



Understanding and Configuring IP Multicast

This chapter describes IP multicast routing on the Catalyst 4006 switch with Supervisor Engine III. It also provides procedures and examples to configure IP multicast routing.



Note

For more information on the syntax and usage for the switch commands used in this chapter, refer to the *Command Reference for the Catalyst 4006 Switch with Supervisor Engine III*.

This chapter contains the following sections:

- [IP Multicast Overview, page 15-1](#)
- [Configuring IP Multicast Routing, page 15-12](#)
- [Configuration Examples, page 15-33](#)

IP Multicast Overview

At one end of the IP communication spectrum is IP unicast, where a source IP host sends packets to a specific destination IP host. In this case, the destination address in the IP packet is the address of a single, unique host in the IP network. These IP packets are forwarded across the network from the source to the destination host by routers. At each point along the path between the source and the destination a router uses a unicast routing table to make unicast forwarding decisions, based on the IP destination address in the packet.

At the other end of the IP communication spectrum is an IP broadcast, where a source host sends packets to all hosts on a network segment. The destination address of an IP broadcast packet has the host portion of the destination IP address set to all ones and the network portion set to the address of the subnet. IP hosts, including routers, understand that packets, which contain an IP broadcast address as the destination address, are addressed to all IP hosts on the subnet. Unless specifically configured otherwise, routers do not forward IP broadcast packets, so IP broadcast communication is normally limited to a local subnet.

IP multicasting falls between IP unicast and IP broadcast communication. IP multicast communication enables a host to send IP packets to a *group* of hosts anywhere within the IP network. To send information to a specific group, IP multicast communication uses a special form of IP destination address called an IP *multicast group address*. The IP multicast group address is specified in the IP destination address field of the packet.

To multicast IP information, Layer 3 switches and routers must forward an incoming IP packet to all output interfaces that lead to *members* of the IP multicast group. In the multicasting process on the Catalyst 4006 switch with Supervisor Engine III, a packet is replicated in the Integrated Switching Engine, forwarded to the appropriate output interfaces, and sent to each member of the multicast group.

It is not uncommon for people to think of IP multicasting and video conferencing as almost the same thing. Although the first application in a network to use IP multicast is often video conferencing, video is only one of many IP multicast applications that can add value to a company's business model. Other IP multicast applications that have potential for improving productivity include: multimedia conferencing, data replication, real-time data multicasts, and simulation applications.

IP Multicast Protocols

The Catalyst 4006 switch with Supervisor Engine III primarily uses the following protocols to implement IP multicast routing:

- Internet Group Management Protocol
- Protocol Independent Multicast
- IGMP Snooping and CGMP




Note

The Catalyst 4006 switch with Supervisor Engine III supports dynamic discovery of Distance Vector Multicast Routing Protocol (DVMRP) routers and can interoperate with them using Ethernet or DVMRP tunnels.

Figure 15-1 shows where these protocols operate within the IP multicast environment.

Figure 15-1 IP Multicast Routing Protocols



Internet Group Management Protocol

Internet Group Management Protocol (IGMP) messages are used by IP multicast hosts to send their local Layer 3 switch or router a request to join a specific multicast group and begin receiving multicast traffic. With some extensions in IGMPv2, IP hosts can also send a request to a Layer 3 switch or router to leave an IP multicast group and not receive the multicast group traffic.

Using the information obtained via IGMP, a Layer 3 switch or router maintains a list of multicast group memberships on a per interface basis. A multicast group membership is active on an interface if at least one host on the interface sends a IGMP request to receive multicast group traffic.

Protocol Independent Multicast

Protocol Independent Multicast (PIM) is *protocol independent* because it can leverage whichever unicast routing protocol is used to populate the unicast routing table, including EIGRP, OSPF, BGP, or static route, to support IP multicast. PIM also uses a unicast routing table to perform the reverse path forwarding (RPF) check function instead of building a completely independent multicast routing table. PIM does not send and receive multicast routing updates between routers like other routing protocols do.

PIM Dense Mode

PIM Dense Mode (PIM-DM) uses a *push* model to flood multicast traffic to every corner of the network. PIM-DM is intended for networks in which most LANs need to receive the multicast, such as LAN TV and corporate or financial information broadcasts. It can be an efficient delivery mechanism if there are active receivers on every subnet in the network.

PIM-DM initially floods multicast traffic throughout the network. It uses reverse path forwarding to send traffic to all Layer 3 switch or router interfaces except the one on which it arrived. Downstream routers that do not need the multicast (either because they have no receivers on their interfaces or because they are already receiving the multicast from another port) reply with a prune message, requesting to be removed from the forwarding list. This process repeats every 3 minutes.

Layer 3 switches and routers create routing state information with the flood and prune mechanism. Routing state is the source and group information that downstream routers use to build their multicast forwarding tables. PIM-DM can support only source trees—(S,G) entries. It cannot be used to build a shared distribution tree.

PIM Sparse Mode

PIM Sparse Mode (PIM-SM) uses a *pull* model to deliver multicast traffic. Only networks with active receivers that have explicitly requested the data will be forwarded the traffic. PIM-SM is intended for networks with several different multicasts, such as desktop video conferencing and collaborative computing, that go to a small number of receivers and are typically in progress simultaneously.

In PIM-SM, senders and receivers first register with a single router, which is designated as a RP. Traffic is sent by the sender to the RP, which then forwards it to the registered receivers.

As intermediate routers see the source and destination of the multicast traffic (it is unlikely that the best path from source to destination goes through the RP), they optimize the paths so that the traffic takes a more direct route (likely bypassing the RP). But traffic is still sent to the RP, in case new receivers want to register.

PIM-SM uses a shared tree to distribute the information about active sources. Depending on the configuration options, the traffic can remain on the shared tree or be changed to an optimized source distribution tree. The latter is the default behavior for PIM-SM on Cisco routers. The traffic starts to flow down the shared tree, and then routers along the path determine whether there is a better path to the source. If a more direct path exists, the designated router (the router closest to the receiver) sends a join message toward the source and then reroutes the traffic along this path.

IGMP Snooping and CGMP

IGMP snooping is used for multicasting in a Layer 2 switching environment. With IGMP snooping, a Layer 3 switch or router examines Layer 3 information in the IGMP packets in transit between hosts and a router. When the switch receives the IGMP Host Report from a host for a particular multicast group, the switch adds the host's port number to the associated multicast table entry. When the switch receives the IGMP Leave Group message from a host, it removes the host's port from the table entry.

Because IGMP control messages are transmitted as multicast packets, they are indistinguishable from multicast data if only the Layer 2 header is examined. A switch running IGMP snooping examines every multicast data packet to determine if it contains any pertinent IGMP control information. If IGMP snooping is implemented on a low end switch with a slow CPU this could have a severe performance impact when data is transmitted at high rates. On the Catalyst 4006 switch with Supervisor Engine III, IGMP snooping is implemented in the forwarding ASIC, so it does not impact the forwarding rate.

**Note**

A Catalyst 4006 switch with Supervisor Engine III can act as a CGMP server for switches that do not support IGMP snooping, such as the Catalyst 4000 family switches with Supervisor Engines I and II. You cannot configure the Catalyst 4006 switch with Supervisor Engine III as a CGMP client. To configure a Catalyst 4006 switch with Supervisor Engine III as a client, use IGMP snooping.

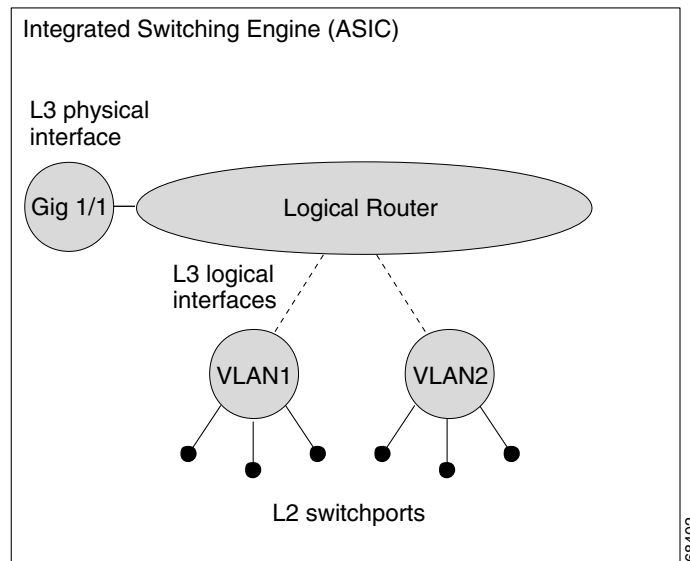
Cisco Group Management Protocol (CGMP) is a proprietary Cisco protocol that allows Catalyst switches to leverage IGMP information on Cisco routers to make Layer 2 forwarding decisions. CGMP is configured on the multicast routers and the Layer 2 switches. As a result, IP multicast traffic is delivered only to those Catalyst switchports with hosts that have requested the traffic. Switchports that have not explicitly requested the traffic will not receive it.

IP Multicast on the Catalyst 4006 Switch with Supervisor Engine III

The Catalyst 4006 switch with Supervisor Engine III supports an ASIC-based Integrated Switching Engine that provides Ethernet bridging at Layer 2 and IP routing at Layer 3. Because the ASIC is specifically designed to forward packets, the Integrated Switching Engine hardware provides very high performance with ACLs and QoS enabled. At wire-speed, forwarding in hardware is significantly faster than the CPU subsystem software, which is designed to handle exception packets.

The Integrated Switching Engine hardware supports interfaces for inter-VLAN routing and switchports for Layer 2 bridging. It also provides a physical Layer 3 interface that can be configured to connect with a host, a switch, or a router.

