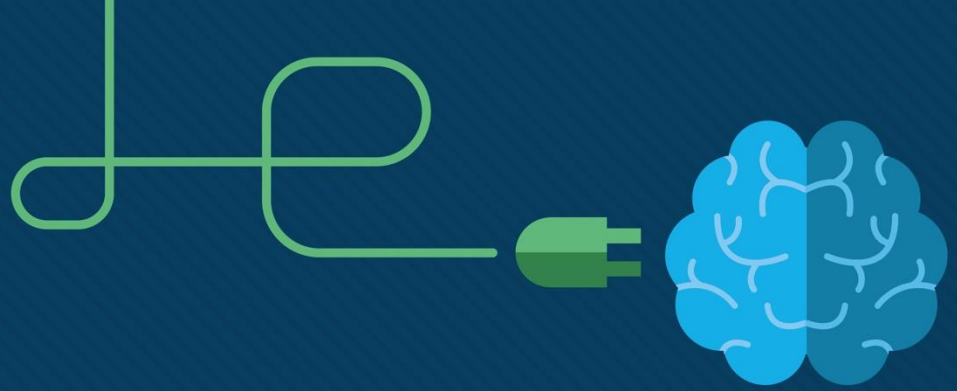

Chapter 13: Multicast

Instructor Materials

CCNP Enterprise: Core Networking

Chapter 13 Content

This chapter covers the following content:

Mulitcast Fundamentals - This section describes multicast concepts, as well as the need for multicast.

Multicast Addressing - This section describes the multicast address scopes used by multicast to operate at Layer 2 and Layer 3.

Internet Group Management Protocol - This section explains how multicast receivers join multicast groups to start receiving multicast traffic using IGMPv2 or IGMPv3. It also describes how multicast flooding on Layer 2 switches is prevented using a feature called IGMP snooping.

Chapter 13 Content (Cont.)

This chapter covers the following content:

Protocol Independent Multicast - This section describes the concepts, operation, and features of PIM. PIM is the protocol used to route multicast traffic across network segments from a multicast source to a group of receivers.

Rendezvous Points - This section describes the purpose, function, and operation of rendezvous points in a multicast network.

Multicast Fundamentals

- Multicast communication is a technology that optimizes network bandwidth utilization and conserves system resources.
- It relies on Internet Group Management Protocol (IGMP) for its operation in Layer 2 networks and Protocol Independent Multicast (PIM) for its operation in Layer 3 networks.

Multicast Architecture

Figure 13-1 illustrates how IGMP operates between the receivers and the local multicast router and how PIM operates between routers.

These two technologies work hand-in-hand to allow multicast traffic to flow from the source to the receivers, and they are explained in this chapter.

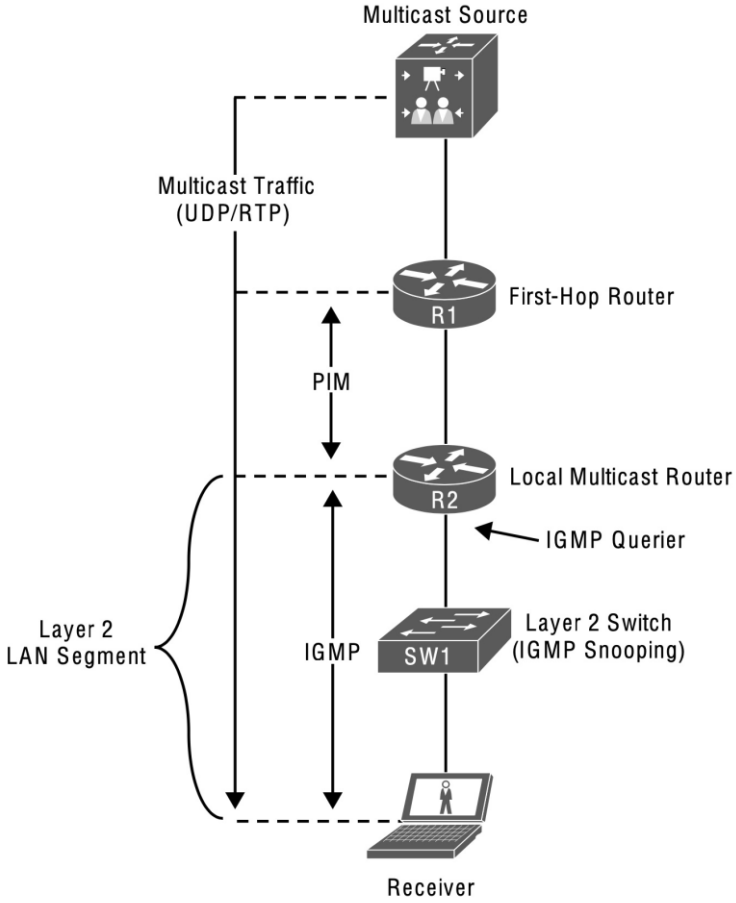

Figure 13-1 *Multicast Architecture*

Traditional IP Communication

Traditional IP communication between network hosts typically uses one of the following transmission methods:

- Unicast (one-to-one)
- Broadcast (one-to-all)
- Multicast (one-to-many)

Unicast Video Feed

Figure 13-2 shows an example where six workstations are watching the same video that is advertised by a server using unicast traffic (one-to-one). Each arrow represents a data stream of the same video going to five different hosts.

If each stream is 10 Mbps, the network link between R1 and R2 needs 50 Mbps of bandwidth. The network link between R2 and R4 requires 30 Mbps of bandwidth, and the link between R2 and R5 requires 20 Mbps of bandwidth. The server must maintain session state information for all the sessions between the hosts.

The bandwidth and load on the server increase as more receivers request the same video feed.

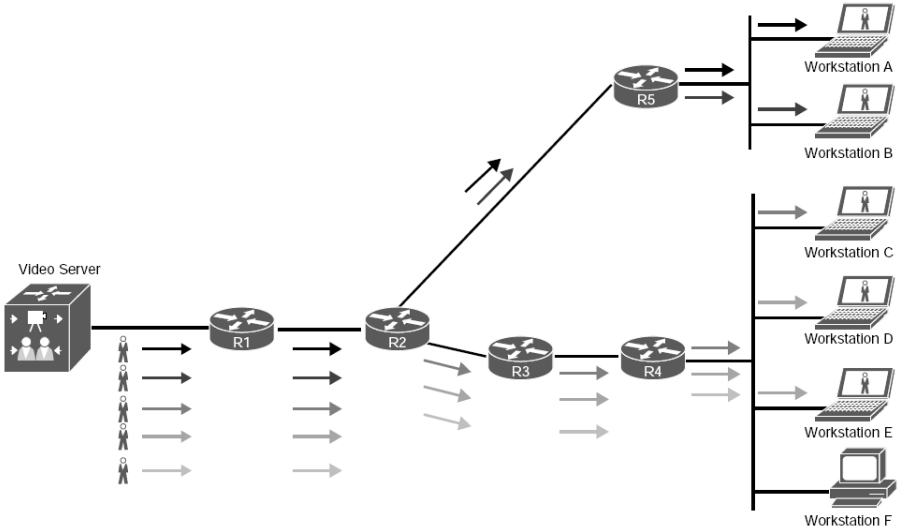

Figure 13-2 *Unicast Video Feed*

Broadcast Video Feed

Figure 13-3 shows an example of how the same video stream is transmitted using IP directed broadcasts. The load on the server is reduced because it needs to maintain only one session state rather than many. The same video stream consumes only 10 Mbps of bandwidth on all network links.

However, this approach does have disadvantages:

- IP directed broadcast functionality is not enabled by default on Cisco routers, and enabling it exposes the router to distributed denial-of-service (DDoS) attacks.
- In Figure 13-3, Workstation F is processing unwanted packets.

