



Fast-Path Multicast Forwarding on Cisco 12000 Series Engine 2 and ISE Line Cards

Feature History

| Release | Modification |
|------------|---|
| 12.0(21)S | This feature was introduced on Cisco 12000 Series Engine 2 line cards in Cisco 12000 series Internet Routers. |
| 12.0(22)S | This feature was introduced on Cisco 12000 Series Packet-over-SONET (POS) and Channelized IP Service Engine (ISE) line cards. |
| 12.0(25)S | This feature was introduced on the Cisco 12000 Series 4-Port OC-12c/STM-4c ATM ISE and 4-Port Gigabit Ethernet ISE line cards. |
| 12.0(26)S | This feature was introduced on the Cisco 12000 Series 4-Port OC-12c/STM-4c Dynamic Packet Transport (DPT) ISE line card. |
| 12.0(27)S1 | <p>The hw-module slotnumber ip multicast hw-accelerate command was enhanced to include the optional [source-table index-number] [offset bit-number] parameters that allow you to reconfigure the maximum number of source IP addresses supported in a group in the source table of a line card and the offset bit in a source IP address used to select packets for fast-path multicast forwarding.</p> <p>The show ip psa-mds command was enhanced to include an optional summary parameter.</p> <p>The show ip hardware-mdfs rx command was enhanced to include an optional summary parameter.</p> <p>The enhanced feature was introduced on the following Cisco 12000 Series Engine 2 line cards:</p> <ul style="list-style-type: none">• 8-Port OC-3c/STM-1c POS/SDH• 16-Port OC-3c/STM-1c POS/SDH• 4-Port OC-12c/STM-4c POS/SDH• 1-Port OC-48c/STM-16c POS/SDH• 1-Port OC-48c/STM-16c DPT• 3-Port Gigabit Ethernet |



Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2004 Cisco Systems, Inc. All rights reserved.

| Release | Modification |
|------------------------|--|
| 12.0(27)S1 (continued) | <p>The enhanced feature was introduced on the following Cisco 12000 Series ISE line cards:</p> <ul style="list-style-type: none"> • 4-Port OC-3c/STM-1c ATM ISE • 4-Port OC12c/STM4c ATM ISE • 1-Port OC48c/STM16c POS ISE • 1-Port Channelized OC-12/STM-4 (DS1/E1) ISE • 4-Port Channelized OC-12/STM-4 (DS3/E3, OC-3c/STM-1c) POS/SDH ISE • 1-Port Channelized OC-48/STM-16 (DS3/E3, OC-3c/STM-1c, OC-12c/STM-4c) POS/SDH ISE • 4-Port OC-12c/STM-4c POS/SDH ISE (Revision B) • 4-Port Gigabit Ethernet ISE • 4-Port OC-12c/STM-4c DPT ISE |
| 12.0(28)S | <p>The hw-module slotnumber ip multicast command was enhanced to include the optional [raw-queue limit <i>queue-limit</i>] parameter that allows you to set the upper threshold for the number of packets that can be queued on the multicast raw queue (MRQ) on an ISE line card.</p> |

This feature module describes the fast-path hardware-based forwarding of IP multicast packets and how to enable it on supported Engine 2 and Channelized IP Service Engine (ISE) line cards in a Cisco 12000 Series Internet Router.

This document includes the following sections:

- [Feature Overview, page 2](#)
- [Supported Platforms, page 6](#)
- [Supported Standards, MIBs, and RFCs, page 7](#)
- [Prerequisites, page 7](#)
- [Configuration Tasks, page 7](#)
- [Monitoring and Maintaining Hardware Engine-Based Multicast Forwarding, page 9](#)
- [Configuration Examples, page 10](#)
- [Command Reference, page 12](#)
- [Glossary, page 21](#)

Feature Overview

The Fast-Path Multicast Forwarding feature allows supported line cards in Cisco 12000 Series Internet Routers to forward packets in the hardware-based Layer 3 forwarding engine (fast path) instead of having multicast traffic sent to the line card's CPU for slower path processing.

Normally IP packets are switched on a line card's CPU using Multicast Distributed Fast Switching (MDFS). MDFS fast-path forwarding reduces the load on the line card's CPU for slower-path processing. Hardware engine-based multicast forwarding, however, provides even higher-speed switching than MDFS.

Because packet forwarding is performed on the line card in the hardware switching engine, higher forwarding rates are possible. Also, more CPU resources are free to perform other features and control plane processes.

The implementation of the Fast-Path Multicast Forwarding feature on Cisco 12000 Series Internet Routers is compatible with multicast routing protocols, such as Protocol Independent Multicast dense mode (PIM-DM), PIM sparse mode (PIM-SM), and source specific multicast (SSM).

By multicasting IP packets, network applications (such as voice and video distribution) use transport resources efficiently by sending only one copy of the original packet from the data source through a multicast distribution tree to hosts on IP subnetworks.

**Note**

Multicast Distributed Fast Switching (MDFS) is also known as Multicast Distributed Switching (MDS).

In IOS Release 12.0(26)S and earlier releases, the implementation of fast-path hardware forwarding of multicast packets used a hash table that supports up to 16 source IP addresses for each multicast group. Starting in IOS Release 12.0(27)S1, to improve the scalability of multicast on the Cisco 12000 Series Internet Router and make it easier for network operators to deploy the multicast virtual private network (MVPN) feature, you can use the **hw-module slotnumber ip multicast hw-accelerate [source-table index-number] [offset bit-number]** command to configure an increased number of sources in a multicast group for a line card.

