



CHAPTER 1

Overview

This chapter describes the multicast features of Cisco NX-OS.

This chapter includes the following sections:

- [Information About Multicast, page 1-1](#)
- [Licensing Requirements for Multicast, page 1-12](#)
- [High-Availability Requirements for Multicast, page 1-13](#)
- [Additional References, page 1-13](#)

Information About Multicast

IP multicast is a method of forwarding the same set of IP packets to a number of hosts within a network. You can use multicast in both IPv4 and IPv6 networks to provide efficient delivery of data to multiple destinations.



Note

Tunnel interfaces do not support Protocol-Independent Multicast (PIM).

Multicast involves both a method of delivery and discovery of senders and receivers of multicast data, which is transmitted on IP multicast addresses called groups. A multicast address that includes a group and source IP address is often referred to as a channel. The Internet Assigned Number Authority (IANA) has assigned 224.0.0.0 through 239.255.255.255 as IPv4 multicast addresses. For more information, see <http://www.iana.org/assignments/multicast-addresses>.

IPv6 multicast addresses begin with 0xFF. The IPv6 addressing architecture is defined by [RFC 4291](#). For more information about the IANA reserved addresses, see <http://www.iana.org/assignments/ipv6-multicast-addresses>.



Note


For a complete list of RFCs related to multicast, see [Appendix A, “IETF RFCs.”](#)

The routers in the network listen for receivers to advertise their interest in receiving multicast data from selected groups. The routers then replicate and forward the data from sources to the interested receivers. Multicast data for a group is transmitted only to those LAN segments with receivers that requested it.

[Figure 1-1](#) shows one source transmitting multicast data that is delivered to two receivers. In the figure, because the center host is on a LAN segment where no receiver requested multicast data, no data is delivered to that receiver.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-1 Multicast Traffic from One Source to Two Receivers



This section includes the following topics:

- [Multicast Distribution Trees, page 1-2](#)
- [Multicast Forwarding, page 1-5](#)
- [Cisco NX-OS PIM and PIM6, page 1-6](#)
- [IGMP and MLD, page 1-10](#)
- [IGMP Snooping, page 1-10](#)
- [Interdomain Multicast, page 1-10](#)
- [MRIB and M6RIB, page 1-11](#)

Multicast Distribution Trees

A multicast distribution tree represents the path that multicast data takes between the routers that connect sources and receivers. The multicast software builds different types of trees to support different multicast methods.

This section includes the following topics:


- [Source Trees, page 1-2](#)
- [Shared Trees, page 1-3](#)
- [Bidirectional Shared Trees, page 1-4](#)

Source Trees

A source tree represents the shortest path that the multicast traffic takes through the network from the sources that transmit to a particular multicast group to receivers that requested traffic from that same group. Because of the shortest path characteristic of a source tree, this tree is often referred to as a shortest path tree (SPT). [Figure 1-2](#) shows a source tree for group 224.1.1.1 that begins at host A and connects to hosts B and C.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-2 Source Tree