[Figure 15-2](#) shows a logical view of Layer 2 and Layer 3 forwarding in the Integrated Switching Engine hardware.

Figure 15-2 Logical View of Layer 2 and Layer 3 Forwarding in Hardware

CEF, MFIB, and Layer 2 Forwarding

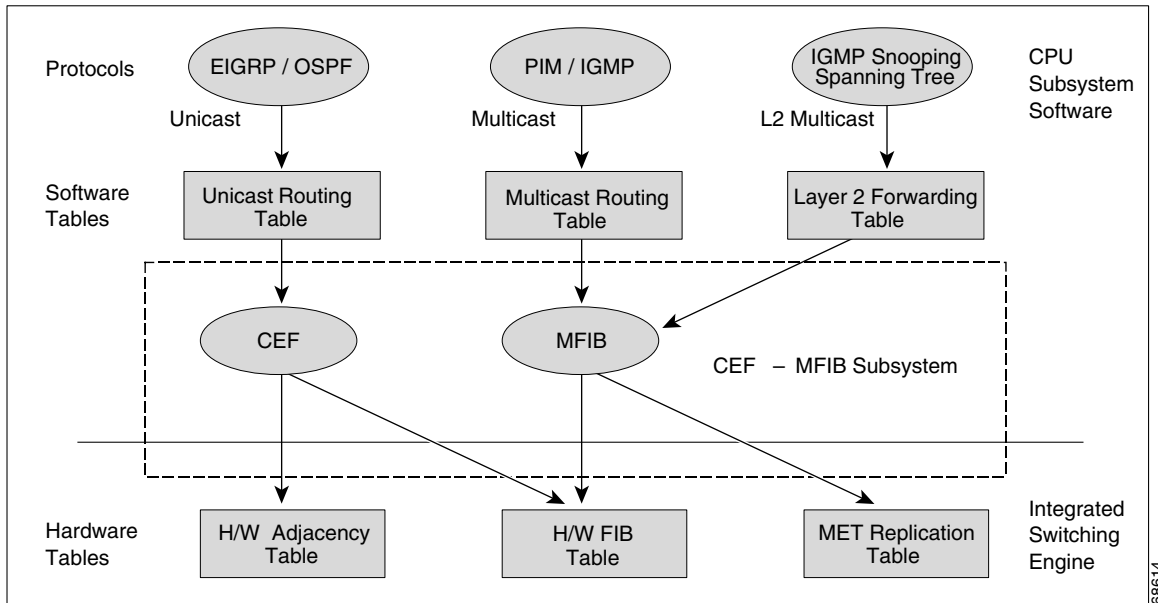
The implementation of IP multicast on the Catalyst 4006 switch with Supervisor Engine III is an extension of centralized Cisco Express Forwarding (CEF). CEF extracts information from the higher-layer unicast routing table, which is created by unicast routing protocols, such as BGP, OSPF, and EIGR and, together with platform-dependent management software, loads it into the hardware Forwarding Information Base (FIB). With the unicast routes in the FIB, when a route is changed in the upper-layer routing table, only one route needs to be changed in the hardware routing state. To forward unicast packets in hardware, the Integrated Switching Engine looks up source and destination routes in ternary content addressable memory (TCAM), takes the adjacency index from the hardware FIB, and gets the Layer 2 rewrite information and next-hop address from the hardware adjacency table.

The new multicast forwarding information base (MFIB) subsystem is the multicast analog of the unicast CEF. The MFIB subsystem extracts the multicast routes that PIM and IGMP create and refines them into a protocol-independent format for forwarding in hardware. The MFIB subsystem removes the protocol-specific information and leaves only the essential forwarding information. Each entry in the MFIB table consists of an (S,G) or (*,G) route, an input RPF VLAN, and a list of Layer 3 output interfaces. The MFIB subsystem, together with platform-dependent management software, loads this multicast routing information into the hardware FIB and hardware multicast expansion table (MET).

The Catalyst 4006 switch with Supervisor Engine III performs Layer 3 routing and Layer 2 bridging at the same time. There can be multiple Layer 2 switchports on any VLAN interface. To determine the set of output switchports on which to forward a multicast packet, the Supervisor Engine III combines the Layer 3 MFIB information with the Layer 2 forwarding information and stores it in the hardware MET for packet replication.

Figure 15-3 shows a functional overview of how the Catalyst 4006 switch with Supervisor Engine III combines unicast routing, multicast routing, and Layer 2 bridging information to forward in hardware.

Figure 15-3 Combining CEF, MFIB, and Layer 2 Forwarding Information in Hardware



Like the CEF unicast routes, the MFIB routes are Layer 3 and must be merged with the appropriate Layer 2 information. The following example shows an MFIB route:

```
(* ,224.1.1.2.3)
  RPF interface is Vlan3
  Output Interfaces are:
    Vlan 1
    Vlan 2
```

The route (*,224.1.1.2.3) is loaded in the hardware FIB table and the list of output interfaces is loaded into the MET. A pointer to the list of output interfaces, the MET index, and the RPF interface are also loaded in the hardware FIB with the (*,224.1.1.2.3) route. With this information loaded in hardware, merging the Layer 2 information can begin. For the output interfaces on VLAN1, the Integrated Switching Engine must send the packet to all switchports in VLAN1 that are in the spanning tree forwarding state. The same process applies to VLAN 2. To determine the set of switchports in VLAN 2, the Layer 2 Forwarding Table is used.

When the hardware routes a packet, in addition to sending it to all of the switchports on all output interfaces, the hardware also sends the packet to all switchports (other than the one it arrived on) in the input VLAN. For example, assume that VLAN 3 has two switchports in it, Gig 3/1 and Gig 3/2. If a host on Gig 3/1 sends a multicast packet, the host on Gig 3/2 might also need to receive the packet. To send a multicast packet to the host on Gig 3/2, all of the switchports in the ingress VLAN must be added to the portset that is loaded in the MET.

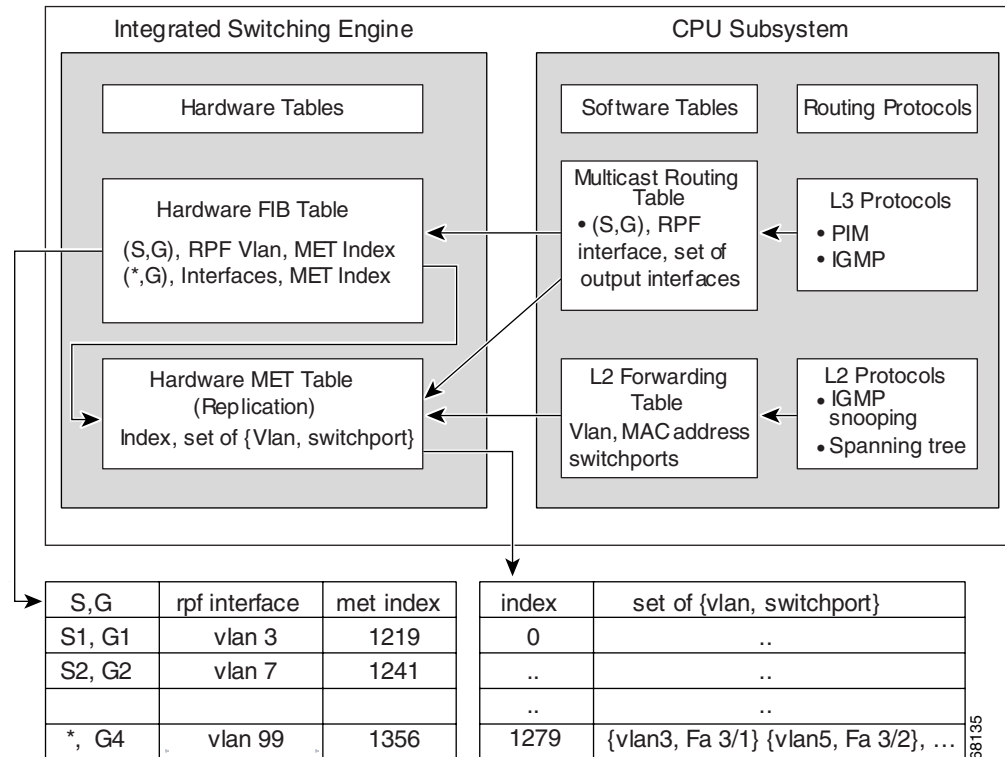
If VLAN 1 contains 1/1 and 1/2, VLAN 2 contains 2/1 and 2/2 and VLAN 3 contains 3/1 and 3/2, the MET chain for this route would contain these switchports: (1/1,1/2,2/1,2/2,3/1, and 3/2).

If IGMP snooping is on, the packet should not be forwarded to all output switchports on VLAN 2. The packet should only be forwarded to switchports where IGMP snooping has determined that there is either a group member or router. For example, if VLAN 1 had IGMP snooping enabled, and IGMP snooping determined that only port 1/2 had a group member on it, then the MET chain would contain these switchports: (1/1,1/2,2/1,2/2,3/1, and 3/2).

IP Multicast Tables

Figure 15-4 shows some key data structures that the Catalyst 4006 switch with Supervisor Engine III uses to forward IP multicast packets in hardware.

Figure 15-4 IP Multicast Tables and Protocols



The Integrated Switching Engine maintains the hardware FIB table to identify individual IP multicast routes. Each entry consists of a destination group IP address and an optional source IP address. Multicast traffic flows on primarily two types of routes: (S,G) and (*,G). The (S,G) routes flow from a source to a group based on the IP address of the multicast source and the IP address of the multicast group destination. Traffic on a (*,G) route flows from the PIM RP to all receivers of group G. Only sparse-mode groups use (*,G) routes. The Integrated Switching Engine hardware contains space for a total of 128,000 routes, which are shared by unicast routes, multicast routes, and multicast fast-drop entries.