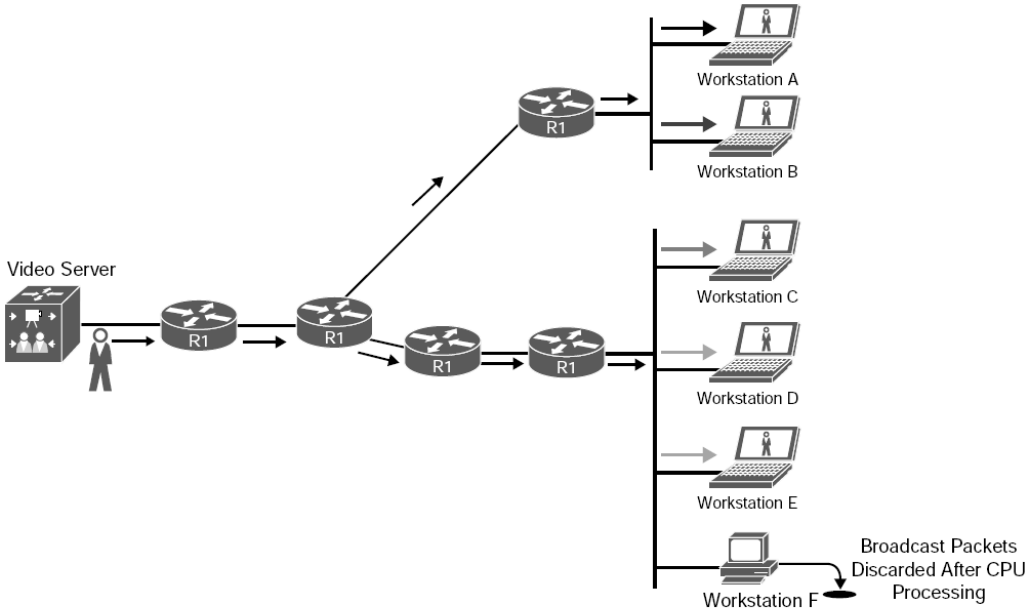

Figure 13-3 *Broadcast Video Feed*

Multicast Video Feed

Figure 13-4 shows an example of the same video feed using multicast. Each of the network links consumes only 10 Mbps of bandwidth, as much as with broadcast traffic, but only receivers that are interested in the video stream process the multicast traffic.

For example, Workstation F would drop the multicast traffic at the NIC level because it would not be programmed to accept the multicast traffic.

Note: Workstation F would not receive any multicast traffic if the switch for that network segment enabled Internet Group Management Protocol (IGMP) snooping, which is covered in a later section.

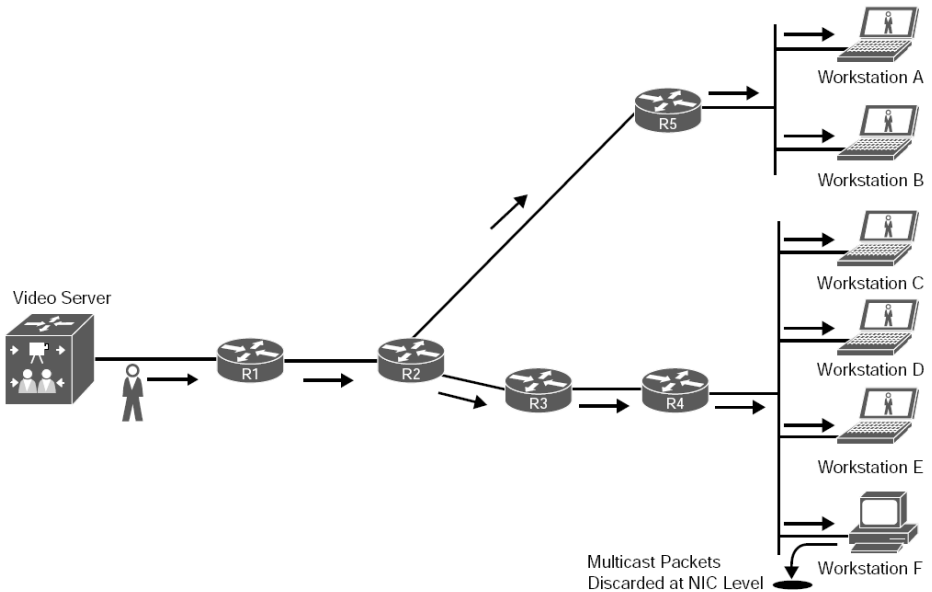

Figure 13-4 *Multicast Video Feed*

Multicast Traffic

Multicast traffic provides one-to-many communication, where only one data packet is sent on a link as needed and then is replicated between links as the data forks (splits) on a network device along the multicast distribution tree (MDT).

The data packets are known as a stream that uses a special destination IP address, known as a group address.

A server for a stream still manages only one session, and network devices selectively request to receive the stream.

Recipient devices of a multicast stream are known as receivers.

Common applications that take advantage of multicast traffic include Cisco TelePresence, real-time video, IPTV, stock tickers, distance learning, video/audio conferencing, music on hold, and gaming.

Multicast Addressing

- The Internet Assigned Number Authority (IANA) assigned the IP Class D address space 224.0.0.0/4 for multicast addressing. It includes addresses ranging from 224.0.0.0 to 239.255.255.255.
- The first 4 bits of this whole range start with 1110.

Class D Addressing

Out of the multicast blocks mentioned in Table 13-2, the most important are discussed in the list that follows:

- **Local network control block** (224.0.0/24) - Addresses in the local network control block are used for protocol control traffic that is not forwarded out a broadcast domain.
- **Internetwork control block** (224.0.1.0/24) - Addresses in the internetwork control block are used for protocol control traffic that may be forwarded through the Internet.

Table 13-2 IP Multicast Addresses Assigned by IANA

Designation	Multicast Address Range
Local network control block	224.0.0.0 to 224.0.0.255
Internetwork control block	224.0.1.0 to 224.0.1.255
Ad hoc block	224.0.2.0 to 224.0.255.255
Reserved	224.1.0.0 to 224.1.255.255
SDP/SAP block	224.2.0.0 to 224.2.255.255
Ad hoc block II	224.3.0.0 to 224.4.255.255
Reserved	224.5.0.0 to 224.255.255.255
Reserved	225.0.0.0 to 231.255.255.255
Source Specific Multicast (SSM) block	232.0.0.0 to 232.255.255.255
GLOP block	233.0.0.0 to 233.251.255.255
Ad hoc block III	233.252.0.0 to 233.255.255.255
Reserved	234.0.0.0 to 238.255.255.255
Administratively scoped block	239.0.0.0 to 239.255.255.255

Well–Known Reserved Multicast Addresses

Table 13-3 lists some of the well-known local network control block and internetwork control block multicast addresses.

Source Specific Multicast (SSM) block (**232.0.0.0/8**): This is the default range used by SSM. SSM is a PIM extension described in RFC 4607.

GLOP block (**233.0.0.0/8**): Addresses in the GLOP block are globally scoped statically assigned addresses. The mapping and assignment are defined in RFC 3180.

IP Multicast Address	Description
224.0.0.0	Base address (reserved)
224.0.0.1	All hosts in this subnet (all-hosts group)
224.0.0.2	All routers in this subnet
224.0.0.5	All OSPF routers (AllSPFRouters)
224.0.0.6	All OSPF DRs (AllDRouters)
224.0.0.9	All RIPv2 routers
224.0.0.10	All EIGRP routers
224.0.0.13	All PIM routers
224.0.0.18	VRRP
224.0.0.22	IGMPv3
224.0.0.102	HSRPv2 and GLBP
224.0.1.1	NTP
224.0.1.39	Cisco-RP-Announce (Auto-RP)
224.0.1.40	Cisco-RP-Discovery (Auto-RP)

Administratively scoped block (**239.0.0.0/8**): These addresses, described in RFC 2365, are limited to a local group or organization. Similar to the reserved IP unicast ranges (such as 10.0.0.0/8) defined in RFC 1918 and will not be assigned by the IANA.

Layer 2 Multicast Addressing

Every multicast group address (IP address) is mapped to a special MAC address that allows Ethernet interfaces to identify multicast packets to a specific group. A LAN segment can have multiple streams, and a receiver knows which traffic to send to the CPU for processing based on the MAC address assigned to the multicast traffic.

- The first 24 bits of a multicast MAC address always start with 01:00:5E.
- The low-order bit of the first byte is the individual/group bit (I/G) bit, also known as the unicast/multicast bit.
- When it is set to 1, it indicates that the frame is a multicast frame, and the 25th bit is always 0.
- The lower 23 bits of the multicast MAC address are copied from the lower 23 bits of the multicast group IP address.

Layer 2 Multicast Addressing (Cont.)

Figure 13-5 shows an example of mapping the multicast IP address 239.255.1.1 into multicast MAC address 01:00:5E:7F:01:01. The first 25 bits are always fixed; the last 23 bits that are copied directly from the multicast IP address vary.

Out of the 9 bits from the multicast IP address that are not copied into the multicast MAC address, the high-order bits 1110 are fixed; that leaves 5 bits that are variable that are not transferred into the MAC address.

