



CHAPTER 13

Configuring CEF for PFC2 and PFC3A

This chapter describes how to configure Cisco Express Forwarding (CEF) for Policy Feature Card 2 (PFC2) and PFC3A on the Catalyst 6500 series switches.

CEF for PFC2 provides IP and Internetwork Packet Exchange (IPX) unicast Layer 3 switching and IP multicast Layer 3 switching for Supervisor Engine 2, PFC2, and Multilayer Switch Feature Card 2 (MSFC2).

CEF for PFC3A provides IP unicast Layer 3 switching and IP multicast Layer 3 switching for Supervisor Engine 720, PFC3A, and Multilayer Switch Feature Card 3 (MSFC3).



Note

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.



Note

For complete information on the syntax and usage information for the supervisor engine commands that are used in this chapter, refer to the *Catalyst 6500 Series Switch Command Reference* publication.

This chapter consists of these sections:

- [Understanding How Layer 3 Switching Works](#), page 13-2
- [Default CEF for PFC2/PFC3A Configuration](#), page 13-12
- [CEF for PFC2/PFC3A Configuration Guidelines and Restrictions](#), page 13-13
- [Configuring CEF for PFC2/PFC3A on the Switch](#), page 13-14
- [Configuring the NetFlow Statistics on the Switch](#), page 13-27
- [Configuring the MLS IP-Directed Broadcasts on the Switch](#), page 13-36



Note

Supervisor Engine 1 with PFC1 and MSFC or MSFC2 provide Layer 3 switching with Multilayer Switching (MLS). See [Chapter 14, “Configuring MLS,”](#) for more information.



Note

To configure MSFC2 to support MLS on a Catalyst 5000 family switch, refer to the *Layer 3 Switching Software Configuration Guide* at this URL:
http://www.cisco.com/en/US/products/hw/switches/ps5304/products_configuration_guide_book09186a00800d67ec.html.

Understanding How Layer 3 Switching Works

These sections describe Layer 3 switching with PFC2:

- [Layer 3 Switching Overview, page 13-2](#)
- [Understanding Layer 3-Switched Packet Rewrite, page 13-2](#)
- [Understanding CEF for PFC2/PFC3A, page 13-5](#)
- [Understanding the NetFlow Statistics, page 13-11](#)

Layer 3 Switching Overview



Note

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.

Layer 3 switching allows a switch, instead of a router, to forward the IP and IPX unicast traffic and the IP multicast traffic between the VLANs. Layer 3 switching is implemented in the hardware and provides wire-speed interVLAN forwarding on the switch, rather than on MSFC2/MSFC3. Layer 3 switching requires minimal support from MSFC2/MSFC3. MSFC2/MSFC3 routes any traffic that cannot be Layer 3 switched.



Note

Layer 3 switching supports the routing protocols that are configured on MSFC2/MSFC3. Layer 3 switching does not replace the routing protocols that are configured on MSFC2/MSFC3. Layer 3 switching uses Protocol Independent Multicast (PIM) for multicast route determination.

Layer 3 switching on Catalyst 6500 series switches provides flow statistics that you can use to identify the traffic characteristics for administration, planning, and troubleshooting. Layer 3 switching uses NetFlow Data Export (NDE) to export the flow statistics. See [Chapter 16, “Configuring NDE”](#) for more information about NDE.



Note

Traffic is Layer 3 switched after being processed by the VLAN access control list (VACL) feature and the quality of service (QoS) feature.

Understanding Layer 3-Switched Packet Rewrite



Note

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.

When a packet is Layer 3 switched from a source in one VLAN to a destination in another VLAN, the switch performs a packet rewrite at the egress port that is based on information learned from MSFC2/MSFC3 so that the packets appear to have been routed by MSFC2/MSFC3.



Note

Rather than just forwarding IP multicast packets, the PFC2/PFC3A replicates them as necessary on the appropriate VLANs.

The packet rewrite alters these five fields:

- Layer 2 (MAC) destination address
- Layer 2 (MAC) source address
- Layer 3 IP Time to Live (TTL) or IPX Transport Control
- Layer 3 checksum
- Layer 2 (MAC) checksum (also called the frame checksum or FCS)



Note

Packets are rewritten with the encapsulation that is appropriate for the next-hop subnet.