When used as a service provider edge (PE) router in a large network, the Cisco 12000 Series Internet Router can support hundreds of multicast sources in a given group. Each PE router in an MVPN becomes a source in the group that represents the MVPN. For more information about how to configure and use multicast VPNs, refer to *Multicast-VPN—IP Multicast Support for MPLS VPNs* at the following location:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ftmltvpn.htm>

Benefits

Multicast Forwarding in the Fast-Path Instead of on the Line Card CPU

Fast-path multicast forwarding is performed on the hardware engine instead of on the line card CPU. As a result, the router forwards IP multicast packets at a faster rate.

Improved Scalability and Easier Deployment of Multicast Virtual Private Networks

Starting in IOS Release 12.0(27)S1, Cisco 12000 series line cards now support hardware-based fast-path multicast forwarding for large numbers of multicast sources in a given group. This enhancement makes the MVPN solution much easier to deploy for large networks.

Restrictions

Supported Line Cards

Starting in IOS Release 12.0(21)S, hardware engine-based fast-path multicast forwarding is supported on the following Cisco 12000 Series Engine 2 line cards:

- Cisco 12000 Series Eight-Port OC-3c/STM-1c POS/SDH line card
- Cisco 12000 Series Sixteen-Port OC-3c/STM-1c POS/SDH line card
- Cisco 12000 Series Four-Port OC-12c/STM-4c POS/SDH line card

- Cisco 12000 Series One-Port OC-48c/STM-16c POS/SDH line card
- Cisco 12000 Series One-Port OC-48c/STM-16c DPT line card
- Cisco 12000 Series Three-Port Gigabit Ethernet line card

In IOS Release 12.0(22)S and later releases, fast-path multicast forwarding is also supported on these Cisco 12000 Series Packet-over-SONET (POS) and Channelized IP Service Engine (ISE) line cards:

- Cisco 12000 Series Sixteen-Port OC-3c/STM-1c POS/SDH ISE line card
- Cisco 12000 Series Four-Port OC-12c/STM-4c POS/SDH ISE line card
- Cisco 12000 Series One-Port OC-48c/STM-16c POS/SDH ISE line card
- Cisco 12000 Series Sixteen-Port Channelized OC-3/STM-1 (DS3/E3) POS/SDH ISE line card
- Cisco 12000 Series Four-Port Channelized OC-12/STM-4 (DS3/E3, OC-3c/STM-1c) POS/SDH ISE line card
- Cisco 12000 Series One-Port Channelized OC-48/STM-16 (DS3/E3, OC-3c/STM-1c) POS/SDH, OC-12c/STM-4c ISE line card

In IOS Release 12.0(25)S and later releases, fast-path multicast forwarding is supported on the Cisco 12000 Series 4-Port OC-12c/STM-4c ATM ISE and 4-Port Gigabit Ethernet ISE line cards.

In IOS Release 12.0(26)S and later releases, fast-path multicast forwarding is supported on the Cisco 12000 Series 4-Port OC-12c/STM-4c Dynamic Packet Transport (DPT) ISE line card.

In IOS Release 12.0(27)S1 and later releases, fast-path multicast forwarding is supported on the following Cisco 12000 Series line cards:

- Engine 2
 - 8-Port OC-3c/STM-1c POS/SDH
 - 16-Port OC-3c/STM-1c POS/SDH
 - 4-Port OC-12c/STM-4c POS/SDH
 - 1-Port OC-48c/STM-16c POS/SDH
 - 1-Port OC-48c/STM-16c DPT
 - 3-Port Gigabit Ethernet
- IP Service Engine
 - 4-Port OC-3c/STM-1c ATM ISE
 - 4-Port OC12c/STM4c ATM ISE
 - 1-Port OC48c/STM16c POS ISE
 - 1-Port Channelized OC-12/STM-4 (DS1/E1) ISE
 - 4-Port Channelized OC-12/STM-4 (DS3/E3, OC-3c/STM-1c) POS/SDH ISE
 - 1-Port Channelized OC-48/STM-16 (DS3/E3, OC-3c/STM-1c, OC-12c/STM-4c) POS/SDH ISE
 - 4-Port OC-12c/STM-4c POS/SDH ISE (Revision B)
 - 4-Port Gigabit Ethernet ISE
 - 4-Port OC-12c/STM-4c DPT ISE

Cisco 12000 Series Engine 2 Line Card-Specific Restrictions

On Cisco 12000 Series Engine 2 line cards, fast-path multicast forwarding is not supported with:

- Label-switched paths (LSPs). Multicast traffic cannot be routed over LSPs in the hardware engine.
- Access control lists (ACLs) in the line card hardware engine.
- Unicast reverse path forwarding (RPF) enabled.
- ATM line cards.

Packets Are Sometimes Punted to the Line Card CPU

If there is a multicast source with a (*,G) entry and only one (S,G) pattern indexed in the source table, multicast packets from the source are normally forwarded in the fast path by the hardware engine. However, MDFS switching in the router may instruct the hardware engine to punt a packet to the line card CPU for forwarding, depending on the packet's (*,G) and (S,G) multicast entries.

If a packet is destined to a group for which there is more than one (S,G) entry, there is also a chance that the packet will be punted to the CPU for further processing. This possibility is reduced as the size of the source table increases.

If no multicast route is usable for a given multicast packet, the packet is punted (by default) to the line card CPU for further processing.

Only Hardware Counters for (S,G) Entries Are Supported

Although hardware counters for (*,G) and (S,G) entries are available in the hardware engine, only (S,G) entries are supported in this implementation. Because (*,G) traffic is handled in the software, (*,G) hardware counters are not provided for fast path forwarding.