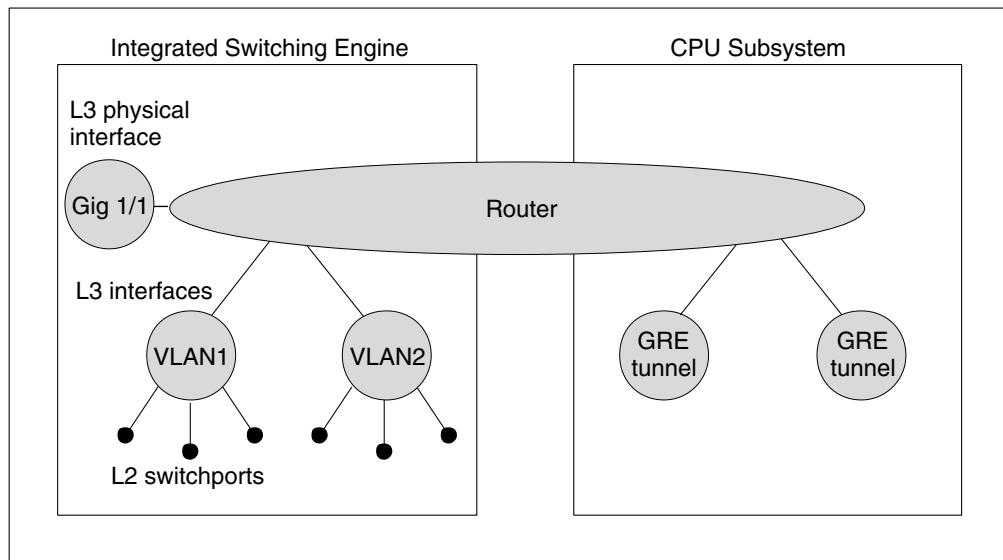
Output interface lists are stored in the multicast expansion table (MET). The MET has room for up to 32,000 output interface lists. The MET resources are shared by both Layer 3 multicast routes and by Layer 2 multicast entries. The actual number of output interface lists available in hardware depends on the specific configuration. If the total number of multicast routes exceed 32,000, multicast packets might not be switched by the Integrated Switching Engine. They would be forwarded by the CPU subsystem at much slower speeds.

Hardware and Software Forwarding

The Integrated Switching Engine forwards the majority of packets in hardware at very high rates of speed. The CPU subsystem forwards exception packets in software. Statistical reports should show that the Integrated Switching Engine is forwarding the vast majority of packets in hardware. The Catalyst 4006 switch with Supervisor Engine III is not designed to forward packets in the CPU subsystem software.

Figure 15-5 shows a logical view of the hardware and software forwarding components.

Figure 15-5 Hardware and Software Forwarding Components



In the normal mode of operation, the Integrated Switching Engine performs inter-VLAN routing in hardware. The CPU subsystem supports Generic Routing Encapsulation (GRE) tunnels for forwarding in software.

Replication is a particular type of forwarding where instead of sending out one copy of the packet, the packet is replicated and multiple copies of the packet are sent out. At Layer 3, replication only occurs for multicast packets; unicast packets are never replicated to multiple Layer 3 interfaces. In IP multicasting, for each incoming IP multicast packet that is received, many replicas of the packet are sent out.

IP multicast packets can be transmitted on the following types of routes:

- Hardware routes
- Software routes
- Partial routes

Hardware routes occur when the Integrated Switching Engine hardware forwards all replicas of a packet. Software routes occur when the CPU subsystem software forwards all replicas of a packet. Partial routes occur when the Integrated Switching Engine forwards some of the replicas in hardware and the CPU subsystem forwards some of the replicas in software.

Partial Routes

The following conditions cause some replicas of a packet for a route to be forwarded by the CPU subsystem:

**Note**

These conditions cause of the replicas to be forwarded by the CPU subsystem software but the performance of the replicas that are forwarded in hardware is not affected.

- The switch is configured with the **ip igmp join-group** command as a member of the IP multicast group on the RPF interface of the multicast source.
- The switch is the first-hop to the source in PIM sparse mode. In this case the switch must send PIM-register messages to the RP.

Software Routes

The following conditions cause all replicas of a packet for a route to be forwarded by the CPU subsystem software:

**Note**

If any one of the following conditions is configured on the RPF interface, then all replication of the output is performed in software. If any one of the following conditions is configured on the output interface, then all replication for that interface is performed in software.

- The interface is configured with multicast helper.
- The interface is a generic routing encapsulation (GRE) or Distance Vector Multicast Routing Protocol (DVMRP) tunnel.
- The interface uses non-ARPA encapsulation.

The following packets are always forwarded in software:


- Packets sent to multicast groups that fall into the range 224.0.0.* (where * is in the range from 0 to 255). This range is used by routing protocols. Layer 3 switching supports all other multicast group addresses.
- Packets with IP options.

Non-RPF Traffic

Traffic that fails an RPF check is called non-RPF traffic. Non-RPF traffic is forwarded by the Integrated Switching Engine by filtering (persistently dropping) or rate limiting the non-RPF traffic.

In a redundant configuration where multiple Layer 3 switches or routers connect to the same LAN segment, only one device forwards the multicast traffic from the source to the receivers on the outgoing interfaces. [Figure 15-6](#) shows how Non-RPF traffic can occur in a common network configuration.

Figure 15-6 Redundant Multicast Router Configuration in a Stub Network



In this kind of topology, only Router A, the PIM designated router (PIM DR), forwards data to the common VLAN. Router B receives the forwarded multicast traffic, but must drop this traffic because it has arrived on the wrong interface and fails the RPF check. Traffic that fails the RPF check is called non-RPF traffic.

Multicast Fast Drop

In IP multicast protocols, such as PIM-SM and PIM-DM, every (S,G) or (*,G) route has an incoming interface associated with it. This interface is referred to as the reverse path forwarding interface. In some cases, when a packet arrives on an interface other than the expected RPF interface, the packet must be forwarded to the CPU subsystem software to allow PIM to perform special protocol processing on the packet. One example of this special protocol processing that PIM performs is the PIM Assert protocol.

By default, the Integrated Switching Engine hardware sends all packets that arrive on a non-RPF interface to the CPU subsystem software. However, processing in software is not necessary in many cases, because these non-RPF packets are often not needed by the multicast routing protocols. The problem is that if no action is taken the non-RPF packets that are sent to the software can overwhelm the CPU.

Use the **ip mfib fastdrop** command to enable or disable MFIB fast drops.

To prevent this from happening, the CPU subsystem software loads fast-drop entries in the hardware when it receives an RPF failed packet that is not needed by the PIM protocols running on the switch. A fast-drop entry is keyed by (S,G, incoming interface). Any packet matching a fast-drop entry is bridged in the ingress VLAN, but is not sent to the software, so the CPU subsystem software is not overloaded by processing these RPF failures unnecessarily.

Protocol events, such as a link going down or a change in the unicast routing table, can impact the set of packets that can safely be fast dropped. A packet that was correctly fast dropped before might, after a topology change, need to be forwarded to the CPU subsystem software so that PIM can process it. The CPU subsystem software handles flushing fast-drop entries in response to protocol events so that the PIM code in IOS can process all the necessary RPF failures.

The use of fast-drop entries in the hardware is critical in some common topologies because it is possible to have persistent RPF failures. Without the fast-drop entries, the CPU would be exhausted by RPF failed packets that it did not need to process.

Multicast Forwarding Information Base

The Multicast Forwarding Information Base (MFIB) subsystem supports IP multicast routing in the Integrated Switching Engine hardware on the Catalyst 4006 switch with Supervisor Engine III. The MFIB logically resides between the IP multicast routing protocols in the CPU subsystem software (PIM, IGMP, MSDP, MBGP, and DVMRP) and the platform-specific code that manages IP multicast routing in hardware. The MFIB translates the routing table information created by the multicast routing protocols into a simplified format that can be efficiently processed and used for forwarding by the Integrated Switching Engine hardware.

To display the information in the multicast routing table, use the **show ip mroute** command. To display the MFIB table information, use the **show ip mfib** command. To display the information in the hardware tables, use the **show platform hardware** command.

The MFIB table contains a set of IP multicast routes. There are several types of IP multicast routes, including (S,G) and (*,G) routes. Each route in the MFIB table can have one or more optional flags associated with it. The route flags indicate how a packet that matches a route should be forwarded. For example, the Internal Copy (IC) flag on an MFIB route indicates that a process on the switch needs to receive a copy of the packet. The following flags can be associated with MFIB routes:

- Internal Copy (IC) flag—set on a route when a process on the router needs to receive a copy of all packets matching the specified route
- Signalling (S) flag—set on a route when a process needs to be notified when a packet matching the route is received. The expected behavior is that the protocol code updates the MFIB state in response to receiving a packet on a signalling interface
- Connected (C) flag—when set on an MFIB route, has the same meaning as the Signalling (S) flag, except that the C flag indicates that only packets sent by directly-connected hosts to the route should be signalled to a protocol process.