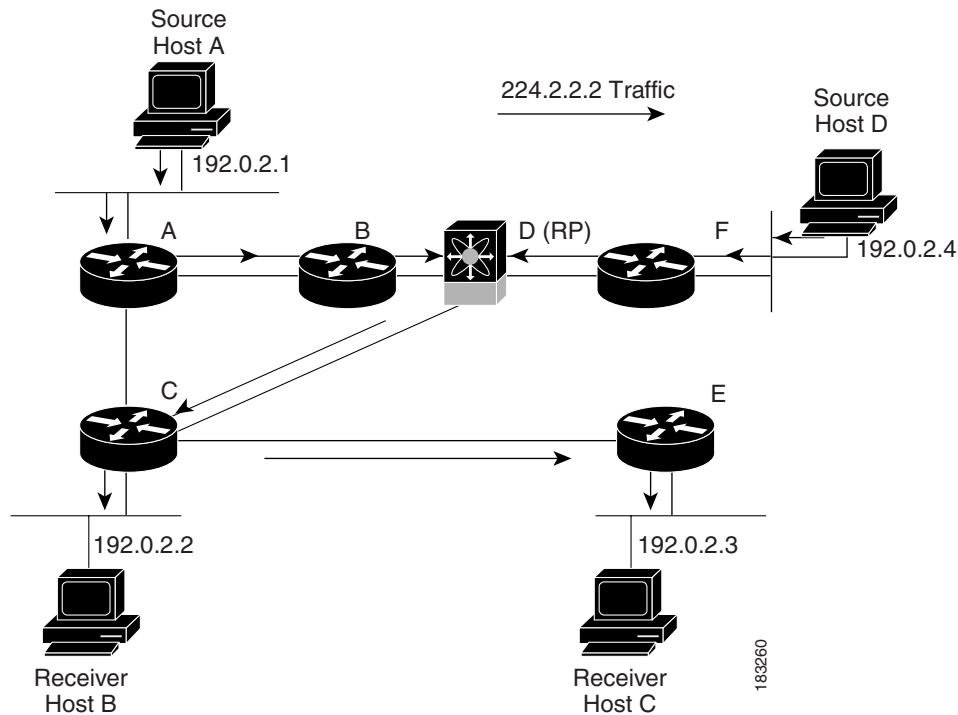
The notation (S, G) represents the multicast traffic from source S on group G. The SPT in [Figure 1-2](#) is written (192.1.1.1, 224.1.1.1). Multiple sources can be transmitting on the same group.

Shared Trees

A shared tree represents the shared distribution path that the multicast traffic takes through the network from a shared root or rendezvous point (RP) to each receiver. (The RP creates an SPT to each source.) A shared tree is also called an RP tree (RPT). [Figure 1-3](#) shows a shared tree for group 224.1.1.1 with the RP at router D. Source hosts A and D send their data to router D, the RP, which then forwards the traffic to receiver hosts B and C.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-3 Shared Tree




The notation (*, G) represents the multicast traffic from any source on group G. The shared tree in [Figure 1-3](#) is written (*, 224.2.2.2).

Bidirectional Shared Trees

A bidirectional shared tree represents the shared distribution path that the multicast traffic takes through the network from a shared root, or rendezvous point (RP), to each receiver. Multicast data is forwarded to receivers encountered on the way to the RP. The advantage of the bidirectional shared tree is shown in [Figure 1-4](#). Multicast traffic flows directly from host A to host B through routers B and C. In a shared tree, the data from source host A is first sent to the RP (router D) and then forwarded to router B for delivery to host B.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-4 Bidirectional Shared Tree



The notation (*, G) represents the multicast traffic from any source on group G. The bidirectional tree in [Figure 1-4](#) is written (*, 224.2.2.2).

Multicast Forwarding

Because multicast traffic is destined for an arbitrary group of hosts, the router uses reverse path forwarding (RPF) to route data to active receivers for the group. When receivers join a group, a path is formed either toward the source (SSM mode) or the RP (ASM or Bidir mode). The path from a source to a receiver flows in the reverse direction from the path that was created when the receiver joined the group.

For each incoming multicast packet, the router performs an RPF check. If the packet arrives on the interface leading to the source, the packet is forwarded out each interface in the outgoing interface (OIF) list for the group. Otherwise, the router drops the packet.



