



CHAPTER 33

Configuring IP Multicast

This chapter describes IP multicast routing on the Catalyst 4500 series switch. It also provides procedures and examples to configure IP multicast routing.



Note

For more detailed information on IP Multicast, refer to this URL:

http://www.cisco.com/en/US/products/ps6552/products_ios_technology_home.html.



Note

For complete syntax and usage information for the switch commands used in this chapter, look at the *Cisco Catalyst 4500 Series Switch Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html>

If the command is not found in the Catalyst 4500 Command Reference, it is located in the larger Cisco IOS library. Refer to the *Cisco IOS Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/ps6350/index.html>

This chapter includes the following major sections:

- [About IP Multicast, page 33-1](#)
- [Configuring IP Multicast Routing, page 33-13](#)
- [Monitoring and Maintaining IP Multicast Routing, page 33-23](#)
- [Configuration Examples, page 33-28](#)

About IP Multicast



Note

Controlling the transmission rate to a multicast group is not supported.

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 IP unicast, 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 on the path between source and 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 4500 series switch, a packet is replicated in the Integrated Switching Engine, forwarded to the appropriate output interfaces, and sent to each member of the multicast group.

We tend to think of IP multicasting and video conferencing as 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.

This section contains the following subsections:

- [IP Multicast Protocols, page 33-2](#)
- [IP Multicast Implementation on the Catalyst 4500 Series Switch, page 33-5](#)
- [Restrictions on Using Bidirectional PIM, page 33-13](#)


IP Multicast Protocols

The Catalyst 4500 series switch primarily uses these protocols to implement IP multicast routing:

- Internet Group Management Protocol (IGMP)
- Protocol Independent Multicast (PIM)
- IGMP snooping and Cisco Group Management Protocol

[Figure 33-1](#) shows where these protocols operate within the IP multicast environment.

Figure 33-1 IP Multicast Routing Protocols



94150

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 by using 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 an IGMP request to receive multicast group traffic.

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.



Note

Catalyst 4900M, Catalyst 4948E, Supervisor Engine 6-E, and Supervisor Engine 6L-E do not increment counters for (*, G) in PIM Dense Mode. (*, G) counters are incremented only when running Bidirectional PIM Mode.

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 active receivers exist on every subnet in the network.

For more detailed information on PIM Dense Mode, refer to this URL:

http://www.cisco.com/en/US/docs/ios-xml/ios/ipmulti_optim/configuration/12-2sx/imc_pim_dense_rfrsh.html.

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 are 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.

Bidirectional PIM Mode

In bidirectional PIM (Bidir-PIM) mode, traffic is routed only along a bidirectional shared tree that is rooted at the rendezvous point (RP) for the group. The IP address of the RP functions as a key enabling all routers to establish a loop-free spanning tree topology rooted in that IP address.

Bidir-PIM is intended for many-to-many applications within individual PIM domains. Multicast groups in bidirectional mode can scale to an arbitrary number of sources without incurring overhead due to the number of sources.

For more detailed information on Bidirectional Mode, refer to this URL:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/ps6592/prod_white_paper0900aecd80310db2.pdf.

Rendezvous Point (RP)

If you configure PIM to operate in sparse mode, you must also choose one or more routers to be rendezvous points (RPs). Senders to a multicast group use RPs to announce their presence. Receivers of multicast packets use RPs to learn about new senders. You can configure Cisco IOS software 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 for the same group. The conditions specified by the access list determine for which groups the router is an RP (as different groups can have different RPs).

IGMP Snooping

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 whether it contains any pertinent IGMP control information. If IGMP snooping is implemented on a low end switch with a slow CPU, performance could be severely impacted when data is transmitted at high rates. On the Catalyst 4500 series switches, IGMP snooping is implemented in the forwarding ASIC, so it does not impact the forwarding rate.


IP Multicast Implementation on the Catalyst 4500 Series Switch

The Catalyst 4500 series switch 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 switch ports 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 33-2 shows a logical view of Layer 2 and Layer 3 forwarding in the integrated switching engine hardware.