A route can also have a set of optional flags associated with one or more interfaces. For example, an (S,G) route with the flags on VLAN 1 indicates how packets arriving on VLAN 1 should be treated, and also indicate whether packets matching the route should be forwarded onto VLAN 1. The per-interface flags supported in the MFIB include:

- Accepting (A)—set on the interface that is known in multicast routing as the RPF interface. A packet that arrives on an interface that is marked as Accepting (A) is forwarded to all Forwarding (F) interfaces.
- Forwarding (F)—used in conjunction with the Accepting (A) flag as described above. The set of Forwarding interfaces that form what is often referred to as the multicast *olist* or output interface list.
- Signalling (S)—set on an interface when some multicast routing protocol process in IOS needs to be notified of packets arriving on that interface.
- Not Platform fast-switched (NP)—used in conjunction with the Forwarding (F) flag. A Forwarding interface is also marked as Not Platform fast switched whenever that output interface cannot be fast switched by the platform. The NP flag is typically used when the Forwarding interface cannot be routed in hardware and requires software forwarding. For example, Catalyst 4006 switch with Supervisor Engine III tunnel interfaces are not hardware switched, so they are marked with the NP flag. If there are any NP interfaces associated with a route, then for every packet arriving on an Accepting interface, one copy of that packet is sent to the software forwarding path for software replication to those interfaces that were not switched in hardware.

**Note**

When PIM-SM routing is in use, the MFIB route might include an interface like in the example: PimTunnel [1.2.3.4]. This is a virtual interface that the MFIB subsystem creates to indicate that packets are being tunneled to the specified destination address. A PimTunnel interface cannot be displayed with the normal **show interface** command.

S/M, 224/4

An (S/M, 224/4) entry is created in the MFIB for every multicast enabled interface. This entry ensures that all packets sent by directly-connected neighbors can be Register-encapsulated to the PIM-SM RP. Typically, only a small number of packets would be forwarded using the (S/M,224/4) route, until the (S,G) route is established by PIM-SM.

For example, on an interface with IP address 10.0.0.1 and netmask 255.0.0.0, a route would be created matching all IP multicast packets in which the source address is anything in the class A network 10. This route can be written in conventional subnet/masklength notation as (10/8,224/4). If an interface has multiple assigned IP addresses, then one route is created for each such IP address.

Unsupported Features

The following IP Multicast features are not supported in this release:

- Controlling the Transmission Rate to a Multicast Group
- Configuring an IP Multicast Boundary
- Load Splitting IP Multicast Traffic Across Equal-Cost Paths

Configuring IP Multicast Routing

The following sections describe IP multicast routing configuration tasks.

These sections describe basic IP multicast routing configuration tasks (Required):

- [Enabling IP Multicast Routing, page 15-14](#)
- [Enabling PIM on an Interface, page 15-14](#)

These sections describe basic IP multicast routing configuration tasks (Optional):

- [Configuring Auto-RP, page 15-16](#)
- [Configuring PIM Version 2, page 15-18](#)

These sections describe advanced IP multicast routing configuration tasks (Optional):

- [Configuring Advanced PIM Features, page 15-23](#)
- [Configuring an IP Multicast Static Route, page 15-27](#)

To see more complete information on IP multicast routing, refer to the *Cisco IOS IP and IP Routing Configuration Guide, Release 12.1*.

Default Configuration

Table 15-1 shows the IP multicast default configuration.

Table 15-1 Default IP Multicast Configuration

Feature	Default Value
Rate limiting of RPF	Enabled globally
IP multicast routing	Disabled globally Note When IP multicast routing is disabled, IP multicast traffic data packets are not forwarded by the Catalyst 4006 switch with Supervisor Engine III. However, IP multicast control traffic will continue to be processed and forwarded. Therefore, IP multicast routes can remain in the routing table even if ip multicast routing is disabled.
PIM	Disabled on all interfaces
IGMP snooping	Enabled on all VLAN interfaces Note If you disable IGMP snooping on an interface, all output ports are forwarded by the Integrated Switching Engine. When IGMP snooping is disabled on an input VLAN interface, multicast packets related to that interface are sent to all forwarding switchports in the VLAN.



Note

Source Specific Multicast and IGMP v3 are supported.

For more information about source specific multicast with IGMPv3 and IGMP, refer to the following url:
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/dtssm5t.htm>

Enabling IP Multicast Routing

Enabling IP multicast routing allows the Catalyst 4006 switch with Supervisor Engine III to forward multicast packets. To enable IP multicast routing on the router, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip multicast-routing	Enables IP multicast routing.

Enabling PIM on an Interface

Enabling PIM on an interface also enables IGMP operation on that interface. An interface can be configured to be in dense mode, sparse mode, or sparse-dense mode. The mode determines how the Layer 3 switch or router populates its multicast routing table and how the Layer 3 switch or router forwards multicast packets it receives from its directly connected LANs. You must enable PIM in one of these modes for an interface to perform IP multicast routing.

In populating the multicast routing table, dense-mode interfaces are always added to the table. Sparse-mode interfaces are added to the table only when periodic join messages are received from downstream routers, or when there is a directly connected member on the interface. When forwarding from a LAN, sparse-mode operation occurs if there is an RP known for the group. If so, the packets are encapsulated and sent toward the RP. When no RP is known, the packet is flooded in a dense-mode fashion. If the multicast traffic from a specific source is sufficient, the receiver's first-hop router can send join messages toward the source to build a source-based distribution tree.

There is no default mode setting. By default, multicast routing is disabled on an interface.

Enabling Dense Mode

To configure PIM on an interface to be in dense mode, enter the following command in interface configuration mode:

Command	Purpose
Switch(config-if)# ip pim dense-mode	Enables dense-mode PIM on the interface.

See the “[PIM Dense Mode Example](#)” section at the end of this chapter for an example of how to configure a PIM interface in dense mode.

Enabling Sparse Mode

To configure PIM on an interface to be in sparse mode, enter the following command in interface configuration mode:

Command	Purpose
Switch(config-if)# ip pim sparse-mode	Enables sparse-mode PIM on the interface.

See the “[PIM Sparse Mode Example](#)” section at the end of this chapter for an example of how to configure a PIM interface in sparse mode.

Enabling Sparse-Dense Mode

When you enter either the **ip pim sparse-mode** or **ip pim dense-mode** command, sparseness or denseness is applied to the interface as a whole. However, some environments might require PIM to run in a single region in sparse mode for some groups and in dense mode for other groups.

An alternative to enabling only dense mode or only sparse mode is to enable sparse-dense mode. In this case, the interface is treated as dense mode if the group is in dense mode; the interface is treated in sparse mode if the group is in sparse mode. If you want to treat the group as a sparse group, and the interface is in sparse-dense mode, you must have an RP.

If you configure sparse-dense mode, the idea of sparseness or denseness is applied to the group on the switch, and the network manager should apply the same concept throughout the network.

Another benefit of sparse-dense mode is that Auto-RP information can be distributed in a dense-mode manner; yet, multicast groups for user groups can be used in a sparse-mode manner. Thus, there is no need to configure a default RP at the leaf routers.

When an interface is treated in dense mode, it is populated in a multicast routing table’s outgoing interface list when either of the following is true:

- There are members or DVMRP neighbors on the interface.
- There are PIM neighbors and the group has not been pruned.

When an interface is treated in sparse mode, it is populated in a multicast routing table’s outgoing interface list when either of the following is true:

- There are members or DVMRP neighbors on the interface.
- An explicit join has been received by a PIM neighbor on the interface.

To enable PIM to operate in the same mode as the group, enter the following command in interface configuration mode:

Command	Purpose
Switch(config-if)# ip pim sparse-dense-mode	Enables PIM to operate in sparse or dense mode, depending on the group.

Configuring a Rendezvous Point

If you configure PIM to operate in sparse mode, you must also choose one or more routers to be a rendezvous point (RP). You need not configure the routers to be RPs; they learn this themselves. RPs are used by senders to a multicast group to announce their existence and by receivers of multicast packets to learn about new senders. The Cisco IOS software can be configured so that packets for a single multicast group can use one or more RPs.

The RP address is used by first-hop routers to send PIM register messages on behalf of a host sending a packet to the group. The RP address is also used by last-hop routers to send PIM join and prune messages to the RP to inform it about group membership. You must configure the RP address on all routers (including the RP router).

A PIM router can be an RP for more than one group. Only one RP address can be used at a time within a PIM domain. The conditions specified by the access list determine for which groups the router is an RP.

To configure the address of the RP, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim rp-address <i>ip-address [access-list-number] [override]</i>	Configures the address of a PIM RP.

Configuring Auto-RP

Auto-RP is a feature that automates the distribution of group-to-RP mappings in a PIM network. This feature has the following benefits:

- It is easy to use multiple RPs within a network to serve different group ranges.
- It allows load splitting among different RPs and arrangement of RPs according to the location of group participants.
- It avoids inconsistent, manual RP configurations that can cause connectivity problems.

Multiple RPs can be used to serve different group ranges or serve as hot backups of each other. To make Auto-RP work, a router must be designated as an *RP-mapping agent*, which receives the RP-announcement messages from the RPs and arbitrates conflicts. The RP-mapping agent then sends the consistent group-to-RP mappings to all other routers. Thus, all routers automatically discover which RP to use for the groups they support.



Note

If you configure PIM in sparse mode or sparse-dense mode and do not configure Auto-RP, you must statically configure an RP as described in the section “[Assigning an RP to Multicast Groups](#)” later in this chapter.



Note

If router interfaces are configured in sparse mode, Auto-RP can still be used if all routers are configured with a static RP address for the Auto-RP groups.

Setting Up Auto-RP in a New Internetwork

If you are setting up Auto-RP in a new internetwork, you do not need a default RP because you configure all the interfaces for sparse-dense mode. Follow the process described in the section “[Adding Auto-RP to an Existing Sparse-Mode Cloud](#),” omitting the first step of choosing a default RP.

Adding Auto-RP to an Existing Sparse-Mode Cloud

The following sections contain some suggestions for the initial deployment of Auto-RP into an existing sparse-mode cloud, to minimize disruption of the existing multicast infrastructure.

