

Chapter 2

IP Multicast Overview

The JUNOS software implements the following protocols to support IP multicast routing:

- Internet Group Management Protocol (IGMP), versions 1 and 2—Used to learn whether group members are present, for IP version 4 (IPv4) routers.
- Multicast Listener Discovery (MLD), versions 1 and 2—Used to learn whether group members are present, for IP version 6 (IPv6) routers.
- Distance Vector Multicast Routing Protocol (DVMRP)—Dense-mode multicast routing protocol.
- Protocol Independent Multicast (PIM)—Multicast routing protocol that routes to multicast groups that might span wide-area and interdomain internetworks. Both dense mode and sparse mode are supported.
- Multicast Source Discovery Protocol (MSDP)—Multicast routing protocol that discovers active sources of multicast messages. PIM sparse mode uses these sources.
- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Handle conference session announcements.

This chapter discusses the following topics:

- IP Multicast Standards on page 20
- Multicast Overview on page 21
- Multicast Addresses on page 22
- Multicast Redundancy on page 22
- Replicating Multicast Packets on page 23

IP Multicast Standards

The protocols related to IP multicast are defined in the following documents:

- RFC 1112, *Host Extensions for IP Multicasting* (defines IGMP Version 1)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2327, *SDP: Session Description Protocol*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- RFC 2365, *Administratively Scoped IP Multicast*
- RFC 2547, *BGP/MPLS VPNs*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 2974, *Session Announcement Protocol*
- RFC 3208, *PGM Reliable Transport Protocol Specification*
- RFC 3376, *Internet Group Management Protocol, Version 3* (source-specific multicast [SSM] include mode only)
- RFC 3446, *Anycast Rendezvous Point (RP) Mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
- RFC 3569, *An Overview of Source-Specific Multicast (SSM)*
- RFC 3590 (SSM include mode only), *Source Address Selection for Multicast Listener Discovery Protocol*
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- Internet draft draft-ietf-pim-sm-bsr-05.txt, *Bootstrap Router (BSR) Mechanism for PIM Sparse Mode* (expired August 2005) (no support for draft's scoping mechanism)
- Internet draft draft-ietf-idmr-dvmrp-v3-11.txt, *Distance Vector Multicast Routing Protocol* (expired April 2004)
- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs, Section 2* (expired December 2004)
- Internet draft draft-rosen-vpn-mcast-06.txt, *Multicast in MPLS/BGP VPNs, Section 2* (expired October 2003)
- Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs, data MDTs only* (expired April 2004)

- Internet draft draft-ietf-pim-sm-v2-new-10.txt, *Protocol Independent Multicast—Sparse Mode (PIM-SM): Protocol Specification (Revised)* (expired January 2005)
- Internet draft draft-ietf-pim-dm-new-v2-05.txt, *Protocol Independent Multicast Version 2 Dense Mode Specification* (expired December 2004)
- Internet draft draft-ietf-ssm-arch-06.txt, *Source-Specific Multicast for IP* (expired March 2005)
- Internet draft draft-ietf-mboned-ssm232-08.txt, *Source-Specific Protocol Independent Multicast in 232/8* (expired September 2004)
- Internet draft draft-holbrook-idmr-igmpv3-ssm-07.txt, *Using IGMPv3 and MLDv2 for Source-Specific Multicast* (expired December 2004)
- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs* (expired December 2004)



NOTE: The implementation of the Distance Vector Multicast Routing Protocol (DVMRP) is based on a series of expired Internet drafts, as are all current implements of DVMRP. None are based on RFC 1075, *Distance Vector Multicast Routing Protocol*.

To access Internet RFCs and drafts, go to the IETF Web site at <http://www.ietf.org>.

Multicast Overview

IPv4 has three fundamental types of addresses: unicast, broadcast, and multicast. A *unicast address* is used to send a packet to a single destination. A *broadcast address* is used to send a datagram to an entire subnetwork. A *multicast address* is used to send a datagram to a set of hosts that can be on different subnetworks and that are configured as members of a multicast group.

A multicast datagram is delivered to destination group members with the same best-effort reliability as a standard unicast IP datagram. This means that multicast datagrams are not guaranteed to reach all members of a group or to arrive in the same order in which they were transmitted. The only difference between a multicast IP packet and a unicast IP packet is the presence of a group address in the IP header destination address field. Multicast addresses use the Class D address format.