Number of Source IP Addresses Supported per Multicast Group

In IOS Release 12.0(26)S and earlier releases, the maximum number of source IP addresses supported for a multicast group is 16. When the number of source addresses exceeds 16, packets from the additional source addresses are forwarded to the line card CPU for processing.

Starting in IOS Release 12.0(27)S1, you can increase the maximum number of source IP addresses supported by reconfiguring the size of the source table used to index multicast (S,G) states.

**Note**

These restrictions for the number of source IP addresses supported for a multicast group apply only to Cisco 12000 Series Engine 2 and ISE line cards. On Cisco 12000 Series Engine 4 and Engine 4 Plus (E4+) line cards, there is no restriction on the number of IP addresses supported.

Related Features and Technologies

- Multicast Distributed Fast Switching (MDFS)
- Multicast Distributed Switching (MDS)
- Multicast Virtual Private Networks (MVPNs)

Related Documents

- [FC: Cisco IOS Release 12.0 Configuration Fundamentals Configuration Guide](#)

This document provides the following information:

- Different methods of entering commands into the router and altering the user environment
- Different types of files you can manipulate on the router, such as configuration files, images, and microcode
- Tasks that allow you to maintain your router after it is configured with the network, routing, and WAN protocols
- *FR: Cisco IOS Release 12.0 Configuration Fundamentals Command Reference*
This document describes how to perform the following tasks:
 - Configure Cisco IOS software using the command-line interface (CLI) and command modes.
 - Use basic Cisco IOS commands, including user interface, file management, and system management commands.
- *Cisco Express Forwarding Overview*
This document describes Cisco Express Forwarding (CEF), an advanced, Layer 3 IP switching technology.
- *Configuring Cisco Express Forwarding*
This chapter describes the procedures for configuring CEF and distributed CEF (dCEF).
- *P1C: Cisco IOS IP and IP Routing Configuration Guide*
This document describes the Internet Protocol (IP) and the IP routing protocols supported by Cisco IOS software.
- *Source Specific Multicast (SSM) Mapping*
This document describes how to configure and use the Source Specific Multicast (SSM) Mapping feature.

Supported Platforms

- Cisco 12000 Series Internet Routers

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

<http://www.cisco.com/go/fn>

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases depends on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

No new or modified MIBs are supported by this feature.

To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

RFCs

No new or modified RFCs are supported by this feature.

Prerequisites

There are no prerequisites for using fast-path multicast forwarding in Cisco 12000 Series Internet Routers.

Configuration Tasks

See the following sections for configuration tasks for the hardware engine-based multicast forwarding feature. Each task in the list is identified as either required or optional.

- [Enabling Hardware Engine-Based IP Multicast Forwarding, page 7](#) (required)
- [Increasing the Number of Source IP Addresses in a Multicast Group, page 8](#) (optional)
- [Configuring the Threshold of the Multicast Raw Queue, page 8](#) (optional)
- [Increasing the Number of Source IP Addresses in a Multicast Group, page 8](#) (optional)

Enabling Hardware Engine-Based IP Multicast Forwarding

By default, Cisco 12000 Series Engine 2 and ISE line cards do not perform fast-path forwarding on ingress interfaces. To enable the forwarding of IP multicast packets on the fast path using the hardware engine of a supported Engine 2 or ISE line card, use the following command in global configuration mode:

| Command | Purpose |
|---|---|
| Router(config)# hw-module slot number ip multicast hw-accelerate | Enables hardware engine-based IP multicast forwarding on the specified line card. |

Configuring the Threshold of the Multicast Raw Queue

This is an optional task and applies only to Cisco 12000 Series ISE line cards.

To optimize performance, you can specify the maximum number of packets supported in the multicast raw queue (MRQ). After this threshold is reached, any packets sent to the MRQ are dropped.

To reconfigure the maximum MRQ threshold, enter the following command in global configuration mode:

| Command | Purpose |
|--|--|
| Router(config)# hw-module slotnumber ip multicast raw-queue limit queue-limit | Configures the maximum number of packets (and therefore the number of buffers used) that can be queued in the MRQ. |

Increasing the Number of Source IP Addresses in a Multicast Group

This is an optional task and applies only to Cisco 12000 Series Engine 2 and ISE line cards.

To reconfigure an increased number of source IP addresses supported in multicast fast-path hardware forwarding, enter the following commands, starting in global configuration mode:

| Command | Purpose |
|---|---|
| Router(config)# no hw-module slot number ip multicast hw-accelerate | Disables hardware engine-based multicast forwarding on the line card. |
| Router(config)# hw-module slotnumber ip multicast hw-accelerate [source-table [size index-number] [offset bit-number]] | Re-enables multicast forwarding and configures the number of source IP address entries allowed in the source table. The optional offset bit-number parameter specifies the offset bit number in source IP addresses used to index entries in the source table. For more information, see the “Example of How to Configure Source Table Parameters” section on page 10. |

Verifying Hardware Engine-Based IP Multicast Forwarding and the MRQ Threshold

Use the **show running-config** command to verify that hardware engine-based IP multicast forwarding is enabled and the currently configured threshold of the multicast raw queue on a supported line card in the Cisco 12000 Series Internet Router. In the following example, the output lines that display the hardware engine-based multicast forwarding and MRQ threshold configured for the line card in slot 1 are shown in **bold**.

```
Router# show running-config
```

```

Current configuration : 1308 bytes
!
version 12.0
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname cisco12000
!
enable password lab
!
monitor event-trace eeprom stacktrace
hw-module slot 4 ip load-sharing per-packet
hw-module slot 7 load-sharing per-packet
hw-module slot 1 ip multicast hw-accelerate source-table 16 offset 0
hw-module slot 1 ip multicast raw-queue limit 6000

