Internet Protocol (IP) Multicast

Background

Internet Protocol (IP) multicast is a routing technique that allows IP traffic to be sent from one source or multiple sources and delivered to multiple destinations. Instead of sending individual packets to each destination, a single packet is sent to a multicast group, which is identified by a single IP destination group address. IP multicast routing arose because unicast and broadcast techniques do not handle the requirements of new applications efficiently. Multicast addressing, for example, supports the transmission of a single IP datagram to multiple hosts. This chapter focuses on the leading multicast routing options. Figure 39-1 illustrates the general nature of a multicast environment

Figure 39-1 IP multicast provides a means to deliver high-bandwidth traffic to multiple destinations.

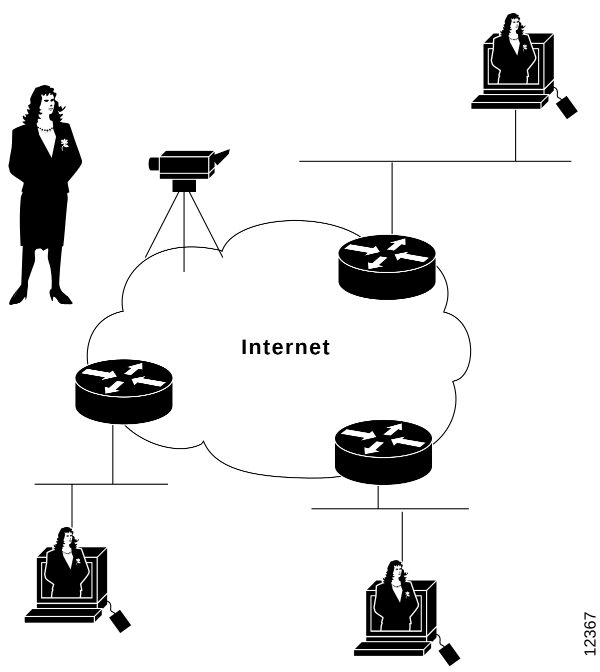

Internet Group-Membership Protocol (IGMP)

A principle component of IP multicast is the Internet Group-Membership Protocol (IGMP). IGMP relies on Class D IP addresses for the creation of multicast groups and is defined in RFC 1112. IGMP is used to dynamically register individual hosts in a multicast group with a Class D address. Hosts identify group memberships by sending IGMP messages, and traffic is sent to all members of that multicast group. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on particular LANs. Routers communicate with each other by using one or more protocols to build multicast routes for each group.

IP Multicast Routing Protocols

Several routing protocols are used to discover multicast groups and to build routes for each group. These include Protocol-Independent Multicast (PIM), Distance-Vector Multicast Routing Protocol (DVMRP), and Multicast Open Shortest Path First (MOSPF). The following table summarizes the unicast requirements needed and flooding algorithms used for each protocol. Table 39-1 summarizes the multicast routing option.

Table 39-1 Summary of Multicast Routing Options

Protocol	Unicast Protocol Requirements	Flooding Algorithm
PIM-dense mode	Any	Reverse path flooding (RPF)
PIM-sparse mode	Any	RPF
DVMRP	Internal, RIP-like routing protocol	RPF
MOSPF	Open Shortest Path First (OSPF)	Shortest-path first (SPF)

Protocol-Independent Multicast (PIM)

Protocol-Independent Multicast (PIM) is addressed in an Internet draft RFC (under discussion by the IETF Multicast Routing Working Group). It includes two different modes of behavior for dense and sparse traffic environments: *dense* mode and *sparse* mode.

The PIM dense mode uses a process of reverse path flooding that is similar to the DVMRP. Differences exist, however, between dense mode PIM and DVMRP. PIM, for example, does not require a particular unicast protocol to determine which interface leads back to the source of a data stream. DVMRP employs its own unicast protocol, while PIM uses whatever unicast protocol the internetwork is using.

The PIM sparse mode is optimized for internetworks with many data streams but relatively few LANs. It defines a rendezvous point that is then used as a registration point to facilitate the proper routing of packets.

When a sender wants to transmit data, the first-hop router (with respect to the source) node sends data to the rendezvous point. When a receiver wants to receive data, the last-hop router (with respect to the receiver) registers with the rendezvous point. A data stream then can flow from the sender to the rendezvous point and to the receiver. Routers in the path optimize the path and automatically remove any unnecessary hops, even at the rendezvous point.

Distance-Vector Multicast Routing Protocol (DVMRP)

The Distance-Vector Multicast Routing Protocol (DVMRP) uses a reverse path-flooding technique and is used as the basis for the Internet's multicast backbone (MBONE). DVMRP is defined in RFC 1075 and has certain some shortcomings. In particular, DVMRP is notorious for poor network scaling, resulting from reflooding, particularly with versions that do not implement pruning. DVMRP's flat unicast routing mechanism also affects its capability to scale.

The reverse path-flooding operation involves a router sending a copy of a packet out to all paths (except the path back to the origin) upon the packet's receipt. Routers then send a prune message back to the source to stop a data stream if the router is attached to a LAN that does not want to receive a particular multicast group.

Reflooding and *DVMRP unicast* are used in DVMRP path-flooding operations. In reflooding, DVMRP routers periodically reflood an attached network to reach new hosts. The flooding mechanism uses an algorithm that takes into account the frequency of flooding and the time required for a new multicast group member to receive the data stream. DVMRP unicast is used to determine which interface leads back to the source of a data stream. It is unique to DVMRP but is similar to RIP in that it is based on hop count. The DVMRP unicast environment permits the use of a different path than the path used for multicast traffic.

Multicast Open Shortest Path First (MOSPF)

The Multicast Open Shortest Path First (MOSPF) is an extension of OSPF. In general, MOSPF employs a unicast routing protocol that requires each router in a network to be aware of all available links.

An MOSPF router calculates routes from the source to all possible group members for a particular multicast group. MOSPF routers include multicast information in OSPF link states. MOSPF calculates the routes for each source/multicast group pair when the router receives traffic for that pair, and routes are cached until a topology change occurs. MOSPF then recalculates the topology.

Several MOSPF implementation issues have been identified and require consideration. First, MOSPF works only in internetworks that use OSPF. In addition, MOSPF is best suited for environments with relatively few active source/group pairs. MOSPF can take up significant router CPU bandwidth in environments that have many active source/group pairs or that are unstable.