Choosing a Default RP

Sparse-mode environments need a default RP; sparse-dense-mode environments do not. If you have sparse-dense mode configured everywhere, you do not choose a default RP.

Adding Auto-RP to a sparse-mode cloud requires a default RP. In an existing PIM sparse-mode region, at least one RP is defined across the network that has good connectivity and availability. That is, the **ip pim rp-address** command is already configured on all routers in this network.

Use that RP for the global groups (for example, 224.x.x.x and other global groups). There is no need to reconfigure the group address range that RP serves. RPs discovered dynamically through Auto-RP take precedence over statically configured RPs. Typically, you would use a second RP for the local groups.

Announcing the RP and the Group Range It Serves

Find another router to serve as the RP for the local groups. The RP-mapping agent can double as an RP itself. Assign the whole range of 239.x.x.x to that RP, or assign a subrange of that (for example, 239.2.x.x).

To designate that a router is the RP, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim send-rp-announce <i>type number scope ttl group-list access-list-number</i>	Configures a router to be the RP.

To change the group ranges this RP optimally serves in the future, change the announcement setting on the RP. If the change is valid, all other routers automatically adopt the new group-to-RP mapping.

The following example advertises the IP address of Ethernet 0 as the RP for the administratively scoped groups:

```
ip pim send-rp-announce ethernet0 scope 16 group-list 1
access-list 1 permit 239.0.0.0 0.255.255.255
```

Assigning the RP Mapping Agent

The RP mapping agent is the router that sends the authoritative Discovery packets notifying other routers which group-to-RP mapping to use. Such a role is necessary in the event of conflicts (such as overlapping group-to-RP ranges).

Find a router for which connectivity is not likely to be interrupted and assign it the role of RP-mapping agent. All routers within the TTL number of hops from the source router receive the Auto-RP Discovery messages. To assign the role of RP mapping agent in that router, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim send-rp-discovery scope <i>ttl</i>	Assigns the RP mapping agent.

Verifying the Group-to-RP Mapping

To learn if the group-to-RP mapping has arrived, enter one of the following commands in EXEC mode on the designated routers:

Command	Purpose
Switch# show ip pim rp mapping	Displays active RPs that are cached with associated multicast routing entries. Information learned by configuration or Auto-RP.
Switch# show ip pim rp [<i>group-name</i> <i>group-address</i>] [mapping]	Displays information actually cached in the routing table.

Preventing Join Messages to False RPs

Note the **ip pim accept-rp** commands previously configured throughout the network. If the **ip pim accept-rp** command is not configured on any router, this problem can be addressed later. In those routers already configured with the **ip pim accept-rp** command, you must specify the command again to accept the newly advertised RP.

To accept all RPs advertised with Auto-RP and reject all other RPs by default, use the **ip pim accept-rp auto-rp** command.

If all interfaces are in sparse mode, a default configured RP to support the two well-known groups 224.0.1.39 and 224.0.1.40. Auto-RP relies on these two well-known groups to collect and distribute RP-mapping information. When this is the case and the **ip pim accept-rp auto-rp** command is configured, another **ip pim accept-rp** command accepting the default RP must be configured. The following is an example of this configuration:

```
ip pim accept-rp default RP address 1
access-list 1 permit 224.0.1.39
access-list 1 permit 224.0.1.40
```

Filtering Incoming RP Announcement Messages

To filter incoming RP announcement messages, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim rp-announce-filter rp-list access-list-number group-list access-list-number	Filters incoming RP announcement messages.

Configuring PIM Version 2

PIM Version 2 includes the following improvements over PIM Version 1:

- A single, active RP exists per multicast group, with multiple backup RPs. This single RP compares to multiple active RPs for the same group in PIM Version 1.
- A bootstrap router (BSR) provides a fault-tolerant, automated RP discovery and distribution mechanism. Thus, routers dynamically learn the group-to-RP mappings.
- Sparse mode and dense mode are properties of a group, as opposed to an interface. We strongly recommend sparse-dense mode, as opposed to either sparse mode or dense mode only.
- PIM join and prune messages have more flexible encodings for multiple address families.
- A more flexible Hello packet format replaces the Query packet to encode current and future capability options.
- Register messages to an RP indicate whether they were sent by a border router or a designated router.
- PIM packets are no longer inside IGMP packets; they are standalone packets.

PIM Version 1, used with the Auto-RP feature, can perform the same tasks as the PIM Version 2 BSR. However, Auto-RP is a standalone protocol, separate from PIM Version 1, and is Cisco proprietary. PIM Version 2 is a standards track protocol in the IETF. We recommend that you use PIM Version 2.

Either the BSR or Auto-RP should be chosen for a given range of multicast groups. If there are PIM Version 1 routers in the network, do not use the BSR.

Cisco's PIM Version 2 implementation allows interoperability and transition between Version 1 and Version 2, although there might be some minor problems. You can upgrade to PIM Version 2 incrementally. PIM Versions 1 and 2 can be configured on different routers within one network. Internally, all routers on a shared media network must run the same PIM version. Therefore, if a PIM Version 2 router detects a PIM Version 1 router, the Version 2 router downgrades itself to Version 1 until all Version 1 routers have been shut down or upgraded.

PIM uses the BSR to discover and announce RP-set information for each group prefix to all the routers in a PIM domain. This is the same function accomplished by Auto-RP, but the BSR is part of the PIM Version 2 specification. The BSR mechanism interoperates with Auto-RP on Cisco routers.

To avoid a single point of failure, you can configure several candidate BSRs in a PIM domain. A BSR is elected among the candidate BSRs automatically; they use bootstrap messages to discover which BSR has the highest priority. This router then announces to all PIM routers in the PIM domain that it is the BSR.

Routers that are configured as candidate RPs then unicast to the BSR the group range for which they are responsible. The BSR includes this information in its bootstrap messages and disseminates it to all PIM routers in the domain. Based on this information, all routers will be able to map multicast groups to specific RPs. As long as a router is receiving the bootstrap message, it has a current RP map.

Prerequisites

When PIM Version 2 routers interoperate with PIM Version 1 routers, Auto-RP should have already been deployed. A PIM Version 2 BSR that is also an Auto-RP mapping agent will automatically advertise the RP elected by Auto-RP. That is, Auto-RP prevails in its single RP being imposed on every router in the group. All routers in the domain refrain from trying to use the PIM Version 2 hash function to select multiple RPs.

Because bootstrap messages are sent hop by hop, a PIM Version 1 router will prevent these messages from reaching all routers in your network. Therefore, if your network has a PIM Version 1 router in it, and only Cisco routers, it is best to use Auto-RP rather than the bootstrap mechanism. If you have a network that includes routers from other vendors, configure the Auto-RP mapping agent and the BSR on a Cisco PIM Version 2 router. Also ensure that no PIM Version 1 router is located on the path between the BSR and a non-Cisco PIM Version 2 router.

Transitioning to PIM Version 2

On each LAN, the Cisco implementation of PIM Version 2 automatically enforces the rule that all PIM messages on a shared LAN are in the same PIM version. To accommodate that rule, if a PIM Version 2 router detects a PIM Version 1 router on the same interface, the Version 2 router downgrades itself to Version 1 until all Version 1 routers have been shut down or upgraded.

Deciding When to Configure a BSR

If there are only Cisco routers in your network (no routers from other vendors), there is no need to configure a BSR. Configure Auto-RP in the mixed PIM Version 1/Version 2 environment.

However, if you have non-Cisco, PIM Version 2 routers that need to interoperate with Cisco routers running PIM Version 1, both Auto-RP and a BSR are required. We recommend that a Cisco PIM Version 2 router be both the Auto-RP mapping agent and the BSR.

Dense Mode

Dense mode groups in a mixed Version 1/Version 2 region need no special configuration; they will interoperate automatically.

Sparse Mode

Sparse mode groups in a mixed Version 1/Version 2 region are possible because the Auto-RP feature in Version 1 interoperates with the RP feature of Version 2. Although all PIM Version 2 routers are also capable of using Version 1, we recommend that the RPs be upgraded to Version 2 (or at least upgraded to PIM Version 1 in the Cisco IOS Release 11.3 software).

To ease the transition to PIM Version 2, we also recommend the following:

- Use Auto-RP throughout the region
- Configure Sparse-dense mode throughout the region

If Auto-RP was not already configured in the PIM Version 1 regions, configure Auto-RP.

PIM Version 2 Configuration Tasks

There are two approaches to using PIM Version 2. You can use Version 2 exclusively in your network, or migrate to Version 2 by employing a mixed PIM version environment.

- If your network is all Cisco routers, you may use either Auto-RP or the bootstrap mechanism (BSR).
- If you have non-Cisco routers in your network, you need to use the bootstrap mechanism.
- If you have PIM Version 1 and PIM Version 2 Cisco routers and routers from other vendors, then you must use both Auto-RP and the bootstrap mechanism.

The tasks to configure PIM Version 2 are described in the following sections:

- [Specifying the PIM Version, page 15-20](#)
- [Configuring PIM Version 2 Only, page 15-21](#)
- [Transitioning to PIM Version 2, page 15-19](#)
- [Monitoring the RP Mapping Information, page 15-23](#)
- [Troubleshooting, page 15-23](#)

Specifying the PIM Version

All systems using Cisco IOS Release 11.3(2)T or later start in PIM Version 2 mode by default. In case you need to reenable PIM Version 2 or specify PIM Version 1 for some reason, you can control the PIM version by entering the following command in interface configuration mode:

Command	Purpose
Switch(interface)# ip pim version [1 2]	Configures the PIM version used.

Configuring PIM Version 2 Only

To configure PIM Version 2 exclusively, perform the tasks in this section. It is assumed that no PIM Version 1 system exists in the PIM domain.