```

Monitoring and Maintaining Hardware Engine-Based Multicast Forwarding

To monitor the hardware engine-based multicast forwarding on a Cisco 12000 Series Engine 2 line card, use the following EXEC command:

| Command | Purpose |
|---|---|
| Router# exec slot number show ip psa-mds group-address | Displays information about hardware engine-based forwarding for all or a particular multicast group (*,G), and all (S,G) patterns of multicast packets being forwarded in the fast path on an Engine 2 line card. Multicast packets with (S,G) patterns that collide are punted to the line card CPU (software) for forwarding and do not appear in the show command output. |

To monitor the hardware engine-based multicast forwarding on a Cisco 12000 Series ISE line card, use the following EXEC command:

| Command | Purpose |
|--|--|
| Router# exec slot number show ip hardware-mdfs rx group-address | Displays information about hardware engine-based forwarding for all or a particular multicast group (*,G), and all (S,G) patterns of multicast packets being forwarded in the fast path on an ISE line card. Multicast packets with (S,G) patterns that collide are punted to the line card CPU (software) for forwarding and do not appear in the show command output. |

For more information about the **show ip mds forwarding** command, refer to:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122sup/122csum/csum2/122csswt/xsfscmd5.htm>

Configuration Examples

Example of How to Enable Fast-Path Multicast Forwarding

The following example shows how to enable hardware engine-based fast-path multicast forwarding on the line card in slot 2:

```
Router(config)# hw-module slot 2 ip multicast hw-accelerate
```



Note

To configure hardware engine-based multicast forwarding on an E2 DPT line card, you must enter the **hw-module slot *number* ip multicast hw-accelerate** command on both side A and side B of the DPT line card.

Example of How to Configure Source Table Parameters

You may sometimes need to increase the number of source addresses that can be forwarded in the fast path on Cisco 12000 Series Engine 2 and ISE line cards, depending on the IP multicast application that you are using. One of the following conditions may apply:

- Most of the IP multicast traffic for a group has one or a few active sources sending packets at a high data rate. For these groups, the default values work well and do not need to be changed.
- In a multicast virtual private network (MVPN) service, the provider edge (PE) and provider (P) routers must transmit IP multicast traffic for default multicast distribution trees (MDTs) and Data MDTs. The number of sources in a default MDT group is equal to the number of PE routers connected to the VPN routing/forwarding (VRF) table using the default MDT group. Therefore, the number of source addresses in a service provider network is often larger than sixteen, the number of source IP addresses supported in a multicast group in IOS Release 12.0(26)S and earlier releases.

In the MVPN architecture, default MDTs create an (S,G) state with a unique source address (S) for each PE in the MVPN network for the same default MDT group address (G). As a result, multiple sources addresses (S) overlap the same default MDT group (G), making it difficult to determine the aggregate rate required on a customer VPN to support traffic forwarding on the default MDT. It is, therefore, recommended that a service provider enable hardware-based forwarding for as many overlapping source addresses as there are PE routers in the largest VPN that the service provider needs to support.

In IOS Release 12.0(27)S1 and later releases, in addition to being able to configure a larger number of sources per group for hardware-based fast-path forwarding, you can also easily determine the source addresses of the PE routers to map to different source table indices.

For example, if you have 200 PE routers, you can configure a source table size of 256 and then ensure that the loopback address on each PE router is configured differently in eight consecutive bits. You can use IP addresses 10.1.1.0 to 10.1.1.255 on PE routers as loopback addresses. Then you can enter the **hw-module slot *number* ip multicast hw-accelerate source-table size 256 offset 0** command to ensure that MVPN traffic from all PE routers is forwarded in the fast path.

This example shows how to use the optional **source-table *index-number* offset *bit-number*** parameters in the **hw-module slot *number* ip multicast hw-accelerate** command to specify the number of entries in the source table of a line card and the offset bit number used to index source IP addresses.

This example assumes a network in which the multicast source IP addresses of three CE routers are used:

```
121.0.1.1
```

122.1.129.2

123.1.131.3

The three CE routers are in the same multicast distribution tree (MDT) group. The (S,G) pattern for the three CE router source addresses is as follows:

(0x79,0x0,0b00000001,0b00000001, G)

(0x7A,0x1,0b10000001,0b00000010, G)

(0x7B,0x1,0b10000011,0b00000011, G)

Each source IP address consists of 32 bits. The following command configures the source table to have 512 entries and to use offset bit number 6 to index the (S,G) patterns of the source IP addresses:

```
Router(config)# hw-module slot 2 ip multicast hw-accelerate source-table 512 offset 6
```

9 bits in the (S,G) pattern of each source address are used to index entries in the source table. This is because for each configured number of source table entries, a unique bit pattern is used to index source IP addresses. The following shows the bit pattern used to index source addresses for each source table size.

| Source Table Size | Bit Pattern Used to Index Source IP Addresses |
|-------------------|---|
| 16 entries | 2 ⁴ 4 bits indexed |
| 256 entries | 2 ⁸ 8 bits indexed |
| 512 entries | 2 ⁹ 9 bits indexed |
| 1024 entries | 2 ¹⁰ 10 bits indexed |
| 4096 entries | 2 ¹² 12 bits indexed |

Starting with offset bit 6, up to 512 source addresses are indexed in the source table by taking bit numbers 6 through 14 in the (S,G) pattern of each 32-bit address as shown here:

```
Bit number:   31       23       15       7       0
Source address: 00000000 00000000 00000000 00000000
Offset bit 6:                                     <0 to 5>
9 bits indexed:                                     <-6 to 14->
```

The configuration command specifies that the following underlined bits are used from each (S,G) pattern for the three CE router source addresses:

(0x79,0x0,0b00000001,0b00000001, G)

(0x7A,0x1,0b10000001,0b00000010, G)

(0x7B,0x1,0b10000011,0b00000011, G)

Note that in this example, the underlined bits in the (S,G) patterns for the first two entries collide because they are identical. When two source entries collide, the first one is forwarded in the fast path; the second one is forwarded on the line card CPU. Since the third source entry does not collide, it is forwarded in the fast path.