Individual hosts can join or leave a multicast group at any time. There are no restrictions on the physical location or the number of members in a multicast group. A host can be a member of more than one multicast group at any time and does not have to belong to a group to send packets to members of a group.

Routers use a group membership protocol to learn about the presence of group members on directly attached subnetworks. When a host joins a multicast group, it transmits a group membership protocol message for the group or groups that it wants to receive and sets its IP process and network interface card to receive frames addressed to the multicast group.

The Internet Multicast Backbone (MBone) is an interconnected set of subnetworks and routers that support the delivery of IP multicast traffic. The MBone is a virtual network that is layered on top of sections of the physical Internet. The MBone is composed of islands of multicast routing capability that are connected to other islands by virtual point-to-point links called tunnels. The tunnels allow multicast traffic to pass undisturbed through the parts of the Internet that are not multicast-capable. Because the MBone and the Internet have different topologies, multicast routers execute a separate routing protocol to decide how to forward multicast packets.

Multicast Addresses

Multicast host group addresses are defined to be the IP addresses whose high-order four bits are **1110**, giving an address range from **224.0.0.0** through **239.255.255.255**, or simply **224.0.0.0/4**. (These addresses also are referred to as Class D addresses.)

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP multicast groups. The base address **224.0.0.0** is reserved and cannot be assigned to any group. The block of multicast addresses from **224.0.0.1** through **224.0.0.255** is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms.

The range **239.0.0.0** through **239.255.255.255** is reserved for administratively scoped addresses. Because packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and because administratively scoped multicast addresses are locally assigned, these addresses do not need to be unique across administrative boundaries.

Multicast Redundancy

The JUNOS software supports nondisruptive, graceful Routing Engine switchover and graceful restart for some routing protocols in the case of Routing Engine or routing process failure. You can configure many routing protocols to continue to forward packets during the Routing Engine switchover or routing process restart. To support graceful Routing Engine switchover, the router must have two Routing Engines installed and be configured properly. For more information about graceful Routing Engine switchover, see the *JUNOS System Basics Configuration Guide*.

Graceful restart of routing protocol processes is required for graceful switchover. By default, graceful restart is disabled and must be enabled at the [edit routing-options graceful-restart] hierarchy level.

PIM sparse mode uses a mechanism called a *generation identifier* to indicate the need for graceful restart. Generation identifiers are included by default in PIM hello messages, as specified in the Internet draft draft-ietf-pim-sm-v2-new-10.txt. An initial generation identifier is created by each PIM neighbor to establish device capabilities. When one of the PIM neighbors restarts, it sends a new generation identifier to its neighbors. All neighbors that support graceful restart and are connected by point-to-point links assist by sending multicast updates to the restarting neighbor.

The restart phase is completed either when the PIM state becomes stable or when the restart interval timer expires. If the neighbors do not support graceful restart or if they connect to each other using multipoint interfaces, the restarting router uses a restart interval timer to define the restart period.

Graceful restart is compatible with the use of all multicast protocols. However, only PIM benefits from this feature. The router does not forward multicast packets for protocols other than PIM during graceful restart or switchover, because all other multicast protocols must completely restart after a routing process failure.

To configure graceful restart for PIM, include the `graceful-restart` statement at the `[edit routing-options]` hierarchy level, and include the `pim` statement at the `[edit protocols pim]` hierarchy level:

```
[edit routing-options]
graceful-restart;

[edit protocols]
pim {...};
```

For more information about graceful restart for PIM sparse mode, see “Configuring PIM Sparse Mode Graceful Restart” on page 223 and the *JUNOS Feature Guide*. For more information about graceful restart and other routing protocols, see the *JUNOS Routing Protocols Configuration Guide*.

Replicating Multicast Packets

The Juniper Networks T-series routing platforms handle extreme multicast packet replication requirements with a minimum of router load. Each memory component replicates a multicast packet twice at most. Even in the worst-case scenario involving maximum fan-out, when one input port and 63 output ports need a copy of the packet, the T-series routing platform only copies a multicast packet only six times. Most multicast distribution trees are much sparser, so in many cases only two or three replications are necessary. In no case does the T-series architecture have an impact on multicast performance, even with the largest multicast fan-out requirements.

