernet or 10-Gigabit Ethernet interface with the specified bridge domain and enters L2VPN bridge group bridge domain attachment circuit configuration mode.

Repeat this step for as many interfaces as you want to associate with the bridge domain.

Step 6 **end** or **commit**

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)# end
```

or

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- 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

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **l2vpn**

Example:

```
RP/0/RP0/CPU0:router(config)# l2vpn
```

Enters L2VPN configuration mode.

Step 3 **bridge group** *bridge-group-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn)# bridge group BG_test
```

Creates a bridge group and enters L2VPN bridge group configuration mode.

Step 4 **bridge-domain** *bridge-domain-name*

Example:

```
RP/0/RP0/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** *identifier*

Example:

```
RP/0/RP0/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/RP0/CPU0:router(config-l2vpn-bg-bd)# end
```

or

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- 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:

show interfaces bvi *identifier* [**accounting** | **brief** | **description** | **detail**]	Displays interface status, line protocol state, and packet counters for the specified BVI.
show adjacency bvi *identifier* [**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.

Additional Information on IRB

Packet Flows Using IRB

This figure 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

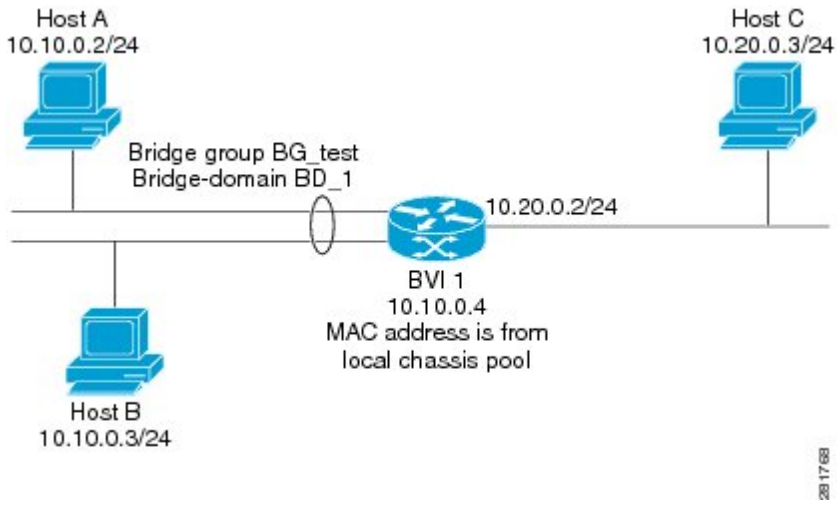

When IRB is configured on a router, the following processing happens:

- 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 this figure, 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 any ARP request the 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:
  - IP source MAC address of Host A (10.10.0.2) is changed to the MAC address of the BVI (10.10.0.4).
  - 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 this figure, 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.

Configuration Examples for IRB

This section provides the following configuration examples:

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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#interface bvi 1
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.10.0.4 255.255.255.0
RP/0/RP0/CPU0:router(config-if))# exit
!
! Configure the Layer 2 AC interface
!
RP/0/RP0/CPU0:router(config)#interface HundredGigE 0/1/0/0 l2transport
RP/0/RP0/CPU0:router(config-if))# exit
!
! Configure the L2VPN bridge group and bridge domain and assign interfaces
!
RP/0/RP0/CPU0:router(config)#l2vpn
RP/0/RP0/CPU0:router(config-l2vpn)#bridge group 10
RP/0/RP0/CPU0:router(config-l2vpn-bg)#bridge-domain 1
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface HundredGigE 0/1/0/0
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-if)# exit
!
! Associate a BVI to the bridge domain
!
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# routed interface bvi 1
RP/0/RP0/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/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#interface bvi 1
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.10.10.4 255.255.255.0
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.20.20.4 255.255.255.0 secondary
RP/0/RP0/CPU0:router(config-if)#ipv4 address 10.30.30.4 255.255.255.0 secondary
RP/0/RP0/CPU0:router(config-if))# commit
```

IRB With BVI and VRRP Configuration: Example

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

Note VRRPv6 is also supported.

```
l2vpn
 bridge group IRB
  bridge-domain IRB-EDGE
   interface TenGigE0/0/0/8
```

```
!
   routed interface BVI 100
!
interface TenGigE0/0/0/8
  l2transport
!
interface BVI 100
 ipv4 address 10.21.1.1 255.255.255.0
!
router vrrp
 interface BVI 100
 address-family ipv4
 vrrp 1
 address 10.21.1.100
 priority 100
 !
```