The first task is recommended, configuring sparse-dense mode. If you configure Auto-RP, none of the other tasks is required to run PIM Version 2. To configure Auto-RP, see the section “[Configuring Auto-RP](#)” earlier in this chapter.

If you want to configure a BSR, complete the tasks in the following sections:

- [Configuring PIM Sparse-Dense Mode, page 15-21](#)
- [Defining the PIM Domain Border, page 15-21](#)
- [Defining the IP Multicast Boundary, page 15-21](#)
- [Configuring Candidate BSRs, page 15-22](#)
- [Configuring Candidate RPs, page 15-22](#)

Configuring PIM Sparse-Dense Mode

To configure PIM sparse-dense mode, enter the following commands on all PIM routers inside the PIM domain, beginning in global configuration mode:

	Command	Purpose
Step 1	Switch(config)# ip multicast-routing	Enables IP multicast routing.
Step 2	Switch(interface)# interface type number	Configures an interface.
Step 3	Switch(config)# ip pim sparse-dense-mode	Enables PIM on the interface. The sparse-dense mode is identical to the implicit interface mode in the PIM Version 2 specification.

Repeat Steps 2 and 3 above for each interface on which you want to run PIM.

Defining the PIM Domain Border

Configure a border for the PIM domain, so that bootstrap messages do not cross this border in either direction. Therefore, different BSRs will be elected on the two sides of the PIM border. Use the following command on the interface of a border router peering with one or more neighbors outside the PIM domain. To configure a PIM domain boundary, enter the following command in interface configuration mode:

Command	Purpose
Switch(config)# ip pim border	Configures a PIM domain boundary.

Defining the IP Multicast Boundary

Defining a IP Multicast Boundary is not supported on this release of the Catalyst 4006 switch with Supervisor Engine III.

Configuring Candidate BSRs

You should configure one or more candidate BSRs. The routers that serve as candidate BSRs should be well connected and be in the backbone portion of the network, as opposed to the dialup portion of the network. On the candidate BSRs, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim bsr-candidate hash-mask-length [priority]	Configure the router to be a candidate bootstrap router.

Configuring Candidate RPs

Configure one or more candidate RPs. Similar to BSRs, the RPs should also be well connected and in the backbone portion of the network. An RP can serve the entire IP multicast address space or a portion of it. Candidate RPs send candidate RP advertisements to the BSR. Consider the following when deciding which routers should be RPs:

- In a network of Cisco routers where only Auto-RP is used, any router can be configured as an RP.
- In a network of routers that includes only Cisco PIM Version 2 routers and routers from other vendors, any router can be used as an RP.
- In a network of Cisco PIM Version 1 routers, Cisco PIM Version 2 routers, and routers from other vendors, only Cisco PIM Version 2 routers should be configured as RPs.

On the candidate RPs, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim rp-candidate type number ttl group-list access-list-number	Configures the router to be a candidate RP.

For examples of configuring PIM Version 2, see the section “[BSR Configuration Example](#)” at the end of this chapter.

Using Auto-RP and a BSR

If you must have one or more BSRs, as described in the prior section “[Deciding When to Configure a BSR](#),” we recommend the following:

- Configure the candidate BSRs as the RP-mapping agents for Auto-RP.
- For group prefixes advertised via Auto-RP, the Version 2 BSR mechanism should not advertise a subrange of these group prefixes served by a different set of RPs. In a mixed Version 1/Version 2 PIM domain, it is preferable to have backup RPs serve the same group prefixes. This prevents the Version 2 DRs from selecting a different RP from those Version 1 DRs, due to longest match lookup in the RP-mapping database.

To verify the consistency of group-to-RP mappings, perform the following tasks in EXEC mode:

	Task	Purpose
Step 1	Switch# show ip pim rp [[<i>group-name</i> <i>group-address</i>] mapping]	Displays the available RP mappings on any router.
Step 2	Switch# show ip pim rp-hash <i>group</i>	Confirms that the same RP appears that a PIM Version 1 system chooses on a PIM Version 2 router.

Monitoring the RP Mapping Information

To monitor the RP mapping information, you can enter the following commands in EXEC mode:

Command	Purpose
Switch# show ip pim bsr	Displays information about the currently elected BSR.
Switch# show ip pim rp-hash <i>group</i>	Displays the RP that was selected for the specified group.
Switch# show ip pim rp [<i>group-name</i> <i>group-address</i> mapping]	Displays how the router learns of the RP (via bootstrap or Auto-RP mechanism).

Troubleshooting

When debugging interoperability problems between PIM Version 1 and Version 2, perform the following tasks:

- | | |
|--------|--|
| Step 1 | Verify RP mapping with the show ip pim rp-hash command, making sure that all systems agree on the same RP for the same group. |
| Step 2 | Verify interoperability between different versions of DRs and RPs. Ensure the RPs are interacting with the DRs properly (by responding with register-stops and forwarding decapsulated data packets from registers). |

Configuring Advanced PIM Features

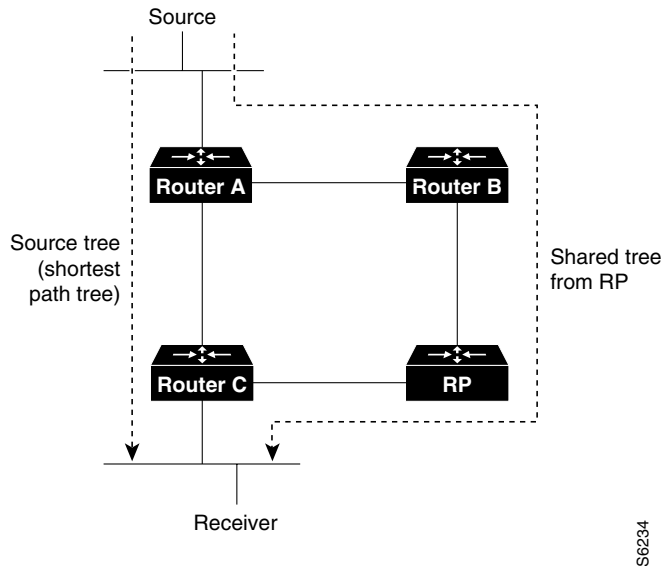
Perform the optional tasks in the following sections to configure PIM features:

- [Understanding PIM Shared Tree and Source Tree \(Shortest Path Tree\)](#), page 15-24
- [Delaying the Use of PIM Shortest Path Tree](#), page 15-25
- [Understanding Reverse-Path Forwarding](#), page 15-25
- [Assigning an RP to Multicast Groups](#), page 15-26
- [Increasing Control over RPs](#), page 15-26
- [Modifying the PIM Router-Query Message Interval](#), page 15-26
- [Enabling PIM Nonbroadcast Multiaccess Mode](#), page 15-27

Understanding PIM Shared Tree and Source Tree (Shortest Path Tree)

By default, members of a group receive data from senders to the group across a single data distribution tree rooted at the RP. This type of distribution tree is called *shared tree*, and is shown in Figure 15-7. Data from senders is delivered to the RP for distribution to group members joined to the shared tree.

Figure 15-7 Shared Tree and Source Tree (Shortest Path Tree)



If the data rate warrants, leaf routers on the shared tree can initiate a switch to the data distribution tree rooted at the source. This type of distribution tree is called a *shortest path tree* or *source tree*. By default, the Cisco IOS software changes to a source tree configuration upon receiving the first data packet from a source.

The following process describes the move from shared tree to source tree:

1. Receiver joins a group; leaf Router C sends a join message toward the RP.
2. RP puts link to Router C in its outgoing interface list.
3. Source sends data; Router A encapsulates data in Register and sends it to the RP.
4. RP forwards data down the shared tree to Router C and sends a join message toward Source. At this point, data may arrive twice at Router C, once encapsulated and once natively.
5. When data arrives natively (unencapsulated) at RP, RP sends a Register-Stop message to Router A.
6. By default, reception of the first data packet prompts Router C to send a join message toward Source.
7. When Router C receives data on (S,G), it sends a prune message for Source up the shared tree.
8. RP deletes the link to Router C from outgoing interface of (S,G). RP triggers a prune message toward Source.

Join and prune messages are sent for sources and RPs. They are sent hop-by-hop and are processed by each PIM router along the path to the source or RP. Register and Register-Stop messages are not sent hop-by-hop. They are sent by the designated router that is directly connected to a source and are received by the RP for the group.

Multiple sources sending to groups used the shared tree.

The network manager can configure the router to stay on the shared tree, as described in the following section, “[Delaying the Use of PIM Shortest Path Tree](#).”

Delaying the Use of PIM Shortest Path Tree

The change from shared to source tree happens upon the arrival of the first data packet at the last hop router (Router C in [Figure 15-7](#)). This change occurs because the `ip pim spt-threshold` command controls that timing, and its default setting is 0 Kbps.

The shortest path tree requires more memory than the shared tree, but reduces delay. You might want to postpone its use. Instead of allowing the leaf router to move to the shortest path tree immediately, you can specify that the traffic must first reach a threshold.

You can configure when a PIM leaf router should join the shortest path tree for a specified group. If a source sends at a rate greater than or equal to the specified *kbps* rate, the router triggers a PIM join message toward the source to construct a source tree (shortest path tree). If **infinity** is specified, all sources for the specified group use the shared tree, never switching to the source tree.

The group list is a standard access list that controls what groups the shortest path tree threshold applies to. If a value of 0 is specified or the group list is not used, the threshold applies to all groups.