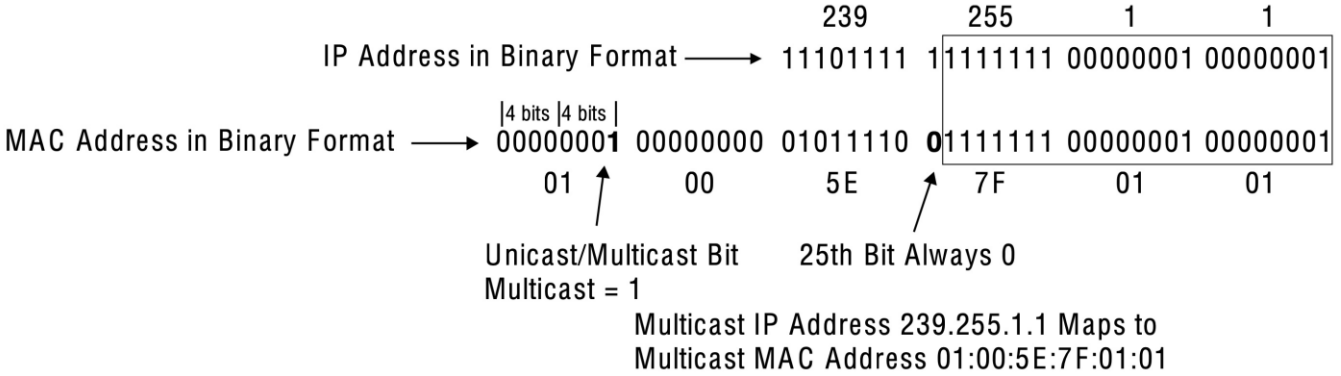

Figure 13-5 *Multicast IP Address-to-Multicast MAC Address Mapping*

Because of this, there are 32 (25) multicast IP addresses that are not universally unique and could correspond to a single MAC address. In other words, they overlap.

Layer 2 Address Mapping Overlap

Figure 13-6 shows an example of two multicast IP addresses that overlap because they map to the same multicast MAC address,

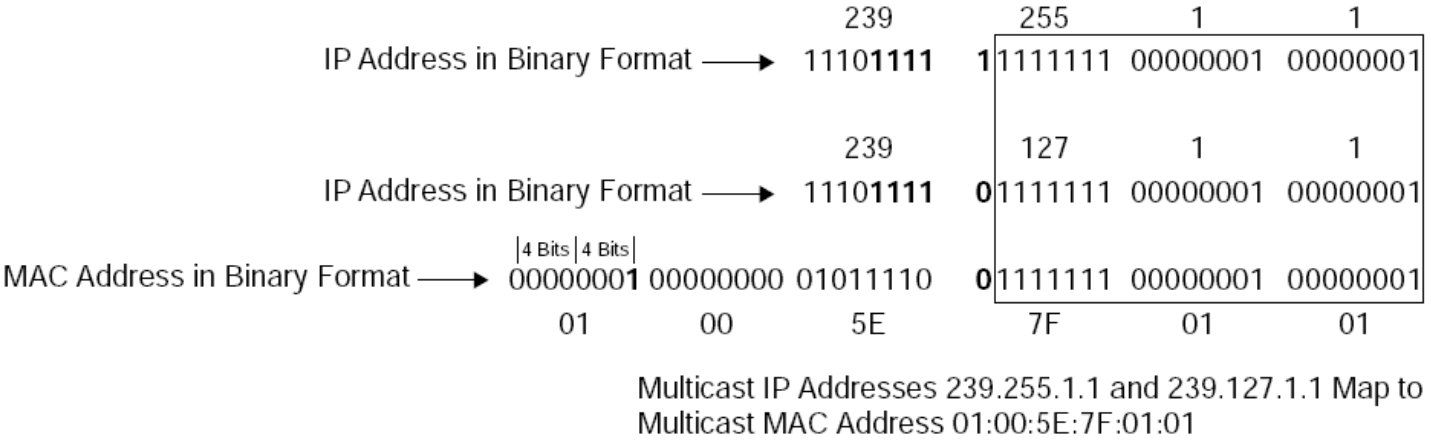

Figure 13-6 *Multicast IP Address to Multicast MAC Address Mapping Overlap*

Internet Group Management Protocol

- Internet Group Management Protocol (IGMP) is the protocol that receivers use to join multicast groups.
- When a receiver wants to receive a specific multicast feed, it sends an IGMP join using the multicast IP group address for that feed. The receiver reprograms its interface to accept the multicast MAC group address that correlates to the group address. For example, a PC could send a join to 239.255.1.1 and would reprogram its NIC to receive 01:00:5E:7F:01:01.
- IGMP must be supported by receivers and the router interfaces facing the receivers.
- Three versions of IGMP exist. RFC 1112 defines IGMPv1, which is old and rarely used. RFC 2236 defines IGMPv2, which is common in most multicast networks, and RFC 3376 defines IGMPv3, which is used by SSM. Only IGMPv2 and IGMPv3 are described in this chapter..

Version 2 Message Format

IGMPv2 uses the message format shown in Figure 13-7. This message is encapsulated in an IP packet with a protocol number of 2. Messages are sent with the IP router alert option set, which indicates that the packets should be examined more closely, and a time-to-live (TTL) of 1. IGMP packets are sent with a TTL of 1 so that packets are processed by the local router and not forwarded by any router.

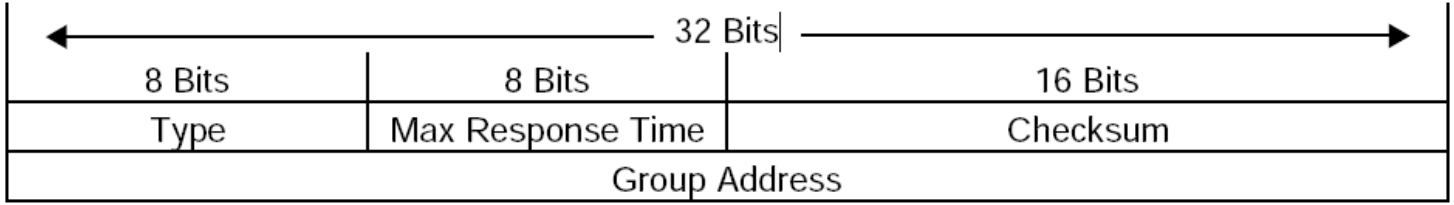

Figure 13-7 *IGMP Message Format*

The Type field describes five different types of IGMP messages used by routers and receivers.

Version 2 Message Format - Type Field Messages

Five different **types** of IGMP messages used by routers and receivers:

- **Version 2 membership report** (type value 0x16) is a message type also commonly referred to as an IGMP join; used by receivers to join a multicast group or to respond to a local router's membership query message.
- **Version 1 membership report** (type value 0x12) is used by receivers for backward compatibility with IGMPv1.
- **Version 2 leave group** (type value 0x17) is used by receivers to indicate they want to stop receiving multicast traffic for a group they joined.
- **General membership query** (type value 0x11) is periodically sent to the all-hosts group address 224.0.0.1 to see whether there are any receivers in the attached subnet. It sets the group address field to 0.0.0.0.
- **Group specific query** (type value 0x11) is sent in response to a leave group message to the group address the receiver requested to leave. The group address is the destination IP address of the IP packet and the group address field.

Version 2 Message Format – Other Fields

Other fields of the IGMP v2 format include:

- **Max response time -** This field is set only in general and group-specific membership query messages (type value 0x11). It specifies the maximum allowed time before sending a responding report in units of one-tenth of a second. In all other messages, it is set to 0x00 by the sender and ignored by receivers.
- **Checksum -** This field is the 16-bit 1s complement of the 1s complement sum of the IGMP message. This is the standard checksum algorithm used by TCP/IP.
- **Group address -** This field is set to 0.0.0.0 in general query messages and is set to the group address in group-specific messages. Membership report messages carry the address of the group being reported in this field; group leave messages carry the address of the group being left in this field.

Version 2 Messages

When a receiver wants to receive a multicast stream, it sends an unsolicited membership report, commonly referred to as an IGMP join, to the local router for the group it wants to join (for example, 239.1.1.1). The local router then sends this request upstream toward the source using a PIM join message. When the local router starts receiving the multicast stream, it forwards it downstream to the subnet where the receiver that requested it resides.

The router then starts periodically sending general membership query messages into the subnet, to the all-hosts group address 224.0.0.1, to see whether any members are in the attached subnet. The general query message contains a max response time field that is set to 10 seconds by default.

In response to this query, receivers set an internal random timer between 0 and 10 seconds. When the timer expires, receivers send membership reports (join message) for each group they belong to. If a receiver receives another receiver's report (join message) for one of the groups it belongs to while it has a timer running, it stops its timer for the specified group and does not send a report (join); this is meant to suppress duplicate reports (join messages).

Version 2 Messages (Cont.)

When a receiver wants to leave a group, if it was the last receiver to respond to a query, it sends a leave group message to the all-routers group address 224.0.0.2. Otherwise, it can leave quietly because there must be another receiver in the subnet.