If Source A and Destination B are on different VLANs and Source A sends a packet to MSFC2/MSFC3 to be routed to Destination B, the switch recognizes that the packet was sent to the Layer 2 (MAC) address of MSFC2/MSFC3.

To perform Layer 3 switching, the switch rewrites the Layer 2 frame header, changing the Layer 2 destination address to the Layer 2 address of Destination B and the Layer 2 source address to the Layer 2 address of MSFC2/MSFC3. The Layer 3 addresses remain the same.

In IP unicast and IP multicast traffic, the switch decrements the Layer 3 TTL value by 1 and recomputes the Layer 3 packet checksum. In IPX traffic, the switch increments the Layer 3 Transport Control value by 1 and recomputes the Layer 3 packet checksum. The switch recomputes the Layer 2 frame checksum and forwards (or for multicast packets, replicates as necessary) the rewritten packet to Destination B's VLAN.

These sections describe how the packets are rewritten:

- [Understanding IP Unicast Rewrite, page 13-3](#)
- [Understanding IPX Unicast Rewrite, page 13-4](#)
- [Understanding IP Multicast Rewrite, page 13-4](#)

Understanding IP Unicast Rewrite

Received IP unicast packets are (conceptually) formatted as follows:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>MSFC2/MSFC3 MAC</i>	<i>Source A MAC</i>	<i>Destination B IP</i>	<i>Source A IP</i>	<i>n</i>	<i>calculation1</i>		

After the switch rewrites an IP unicast packet, it is (conceptually) formatted as follows:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Destination B MAC</i>	<i>MSFC2/MSFC3 MAC</i>	<i>Destination B IP</i>	<i>Source A IP</i>	<i>n-1</i>	<i>calculation2</i>		

Understanding IPX Unicast Rewrite

Received IPX packets are (conceptually) formatted as follows:

Layer 2 Frame Header		Layer 3 IPX Header			Data	FCS
Destination	Source	Checksum/ IPX Length/ Transport Control	Destination Net/ Node/ Socket	Source Net/ Node/ Socket		
<i>MSFC2 MAC</i>	<i>Source A MAC</i>	<i>n</i>	<i>Destination B IPX</i>	<i>Source A IPX</i>		

After the switch rewrites an IPX packet, it is (conceptually) formatted as follows:

Layer 2 Frame Header		Layer 3 IPX Header			Data	FCS
Destination	Source	Checksum/ IPX Length/ Transport Control	Destination Net/ Node/ Socket	Source Net/ Node/ Socket		
<i>Destination B MAC</i>	<i>MSFC2 MAC</i>	<i>n+1</i>	<i>Destination B IPX</i>	<i>Source A IPX</i>		

Understanding IP Multicast Rewrite

Received IP multicast packets are (conceptually) formatted as follows:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Group G1 MAC¹</i>	<i>Source A MAC</i>	<i>Group G1 IP</i>	<i>Source A IP</i>	<i>n</i>	<i>calculation1</i>		

1. In this example, Destination B is a member of Group G1.

After the switch rewrites an IP multicast packet, it is (conceptually) formatted as follows:

Frame Header		IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Group G1 MAC</i>	<i>MSFC2/MSFC3 MAC</i>	<i>Group G1 IP</i>	<i>Source A IP</i>	<i>n-1</i>	<i>calculation2</i>		

Understanding CEF for PFC2/PFC3A

**Note**

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.

These sections describe CEF for PFC2:

- [CEF for PFC2/PFC3A Overview, page 13-5](#)
- [Understanding the Forwarding Decisions, page 13-6](#)
- [Understanding the FIB, page 13-6](#)
- [Understanding the Adjacency Table, page 13-7](#)
- [Partially and Completely Switched Multicast Flows, page 13-9](#)
- [CEF for PFC2/PFC3A Examples, page 13-10](#)

CEF for PFC2/PFC3A Overview

Supervisor Engine 2, PFC2, and MSFC2 provide Layer 3 switching with CEF for PFC2. CEF for PFC2 is permanently enabled on Supervisor Engine 2. Cisco IOS CEF is permanently enabled on MSFC2 in support of CEF for PFC2.

Supervisor Engine 720, PFC3A, and MSFC3 provide Layer 3 switching with CEF for PFC3A. CEF for PFC3A is permanently enabled on Supervisor Engine 720. Cisco IOS CEF is permanently enabled on MSFC3 in support of CEF for PFC3A.

