



Configuring IRB

This chapter describes how to configure integrated routing and bridging (IRB) for the ML-Series card. For more information about the Cisco IOS commands used in this chapter, refer to the *Cisco IOS Command Reference* publication.

This chapter includes the following major sections:

- [Understanding Integrated Routing and Bridging, page 12-1](#)
- [Configuring IRB, page 12-2](#)
- [IRB Configuration Example, page 12-3](#)
- [Monitoring and Verifying IRB, page 12-4](#)



Caution

Cisco Inter-Switch Link (ISL) and Cisco Dynamic Trunking Protocol (DTP) are not supported by the ML-Series, but the ML-Series broadcast forwards these formats. Using ISL or DTP on connecting devices is not recommended. Some Cisco devices attempt to use ISL or DTP by default.

Understanding Integrated Routing and Bridging

Your network might require you to bridge local traffic within several segments and have hosts on the bridged segments reach the hosts or ML-Series card on routed networks. For example, if you are migrating bridged topologies into routed topologies, you might want to start by connecting some of the bridged segments to the routed networks.

Using the integrated routing and bridging (IRB) feature, you can route a given protocol between routed interfaces and bridge groups within a single ML-Series card. Specifically, local or unroutable traffic is bridged among the bridged interfaces in the same bridge group, while routable traffic is routed to other routed interfaces or bridge groups.

Because bridging is in the data link layer and routing is in the network layer, they have different protocol configuration models. With IP, for example, bridge group interfaces belong to the same network and have a collective IP network address. In contrast, each routed interface represents a distinct network and has its own IP network address. Integrated routing and bridging uses the concept of a Bridge Group Virtual Interface (BVI) to enable these interfaces to exchange packets for a given protocol.

A BVI is a virtual interface within the ML-Series card that acts like a normal *routed* interface. A BVI does not support bridging but actually represents the corresponding bridge group to routed interfaces within the ML-Series card. The interface number is the link between the BVI and the bridge group.

Before configuring IRB, consider the following:

- The default routing/bridging behavior in a bridge group (when IRB is enabled) is to bridge all packets. Make sure that you explicitly configure routing on the BVI for IP traffic.
- Packets of unroutable protocols such as local-area transport (LAT) are always bridged. You cannot disable bridging for the unroutable traffic.
- Protocol attributes should not be configured on the bridged interfaces when you are using IRB to bridge and route a given protocol. You can configure protocol attributes on the BVI, but you cannot configure bridging attributes on the BVI.
- A bridge links several network segments into one large, flat network. To bridge a packet coming from a routed interface among bridged interfaces, the bridge group should be represented by one interface.
- All ports in a BVI group must have matching maximum transmission unit (MUTT) settings.

Configuring IRB

The process of configuring integrated routing and bridging consists of the following tasks:

1. Configure bridge groups and routed interfaces.
 - a. Enable bridging.
 - b. Assign interfaces to the bridge groups.
 - c. Configure the routing.
2. Enable IRB.
3. Configure the BVI.
 - a. Enable the BVI to accept routed packets.
 - b. Enable routing on the BVI.
4. Configure IP addresses on the routed interfaces.
5. Verify the IRB configuration.

When you configure the BVI and enable routing on it, packets that come in on a routed interface destined for a host on a segment that is in a bridge group are routed to the BVI and forwarded to the bridging engine. From the bridging engine, the packet exits through a bridged interface. Similarly, packets that come in on a bridged interface but are destined for a host on a routed interface go first to the BVI. The BVI forwards the packets to the routing engine that sends them out on the routed interface.

To configure a bridge group and an interface in the bridge group, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# bridge <i>bridge-group</i> protocol {ieee rstp}	Defines one or more bridge groups.
Step 2	Router(config)# interface <i>type number</i>	Enters interface configuration mode.
Step 3	Router(config-if)# bridge-group <i>bridge-group</i>	Assigns the interface to the specified bridge group.
Step 4	Router(config-if)# end	Returns to privileged EXEC mode.


To enable and configure IRB and BVI, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# bridge irb	Enables IRB. Allows bridging of traffic.
Step 2	Router(config)# interface bvi <i>bridge-group</i>	Configures the BVI by assigning the number of the corresponding bridge group to the BVI. Each bridge group can have only one corresponding BVI.
Step 3	Router(config-if)# ip address <i>ip-address ip-address-subnet-mask</i>	Configures IP addresses on routed interfaces.
Step 4	Router(config-if)# exit	Exits the interface configuration mode.
Step 5	Router(config)# bridge bridge-group route protocol	Enables a BVI to accept and route routable packets received from its corresponding bridge group. Enter this command for each protocol that you want the BVI to route from its corresponding bridge group to other routed interfaces.
Step 6	Router(config)# end	Returns to the privileged EXEC mode.
Step 7	Router# copy running-config startup-config	(Optional) Saves your configuration changes to NVRAM.

IRB Configuration Example

Figure 12-1 shows an example of IRB configuration. Example 12-1 shows the configuration code for Router A, and Example 12-2 shows the configuration code for Router B.