When the leave group message is received by the router, it follows with a specific membership query to the group multicast address to determine whether there are any receivers interested in the group remaining in the subnet. If there are none, the router removes the IGMP state for that group.

If there is more than one router in a LAN segment, an IGMP querier election takes place to determine which router will be the querier. IGMPv2 routers send general membership query messages with their interface address as the source IP address and destined to the 224.0.0.1 multicast address.

When an IGMPv2 router receives such a message, it checks the source IP address and compares it to its own interface IP address. The router with the lowest interface IP address in the LAN subnet is elected as the IGMP querier.

Internet Group Management Protocol Version 3

In IGMPv2, when a receiver sends a membership report to join a multicast group. It does not specify which source it would like to receive multicast traffic from. IGMPv3 adds support for multicast source filtering, giving the receivers the capability to pick the source they wish to accept multicast traffic from.

IGMPv3 supports all IGMPv2's IGMP message types and is backward compatible with IGMPv2.

IGMPv3 added new fields to the IGMP membership query and introduced a new IGMP message type called Version 3 membership report to support source filtering in the following two modes:

Include mode - The receiver announces membership to a multicast group address and provides a list of source addresses (the include list) from which it wants to receive traffic.

Exclude mode - The receiver announces membership to a multicast group address and provides a list of source addresses (the exclude list) from which it does not want to receive traffic. The receiver then receives traffic only from sources whose IP addresses are not listed on the exclude list.

IGMP Snooping

In the case of multicast traffic, a multicast MAC address is never used as a source MAC address. Switches treat multicast MAC addresses as unknown frames and flood them out all ports. All workstations then process these frames. It is then up to the workstations to select interested frames for processing and select the frames that should be discarded.

The flooding of multicast traffic on a switch wastes bandwidth utilization on each LAN segment.

Cisco switches use two methods to reduce multicast flooding on a LAN segment:

- IGMP snooping
- Static MAC address entries

IGMP Snooping

IGMP snooping, defined in RFC 4541, is the most widely used method and works by examining IGMP joins sent by receivers and maintaining a table of interfaces to IGMP joins. When the switch receives a multicast frame destined for a multicast group, it forwards the packet only out the ports where IGMP joins were received for that specific multicast group.

Figure 13-10 illustrates Workstation A and Workstation C sending IGMP joins to 239.255.1.1, which translates to the multicast MAC address 01:00:5E:7F:01:01. Switch 1 has IGMP snooping enabled and populates the MAC address table with this information.

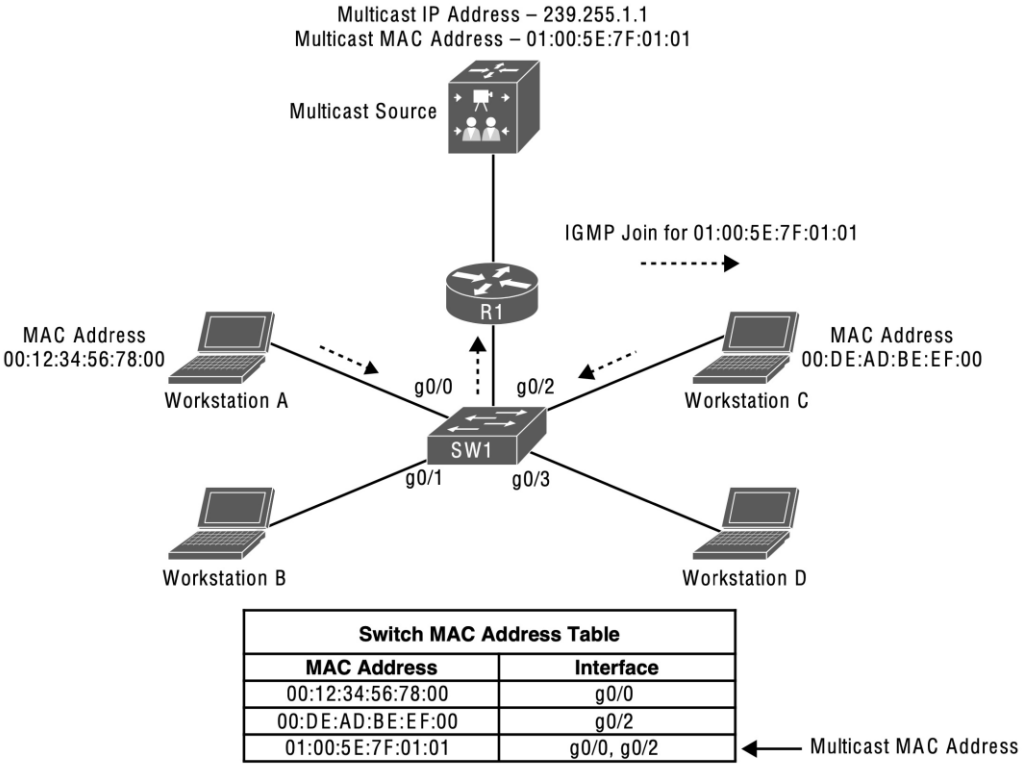

Figure 13-10 *IGMP Snooping Example*

IGMP Snooping

Figure 13-11 illustrates the source sending traffic to 239.255.1.1(01:00:5E:7F:01:01).

Switch 1 receives this traffic, and it forwards it out only the g0/0 and g0/2 interfaces because those are the only ports that received IGMP joins for that group.

A multicast static entry can also be manually programmed into the MAC address table, but this is not a scalable solution because it cannot react dynamically to changes. For this reason, it is not a recommended approach.

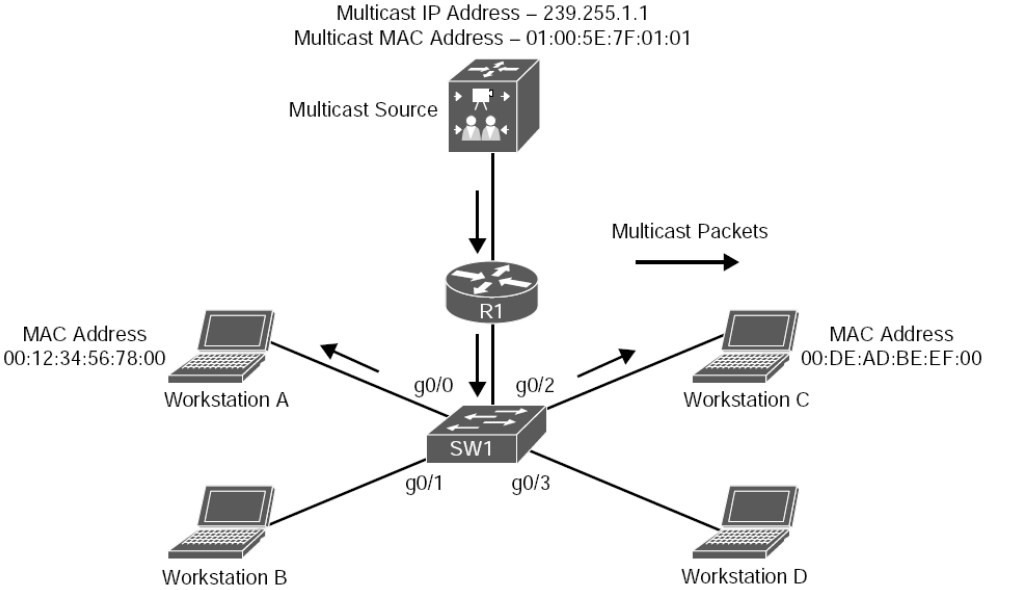

Figure 13-11 *No Flooding with IGMP Snooping*

Protocol Independent Multicast

- A multicast routing protocol is necessary to route the multicast traffic throughout the network so that routers can locate and request multicast streams from other routers. Multiple multicast routing protocols exist, but Cisco fully supports only Protocol Independent Multicast (PIM).
- PIM is a multicast routing protocol that routes multicast traffic between network segments.
- PIM can use any of the unicast routing protocols to identify the path between the source and receivers.
- Multicast routers create distribution trees that define the path that IP multicast traffic follows through the network to reach the receivers. The two basic types of multicast distribution trees are:
  - source trees, also known as shortest path trees (SPTs)
  - shared trees, also known as rendezvous point trees (RPTs)

Source Tree

A source tree is a multicast distribution tree where the source is the root of the tree, and branches form a distribution tree through the network all the way down to the receivers.

When this tree is built, it uses the shortest path through the network from the source to the leaves of the tree. For this reason, it is also referred to as a shortest path tree (SPT).

The forwarding state of the SPT is known by the notation (S,G), pronounced “S comma G,” where S is the source of the multicast stream and G is the multicast group address.