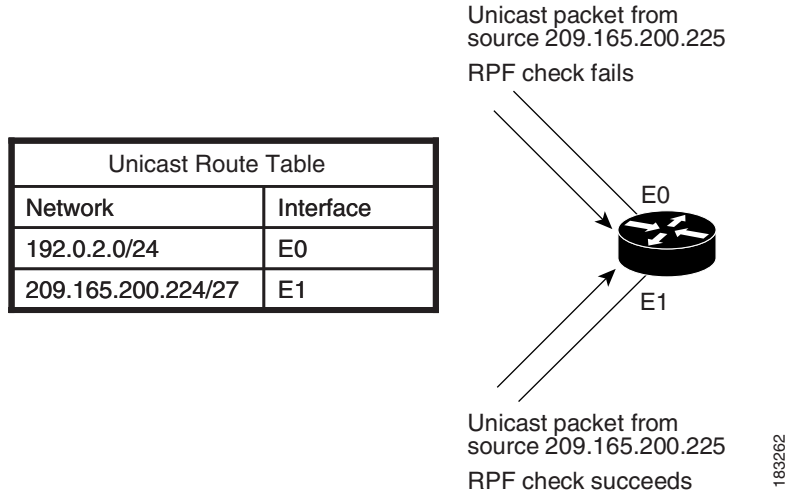
Note

In Bidir mode, if a packet arrives on a non-RPF interface, and the interface was elected as the designated forwarder (DF), then the packet is also forwarded in the upstream direction toward the RP. For more information about DFs, see the [“Designated Forwarders”](#) section on page 3-7.

[Figure 1-5](#) shows an example of RPF checks on packets coming in from different interfaces. The packet that arrives on E0 fails the RPF check because the unicast route table lists the source of the network on interface E1. The packet that arrives on E1 passes the RPF check because the unicast route table lists the source of that network on interface E1.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-5 RPF Check Example



Cisco NX-OS PIM and PIM6

Cisco NX-OS supports multicasting with Protocol Independent Multicast (PIM) sparse mode. PIM is IP routing protocol independent and can leverage whichever unicast routing protocols are used to populate the unicast routing table. In PIM sparse mode, multicast traffic is sent only to locations of the network that specifically request it. PIM dense mode is not supported by Cisco NX-OS.



Note

In this publication, the term “PIM” is used for PIM sparse mode version 2.

To access multicast commands, you must enable the PIM or PIM6 feature. Multicast is enabled only after you enable PIM or PIM6 on an interface of each router in a domain. You configure PIM for an IPv4 network and PIM6 for an IPv6 network. By default, IGMP and MLD are running on the system.

PIM, which is used between multicast-capable routers, advertises group membership across a routing domain by constructing multicast distribution trees. PIM builds shared distribution trees on which packets from multiple sources are forwarded, as well as source distribution trees, on which packets from a single source are forwarded.

The distribution trees change automatically to reflect the topology changes due to link or router failures. PIM dynamically tracks both multicast-capable sources and receivers, although the source state is not created in Bidir mode.

The router uses the unicast routing table and RPF routes for multicast to create multicast routing information. In Bidir mode, additional routing information is created.



Note

In this publication, “PIM for IPv4” and “PIM6 for IPv6” refer to the Cisco NX-OS implementation of PIM sparse mode. A PIM domain can include both an IPv4 and an IPv6 network.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-6 shows two PIM domains in an IPv4 network.

Figure 1-6 PIM Domains in an IPv4 Network




Figure 1-6 shows the following elements of PIM:

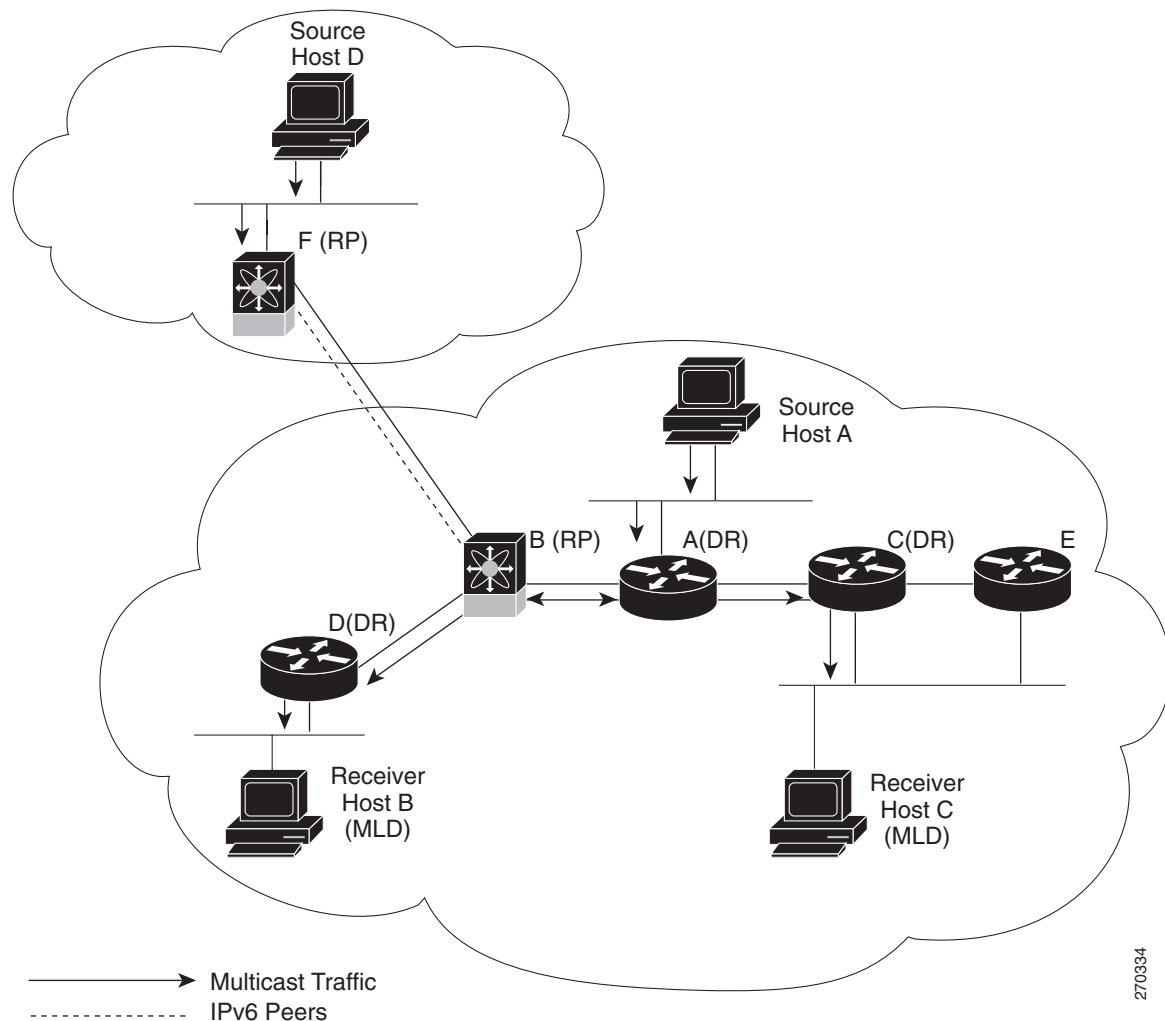
- The lines with arrows show the path of the multicast data through the network. The multicast data originates from the sources at hosts A and D.
- The dashed line connects routers B and F, which are Multicast Source Discovery Protocol (MSDP) peers. MSDP supports the discovery of multicast sources in other PIM domains.
- Hosts B and C receive multicast data by using Internet Group Management Protocol (IGMP) to advertise requests to join a multicast group.
- Routers A, C, and D are designated routers (DRs). When more than one router is connected to a LAN segment, such as C and E, the PIM software chooses one router to be the DR so that only one router is responsible for putting multicast data on the segment.