Figure 33-2 Logical View of Layer 2 and Layer 3 Forwarding in Hardware



This section contains the following subsections:

- [CEF, MFIB, and Layer 2 Forwarding, page 33-5](#)
- [IP Multicast Tables, page 33-7](#)
- [Hardware and Software Forwarding, page 33-9](#)
- [Non-Reverse Path Forwarding Traffic, page 33-10](#)
- [Multicast Fast Drop, page 33-11](#)
- [Multicast Forwarding Information Base, page 33-12](#)
- [S/M, 224/4, page 33-13](#)

CEF, MFIB, and Layer 2 Forwarding

The implementation of IP multicast on the Catalyst 4500 series switch is an extension of centralized Cisco Express Forwarding (CEF). CEF extracts information from the unicast routing table, which is created by unicast routing protocols, such as BGP, OSPF, and EIGR and 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).

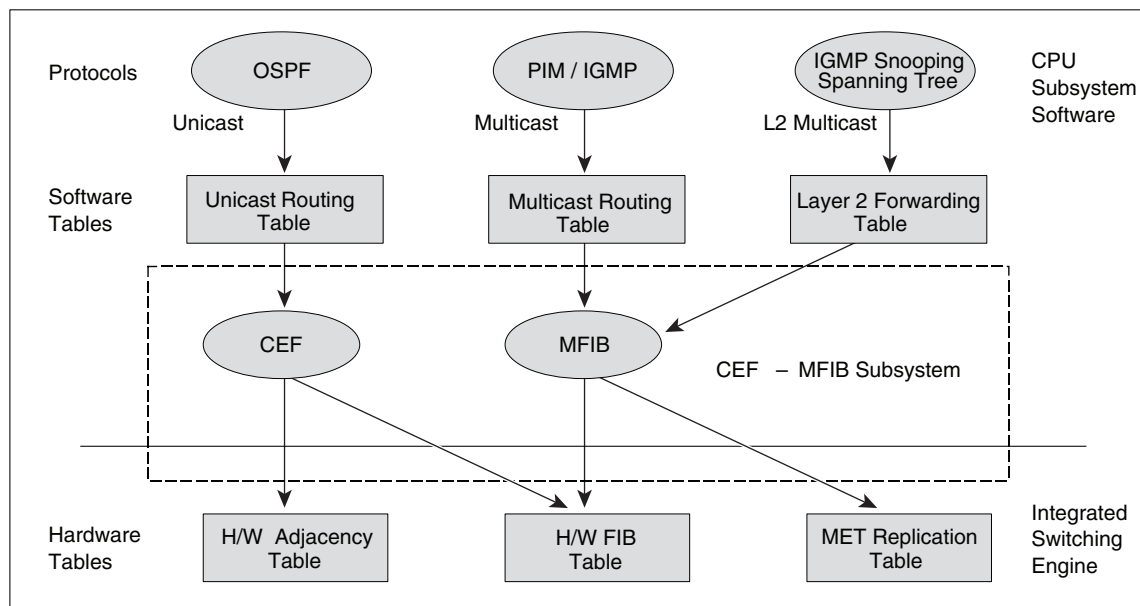
**Note**

On the Supervisor Engine 6-E and 6L-E, MET has been replaced by the RET (Replica Expansion Table).

The Catalyst 4500 series switch performs Layer 3 routing and Layer 2 bridging at the same time. There can be multiple Layer 2 switch ports on any VLAN interface. To determine the set of output switch ports on which to forward a multicast packet, the Supervisor Engine III combines Layer 3 MFIB information with Layer 2 forwarding information and stores it in the hardware MET for packet replication.

Figure 33-3 shows a functional overview of how the Catalyst 4500 series switch combines unicast routing, multicast routing, and Layer 2 bridging information to forward in hardware.

Figure 33-3 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 (*,224.1.2.3) 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 (*,224.1.2.3) route. With this information loaded in hardware, merging of the Layer 2 information can begin. For the output interfaces on VLAN1, the integrated switching engine must send the packet to all switch ports in VLAN1 that are in the spanning tree forwarding state. The same process applies to VLAN 2. To determine the set of switch ports in VLAN 2, the Layer 2 Forwarding Table is used.