Figure 13-12 shows the SG notation as (10.1.1.2, 239.1.1.1)

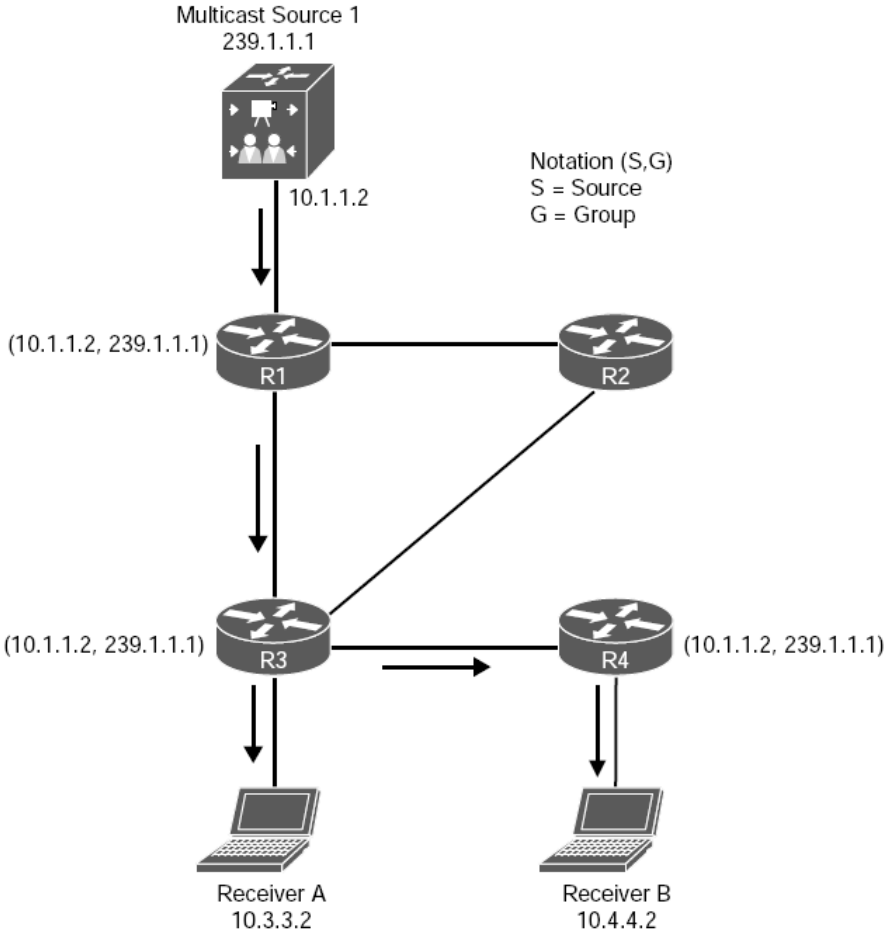

Figure 13-12 *Source Tree Example*

Shared Tree

A shared tree is a multicast distribution tree where the root of the shared tree is not the source but a router designated as the rendezvous point (RP). Shared trees are also referred to as RP trees (RPTs).

Multicast traffic is forwarded down the shared tree according to the group address G that the packets are addressed to, regardless of the source address. The forwarding state on the shared tree is referred to by the notation (*,G), pronounced “star comma G.”

Figure 13-13 illustrates a shared tree where R2 is the RP, and the (*,G) is (*,239.1.1.1).

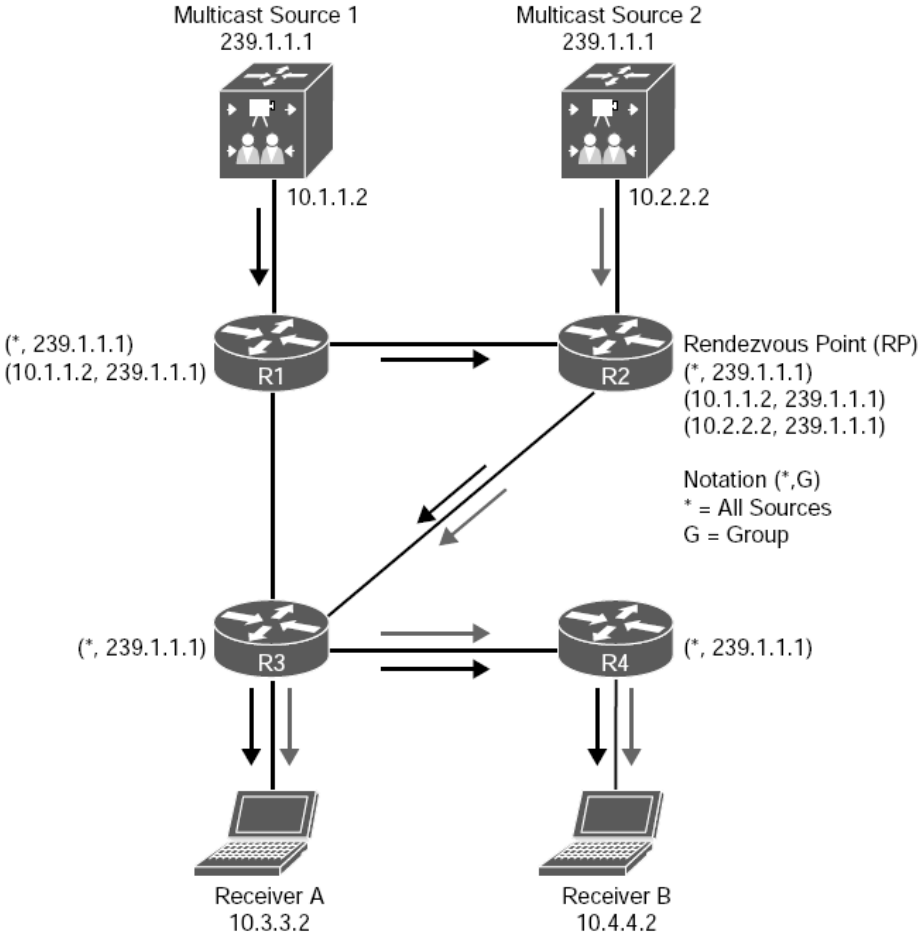

Figure 13-13 *Shared Tree Between RP and LHRs*

Protocol Independent Multicast
Terminology

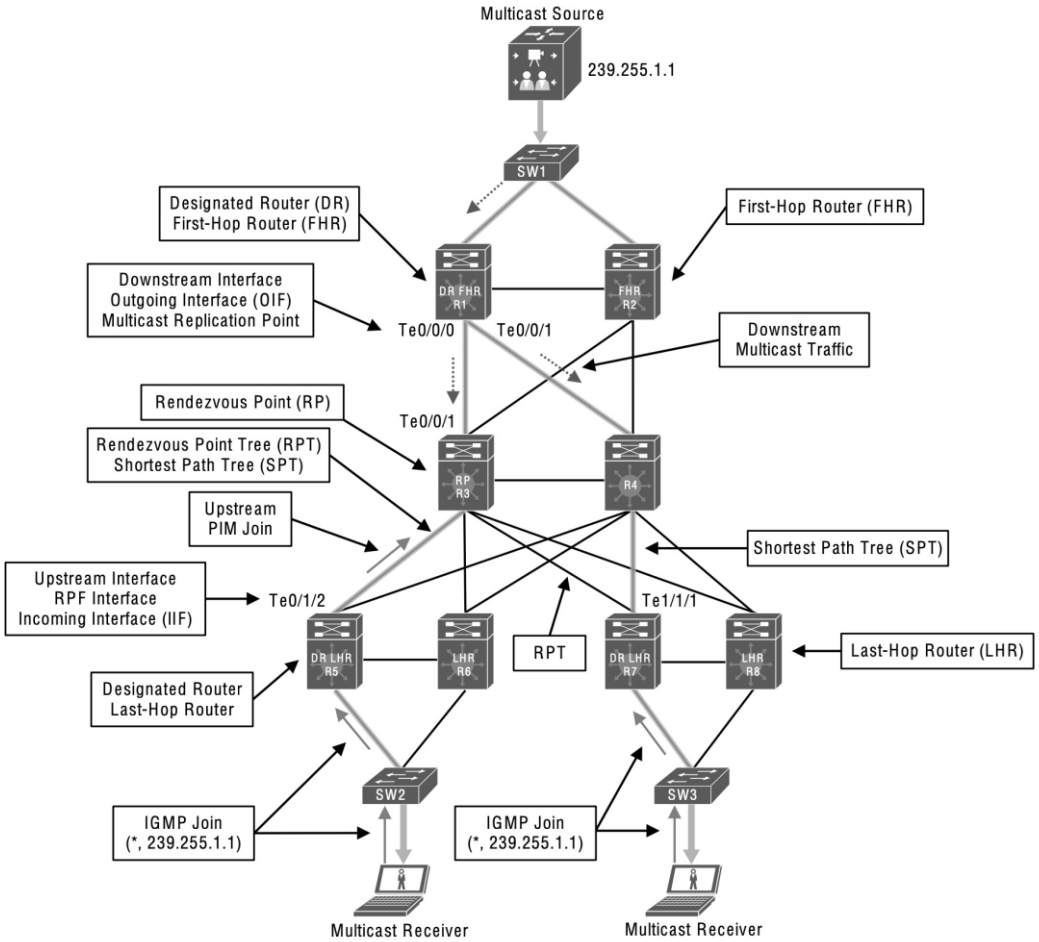

Figure 13-14 *PIM Terminology Illustration*

Operating Modes

There are currently five PIM operating modes:

- PIM Dense Mode (PIM-DM)
- PIM Sparse Mode (PIM-SM)
- PIM Sparse Dense Mode
- PIM Source Specific Multicast (PIM-SSM)
- PIM Bidirectional Mode (Bidir-PIM)

PIM-DM and PIM-SM are also commonly referred to as any-source multicast (ASM).