This example shows why it is important to carefully select an offset bit number that will index entries in the source table so that the number of collisions is minimized.

If you specify 0 as the offset bit 0, as shown in the following command, different (S,G) patterns are selected and no collision occurs:

```
Router(config)# hw-module slot 2 ip multicast hw-accelerate source-table 512 offset 0
```

```
(0x79,0x0,0b00000001,0b00000001, G)
```

```
(0x7A,0x1,0b10000001,0b00000010, G)
```

```
(0x7B,0x1,0b10000011,0b00000011, G)
```



Note Before you can reconfigure the number of IP addresses supported in the source table or the offset bit value used to index source IP addresses, you must first disable fast-path multicast forwarding on the line card.

```
Router(config)# no hw-module slot 2 ip multicast hw-accelerate
```

```
Router(config)# hw-module slot 2 ip multicast hw-accelerate source-table 512 offset 6
```

Command Reference

This section documents only new or modified commands used in the Fast-Path Multicast Forwarding feature on Cisco 12000 Series Engine 2 and ISE line cards.

- [hw-module slot ip multicast, page 13](#)
- [show ip psa-mds, page 16](#)
- [show ip hardware-mdfs rx, page 18](#)

hw-module slot ip multicast

Use the **hw-module slotnumber ip multicast** command in global configuration mode to:

- Enable fast-path IP multicast packet forwarding on the hardware engine of a supported Engine 2 or ISE line card and configure the number of source addresses supported in a multicast group or the offset bit used to index source addresses.
- Configure the maximum buffer threshold for the multicast raw queue (MRQ) on an ISE line card.

To disable IP multicast packet forwarding and the other settings, use the **no** form of this command.

hw-module slot *number* **ip multicast** {**hw-accelerate source-table** [**size** *index-number*] [**offset** *bit-number*] | **raw-queue limit** *queue-limit*}

no hw-module slot *number* **ip multicast** {**hw-accelerate source-table** [**size** *index-number*] [**offset** *bit-number*] | **raw-queue limit** *queue-limit*}

| Syntax Description | | |
|---|--|---|
| <i>number</i> | | Slot number of the line card to be configured for IP multicast forwarding. |
| hw-accelerate source-table <i>index-number</i> | | (Optional) Number of source IP addresses that are indexed in the source table and supported for fast-path hardware engine-based multicast forwarding. Valid values are 16, 256, 512, 1024, and 4096. The default number is 16. |
| offset <i>bit-number</i> | | (Optional) Offset bit number in a source IP address used to start indexing source addresses in the source table for multicast forwarding. Valid values are from 0 to 27. (The maximum valid value is less than 27 for source tables with more than 16 entries.) The default offset bit number is 0. |
| raw-queue limit <i>queue-limit</i> | | (Optional) Maximum number of packets supported in the multicast raw queue on an ISE line card. Valid values are from 2048 to 58981 bytes. |

Defaults Fast-path multicast forwarding is not enabled by default on supported Engine 2 and ISE line cards. On ISE line cards, the default MRQ threshold is set to the total number of buffers carved divided by 4.

Command Modes Global configuration mode

| Command History | Release | Modification |
|-----------------|------------|---|
| | 12.0(21)S | This command was introduced on Cisco 12000 series Internet Routers. |
| | 12.0(27)S1 | This command was enhanced to include the optional [source-table <i>index-number</i>] [offset <i>bit-number</i>] parameters. |
| | 12.0(29)S | This command was enhanced to include the raw-queue limit <i>queue-limit</i> parameter. |

Usage Guidelines

By default, Cisco 12000 Series Engine 2 and ISE line cards do not perform fast-path forwarding on ingress interfaces. You must use the **hw-module slot *number* ip multicast hw-accelerate** command to enable the forwarding of IP multicast packets on the fast path using the hardware engine of a supported Engine 2 or ISE line card.

Fast-path forwarding on an Engine 2 or ISE line card is performed through a hardware lookup procedure for each IP multicast (S,G) packet that consists of two steps:

1. The router locates the table of source IP addresses that matches the group address (G) in the packet.
2. The router locates an IP source address indexed in the table that matches the source address (S) in the packet.

If an indexed entry is found in the table, the source address is used for fast-path hardware forwarding. If an entry is not found, the (S,G) packet is punted to the line card CPU for forwarding.

Although an Engine 2 or ISE line card can only forward one source address per source table in the fast path, different source addresses in the source tables of different groups can be fast-path forwarded because individual groups are not related.

Source addresses are indexed in a source table on a first-come, first-served basis. For example, if the router receives the (10.1.15.0, G1) packet, the (S,G) pattern can be indexed from 0 to 15 in the hash table. If the index calculated by the hash algorithm for the packet is 0 and if another packet's source address for group G1 is received (for example, 9.1.2.0, G1) and the index calculated for the second (S, G) pattern is also 0, the new (S,G) pattern - (9.1.2.0, G1) - is forwarded on the line card CPU. This is called a hash collision between (10.1.15.0) and (9.1.2.0) for group G1.

In IOS Release 12.0(26)S and earlier releases, the size of the source table for each multicast group is fixed at 16. An index is calculated for the source address in a packet using the same algorithm used by Cisco Express Forwarding (CEF) for multipath selection.