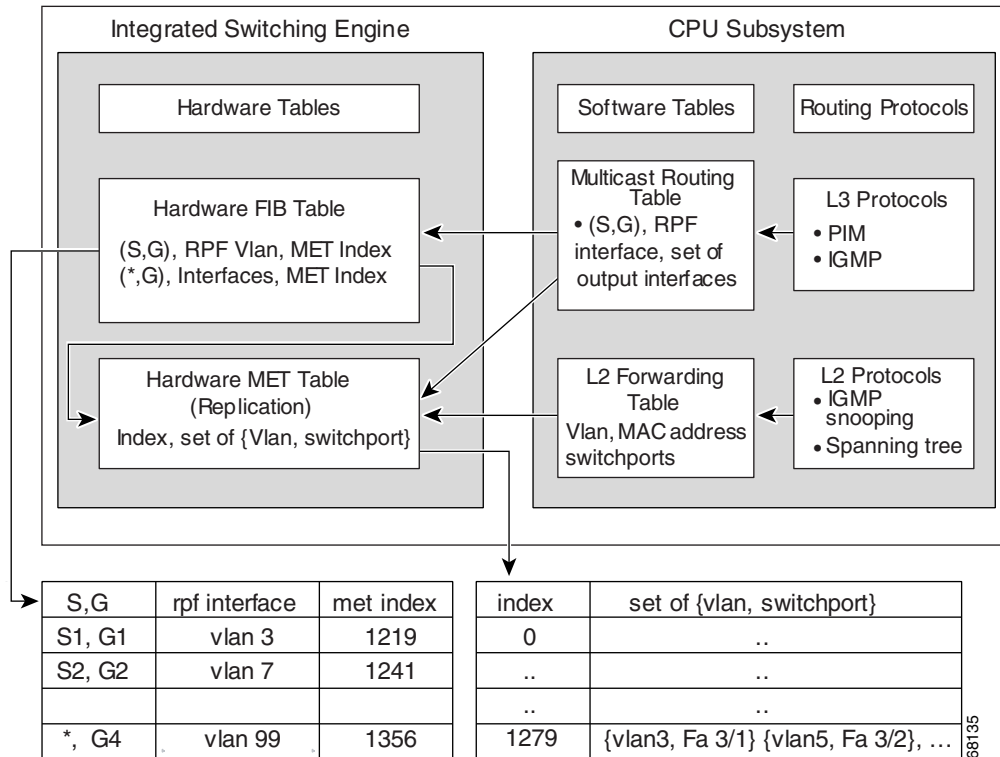
When the hardware routes a packet, in addition to sending it to all of the switch ports on all output interfaces, the hardware also sends the packet to all switch ports (other than the one it arrived on) in the input VLAN. For example, assume that VLAN 3 has two switch ports in it, Gig 3/1 and Gig 3/2. 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 switch ports in the ingress VLAN must be added to the port set 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 switch ports: (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 switch ports on VLAN 2. The packet should be forwarded only to switch ports 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 switch ports: (1/1,1/2, 2/1, 2/2, 3/1, and 3/2).

IP Multicast Tables

Figure 33-4 shows some key data structures that the Catalyst 4500 series switch uses to forward IP multicast packets in hardware.

Figure 33-4 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. (For RET, we can have up to 102 K entries (32 K used for floodsets, 70,000 used for multicast entries)). 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.

**Note**

For RET, a maximum of 102 K entries is supported (32 K used for floodsets, 70 K used for multicast entries).

**Note**


Partial routing is not supported on Catalyst 4900M, Catalyst 4948E, Supervisor Engine 6-E, and Supervisor Engine 6L-E; only hardware and software routing are supported.

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.

Figure 33-5 shows a logical view of the hardware and software forwarding components.

Figure 33-5 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 occurs only 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



Note

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

- 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. The switch must send PIM-register messages to the RP.

Software Routes



Note

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

The following conditions cause all replicas of a packet for a route to be forwarded by the CPU subsystem 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-Advanced Research Products Agency (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-Reverse Path Forwarding Traffic

Traffic that fails an Reverse Path Forwarding (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 33-6](#) shows how non-RPF traffic can occur in a common network configuration.

Figure 33-6 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.

To prevent this situation 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 you may have persistent RPF failures. Without the fast-drop entries, the CPU is 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 4500 series switch. 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.

The MFIB table contains a set of IP multicast routes. IP multicast routes include (S,G) and (*,G). 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—Sets on a route when a process on the router needs to receive a copy of all packets matching the specified route.
- Signalling (S) flag—Sets 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 Signaling (S) flag, except that the C flag indicates that only packets sent by directly connected hosts to the route should be signaled 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 handled, and they also indicate whether packets matching the route should be forwarded onto VLAN 1. The per-interface flags supported in the MFIB include the following:

- Accepting (A)—Sets 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.
- Signaling (S)—Sets on an interface when some multicast routing protocol process in Cisco IOS needs to be notified of packets arriving on that interface.