To configure a traffic rate threshold that must be reached before multicast routing is switched from the source tree to the shortest path tree, enter the following command in interface configuration mode:

Command	Purpose
Switch(config)# <code>ip pim spt-threshold {kbps infinity} [group-list access-list-number]</code>	Specifies the threshold that must be reached before moving to shortest path tree (spt).

Understanding Reverse-Path Forwarding

Reverse-Path Forwarding (RPF) is an algorithm used for forwarding multicast datagrams. It functions as follows:

- If a router receives a datagram on an interface it uses to send unicast packets to the source, the packet has arrived on the RPF interface.
- If the packet arrives on the RPF interface, a router forwards the packet out the interfaces present in the outgoing interface list of a multicast routing table entry.
- If the packet does not arrive on the RPF interface, the packet is silently discarded to prevent loops.

PIM uses both source trees and RP-rooted shared trees to forward datagrams; the RPF check is performed differently for each, as follows:

- If a PIM router has source-tree state (that is, an (S,G) entry is present in the multicast routing table), the router performs the RPF check against the IP address of the source of the multicast packet.
- If a PIM router has shared-tree state (and no explicit source-tree state), it performs the RPF check on the RP's address (which is known when members join the group).

Sparse-mode PIM uses the RPF lookup function to determine where it needs to send joins and prunes. (S,G) joins (which are source-tree states) are sent toward the source. (*,G) joins (which are shared-tree states) are sent toward the RP.

DVMRP and dense-mode PIM use only source trees and use RPF as previously described.

Assigning an RP to Multicast Groups

If you have configured PIM sparse mode, you must configure a PIM RP for a multicast group. An RP can either be configured statically in each box, or learned through a dynamic mechanism. This section explains how to statically configure an RP. If the RP for a group is learned through a dynamic mechanism (such as Auto-RP), you need not perform this task for that RP. You should use Auto-RP, which is described in the section “[Configuring Auto-RP](#)” earlier in this chapter.

PIM designated routers forward data from directly connected multicast sources to the RP for distribution down the shared tree.

Data is forwarded to the RP in one of two ways. It is encapsulated in Register packets and unicast directly to the RP, or, if the RP has itself joined the source tree, it is multicast forwarded per the RPF forwarding algorithm described in the preceding section, “[Understanding Reverse-Path Forwarding](#).” Last-hop routers directly connected to receivers can join themselves to the source tree and prune themselves from the shared tree.

A single RP can be configured for multiple groups defined by an access list. If there is no RP configured for a group, the router treats the group as dense using the dense-mode PIM techniques.

If a conflict exists between the RP configured with this command and one learned by Auto-RP, the Auto-RP information is used, unless the **override** keyword is configured.

To assign an RP to one or more multicast groups, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim rp-address <i>ip-address</i> [<i>group-access-list-number</i>] [override]	Assigns an RP to multicast groups.

Increasing Control over RPs

You can take a defensive measure to prevent a misconfigured leaf router from interrupting PIM service to the remainder of a network. To do so, configure the local router to accept join messages only if they contain the RP address specified, when the group is in the group range specified by the access list. To configure this feature, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip pim accept-rp { <i>address</i> auto-rp } [<i>access-list-number</i>]	Controls which RPs the local router will accept join messages to.

Modifying the PIM Router-Query Message Interval

Router-query messages are used to elect a PIM designated router. The designated router is responsible for sending IGMP host-query messages. By default, multicast routers send PIM router-query messages every 30 seconds. To modify this interval, enter the following command in interface configuration mode:

Command	Purpose
Switch(interface)# ip pim query-interval <i>seconds</i>	Configures the frequency at which multicast routers send PIM router-query messages.

Enabling PIM Nonbroadcast Multiaccess Mode


Enabling PIM Nonbroadcast Multiaccess Mode is not supported on this release of the Catalyst 4006 switch with Supervisor Engine III.

Configuring an IP Multicast Static Route

IP multicast static routes (mroutes) allow you to have multicast paths diverge from the unicast paths. When using PIM, the router expects to receive packets on the same interface where it sends unicast packets back to the source. This expectation is beneficial if your multicast and unicast topologies are congruent. However, you might want unicast packets to take one path and multicast packets to take another.

The most common reason for using separate unicast and multicast paths is tunneling. When a path between a source and a destination does not support multicast routing, a solution is to configure two routers with a GRE tunnel between them. In [Figure 15-8](#), each unicast router (UR) supports unicast packets only; each multicast router (MR) supports multicast packets.


Figure 15-8 Tunnel for Multicast Packets



As shown in [Figure 15-8](#), Source delivers multicast packets to Destination by using MR 1 and MR 2. MR 2 accepts the multicast packet only if it thinks it can reach Source over the tunnel. If this is true, when Destination sends unicast packets to Source, MR 2 sends them over the tunnel. This could be slower than natively sending the unicast packet through UR 2, UR 1, and MR 1.

Prior to multicast static routes, the configuration shown in [Figure 15-9](#) was used to overcome the problem of both unicasts and multicasts using the tunnel. In this figure, MR 1 and MR 2 are used as multicast routers only. When Destination sends unicast packets to Source, it uses the (UR 3,UR 2,UR 1) path. When Destination sends multicast packets, the UR routers do not understand or forward them. However, the MR routers forward the packets.

Figure 15-9 Separate Paths for Unicast and Multicast Packets



To make the configuration shown in [Figure 15-9](#) work, MR 1 and MR 2 must run another routing protocol (typically a different instantiation of the same protocol running in the UR routers), so that paths from sources are learned dynamically.

A multicast static route allows you to use the configuration shown in [Figure 15-8](#) by configuring a static multicast source. The Cisco IOS software uses the configuration information instead of the unicast routing table. Therefore, multicast packets can use the tunnel without having unicast packets use the tunnel. Static mroutes are local to the router they are configured on and not advertised or redistributed in any way to any other router.

To configure a multicast static route, enter the following command in global configuration mode:

Command	Purpose
Switch(config)# ip mroute <i>source mask</i> [<i>protocol as-number</i>] [<i>rpf-address</i> <i>type number</i>] [<i>distance</i>]	Configures an IP multicast static route.

Monitoring and Maintaining IP Multicast Routing

You can remove all contents of a particular cache, table, or database. You also can display specific statistics. The following sections describe how to monitor and maintain IP multicast:

- [Displaying System and Network Statistics, page 15-28](#)
- [Displaying the Multicast Routing Table, page 15-29](#)
- [Displaying IP MFIB, page 15-31](#)
- [Displaying IP MFIB Fast Drop, page 15-32](#)
- [Displaying PIM Statistics, page 15-32](#)
- [Clearing Tables and Databases, page 15-33](#)

Displaying System and Network Statistics

You can display specific statistics, such as the contents of IP routing tables and databases. Information provided can be used to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path your device's packets are taking through the network.

To display various routing statistics, you can enter any of the following commands in EXEC mode:

Command	Purpose
Switch# ping [<i>group-name</i> <i>group-address</i>]	Sends an ICMP Echo Request to a multicast group address.
Switch# show ip mroute [<i>hostname</i> <i>group_number</i>]	Displays the contents of the IP multicast routing table.
Switch# show ip pim interface [<i>type number</i>] [<i>count</i>]	Displays information about interfaces configured for PIM.
Switch# show ip interface	Displays PIM information for all interfaces.

Displaying the Multicast Routing Table

The following is sample output from the **show ip mroute** command for a router operating in dense mode. This command displays the contents of the IP multicast FIB table for the multicast group named *cbone-audio*.

```
Switch# show ip mroute cbone-audio

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.1), uptime 0:57:31, expires 0:02:59, RP is 0.0.0.0, flags: DC
  Incoming interface: Null, RPF neighbor 0.0.0.0, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Dense, 0:57:31/0:02:52
    Tunnel0, Forward/Dense, 0:56:55/0:01:28

(198.92.37.100/32, 224.0.255.1), uptime 20:20:00, expires 0:02:55, flags: C
  Incoming interface: Tunnel0, RPF neighbor 10.20.37.33, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Dense, 20:20:00/0:02:52
```

The following is sample output from the **show ip mroute** command for a router operating in sparse mode:

```
Switch# show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
  Incoming interface: Tunnel0, RPF neighbor 10.3.35.1, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.0/24, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
  Incoming interface: Tunnel0, RPF neighbor 10.3.35.1
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57
```



Note

Output interface timers are not updated for hardware-forwarded packets. Entry timers are updated approximately every five seconds.