Starting in IOS Release 12.0(27)S1, the index is calculated for a source address as follows:

$$\text{index number} = (\text{bit-number})\% (\text{index-number}).$$

The offset bit number (*bit-number*) is used as a percentage amount and multiplied by the number of source IP addresses that can be indexed in the source table. The default value for the offset bit number is 0; the default value for the number of source table entries is 16. You can configure new values for these arguments using the **hw-module slot *number* ip multicast hw-accelerate [source-table [size *index-number*] [offset *bit-number*]]** command.

For example, if you enter the following command:

```
Router(config)# hw-module slot 2 ip multicast hw-accelerate source-table size 256 offset 8
```

the index is the third byte of the source address; for example, for 10.2.33.5 or 192.2.33.7, the index would be 33.

When using the **hw-module slot *number* ip multicast hw-accelerate** command, note the following guidelines:

- If you configure a larger value for *index-number*, the total number of groups that can be fast-path forwarded by the line card decreases.
- If multicast error debugging is switched on for a line card, hash collisions are reported in the log. To enable multicast error debugging:
 - On an Engine 2 line card, use the **debug ip psa-mds errors** command.
 - On an ISE line card, use the **debug lc mcast event** command.
- When you downgrade the Cisco 12000 Series Internet Router from IOS Release 12.0(27)S1 or later to IOS Release 12.0(26)S or earlier, the **hw-module slot *number* ip multicast hw-accelerate** command is removed and must be added in the configuration.

- You must first disable hardware engine-based IP packet multicast forwarding on a line card before you can reconfigure the number of IP addresses supported in the source table or the offset bit value used to index source IP addresses.
- On Cisco 12000 Series Engine 2 line cards, fast-path multicast forwarding is not supported with the following:
 - PSA access control lists (ACLs)
 - Unicast reverse path forwarding (RPF) enabled
 - ATM line cards

Starting in IOS Release 12.0(29)S, you can change the default threshold of the multicast raw queue (MRQ) on an ISE line card by using the **hw-module slot *number* ip multicast raw-queue limit *queue-limit*** command. This is necessary if you find that the default MRQ threshold is too large. This condition occurs when there is a high rate of multicast traffic and packets start to build up in the MRQ, resulting in the use of too many buffers before the multicast packets are tail dropped.

All multicast packets are queued in the MRQ while waiting to be processed by the multicast raw queue server (MQS). The MQS performs the replications to different egress interfaces.

To verify the threshold size of the MRQ, use the **show running-config** command in privileged EXEC mode (see [Verifying Hardware Engine-Based IP Multicast Forwarding and the MRQ Threshold](#), page 8).

Examples

This example shows how to enable and disable hardware engine-based IP packet multicast forwarding on the Cisco 12000 Series line card in slot 1:

```
Router(config)# hw-module slot 1 ip multicast hw-accelerate
Router(config)# no hw-module slot1 ip multicast hw-accelerate
```

The next example shows how to configure support for 256 source IP addresses and specify that multicast source addresses are selected using offset bit 4 in a source multicast (S,G) pattern for the line card in slot 1.

```
Router(config)# no hw-module slot 1 ip multicast hw-accelerate
Router(config)# hw-module slot 1 ip multicast hw-accelerate source-table 256 offset 4
```

The following example shows how to configure the maximum MRQ size to 5000 bytes so that the MRQ uses fewer memory buffers when an ISE line card in slot 2 processes multicast packets.

```
Router(config)# hw-module slot 2 ip multicast raw-queue limit 5000
```

Related Commands

| Command | Description |
|---|---|
| ip multicast-routing distributed | Enables multicast distributed switching (MDS) globally on all line cards in the Cisco 12000 Series Internet Router. |

show ip psa-mds

To display information about E2 hardware engine-based forwarding for all IP multicast traffic or a particular multicast group (*,G) and all (S,G) in the Cisco 12000 Series Internet Router, use the **show ip psa-mds** command in privileged EXEC mode.

```
show ip psa-mds [group-address] [summary]
```

| Syntax Description | |
|----------------------|---|
| <i>group-address</i> | (Optional) Address of a multicast group for which hardware engine-based forwarding is displayed. |
| summary | (Optional) Displays a summary of the information about fast-path multicast forwarding on an Engine 2 line card. |

Defaults No default behavior or values.

Command Modes Privileged EXEC mode

| Command History | Release | Modification |
|-----------------|------------|--|
| | 12.0(21)S | This command was introduced on Engine 2 line cards in Cisco 12000 series Internet Routers. |
| | 12.0(27)S1 | This command was enhanced to include the optional summary parameter. |

Usage Guidelines Use this command to monitor the hardware engine structures on a supported line card that forwards IP multicast packets.

The **show ip psa-mds** command is supported only on Cisco 12000 Series Engine 2 line cards. To display information about hardware engine-based forwarding on Cisco 12000 Series ISE line cards, use the **show ip hardware-mdfs rx** command.