Control Messages

All PIM control messages use the IP protocol number 103; they are either unicast (that is, register and register stop messages) or multicast, with a TTL of 1 to the all PIM routers address 224.0.0.13.

PIM hello messages are sent by default every 30 seconds out each PIM-enabled interface to learn about the neighboring PIM routers on each interface to the all PIM routers address shown in Table 13-4.

Hello messages are also the mechanism used to elect a designated router (DR) and to negotiate additional capabilities.

All PIM routers must record the hello information received from each PIM neighbor. PIM Dense Mode (PIM-DM)

Table 13-4 PIM Control Message Types

Type	Message Type	Destination	PIM Protocol
0	Hello	224.0.0.13 (all PIM routers)	PIM-SM, PIM-DM, Bidir-PIM and SSM
1	Register	RP address (unicast)	PIM-SM
2	Register stop	First-hop router (unicast)	PIM SM
3	Join/prune	224.0.0.13 (all PIM routers)	PIM-SM, Bidir-PIM and SSM
4	Bootstrap	224.0.0.13 (all PIM routers)	PIM-SM and Bidir-PIM
5	Assert	224.0.0.13 (all PIM routers)	PIM-SM, PIM-DM, and Bidir-PIM
8	Candidate RP advertisement	Bootstrap router (BSR) address (unicast to BSR)	PIM-SM and Bidir-PIM
9	State refresh	224.0.0.13 (all PIM routers)	PIM-DM
10	DF election	224.0.0.13 (all PIM routers)	Bidir-PIM

PIM Dense Mode

Figure 13-15 shows the flood and prune operation of Dense Mode.

The multicast traffic from the source is flooding throughout the entire network. As each router receives the multicast traffic from its upstream neighbor via its RPF interface, it forwards the multicast traffic to all its PIM-DM neighbors.

This results in some traffic arriving via a non-RPF interface, as in the case of R3 receiving traffic from R2 on its non-RPF interface.

Packets arriving via the non-RPF interface are discarded.

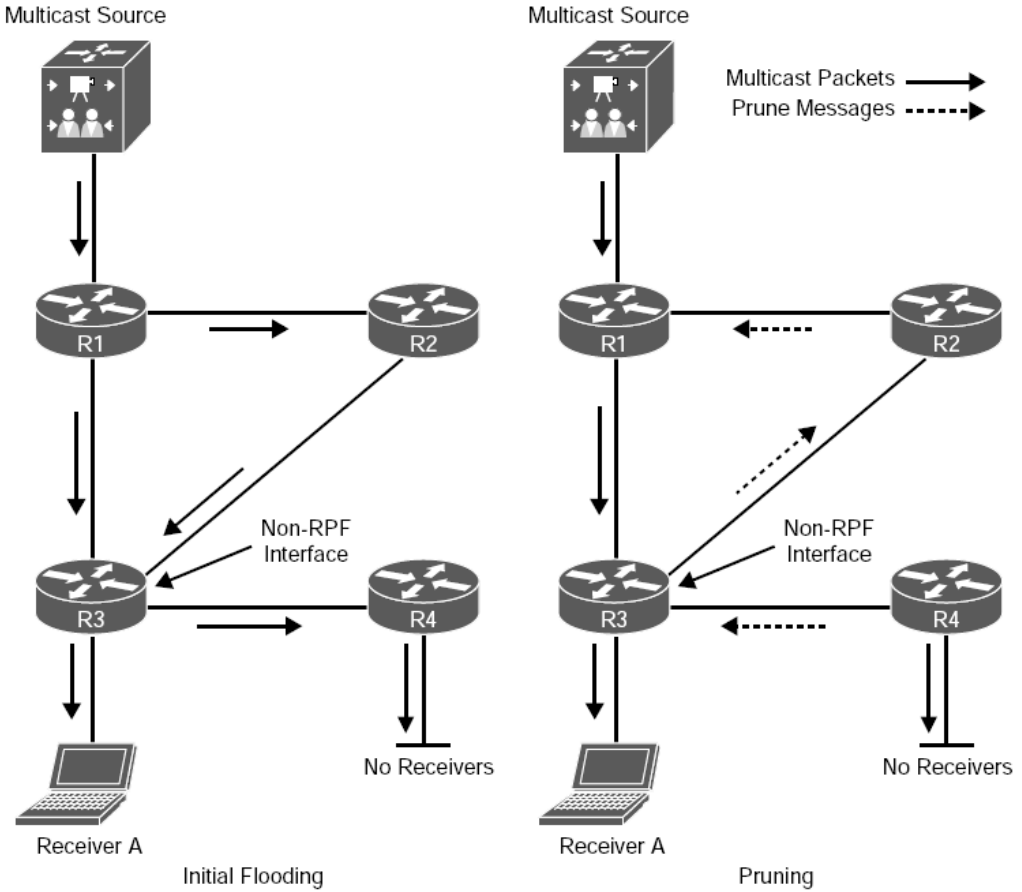

Figure 13-15 *PIM-DM Flood and Prune Operation*

PIM Dense Mode (Cont.)

Figure 13-16 illustrates the resulting topology after all unnecessary links have been pruned off. This results in an SPT from the source to the receiver.

Even though the flow of multicast traffic is no longer reaching most of the routers in the network, the (S,G) state still remains in all routers until the source stops transmitting.

PIM-DM is applicable to small networks where there are active receivers on every subnet of the network. Because this is rarely the case, and the flood and prune behavior, PIM-DM is not generally recommended for production environments

However, it can be useful for a lab environment because it is easy to set up.

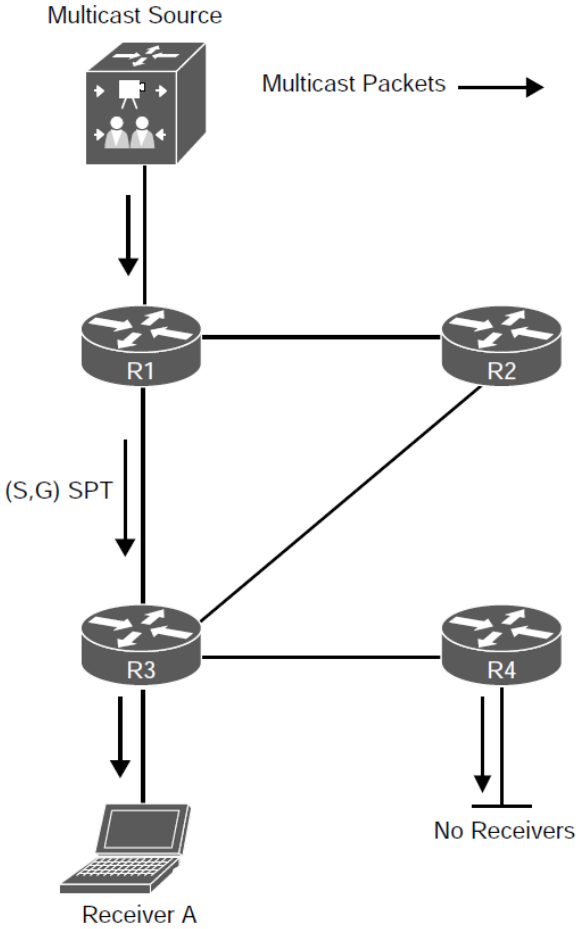

Figure 13-16 *PIM-DM Resulting Topology After Pruning*

PIM Sparse Mode

PIM-SM was designed for receivers scattered throughout the network but works well in densely populated networks.

