

Overview

This chapter describes the Cisco NX-OS multicast features for Cisco Nexus 3600 platform switches.

This chapter includes the following sections:

- Licensing Requirements, on page 1
- Supported Platforms, on page 1
- About Multicast, on page 1
- General Multicast Restrictions, on page 8
- Troubleshooting Inconsistency Between SW and HW Multicast Routes , on page 8
- Additional References, on page 9

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide* and the *Cisco NX-OS Licensing Options Guide* .

Supported Platforms

Use the Nexus Switch Platform Support Matrix to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.

About Multicast

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

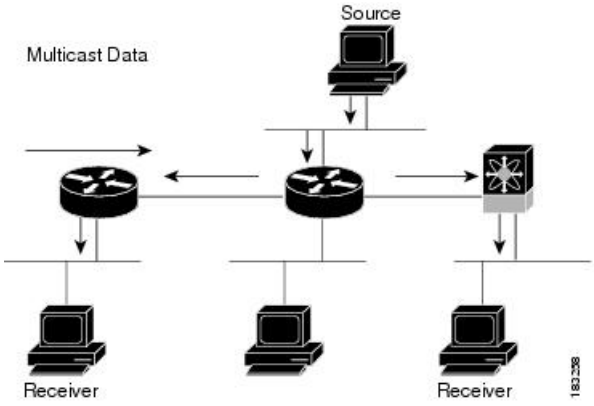
Multicast involves both a method of delivery and discovery of senders and receivers of multicast data, which is transmitted on IP multicast addresses called groups. A multicast address that includes a group and source IP address is often referred to as a channel. The Internet Assigned Number Authority (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.

Note For a complete list of RFCs related to multicast, see IETF RFCs for IP Multicast.

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.

The following figure 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.

Figure 1: Multicast Traffic from One Source to Two Receivers



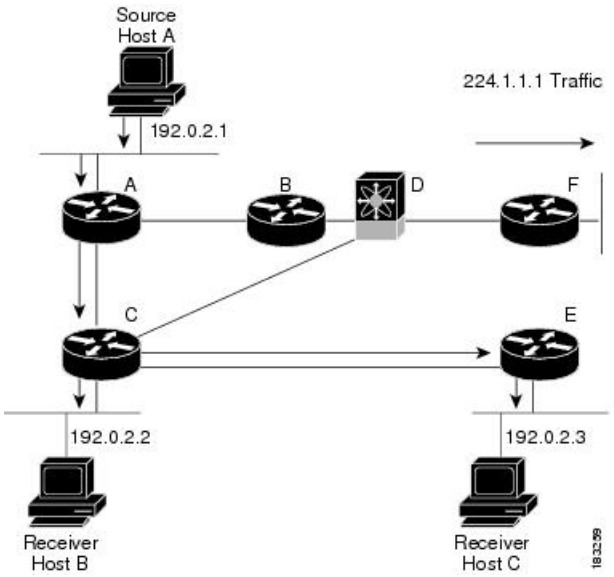
Multicast Distribution Trees

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.

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). The following figure shows a source tree for group 224.1.1.1 that begins at host A and connects to hosts B and C.

Figure 2: Source Tree

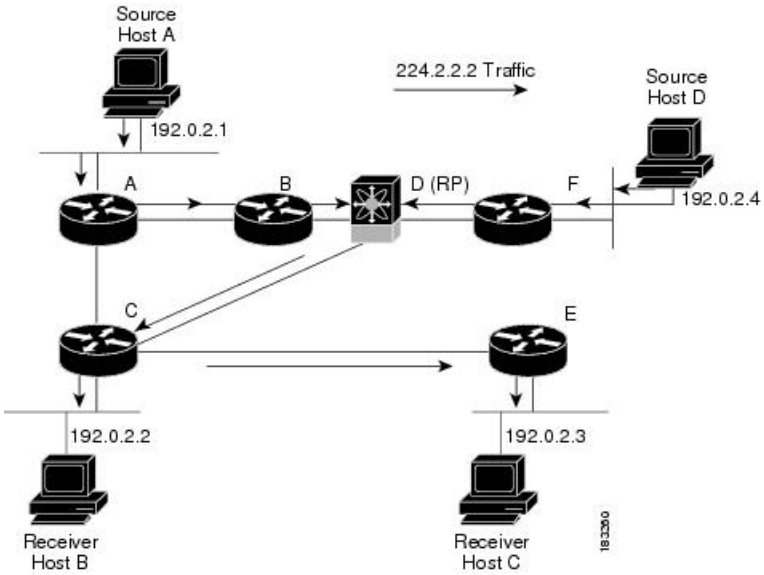


The notation (S, G) represents the multicast traffic from source S on group G. The SPT in this figure is written (192.0.2.1, 224.1.1.1). Multiple sources can be transmitting on the same group.

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). The following figure 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.

Figure 3: Shared Tree



The notation (*, G) represents the multicast traffic from any source on group G. The shared tree in Figure above is written (*, 224.2.2.2).