Note

When PIM-SM routing is in use, the MFIB route might include an interface as in this example:

```
PimTunnel [1.2.3.4].
```

it is a virtual interface that the MFIB subsystem creates to indicate that packets are being tunnelled 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 are 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 is 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.

Restrictions on Using Bidirectional PIM

Restrictions include the following:

- IPv4 Bidirectional (Bidir) PIM is supported on Catalyst 4900M, Catalyst 4948E, Supervisor Engine 6-E, and Supervisor Engine 6L-E. IPv6 Bidir PIM is not supported.
- The Catalyst 4500 series switch allows you to forward Bidir PIM traffic in hardware up to a maximum of seven RPs. If you configure more than seven Bidir RPs, only the first seven RPs can forward traffic in hardware. The traffic directed to the remaining RPs is forwarded in software.

Configuring IP Multicast Routing

The following sections describe IP multicast routing configuration tasks:

- [Default Configuration in IP Multicast Routing, page 33-13](#)
- [Enabling IP Multicast Routing, page 33-14](#)
- [Enabling PIM on an Interface, page 33-14](#)
- [Enabling Bidirectional Mode, page 33-16](#)
- [Enabling PIM-SSM Mapping, page 33-17](#)
- [Configuring a Rendezvous Point, page 33-17](#)
- [Configuring a Single Static RP, page 33-20](#)
- [Load Splitting of IP Multicast Traffic, page 33-22](#)

For more detailed information on IP multicast routing, such as Auto-RP, PIM Version 2, and IP multicast static routes, refer to the *Cisco IOS IP and IP Routing Configuration Guide, Cisco IOS Release 12.3*.

Default Configuration in IP Multicast Routing

[Table 33-1](#) shows the IP multicast default configuration.

Table 33-1 **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 4500 series switch. However, IP multicast control traffic continues to be processed and forwarded. 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 switch ports in the VLAN.

**Note**

Source-specific multicast and IGMP v3 are supported.

For more information about source-specific multicast with IGMPv3 and IGMP, see the following URL:
http://www.cisco.com/en/US/docs/ios/ipmulti/configuration/guide/imc_cfg_ssm.html

Enabling IP Multicast Routing

Enabling IP multicast routing allows the Catalyst 4500 series switch to forward multicast packets. To enable IP multicast routing on the router, enter this command:

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.

When the switch populates 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 this command:

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

For an example of how to configure a PIM interface in dense mode, see the “[PIM Dense Mode Example](#)” section.

Enabling Sparse Mode

To configure PIM on an interface to be in sparse mode, enter this command:

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

For an example of how to configure a PIM interface in sparse mode, see the “[PIM Sparse Mode Example](#)” section.

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. 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. You do not 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:

- When members or DVMRP neighbors exist on the interface
- When PIM neighbors exist 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:

- When members or DVMRP neighbors exist on the interface
- When 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 this command:

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

Enabling Bidirectional Mode

Most of the configuration requirements for Bidir-PIM are the same as those for configuring PIM-SM. You need not enable or disable an interface for carrying traffic for multicast groups in bidirectional mode. Instead, you configure which multicast groups you want to operate in bidirectional mode. Similar to PIM-SM, you can perform this configuration with Auto-RP, static RP configurations, or the PIM Version 2 bootstrap router (PIMv2 BSR) mechanism.

To enable Bidir-PIM, perform this task in global configuration mode:

Command	Purpose
Switch(config)# ip pim bidir-enable	Enables bidir-PIM on a switch.

To configure Bidir-PIM, enter one of these commands, depending on which method you use to distribute group-to-RP mappings:

Command	Purpose
Switch(config)# ip pim rp-address rp-address [access-list] [override] bidir	Configures the address of a PIM RP for a particular group, and specifies bidirectional mode. Use this command when you are not distributing group-to-RP mappings using either Auto-RP or the PIMv2 BSR mechanism
Switch(config)# ip pim rp-candidate type number [group-list access-list] bidir	Configures the router to advertise itself as a PIM Version 2 candidate RP to the BSR, and specifies bidirectional mode. Use this command when you are using the PIMv2 BSR mechanism to distribute group-to-RP mappings.
Switch(config)# ip pim send-rp-address type number scope ttl-value [group-list access-list] [interval seconds] bidir	Configures the router to use Auto-RP to configure the groups the router is willing to act as RP, and specifies bidirectional mode. Use this command when you are using Auto-RP to distribute group-to-RP mappings.