The following is sample output from the **show ip mroute** command with the **summary** keyword:

```
Switch# show ip mroute summary

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
       R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.255.255.255), 2d16h/00:02:30, RP 171.69.10.13, flags: SJPC

(*, 224.2.127.253), 00:58:18/00:02:00, RP 171.69.10.13, flags: SJC
```

```
(*, 224.1.127.255), 00:58:21/00:02:03, RP 171.69.10.13, flags: SJC

(*, 224.2.127.254), 2d16h/00:00:00, RP 171.69.10.13, flags: SJCL
(128.9.160.67/32, 224.2.127.254), 00:02:46/00:00:12, flags: CLJT
(129.48.244.217/32, 224.2.127.254), 00:02:15/00:00:40, flags: CLJT
(130.207.8.33/32, 224.2.127.254), 00:00:25/00:02:32, flags: CLJT
(131.243.2.62/32, 224.2.127.254), 00:00:51/00:02:03, flags: CLJT
(140.173.8.3/32, 224.2.127.254), 00:00:26/00:02:33, flags: CLJT
(171.69.60.189/32, 224.2.127.254), 00:03:47/00:00:46, flags: CLJT
```

The following is sample output from the **show ip mroute** command with the **active** keyword:

```
Switch# show ip mroute active

Active IP Multicast Sources - sending >= 4 kbps

Group: 224.2.127.254, (sdr.cisco.com)
  Source: 146.137.28.69 (mbone.ipd.anl.gov)
  Rate: 1 pps/4 kbps(1sec), 4 kbps(last 1 secs), 4 kbps(life avg)

Group: 224.2.201.241, ACM 97
  Source: 130.129.52.160 (webcast3-e1.acm97.interop.net)
  Rate: 9 pps/93 kbps(1sec), 145 kbps(last 20 secs), 85 kbps(life avg)

Group: 224.2.207.215, ACM 97
  Source: 130.129.52.160 (webcast3-e1.acm97.interop.net)
  Rate: 3 pps/31 kbps(1sec), 63 kbps(last 19 secs), 65 kbps(life avg)
```

The following is sample output from the **show ip mroute** command with the **count** keyword:

```
Switch# show ip mroute count

IP Multicast Statistics - Group count: 8, Average sources per group: 9.87
Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second

Group: 224.255.255.255, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.2.127.253, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.1.127.255, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.2.127.254, Source count: 9, Group pkt count: 14
  RP-tree: 0/0/0/0
  Source: 128.2.6.9/32, 2/0/796/0
  Source: 128.32.131.87/32, 1/0/616/0
  Source: 128.125.51.58/32, 1/0/412/0
  Source: 130.207.8.33/32, 1/0/936/0
  Source: 131.243.2.62/32, 1/0/750/0
  Source: 140.173.8.3/32, 1/0/660/0
  Source: 146.137.28.69/32, 1/0/584/0
  Source: 171.69.60.189/32, 4/0/447/0
  Source: 204.162.119.8/32, 2/0/834/0

Group: 224.0.1.40, Source count: 1, Group pkt count: 3606
  RP-tree: 0/0/0/0
  Source: 171.69.214.50/32, 3606/0/48/0, RPF Failed: 1203

Group: 224.2.201.241, Source count: 36, Group pkt count: 54152
  RP-tree: 7/0/108/0
  Source: 13.242.36.83/32, 99/0/123/0
```

```

Source: 36.29.1.3/32, 71/0/110/0
Source: 128.9.160.96/32, 505/1/106/0
Source: 128.32.163.170/32, 661/1/88/0
Source: 128.115.31.26/32, 192/0/118/0
Source: 128.146.111.45/32, 500/0/87/0
Source: 128.183.33.134/32, 248/0/119/0
Source: 128.195.7.62/32, 527/0/118/0
Source: 128.223.32.25/32, 554/0/105/0
Source: 128.223.32.151/32, 551/1/125/0
Source: 128.223.156.117/32, 535/1/114/0
Source: 128.223.225.21/32, 582/0/114/0
Source: 129.89.142.50/32, 78/0/127/0
Source: 129.99.50.14/32, 526/0/118/0
Source: 130.129.0.13/32, 522/0/95/0
Source: 130.129.52.160/32, 40839/16/920/161
Source: 130.129.52.161/32, 476/0/97/0
Source: 130.221.224.10/32, 456/0/113/0
Source: 132.146.32.108/32, 9/1/112/0

```

**Note**

Multicast route byte and packet statistics are supported only for the first 1024 multicast routes. Output interface statistics are not maintained.

Displaying IP MFIB

You can display all routes in the MFIB, including routes that might not exist directly in the upper-layer routing protocol database, but that are used to accelerate fast switching. These routes appear in the MFIB even if dense-mode forwarding is in use.

To display various MFIB routing routes, you can enter the following commands in EXEC mode:

Command	Purpose
Switch# show ip mfib	Displays the (S,G) and (*,G) routes that are used for packet forwarding. Displays counts for fast, slow, and partially-switched packets for every multicast route.
Switch# show ip mfib all	Displays all routes in the MFIB, including routes that may not exist directly in the upper-layer routing protocol database, but that are used to accelerate fast switching. These routes include the (S/M,224/4) routes.
Switch# show ip mfib log [n]	Displays a log of the most recent n MFIB related events, most recent first.
Switch# show ip mfib counters	Displays counts of MFIB related events. Only non-zero counters are shown.

The following is sample output from the **show ip mfib** command.

```

IP Multicast Forwarding Information Base
Entry Flags: C - Directly Connected, S - Signal,
             IC - Internal Copy
Interface Flags: A - Accept, F - Forward, S - Signal,
                NP - Not platform switched
Packets: Fast/Partial/Slow Bytes: Fast/Partial/Slow:
(171.69.10.13, 224.0.1.40), flags (IC)

```

```

Packets: 2292/2292/0, Bytes: 518803/0/518803
Vlan7 (A)
Vlan100 (F NS)
Vlan105 (F NS)
(*, 224.0.1.60), flags ()
  Packets: 2292/0/0, Bytes: 518803/0/0
  Vlan7 (A NS)
(*, 224.0.1.75), flags ()
  Vlan7 (A NS)
(10.34.2.92, 239.192.128.80), flags ()
  Packets: 24579/100/0, 2113788/15000/0 bytes
  Vlan7 (F NS)
  Vlan100 (A)
(*, 239.193.100.70), flags ()
  Packets: 1/0/0, 1500/0/0 bytes
  Vlan7 (A)
..

```

The fast-switched packet count represents the number of packets that were switched in hardware on the corresponding route.

The partially-switched packet counter represents the number of times that a fast-switched packet was also copied to the CPU for software processing or for forwarding to one or more non-platform switched interfaces (such as a PimTunnel interface).

The slow-switched packet count represents the number of packets that were switched completely in software on the corresponding route.

Displaying IP MFIB Fast Drop

To display fast-drop entries, enter the following command in EXEC mode:

Command	Purpose
Switch# show ip mfib fastdrop	Displays all currently-active fast-drop entries and whether or not fastdrop is enabled.

The following is sample output from the **show ip mfib fastdrop** command.

```

Switch> show ip mfib fasttdrop
MFIB fastdrop is enabled.
MFIB fast-dropped flows:
(10.0.0.1, 224.1.2.3, Vlan9 ) 00:01:32
(10.1.0.2, 224.1.2.3, Vlan9 ) 00:02:30
(1.2.3.4, 225.6.7.8, Vlan3) 00:01:50

```

The full (S,G) flow and the ingress interface on which incoming packets are dropped is shown. The timestamp indicates the age of the entry.

Displaying PIM Statistics

The following is sample output from the **show ip pim interface** command:

```

Switch# show ip pim interface

Address          Interface      Mode   Neighbor  Query   DR
                Count         Interval
198.92.37.6     Ethernet0     Dense  2         30     198.92.37.33
198.92.36.129   Ethernet1     Dense  2         30     198.92.36.131
10.1.37.2       Tunnel0       Dense  1         30     0.0.0.0

```

The following is sample output from the **show ip pim interface** command with a **count**:

```
Switch# show ip pim interface count

Address          Interface      FS  Mpackets In/Out
171.69.121.35    Ethernet0     *   548305239/13744856
171.69.121.35    Serial0.33    *   8256/67052912
198.92.12.73     Serial0.1719  *   219444/862191
```

The following is sample output from the **show ip pim interface** command with a **count** when IP multicast is enabled. The examples lists the PIM interfaces that are fast switched and process switched, and the packet counts for these. The H is added to interfaces where IP multicast is enabled.

```
Switch# show ip pim interface count

States: FS - Fast Switched, H - Hardware Switched
Address          Interface      FS  Mpackets In/Out
192.1.10.2       Vlan10        * H 40886/0
192.1.11.2       Vlan11        * H 0/40554
192.1.12.2       Vlan12        * H 0/40554
192.1.23.2       Vlan23        *   0/0
192.1.24.2       Vlan24        *   0/0
```

Clearing Tables and Databases

You can remove all contents of a particular cache, table, or database. Clearing a cache, table, or database might be necessary when the contents of the particular structure have become, or are suspected to be, invalid.

To clear IP multicast caches, tables, and databases, enter the following commands in EXEC mode:

Command	Purpose
Switch# clear ip mroute	Deletes entries from the IP routing table.
Switch# clear ip mfib counters	Deletes all per-route and global MFIB counters.
Switch# clear ip mfib fastdrop	Deletes all clear all fast-drop entries.



Note

IP multicast routes can be regenerated in response to protocol events and as data packets arrive.

Configuration Examples

The following sections provide IP multicast routing configuration examples:

- [PIM Dense Mode Example, page 15-34](#)
- [PIM Sparse Mode Example, page 15-34](#)
- [BSR Configuration Example, page 15-34](#)

PIM Dense Mode Example

This example is a configuration of dense-mode PIM on an Ethernet interface:

```
ip multicast-routing
interface ethernet 0
 ip pim dense-mode
```

PIM Sparse Mode Example

This example is a configuration of sparse-mode PIM. The RP router is the router with the address 10.8.0.20.

```
ip multicast-routing
 ip pim rp-address 10.8.0.20 1
interface ethernet 1
 ip pim sparse-mode
```

BSR Configuration Example

This example is a configuration of a candidate BSR, which also happens to be a candidate RP:

```
version 11.3
!
ip multicast-routing
!
interface Ethernet0
 ip address 171.69.62.35 255.255.255.240
!
interface Ethernet1
 ip address 172.21.24.18 255.255.255.248
 ip pim sparse-dense-mode
!
interface Ethernet2
 ip address 172.21.24.12 255.255.255.248
 ip pim sparse-dense-mode
!
router ospf 1
 network 172.21.24.8 0.0.0.7 area 1
 network 172.21.24.16 0.0.0.7 area 1
!
ip pim bsr-candidate Ethernet2 30 10
ip pim rp-candidate Ethernet2 group-list 5
access-list 5 permit 239.255.2.0 0.0.0.255
```