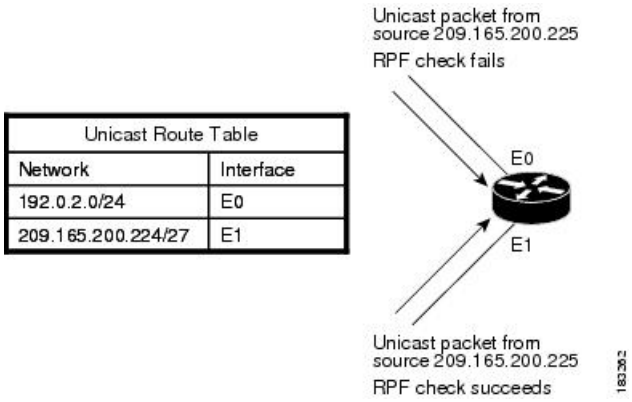
Multicast Forwarding

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 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.

The following figure 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.

Figure 4: RPF Check Example



PIM

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 feature. Multicast is enabled only after you enable PIM on an interface of each router in a domain. You configure PIM for an IPv4 network. By default, IGMP runs 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.

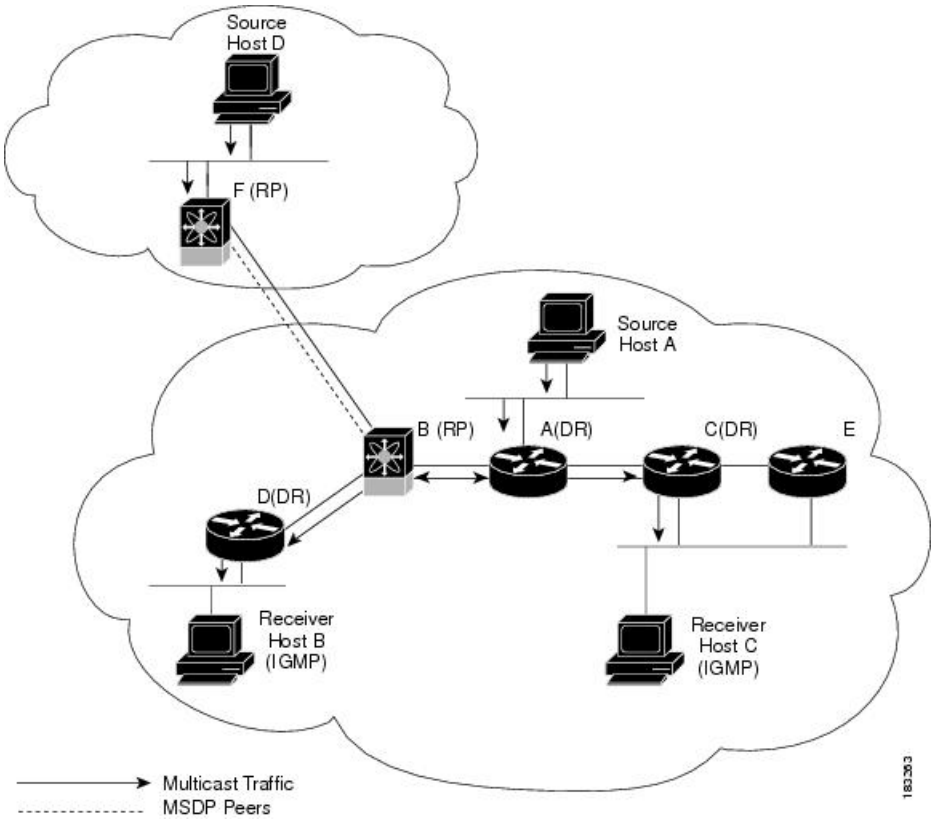
The router uses the unicast routing table and RPF routes for multicast to create multicast routing information.

The following figure shows two PIM domains in an IPv4 network.

Note In this publication, "PIM for IPv4" refer to the Cisco NX-OS implementation of PIM sparse mode. A PIM domain can include an IPv4 network.

The following figure shows two PIM domains in an IPv4 network.

Figure 5: PIM Domains in an IPv4 Network



- 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.
- 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 F is the RP for the other PIM domain. The RP provides a common point for connecting sources and receivers within a PIM domain.

PIM supports these multicast modes for connecting sources and receivers:

- Any source multicast (ASM)
- Source-Specific Multicast (SSM)

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

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.

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

For information about configuring ASM, see the Configuring ASM section.

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.

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.

IGMP

By default, the Internet Group Management Protocol (IGMP) for PIM is running on the system.

The IGMP protocol is 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.

For information about configuring IGMP, see Configuring IGMP.

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 Configuring IGMP Snooping.

Interdomain Multicast

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

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 mode cannot access sources from another PIM domain without the use of another protocol.

Once you enable PIM 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 (PIM) section.