For an example of how to configure bidir-PIM, see the “[Bidirectional PIM Mode Example](#)” section on page 33-29.

Enabling PIM-SSM Mapping

The Catalyst 4500 series switch supports SSM mapping, enabling an SSM transition in cases either where neither URD nor IGMP v3-lite is available, or when supporting SSM on the end system is impossible or unwanted due to administrative or technical reasons. With SSM mapping, you can leverage SSM for video delivery to legacy set-top boxes (STBs) that do not support IGMPv3 or for applications that do not take advantage of the IGMPv3 host stack.

For more details, refer to this URL:

http://www.cisco.com/en/US/docs/ios/12_3t/12_3t2/feature/guide/gtssmma.html

Configuring a Rendezvous Point

A rendezvous point (RP) is required in networks running Protocol Independent Multicast sparse mode (PIM-SM). In PIM-SM, traffic is forwarded only to network segments with active receivers that have explicitly requested multicast data.

The most commonly used methods to configure a rendezvous point (described here) are the use of Static RP and the use of the Auto-RP protocol. Another method (not described here) is the use of the Bootstrap Router (BSR) protocol.

Configuring Auto-RP

Auto-rendezvous point (Auto-RP) automates the distribution of group-to-rendezvous point (RP) mappings in a PIM network. 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 by way of dense mode flooding.

All routers automatically discover which RP to use for the groups they support. The Internet Assigned Numbers Authority (IANA) has assigned two group addresses, 224.0.1.39 and 224.0.1.40, for Auto-RP.

The mapping agent receives announcements of intention to become the RP from Candidate-RPs. The mapping agent then announces the winner of the RP election. This announcement is made independently of the decisions by the other mapping agents.

To configure a rendezvous point, perform this task:

	Command or Action	Purpose
Step 1	Switch> enable	Enables privileged EXEC mode.
Step 2	Switch# configure terminal	Enters global configuration mode.
Step 3	Switch(config)# ip multicast-routing	Enables IP multicast routing.
Step 4	Switch(config)# interface [FastEthernet GigabitEthernet Loopback Null Port-channel TenGigabitEthernet Tunnel Vlan] <i>number</i>	Selects an interface that is connected to hosts on which PIM can be enabled.

	Command or Action	Purpose
Step 5	Switch(config-if)# ip pim [sparse-mode sparse-dense-mode]	Enables PIM sparse or sparse-dense mode on an interface. When configuring Auto-RP in sparse mode, you must also configure the Auto-RP listener in the next step.
Step 6	Switch(config-if)# exit	Returns to global configuration mode.
Step 7	Repeat Steps 4 and 5 on all PIM interfaces.	—
Step 8	Switch(config)# ip pim autorp listener	Causes IP multicast traffic for the two Auto-RP groups 224.0.1.39 and 224.0.1.40 to be PIM dense mode flooded across interfaces operating in PIM sparse mode. <ul style="list-style-type: none"> • Skip this step if you are configuring sparse-dense mode in Step 8.
Step 9	Switch(config)# ip pim send-rp-announce { <i>interface-type interface-number</i> <i>ip-address</i> } scope <i>tvl-value</i> [group-list <i>access-list</i>] [interval <i>seconds</i>] [bidir]	Sends RP announcements out all PIM-enabled interfaces. <ul style="list-style-type: none"> • Perform this step on the RP router only. • Use the <i>interface-type</i> and <i>interface-number</i> arguments to define which IP address is to be used as the RP address. • Use the <i>ip-address</i> argument to specify a directly connected IP address as the RP address. <p>Note If the <i>ip-address</i> argument is configured for this command, the RP-announce message is sourced by the interface to which this IP address is connected (that is, the source address in the IP header of the RP-announce message is the IP address of that interface).</p> <ul style="list-style-type: none"> • This example shows that the interface is enabled with a maximum of 31 hops. The IP address by which the router wants to be identified as RP is the IP address associated with loopback interface 0. Access list 5 describes the groups for which this router serves as RP.