It also assumes that no receivers are interested in multicast traffic unless they explicitly request it.

Just like PIM-DM, PIM-SM uses the unicast routing table to perform RPF checks, and it does not care which routing protocol (including static routes) populates the unicast routing table

PIM-SM uses an explicit join model where the receivers send an IGMP join to their locally connected router, which is also known as the last-hop router (LHR) and this join causes the LHR to send a PIM join in the direction of the root of the tree, which is either the RP in the case of a shared tree (RPT) or the first-hop router (FHR) where the source transmitting the multicast streams is connected in the case of an SPT.

A multicast forwarding state is created as the result of these explicit joins; it is very different from the flood and prune behavior of PIM-DM.

PIM Sparse Mode

Figure 13-17 illustrates a multicast source sending multicast traffic to the FHR. The FHR then sends this multicast traffic to the RP, which makes the multicast source known to the RP. It also illustrates a receiver sending an IGMP join to the LHR to join the multicast group.

The LHR then sends a PIM join (*,G) to the RP, and this forms a shared tree from the RP to the LHR.

The RP then sends a PIM join (S,G) to the FHR, forming a source tree between the source and the RP.

In essence, two trees are created: an SPT from the FHR to the RP (S,G) and a shared tree from the RP to the LHR (*,G).

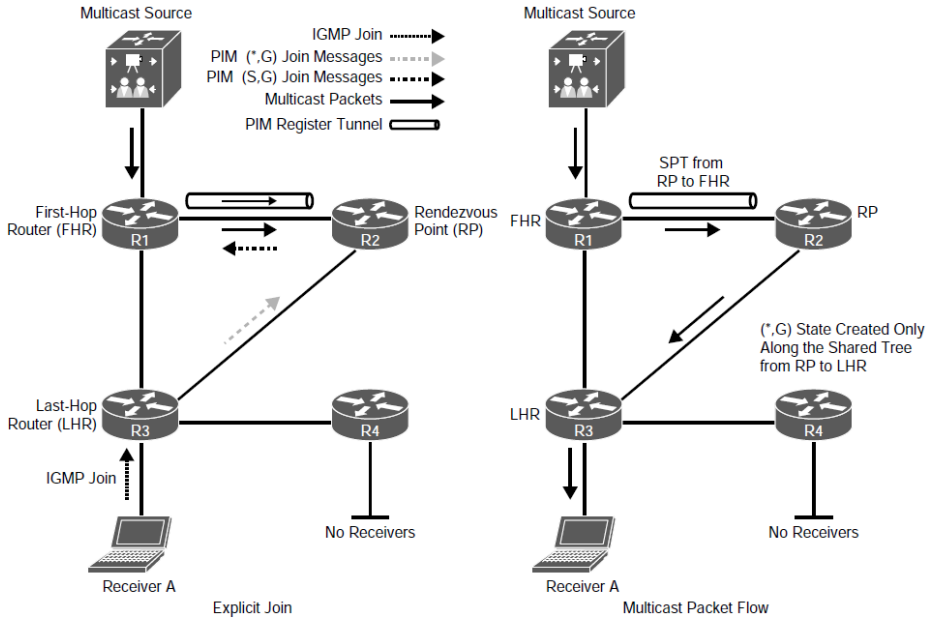

Figure 13-17 *PIM-SM Multicast Distribution Tree Building*

PIM Sparse Switchover Mode

PIM-SM allows the LHR to switch from the shared tree to an SPT for a specific source. In Cisco routers, this is the default behavior, and it happens immediately after the first multicast packet is received from the RP via the shared tree, even if the shortest path to the source is through the RP.

Figure 13-18 illustrates the SPT switchover concept. When the LHR receives the first multicast packet from the RP, it becomes aware of the IP address of the multicast source.

In Figure 13-18, the shortest path to the source is between R1 and R3; if that link were shut down or not present, the shortest path would be through the RP, in which case an SPT switchover would still take place.

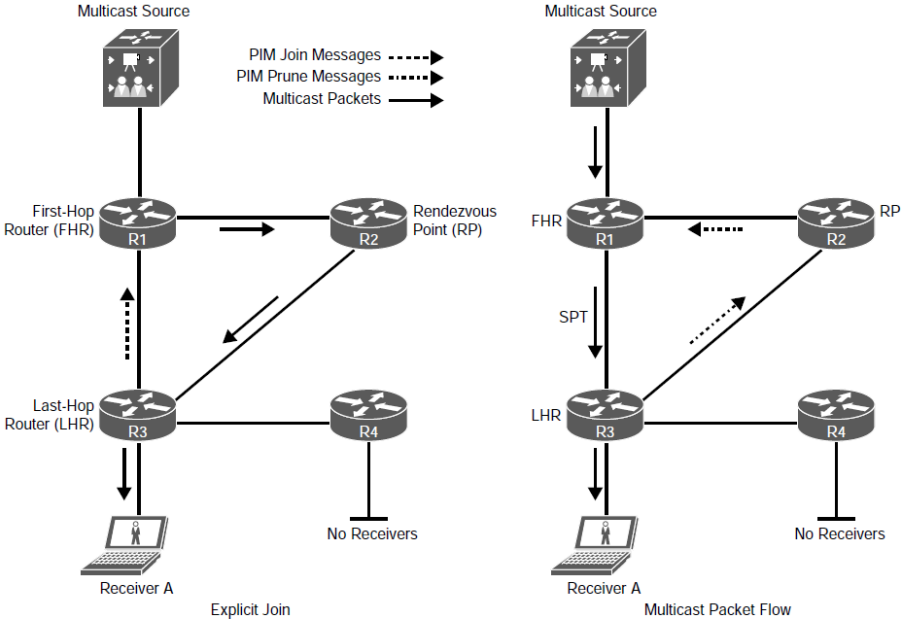

Figure 13-18 *PIM-SM SPT Switchover Example*

Designated Routers

When multiple PIM-SM routers exist on a LAN segment, PIM hello messages are used to elect a designated router (DR) to avoid sending duplicate multicast traffic into the LAN or the RP.

By default, the DR priority value of all PIM routers is 1, and it can be changed to force a router to become the DR.

If all routers have the same priority value, the highest IP address in the subnet is used as a tiebreaker.

Without DRs, all LHRs on the same LAN segment would be capable of sending PIM joins upstream, which could result in duplicate multicast traffic arriving on the LAN.

The default DR hold time is 3.5 times the hello interval, or 105 seconds. If there are no hellos after this interval, a new DR is elected.

To reduce DR failover time, the hello query interval can be reduced

Reverse Path Forwarding

Reverse Path Forwarding (RPF) is an algorithm used to prevent loops and ensure that multicast traffic is arriving on the correct interface. RPF functions as follows:

- If a router receives a multicast packet 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 (OIL) of a multicast routing table entry.
- If the packet does not arrive on the RPF interface, the packet is discarded to prevent loops.

PIM-SM uses the RPF lookup function to determine where it needs to send joins and prunes.

- (S,G) joins are sent toward the source.
- (*,G) joins (shared tree states) are sent toward the RP.

PIM Forwarder

In Figure 13-20, PIM-DM would send duplicate flows into the LAN.

For example, assuming that R1 is the RP, when R4 sends a PIM join message upstream toward it, it sends it to the all PIM routers address 224.0.0.13, and R2 and R3 receive it.

One of the fields of the PIM join message includes the IP address of the upstream neighbor (RPF neighbor).

Assuming that R3 is the RPF neighbor, R3 is the only one that will send a PIM join to R1.

R2 will not because the PIM join was not meant for it. At this point, a shared tree exists between R1, R3, and R4, and no traffic duplication exists.