Router B is the rendezvous point (RP) for one PIM domain and router E is the RP for the other PIM domain. The RP provides a common point for connecting sources and receivers within a PIM domain.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-7 shows two PIM6 domains in an IPv6 network. In an IPv6 network, receivers that want to receive multicast data use the Multicast Listener Discovery (MLD) protocol to advertise requests to join a multicast group. MSDP, which allows for discovery of multicast sources in other PIM domains, is not supported for IPv6. You can configure IPv6 peers and use Source-Specific Multicast (SSM) and Multiprotocol BGP (MBGP) to forward multicast data between PIM6 domains. For more information, see the “Interdomain Multicast” section on page 1-10.

Figure 1-7 PIM6 Domains in an IPv6 Network



PIM supports three multicast modes for connecting sources and receivers:

- Any source multicast (ASM)
- Source-specific multicast (SSM)
- Bidirectional shared trees (Bidir)

Cisco NX-OS supports a combination of these modes for different ranges of multicast groups. You can also define RPF routes for multicast.

This section includes the following topics:

- [ASM, page 1-9](#)

Send document comments to nexus7k-docfeedback@cisco.com

- [SSM, page 1-9](#)
- [Bidir, page 1-9](#)
- [RPF Routes for Multicast, page 1-9](#)

ASM

Any Source Multicast (ASM) is a PIM tree building mode that uses shared trees to discover new sources and receivers as well as source trees to form shortest paths from receivers to sources. The shared tree uses a network node as the root, called the rendezvous point (RP). The source tree is rooted at first-hop routers, directly attached to each source that is an active sender. The ASM mode requires an RP for a group range. An RP can be configured statically or learned dynamically by the Auto-RP or BSR group-to-RP discovery protocols. If an RP is learned and is not known to be a Bidir-RP, then the group operates in ASM mode.

The ASM mode is the default mode when you configure RPs.

For information about configuring ASM, see the [“Configuring ASM and Bidir” section on page 3-16](#).

Bidir

Bidirectional shared trees (Bidir) is a PIM mode that, like the ASM mode, builds a shared tree between receivers and the RP, but does not support switching over to a source tree when a new receiver is added to a group. In the Bidir mode, the router that is connected to a receiver is called the designated forwarder because multicast data can be forwarded directly from the DR to the receiver without first going to the RP. The Bidir mode requires that you configure an RP.

The Bidir mode can reduce the amount of resources required on a router when there are many multicast sources and can continue to operate whether or not the RP is operational or connected.

For information about configuring Bidir, see the [“Configuring ASM and Bidir” section on page 3-16](#).

SSM

Source-Specific Multicast (SSM) is a PIM mode that builds a source tree that originates at the designated router on the LAN segment that receives a request to join a multicast source. Source trees are built by sending PIM join messages in the direction of the source. The SSM mode does not require you to configure RPs.

The SSM mode allows receivers to connect to sources outside the PIM domain.

For information about configuring SSM, see the [“Configuring SSM” section on page 3-28](#).

RPF Routes for Multicast

You can configure static multicast RPF routes to override what the unicast routing table uses. This feature is used when the multicast topology is different than the unicast topology.

For information about configuring RPF routes for multicast, see the [“Configuring RPF Routes for Multicast” section on page 3-30](#).

Send document comments to nexus7k-docfeedback@cisco.com

IGMP and MLD

By default, the Internet Group Management Protocol (IGMP) for PIM and Multicast Listener Discovery (MLD) for PIM6 are running on the system.

IGMP and MLD protocols are used by hosts that want to receive multicast data to request membership in multicast groups. Once the group membership is established, multicast data for the group is directed to the LAN segment of the requesting host.

You can configure IGMPv2 or IGMPv3 on an interface. You will usually configure IGMPv3 to support SSM mode. By default, the software enables IGMPv2.

You can configure MLDv1 or MLDv2 on an interface. You will usually configure MLDv2 to support SSM mode. By default, the software enables MLDv2.

For information about configuring IGMP and MLD, see [Chapter 2, “Configuring IGMP and MLD.”](#)

IGMP Snooping

IGMP snooping is a feature that limits multicast traffic on VLANs to the subset of ports that have known receivers. By examining (snooping) IGMP membership report messages from interested hosts, multicast traffic is sent only to VLAN ports that interested hosts reside on. By default, IGMP snooping is running on the system.