	Command or Action	Purpose
Step 10	Switch(config)# ip pim send-rp-discovery [<i>interface-type interface-number</i>] scope <i>ttl-value [interval seconds]</i>	<p>Configures the router to be an RP mapping agent.</p> <ul style="list-style-type: none"> Perform this step on the RP router only. Use the optional <i>interface-type</i> and <i>interface-number</i> arguments to define which IP address is to be used as the source address of the RP mapping agent. Use the scope keyword and <i>ttl-value</i> argument to specify the Time-to-Live (TTL) value in the IP header of Auto-RP discovery messages. Use the optional interval keyword and <i>seconds</i> argument to specify the interval at which Auto-RP discovery messages are sent. <p>Note Lowering the interval at which Auto-RP discovery messages are sent from the default value of 60 seconds results in more frequent floodings of the group-to-RP mappings. In some network environments, the disadvantages of lowering the interval (more control packet overhead) may outweigh the advantages (more frequent Requirementsgroup-to-RP mapping updates).</p> <ul style="list-style-type: none"> The example shows limiting the Auto-RP discovery messages to 31 hops on loopback interface 1.
Step 11	Switch(config)# ip pim rp-announce-filter rp-list access-list group-list access-list	<p>Filters incoming Auto-RP announcement messages coming from the RP.</p> <ul style="list-style-type: none"> Perform this step on the RP router only. Two example access lists that apply to this step could be: <pre>access-list 1 permit 10.0.0.1 access-list 1 permit 10.0.0.2 access-list 2 permit 224.0.0.0 15.255.255.255</pre>
Step 12	Switch(config)# interface <i>type number</i>	Selects an interface that is connected to hosts on which PIM can be enabled.
Step 13	Switch(config-if)# interface ethernet 1 ip multicast boundary access-list [filter-autorp]	<p>Configures an administratively scoped boundary.</p> <ul style="list-style-type: none"> Perform this step on the interfaces that are boundaries to other routers. The access list is not shown in this task. An access list entry that uses the deny keyword creates a multicast boundary for packets that match that entry.
Step 14	Switch(config-if)# end	Returns to EXEC mode.
Step 15	Switch# show ip pim autorp	(Optional) Displays the Auto-RP information.
Step 16	Switch# show ip pim rp [mapping] [rp-address]	(Optional) Displays RPs known in the network and shows how the router learned about each RP.

	Command or Action	Purpose
Step 17	Switch# show ip igmp groups [<i>group-name</i> <i>group-address</i> <i>interface-type</i> <i>interface-number</i>] [detail]	(Optional) Displays the multicast groups having receivers that are directly connected to the router and that were learned through Internet Group Management Protocol (IGMP). <ul style="list-style-type: none"> A receiver must be active on the network at the time that this command is issued to ensure the presence of receiver information on the resulting display.
Step 18	Switch# show ip mroute [<i>group-address</i> <i>group-name</i>] [<i>source-address</i> <i>source-name</i>] [<i>interface-type</i> <i>interface-number</i>] [summary] [count] [active <i>kbps</i>]	(Optional) Displays the contents of the IP multicast routing (mroute) table.

This example illustrates how to configure Auto-RP:

```
Switch> enable
Switch# configure terminal
Switch(config)# ip multicast-routing
Switch(config)# interface ethernet 1
Switch(config-if)# ip pim sparse-mode
Switch(config-if)# end
Switch(config)# ip pim autorp listener
Switch(config)# ip pim send-rp-announce loopback0 scope 31 group-list 5
Switch(config)# ip pim send-rp-discovery loopback 1 scope 31
Switch(config)# ip pim rp-announce-filter rp-list 1 group-list 2
Switch(config)# interface ethernet 1
Switch(config-if)# ip multicast boundary 10 filter-autorp
Switch(config-if)# end
Switch# show ip pim autorp
Switch# show ip pim rp mapping
Switch# show ip igmp groups
Switch# show ip mroute cbone-audio
```

Configuring a Single Static RP

If you are configuring PIM sparse mode, you must configure a PIM RP for a multicast group. An RP can either be configured statically in each device, or learned through a dynamic mechanism. This task explains how to statically configure an RP, as opposed to the router learning the RP through a dynamic mechanism such as Auto-RP.

PIM designated routers (DRs) 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. Last hop routers directly connected to receivers may, at their discretion, join themselves to the source tree and prune themselves from the shared tree.

A single RP can be configured for multiple groups that are defined by an access list. If no RP is configured for a group, the router treats the group as dense using the PIM dense mode techniques. (You can prevent this occurrence by configuring the **no ip pim dm-fallback** command.)