Examples The following example shows sample output from the **show ip psa-mds** command for the Cisco 12000 Series E2 line card installed in slot 6.

```
Router# execute-on slot 6 show ip psa-mds
```

```
===== Line Card (Slot 6) =====
```

```
224.0.1.40: mcast hw hash (0x78003A80), max_states:F pad: EEEEEEEE
*: u/c loq:0x9840 oq:0x4040 on 0x0, if: 0x0, mtu: 0x2, mgid:0x84000000
```

The next example shows sample output from the **show ip psa-mds summary** command for the Cisco 12000 Series E2 line card.

```
Router# execute-on slot 6 show ip psa-mds summary
```

```
[2] (232.20.20.20,20.20.0.2): no collisions
Group 232.20.20.20 has total of 1 hardware entry with 0 collided
```

```
[0] (232.5.5.5,30.30.52.0): no collisions
[513] (232.5.5.5,30.30.50.1): no collisions
[514] (232.5.5.5,30.30.50.2): no collisions
[515] (232.5.5.5,30.30.50.3): no collisions
[516] (232.5.5.5,30.30.50.4): no collisions
[517] (232.5.5.5,30.30.50.5): no collisions
[518] (232.5.5.5,30.30.50.6): no collisions
[519] (232.5.5.5,30.30.50.7): no collisions
[520] (232.5.5.5,30.30.50.8): no collisions
[521] (232.5.5.5,30.30.50.9): no collisions
[522] (232.5.5.5,30.30.50.10): no collisions
[523] (232.5.5.5,30.30.50.11): no collisions
[524] (232.5.5.5,30.30.50.12): no collisions
[525] (232.5.5.5,30.30.50.13): no collisions
[526] (232.5.5.5,30.30.50.14): no collisions
[527] (232.5.5.5,30.30.50.15): no collisions
[528] (232.5.5.5,30.30.50.16): no collisions
[529] (232.5.5.5,30.30.50.17): no collisions
Group 232.5.5.5 has total of 18 hardware entry with 0 collided
```

Related Commands

| Command | Description |
|---|---|
| execute-on slot <i>number</i> show ip hardware-mdfs rx | Displays information about hardware engine-based forwarding on Cisco 12000 Series ISE line cards |
| execute-on slot <i>number</i> show ip mds forwarding | Displays the Multicast Forwarding Information Base (MFIB) table, forwarding information, related flags, and counters for all multicast traffic, or for a specific multicast group or source IP address. |

show ip hardware-mdfs rx

To display information about ISE hardware engine-based forwarding for all IP multicast traffic or a particular multicast group (*,G) and all (S,G) in the Cisco 12000 Series Internet Router, use the **show ip hardware-mdfs rx** command in privileged EXEC mode.

```
show ip hardware-mdfs rx [group-address] [summary]
```

| Syntax Description | |
|----------------------|---|
| <i>group-address</i> | (Optional) Address of a multicast group for which hardware engine-based forwarding is displayed. |
| summary | (Optional) Displays a summary of the information about fast-path multicast forwarding received on an ISE line card. |

Defaults No default behavior or values.

Command Modes Privileged EXEC mode

| Command History | Release | Modification |
|-----------------|------------|---|
| | 12.0(21)S | This command was introduced on ISE line cards in Cisco 12000 series Internet Routers. |
| | 12.0(27)S1 | This command was enhanced to include the optional summary parameter. |

Usage Guidelines Use this command to monitor the hardware engine structures on a supported line card that forwards IP multicast packets.

The **show ip hardware-mdfs rx** command is supported only on Cisco 12000 Series ISE line cards. To display information about hardware engine-based forwarding on Cisco 12000 Series E2 line cards, use the **show ip psa-mds** command.

Examples The following example shows sample output from the **show ip hardware-mdfs rx** command for the Cisco 12000 Series ISE line card installed in slot 2.

```
Router# exec slot 2 show ip hardware-mdfs rx

The vrf is Default
RX side
(*, 0.0.0.0) [MGID 0x0] [DEF] [Addr - 0x702125C0] [Def_SG 0x40020B3D]
  [max_st 0x3FF] [RPF 0xFFFF] [slotmask 0x0] [ipmtu=2]
[LOQ 0xC000] [OQ 0x801F] - Hardcoded - Not stored in Adjacency
Counters: Pkt 0, Byte 0
DEF:(0.0.0.0,G) [MGID 0x0] [DEF] [Addr - 0x7020B3D0] [Coll Ctr Addr 0x7819E110]
  RPF is tunnel = FALSE
  [      ] [RPF 0xFFFF] [slotmask 0x0] [ipmtu=2] [SPD Non RPF Intrfaces
Bit Drop Vector] 0
  [LOQ 0xC000] [OQ 0x801F] - Hardcoded - Not stored in Adjacency
  [DMA_ADDR 0x67844]
```

```

Counters: Pkt 30023, Byte 3680644
Collison: Pkt 30023, Bytes 3560552  RP_Swt: Pkt 0, Bytes 0
-----
RX side
(*, 232.200.200.1) [MGID 0x63] [Addr - 0x7020F500] [Def_SG 0x40020744]
  [max_st 0x3FF] [RPF 0x1C] [slotmask 0x0] [ipmtu=2]
  [LOQ 0xC000] [OQ 0x801F] - Hardcoded - Not stored in Adjacency
Counters: Pkt 0, Byte 0
DEF:(0.0.0.0,G) [MGID 0x0] [DEF] [Addr - 0x70207440] [Coll Ctr Addr 0x7819CBE0]
  RPF is tunnel = FALSE
  [      ] [RPF 0xFFFF] [slotmask 0x0] [ipmtu=2] [SPD Non RPF Intrfaces
Bit Drop Vector] 0
  [LOQ 0xC000] [OQ 0x801F] - Hardcoded - Not stored in Adjacency
  [DMA_ADDR 0x672F8]
Counters: Pkt 1, Byte 56
Collison: Pkt 1, Bytes 52  RP_Swt: Pkt 0, Bytes 0
261. (5.5.5.5,G) [MGID 0x71] [Addr - 0x703FBC90] [Coll Ctr Addr 0x781B2050]
  RPF is tunnel = FALSE
  [V      ] [RPF 0x1C] [slotmask 0x10] [ipmtu=1514] [SPD Non RPF
Intrfaces Bit Drop Vector] 0
  [LOQ 0xC000] [OQ 0x801F] - Hardcoded - Not stored in Adjacency
  [DMA_ADDR 0x6C814]
Counters: Pkt 6415, Byte 417434
Collison: Pkt 1, Bytes 70  RP_Swt: Pkt 0, Bytes 0
771. (3.3.3.3,G) [MGID 0x6D] [Addr - 0x703FB9C0] [Coll Ctr Addr 0x781B1F60]
  RPF is tunnel = FALSE
  [V      ] [RPF 0x1B] [slotmask 0x2] [ipmtu=4470] [SPD Non RPF
Intrfaces Bit Drop Vector] 0
  [LOQ 0xC000] [OQ 0x801F] - Hardcoded - Not stored in Adjacency
  [DMA_ADDR 0x6C7D8]
Counters: Pkt 0, Byte 0
Collison: Pkt 0, Bytes 0  RP_Swt: Pkt 0, Bytes 0