For information about configuring IGMP snooping, see [Chapter 4, “Configuring IGMP Snooping.”](#)

Interdomain Multicast

Cisco NX-OS provides several methods that allow multicast traffic to flow between PIM domains.

This section includes the following topics:

- [SSM, page 1-10](#)
- [MSDP, page 1-10](#)
- [MBGP, page 1-11](#)

SSM

The PIM software uses SSM to construct a shortest path tree from the designated router for the receiver to a known source IP address, which may be in another PIM domain. The ASM and Bidir modes cannot access sources from another PIM domain without the use of another protocol.

Once you enable PIM or PIM6 in your networks, you can use SSM to reach any multicast source that has an IP address known to the designated router for the receiver.

For information about configuring SSM, see the [“Configuring SSM” section on page 3-28](#).

MSDP

Multicast Source Discovery Protocol (MSDP) is a multicast routing protocol that is used with PIM to support the discovery of multicast sources in different PIM domains.

Send document comments to nexus7k-docfeedback@cisco.com

**Note**

Cisco NX-OS supports the PIM Anycast-RP, which does not require MSDP configuration. For information about PIM Anycast-RP, see the [“Configuring a PIM Anycast-RP Set” section on page 3-24](#).

For information about MSDP, see [Chapter 5, “Configuring MSDP”](#).

MBGP

Multiprotocol BGP (MBGP) defines extensions to BGP4 that enable routers to carry multicast routing information. PIM and PIM6 can use this multicast information to reach sources in external BGP Autonomous Systems.

For information about MBGP, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0*.

MRIB and M6RIB

The Cisco NX-OS IPv4 Multicast Routing Information Base (MRIB) is a repository for route information that is generated by multicast protocols such as PIM and IGMP. The MRIB does not affect the route information itself. The MRIB maintains independent route information for each virtual routing and forwarding (VRF) instance in a virtual device context (VDC). For more information about VDCs, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0*.

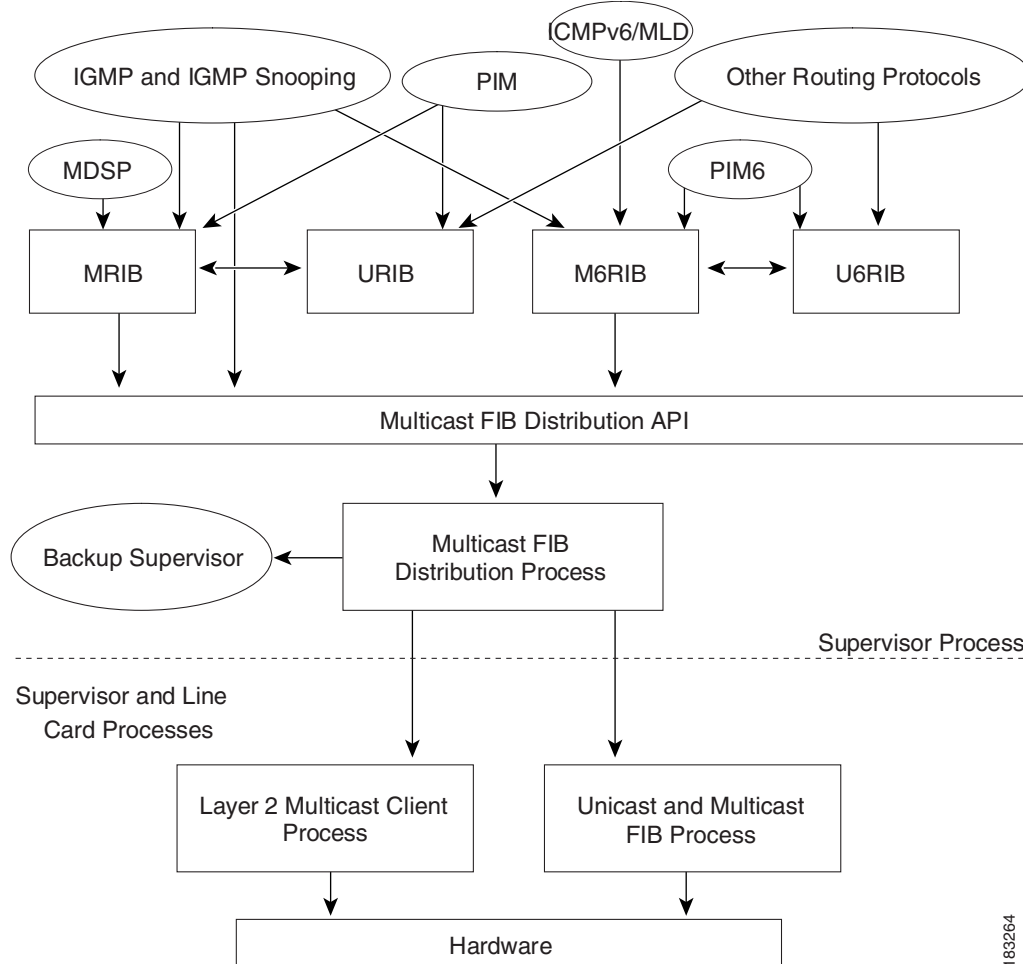
Similar to the MRIB for IPv4 routing information, the M6RIB maintains IPv6 routing information that is generated by protocols such as PIM6 and MLD.