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

To configure a single static RP, perform this task:

	Command or Action	Purpose
Step 1	Switch> enable	Enables privileged EXEC mode.
Step 2	Switch# configure terminal	Enters global configuration mode.
Step 3	Switch(config)# ip multicast-routing	Enables IP multicast routing.
Step 4	Switch(config)# interface <i>type number</i>	Selects an interface that is connected to hosts on which PIM can be enabled.
Step 5	Switch(config-if)# ip pim [sparse-mode sparse-dense-mode]	Enables PIM on an interface. You must use sparse mode.
Step 6	Repeat Steps 4 and 5 on every interface that uses IP multicast.	—
Step 7	Switch(config-if)# exit	Returns to global configuration mode.
Step 8	Switch(config)# ip pim rp-address <i>rp-address</i> [<i>access-list</i>] [override]	Configures the address of a PIM RP for a particular group. <ul style="list-style-type: none"> Perform this step on any router. The <i>access-list</i> argument specifies the number or name of an access list that defines for which multicast groups the RP should be used. The override keyword specifies that if there is a conflict between the RP configured with this command and one learned by Auto-RP, the RP configured with this command prevails.
Step 9	Switch(config)# end	Ends the current configuration session and returns to EXEC mode.
Step 10	Switch# show ip pim rp [mapping] [<i>rp-address</i>]	(Optional) Displays RPs known in the network and shows how the router learned about each RP.
Step 11	Switch# show ip igmp groups [<i>group-name</i> <i>group-address</i> <i>interface-type interface-number</i>] [detail]	(Optional) Displays the multicast groups having receivers that are directly connected to the router and that were learned through IGMP. <ul style="list-style-type: none"> A receiver must be active on the network at the time that this command is issued to ensure that receiver information is present on the resulting display.
Step 12	Switch# show ip mroute [<i>group-address</i> <i>group-name</i>] [<i>source-address</i> <i>source-name</i>] [<i>interface-type interface-number</i>] [summary] [count] [active <i>kbps</i>]	(Optional) Displays the contents of the IP multicast routing (mroute) table.

This example shows how to configure a single-static RP:

```
Switch> enable
Switch# configure terminal
Switch(config)# ip multicast-routing
Switch(config)# interface ethernet 1
Switch(config-if)# ip pim sparse-mode
Switch(config-if)# exit
Switch(config)# ip pim rp-address 192.168.0.0
Switch(config)# end
Switch# show ip pim rp mapping
Switch# show ip igmp groups
Switch# show ip mroute cbone-audio
```

Load Splitting of IP Multicast Traffic

If two or more equal-cost paths from a source are available, unicast traffic is load split across those paths. However, by default, multicast traffic is not load split across multiple equal-cost paths. In general, multicast traffic flows down from the reverse path forwarding (RPF) neighbor. According to the Protocol Independent Multicast (PIM) specifications, this neighbor must have the highest IP address if more than one neighbor has the same metric.

Use the **ip multicast multipath** command to enable load splitting of IP multicast traffic across multiple equal-cost paths.



Note

The **ip multicast multipath** command does not work with bidirectional Protocol Independent Multicast (PIM).

To enable IP multicast multipath, perform this task:

	Command	Purpose
Step 1	Switch# config t	Enters configuration mode.
Step 2	Switch(config)# ip multicast multipath	Enables IP multicast multipath.
Step 3	Switch(config)# end	Exits configuration mode.



Note

The **ip multicast multipath** command load splits the traffic but does not load balance the traffic. Traffic from a source uses only one path, even if the traffic far outweighs traffic from other sources.

Configuring load splitting with the **ip multicast multipath** command causes the system to load split multicast traffic across multiple equal-cost paths based on source address using the S-hash algorithm. When the **ip multicast multipath** command is configured and multiple equal-cost paths exist, the path in which multicast traffic travel is selected based on the source IP address. Multicast traffic from different sources is load split across the different equal-cost paths. Load splitting does not occur across equal-cost paths for multicast traffic from the same source sent to different multicast groups.

The following example shows how to enable ECMP multicast load splitting on a router based on a source address using the S-hash algorithm:

```
Switch(config)# ip multicast multipath
```

The following example shows how to enable ECMP multicast load splitting on a router based on a source and group address using the basic S-G-hash algorithm:

```
Switch(config)# ip multicast multipath s-g-hash basic
```

The following example shows how to enable ECMP multicast load splitting on a router based on a source, group, and next-hop address using the next-hop-based S-G-hash algorithm:

```
Switch(config)# ip multicast multipath s-g-hash next-hop-based
```

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 33-23](#)
- [Displaying the Multicast Routing Table, page 33-23](#)
- [Displaying IP MFIB, page 33-26](#)
- [Displaying Bidirectional PIM Information, page 33-27](#)
- [Displaying PIM Statistics, page 33-27](#)
- [Clearing Tables and Databases, page 33-28](#)

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 using the network.