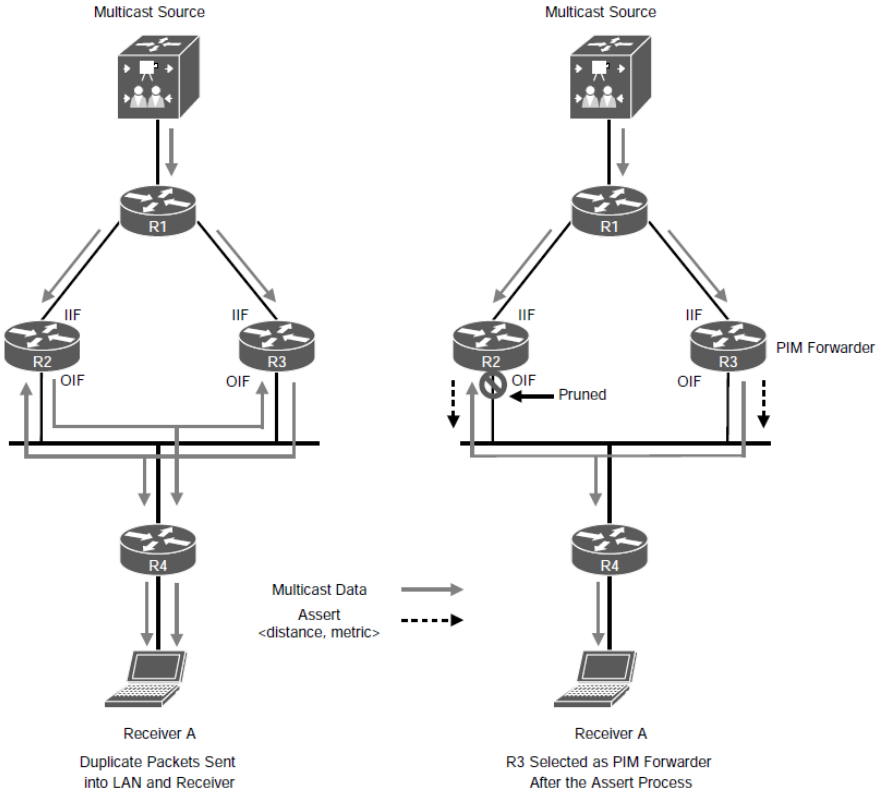

Figure 13-20 *PIM Forwarder Example*

Rendezvous Points

- In PIM-SM, it is mandatory to choose one or more routers to operate as rendezvous points (RPs). An RP is a single common root placed at a chosen point of a shared distribution tree, as described earlier.
- An RP can be either configured statically in each router or learned through a dynamic mechanism.
- A PIM router can be configured to function as an RP either statically in each router in the multicast domain or dynamically by configuring Auto-RP or a PIM bootstrap router (BSR), as described in the following sections.

Static RP

It is possible to statically configure RP for a multicast group range by configuring the address of the RP on every router in the multicast domain. Configuring static RPs is relatively simple and can be achieved with one or two lines of configuration on each router.

If the network does not have many different RPs defined or if the RPs do not change very often, this could be the simplest method for defining RPs. It can also be an attractive option if the network is small.

However, static configuration can increase administrative overhead in a large and complex network. Every router must have the same RP address. Changing the RP address requires reconfiguring every router. If several RPs are active for different groups, information about which RP is handling which multicast group must be known by all routers. To ensure this information is complete, multiple configuration commands may be required.

If a manually configured RP fails, there is no failover procedure for another router to take over the function performed by the failed RP, and this method by itself does not provide any kind of load splitting.

Auto- RP

Auto-RP is a Cisco proprietary mechanism that automates the distribution of group-to-RP mappings in a PIM network. Auto-RP 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.
- It simplifies RP placement according to the locations of group participants.
- It prevents inconsistent manual static RP configurations that might cause connectivity problems.
- Multiple RPs can be used to serve different group ranges or to serve as backups for each other.
- The Auto-RP mechanism operates using two basic components, candidate RPs (C-RPs)and RP mapping agents (MAs).

Candidate RPs and RP Mapping Agents

A C-RP advertises its willingness to be an RP via RP announcement messages to the reserved well-known multicast group 224.0.1.39 (Cisco-RP-Announce).

The RP announcements contain the default group range 224.0.0.0/4, the C-RP's address, and the hold time, which is three times the RP announce interval.

If there are multiple C-RPs, the C-RP with the highest IP address is preferred.

RP MAs join group 224.0.1.39 to receive the RP announcements. They store the information contained in the announcements in a group-to-RP mapping cache, along with hold times. If multiple RPs advertise the same group range, the C-RP with the highest IP address is elected.

The RP MAs advertise the RP mappings to another well-known multicast group address, 224.0.1.40 (Cisco-RP-Discovery). These messages are advertised by default every 60 seconds or when changes are detected. The MA announcements contain the elected RPs and the group-to-RP mappings. All PIM-enabled routers join 224.0.1.40 and store the RP mappings in their private cache.

PIM Bootstrap Router

The bootstrap router (BSR) mechanism, described in RFC 5059, is a nonproprietary mechanism that provides a fault-tolerant, automated RP discovery and distribution mechanism.

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 RP set is a group-to-RP mapping that contains the following components:

- Multicast group range
- RP priority
- RP address
- Hash mask length
- SM/Bidir flag

Candidate RPs

Figure 13-23 illustrates the BSR mechanism, where the elected BSR receives candidate RP advertisement messages from all candidate RPs in the domain, and it then sends BSR messages with RP set information out all PIM-enabled interfaces, which are flooded hop-by-hop to all routers in the network.

The active BSR stores all incoming C-RP advertisements in its group-to-RP mapping cache. The BSR then sends the entire list of C-RPs from its group-to-RP mapping cache in BSR messages every 60 seconds by default to all PIM routers in the entire network. As the routers receive copies of these BSR messages, they update the information in their local group-to-RP mapping caches, and this allows them to have full visibility into the IP addresses of all C-RPs in the network.

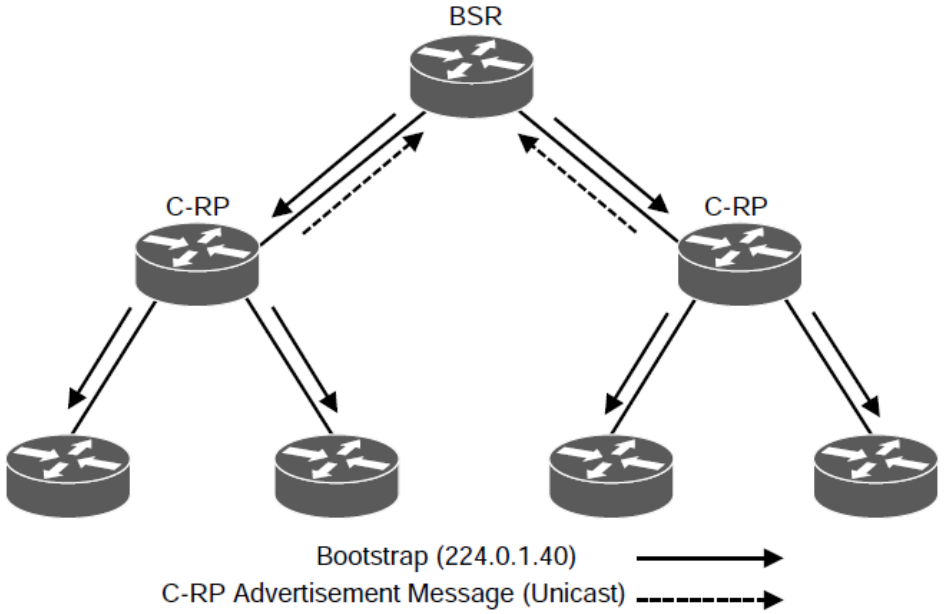

Figure 13-23 *BSR Mechanism*

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 13

Description
Multicast fundamentals
IP Multicast Addresses Assigned by IANA
Well-Known Reserved Multicast Addresses
Layer 2 multicast addresses
IGMP description
IGMPv2
IGMP message format field definitions
IGMPv2 operation

Prepare for the Exam

Key Topics for Chapter 13 (Cont.)

Description
IGMPv3 definition
IGMP snooping
PIM definition
PIM source tree definition
PIM shared tree definition
PIM terminology
PIM operating modes
PIM Control Message Types

Prepare for the Exam

Key Topics for Chapter 13 (Cont.)

Description
PIM-DM definition
PIM-SM definition
PIM-SM shared tree operation
PIM-SM source registration
PIM-SM SPT switchover
PIM-SM designated routers
RPF definition
PIM forwarder

Prepare for the Exam

Key Topics for Chapter 13 (Cont.)

Description
Rendezvous point definition
Static RP definition
Auto-RP definition
Auto-RP C-RP definition
Auto-RP mapping agent definition
PIM BSR definition
PIM BSR C-RP definition

Prepare for the Exam

Key Terms for Chapter 13

Key Terms		
designated router (DR) (Context of PIM)	Multicast Forwarding Information Base (MFIB)	ren
downstream interface	Multicast Routing Information Base (MRIB)	rendezvous point tree (RPT)
first-hop router (FHR)	multicast state	Reverse Path Forwarding (RPF) interface
incoming interface (IIF)	outgoing interface (OIF)	RPF neighbor
IGMP snooping	outgoing interface list (OIL)	shortest path tree (SPT)
Internet Group Management Protocol (IGMP)	Protocol Independent Multicast (PIM)	upstream
last-hop router (LHR)	dezvous point (RP)	upstream interface


CISCO