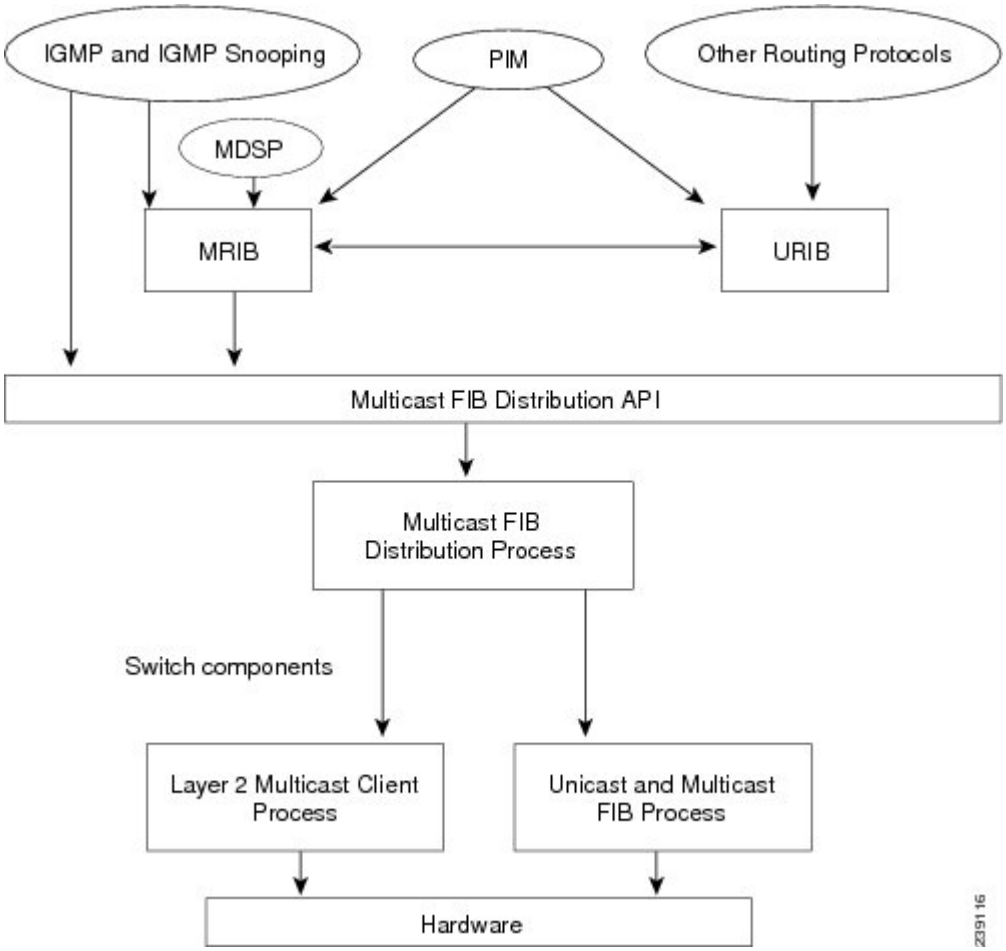
MRIB

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.

The following figure shows the major components of the Cisco NX-OS multicast software architecture:

- The Multicast FIB (MFIB) Distribution (MFDM) API defines an interface between the multicast Layer 2 and Layer 3 control plane modules, including the MRIB, 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 the switch.
- The Layer 2 multicast client process sets up the Layer 2 multicast hardware forwarding path.
- The unicast and multicast FIB process manages the Layer 3 hardware forwarding path.

Figure 6: Cisco NX-OS Multicast Software Architecture



General Multicast Restrictions

The following are the guidelines and limitations for multicast on Cisco Nexus 3600 platform switches:

- Cisco Nexus 3600 platform switches do not support Pragmatic General Multicast (PGM).
- Layer 3 IPv6 multicast routing is not supported.
- Layer 2 IPv6 multicast packets are flooded on the incoming VLAN.
- Multicast counters are not supported on the Cisco Nexus 3600 platform switches.

Troubleshooting Inconsistency Between SW and HW Multicast Routes

Symptom

This section provides symptoms, possible causes, and recommended actions for when *, G, or S,G entries that are seen in the MRIB with active flow, but are not programmed in MFIB.

Possible Cause

The issue can be seen when numerous active flows are received beyond the hardware capacity. This causes some of the entries not to be programmed in hardware while there is no free hardware index.

If the number of active flows are significantly reduced to free up the hardware resource, inconsistency may be seen between MRIB and MFIB for flows that were previously affected when the hardware table was full until the entry, times out, repopulates, and triggers programming.

There is currently no mechanism to walk the MRIB table and reprogram missing entries in HW after hardware resource is freed.

Corrective Action

To ensure reprogramming of the entries, use the **clear ip mroute *** command.

Additional References

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

- IETF RFCs for IP Multicast

MIBs

MIBs	MIBs Link
IP Multicast	To locate and download MIBs, go to the following: MIB