To display various routing statistics, enter any of these commands:

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

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, enter one of these commands:

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 <i>n</i> MFIB-related events, the most recent first. <i>n</i> represents the number of events.
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 Bidirectional PIM Information

To display Bidir-PIM information, enter one of these commands:

Command	Purpose
Switch(config)# show ip pim interface [<i>type number</i>] [<i>df</i> <i>count</i>] [<i>rp-address</i>]	Displays information about the elected designated forward (DF) for each RP of an interface, along with the unicast routing metric associated with the DF.
Switch(config)# show ip pim rp [<i>mapping</i> <i>metric</i>] [<i>rp-address</i>]	Displays information about configured RPs, learned by using Auto-RP or BSR, along with their unicast routing metric.

Displaying PIM Statistics

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

```
Switch# show ip pim interface
```

Address	Interface	Mode	Neighbor Count	Query Interval	DR
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 example 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 one of these commands:

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.



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 33-28](#)
- [PIM Sparse Mode Example, page 33-29](#)
- [Bidirectional PIM Mode Example, page 33-29](#)
- [Sparse Mode with a Single Static RP Example, page 33-29](#)
- [Sparse Mode with Auto-RP: Example, page 33-30](#)

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
```

Bidirectional PIM Mode Example

By default, a bidirectional RP advertises all groups as bidirectional. Use an access list on the RP to specify a list of groups to be advertised as bidirectional. Groups with the **deny** keyword operate in dense mode. A different, nonbidirectional RP address is required for groups that operate in sparse mode, because a single access list only allows either a **permit** or **deny** keyword.

The following example shows how to configure an RP for both sparse and bidirectional mode groups. 224/8 and 227/8 are bidirectional groups, 226/8 is sparse mode, and 225/8 is dense mode. The RP must be configured to use different IP addresses for sparse and bidirectional mode operations. Two loopback interfaces are used to allow this configuration and the addresses of these interfaces must be routed throughout the PIM domain so that the other routers in the PIM domain can receive Auto-RP announcements and communicate with the RP:

```
ip multicast-routing !Enable IP multicast routing
ip pim bidir-enable !Enable bidir-PIM
!
interface loopback 0
 description One Loopback address for this routers Bidir Mode RP function
 ip address 10.0.1.1 255.255.255.0
 ip pim sparse-dense-mode
!
interface loopback 1
 description One Loopback address for this routers Sparse Mode RP function
 ip address 10.0.2.1 255.255.255.0
 ip pim sparse-dense-mode
ip pim send-rp-announce Loopback0 scope 10 group-list 45 bidir
ip pim send-rp-announce Loopback1 scope 10 group-list 46
ip pim send-rp-discovery scope 10
access-list 45 permit 224.0.0.0 0.255.255.255
access-list 45 permit 227.0.0.0 0.255.255.255
access-list 45 deny 225.0.0.0 0.255.255.255
access-list 46 permit 226.0.0.0 0.255.255.255
```

Sparse Mode with a Single Static RP Example

The following example sets the PIM RP address to 192.168.1.1 for all multicast groups and defines all groups to operate in sparse mode:

```
ip multicast-routing
interface ethernet 1
 ip pim sparse-mode
ip pim rp-address 192.168.1.1
no ip pim dm-fallback
```



Note

The same RP cannot be used for both bidirectional and sparse mode groups.

The following example sets the PIM RP address to 172.16.1.1 for the multicast group 225.2.2.2 only:

```
access-list 1 225.2.2.2 0.0.0.0
ip pim rp-address 172.17.1.1
```

Sparse Mode with Auto-RP: Example

The following example configures sparse mode with Auto-RP:

```
ip multicast-routing
ip pim autorp listener
ip pim send-rp-announce Loopback0 scope 16 group-list 1
ip pim send-rp-discovery Loopback1 scope 16
no ip pim dm-fallback
access-list 1 permit 239.254.2.0 0.0.0.255
access-list 1 permit 239.254.3.0 0.0.0.255
.
.
.
access-list 10 permit 224.0.1.39
access-list 10 permit 224.0.1.40
access-list 10 permit 239.254.2.0 0.0.0.255
access-list 10 permit 239.254.3.0 0.0.0.255
```