```

The next example shows sample output from the **show ip hardware-mdfs rx summary** command for the fast-path multicast traffic received on a Cisco 12000 Series ISE line card installed in slot 2.

```
Router# exec slot 2 show ip hardware-mdfs rx summary
```

```
The vrf is Default
```

```
RX side
```

```
RX side
```

```
[261] (5.5.5.5,232.200.200.1): no collisions
[771] (3.3.3.3,232.200.200.1): no collisions
Group 232.200.200.1 has total of 2 hardware entry with 0 collided
```

```
RX side
```

```
[261] (5.5.5.5,232.100.100.1): no collisions
[771] (3.3.3.3,232.100.100.1): no collisions
Group 232.100.100.1 has total of 2 hardware entry with 0 collided
```

```
RX side
```

```
[771] (3.3.3.3,232.111.111.9): no collisions
Group 232.111.111.9 has total of 1 hardware entry with 0 collided
```

```
RX side
```

```
[771] (3.3.3.3,232.111.111.8): no collisions
Group 232.111.111.8 has total of 1 hardware entry with 0 collided
```

```
RX side
```

```
[771] (3.3.3.3,232.111.111.11): no collisions
Group 232.111.111.11 has total of 1 hardware entry with 0 collided
```

```
RX side
```

```
[261] (5.5.5.5,232.111.111.10): no collisions
```

■ show ip hardware-mdfs rx

```
[771] (3.3.3.3,232.111.111.10): no collisions
Group 232.111.111.10 has total of 2 hardware entry with 0 collided
```

Related Commands

| Command | Description |
|---|---|
| execute-on slot <i>number</i> show ip psa-mds | Displays information about hardware engine-based forwarding on Cisco 12000 Series Engine 2 line cards |
| execute-on slot <i>number</i> show ip psa-mds forwarding | Displays the Multicast Forwarding Information Base (MFIB) table, forwarding information, related flags, and counters for all multicast traffic, or for a specific multicast group or source IP address. |

Glossary

ACL—Access control list. A roster of users and groups of users, along with their access rights.

ALPHA—Packet-switching ASIC on Cisco 12000 Series Engine 3 line cards.

ASIC—application specific integrated circuit. Circuit designs used by manufacturers to consolidate many chips into a single package, reducing board size and power consumption.

CEF—Cisco Express Forwarding, a high-performance IP forwarding path. Forwarding can be done in software or hardware.

DM—Dense mode. Dense mode actively attempts to send multicast data to all potential receivers (flooding) and relies upon their self-pruning (removal from group) to achieve desired distribution.

DPT—Dynamic Packet Transport. DPT rings are dual, counter-rotating fiber rings. Both fibers are used concurrently to transport both data and control traffic, and use Intelligent Protection Switching (IPS) that provides proactive performance monitoring, event detection, and rapid self-healing, and restores IP service after fiber facility or node failures.

ISE—IP Services Engine. ISE line cards for Cisco 12000 series Internet Routers provide enhanced Layer 3 capabilities for high-speed customer aggregation, backbone connectivity, and peering solutions. These line cards are available in both concatenated and channelized versions.

MDFS—Multicast Distributed Fast Switching.

MDS—Multicast Distributed Switching.

MFIB—Multicast Forwarding Information Base.

MIB—Management Information Base. Database of network management information that is used and maintained by a network management protocol such as SNMP. The value of a MIB object can be changed or retrieved using SNMP commands, usually through a network management system (NMS). MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.

MPPS—million packets per second.

MRQ—multicast raw queue.

multicast—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 group of destinations known as a multicast group, which is identified by a single IP destination group address. Multicast addressing supports the transmission of a single IP datagram to multiple hosts.

MVPN—Multicast Virtual Private Network.

PIM—Protocol Independent Multicast. A multicast routing architecture defined by the IETF that enables IP multicast routing on existing IP networks. Its key point is its independence from any underlying Unicast protocol such as OSPF or BGP.

POS—Packet-over-SONET

PSA—Packet-switching ASIC on Cisco 12000 Series Engine 2 line cards.

RP—Route processor.

SM—Spare mode. Spare mode relies upon an explicitly joining method before attempting to send multicast data to receivers of a multicast group.

VPN—Virtual private network. VPNs connect branch offices and remote users through a shared or public network, such as the Internet, and provide the same security and availability as a private network. Because VPNs use an existing shared WAN infrastructure, costs are lower and deployment is faster than

is the case with traditional private networks. A VPN can consist of sites (or systems) that are all from the same enterprise (intranet) or from different enterprises (extranet); it can consist of sites (or systems) that all attach to the same service provider backbone or to different service provider backbones.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)

Copyright © 2004 Cisco Systems, Inc. All rights reserved.