Figure 12-1 Configuring IRB



Example 12-1 Configuring Router A

```
bridge irb
bridge 1 protocol ieee
  bridge 1 route ip
!
!
interface FastEthernet0
  ip address 192.168.2.1 255.255.255.0
```

```

!
interface POS0
  no ip address
  crc 32
bridge-group 1
  pos flag c2 1
!
interface POS1
  no ip address
  crc 32
bridge-group 1
  pos flag c2 1
!
interface BVI1
  ip address 192.168.1.1 255.255.255.0
!
router ospf 1
  log-adjacency-changes
  network 192.168.1.0 0.0.0.255 area 0
  network 192.168.2.0 0.0.0.255 area 0

```

Example 12-2 Configuring Router B

```

bridge irb
bridge 1 protocol ieee
  bridge 1 route ip
!
!
interface FastEthernet0
  ip address 192.168.3.1 255.255.255.0
!
interface POS0
  no ip address
  crc 32
bridge-group 1
  pos flag c2 1
!
interface POS1
  no ip address
  crc 32
bridge-group 1
  pos flag c2 1
!
interface BVI1
  ip address 192.168.1.2 255.255.255.0
!
router ospf 1
  log-adjacency-changes
  network 192.168.1.0 0.0.0.255 area 0
  network 192.168.3.0 0.0.0.255 area 0

```

Monitoring and Verifying IRB

Table 12-1 shows the privileged EXEC commands for monitoring and verifying IRB.

Table 12-1 Commands for Monitoring and Verifying IRB

Command	Purpose
Router# show interfaces bvi bvi-interface-number	Shows BVI information, such as the BVI MAC address and processing statistics. The bvi-interface-number is the number of the bridge group assigned to the BVI interface.
Router# show interfaces [type-number] irb	Shows BVI information for the following: <ul style="list-style-type: none"> • Protocols that this bridged interface can route to the other routed interface (if this packet is routable). • Protocols that this bridged interface bridges

The following is sample output from the **show interfaces bvi** and **show interfaces irb** commands:

Example 12-3 Monitoring and Verifying IRB

```

Router# show interfaces bvi1
BVI1 is up, line protocol is up
  Hardware is BVI, address is 0011.2130.b340 (bia 0000.0000.0000)
  Internet address is 100.100.100.1/24
  MTU 1500 bytes, BW 145152 Kbit, DLY 5000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 03:35:28, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    0 packets input, 0 bytes, 0 no buffer
    Received 0 broadcasts (0 IP multicast)
    0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    1353 packets output, 127539 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out

Router# show interfaces irb
BVI1
Software MAC address filter on BVI1
  Hash Len  Address      Matches Act   Type
  0x00:  0 ffff.ffff.ffff      0 RCV Physical broadcast
GigabitEthernet0
Bridged protocols on GigabitEthernet0:
  clns      ip
Software MAC address filter on GigabitEthernet0
  Hash Len  Address      Matches Act   Type
  0x00:  0 ffff.ffff.ffff      0 RCV Physical broadcast
  0x58:  0 0100.5e00.0006      0 RCV IP multicast
  0x5B:  0 0100.5e00.0005      0 RCV IP multicast
  0x65:  0 0011.2130.b344      0 RCV Interface MAC address
  0xC0:  0 0100.0ccc.cccc      0 RCV CDP
  0xC2:  0 0180.c200.0000      0 RCV IEEE spanning tree
POS0

```

```

Routed protocols on POS0:
  ip
Bridged protocols on POS0:
  clns      ip
Software MAC address filter on POS0
  Hash Len   Address      Matches  Act    Type
0x00:  0  ffff.ffff.ffff      9  RCV  Physical broadcast
0x58:  0  0100.5e00.0006      0  RCV  IP multicast
0x5B:  0  0100.5e00.0005    1313  RCV  IP multicast
0x61:  0  0011.2130.b340     38  RCV  Interface MAC address
0x61:  1  0011.2130.b340      0  RCV  Bridge-group Virtual Interface
0x65:  0  0011.2130.b344      0  RCV  Interface MAC address
0xC0:  0  0100.0ccc.cccc    224  RCV  CDP
0xC2:  0  0180.c200.0000      0  RCV  IEEE spanning tree
POS1
SPR1
Bridged protocols on SPR1:
  clns      ip
Software MAC address filter on SPR1
  Hash Len   Address      Matches  Act    Type
0x00:  0  ffff.ffff.ffff      0  RCV  Physical broadcast
0x60:  0  0011.2130.b341      0  RCV  Interface MAC address
0x65:  0  0011.2130.b344      0  RCV  Interface MAC address
0xC0:  0  0100.0ccc.cccc      0  RCV  CDP
0xC2:  0  0180.c200.0000      0  RCV  IEEE spanning tree

```

Table 12-1 describes significant fields shown in the display.

Table 12-2 *show interfaces irb Field Descriptions*

Field	Description
Routed protocols on...	List of the routed protocols configured for the specified interface.
Bridged protocols on...	List of the bridged protocols configured for the specified interface.
Software MAC address filter on...	Table of software MAC address filter information for the specified interface.
Hash	Hash key/relative position in the keyed list for this MAC-address entry.
Len	Length of this entry to the beginning element of this hash chain.
Address	Canonical (Ethernet ordered) MAC address.
Matches	Number of received packets matched to this MAC address.
Routed protocols on...	List of the routed protocols configured for the specified interface.
Bridged protocols on...	List of the bridged protocols configured for the specified interface.