[Figure 1-8](#) shows the major components of the Cisco NX-OS multicast software architecture:

- The Multicast FIB (MFIB and M6FIB) Distribution (MFDM) API defines an interface between the multicast Layer 2 and Layer 3 control plane modules, including the MRIB and M6RIB, and the platform forwarding plane. The control plane modules send the Layer 3 route update and Layer 2 lookup information using the MFDM API.
- The multicast FIB distribution process distributes the multicast update messages to all the relevant modules and the standby supervisor. It runs only on the supervisor.
- The Layer 2 multicast client process sets up the Layer 2 multicast hardware forwarding path. It runs on both the supervisor and the modules.
- The unicast and multicast FIB process manages the Layer 3 hardware forwarding path. It runs on both the supervisor and the modules.

Send document comments to nexus7k-docfeedback@cisco.com

Figure 1-8 Cisco NX-OS Multicast Software Architecture



Licensing Requirements for Multicast

The multicast features that require a license are as follows:

- PIM and PIM6
- MSDP

For information about multicast licensing, see the “[Licensing Requirements for PIM and PIM6](#)” section on page 3-8 and the “[Licensing Requirements for MSDP](#)” section on page 5-3.

The multicast features that require no license are as follows:

- IGMP
- MLD
- IGMP snooping

For a complete explanation of the NX-OS licensing scheme, see the *Cisco Nexus 7000 Series NX-OS Licensing Guide, Release 4.0*.

Send document comments to nexus7k-docfeedback@cisco.com

High-Availability Requirements for Multicast

After a multicast routing protocol is restarted, its state is recovered from the MRIB process. When a supervisor switchover occurs, the MRIB recovers its state from the hardware, and the multicast protocols recover their state from periodic message activity. For more information about high availability, see the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 4.0*.

Additional References

For additional information related to implementing multicast, see the following sections:

- [Related Documents, page 1-13](#)
- [Appendix A, “IETF RFCs”](#)
- [Technical Assistance, page 1-13](#)

Related Documents

Related Topic	Document Title
VDCs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference, Release 4.0</i>
CLI Commands	<i>Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference, Release 4.0</i>

Technical Assistance

Description	Link
Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/public/support/tac/home.shtml

Send document comments to nexus7k-docfeedback@cisco.com