CEF for PFC2/PFC3A works with CEF (for unicast traffic) and PIM (for multicast traffic) on MSFC2/MSFC3 to support IP, IP multicast, and IPX traffic. CEF and PIM on MSFC2/MSFC3 are enhanced to support CEF for PFC2/PFC3A. CEF for PFC2/PFC3A generates flow statistics for Layer 3-switched traffic that can be displayed at the CLI or used for NDE.

CEF for PFC2/PFC3A provides Layer 3 switching for all packets that match a complete forwarding information base (FIB) entry (see the [“Understanding the FIB” section on page 13-6](#)). CEF for PFC2/PFC3A sends all packets that match an incomplete FIB entry (one where the MAC address has not been resolved) to MSFC2/MSFC3 to be routed until MSFC2/MSFC3 resolves the MAC address.

**Note**

CEF for PFC2/PFC3A sends bridge traffic that is addressed at Layer 2 to MSFC2/MSFC3 to be processed.

**Note**

Access control lists (ACLs) and policy-based routing can cause CEF for PFC2/PFC3A to ignore the FIB when making a forwarding decision (see the [“Understanding the Forwarding Decisions” section on page 13-6](#)).

Understanding the Forwarding Decisions

CEF for PFC2/PFC3A provides Layer 3 switching that is based on the following:

- Entries in the ACL ternary content addressable memory (TCAM) for policy-based routing decisions
- Entries in the NetFlow table for TCP intercept and reflexive ACL forwarding decisions (see the “[Understanding the NetFlow Statistics](#)” section on page 13-11)
- Entries in the FIB and adjacency table for all other forwarding decisions

Enter the **show mls entry** command to display information about the entries that are used to make forwarding decisions. CEF for PFC2/PFC3A makes a forwarding decision for each packet and sends the rewrite information for each packet to the egress port, where the rewrite occurs when the packet is transmitted from the switch.

Understanding the FIB

The FIB resides in a separate TCAM. The adjacency table is stored separately in DRAM. The NetFlow table is stored separately in DRAM. The FIB, the adjacency table, and the NetFlow table do not compete with any other features for storage space.

The FIB is conceptually similar to a routing table. It maintains a mirror image of the forwarding information that is contained in the unicast and multicast routing tables on MSFC2/MSFC3. When routing or topology changes occur in the network, the unicast and multicast routing tables on MSFC2/MSFC3 are updated and those changes are reflected in the FIB. The FIB maintains next-hop address information that is based on the information in the routing tables on MSFC2/MSFC3. The FIB supports 256,000 entries, which includes 16,000 IP multicast entries (128,000 IP multicast entries on MSFC3). With reverse path forwarding (RPF) check enabled, the number of IP entries doubles (with Supervisor Engine 720, the number of IP entries remain the same).

FIB lookup uses the following criteria:

- Destination IP address for IP unicast
- Destination IPX network for IPX unicast
- Source and destination IP address for IP unicast with RPF check
- Source and destination IP address for IP multicast with RPF check



Note

Because the FIB mirrors the unicast and multicast routing tables on MSFC2/MSFC3, any commands on MSFC2/MSFC3 that change the unicast or multicast routing tables affect the FIB. Forwarding entries cannot be cleared from the Supervisor Engine 2 or Supervisor Engine 720 command-line interface (CLI).

In switches with redundant supervisor engines and MSFC2s/MSFC3s, the designated MSFC2/MSFC3 supports the FIB on the active Supervisor Engine 2 or Supervisor Engine 720. The routing protocols on the nondesignated MSFC2/MSFC3 send information to the routing protocols on the designated MSFC2/MSFC3.

Enter the **show mls entry cef** command to display the following:

- Module number of the MSFC that is supporting the FIB
- FIB entry type (receive, connected, resolved, drop, wildcard, or default)
- Destination address (IP address or IPX network)
- Destination mask
- Next-hop address (IP address or IPX network)

- Next-hop mask
- Next-hop load-sharing weight

```

Console> (enable) show mls entry cef
Mod FIB-Type Destination-IP Destination-Mask NextHop-IP Weight
-----
15 receive 0.0.0.0 255.255.255.255
15 receive 255.255.255.255 255.255.255.255
15 receive 127.0.0.0 255.255.255.255
15 receive 127.0.0.52 255.255.255.255
15 receive 127.255.255.255 255.255.255.255
15 receive 10.1.1.2 255.255.255.255
15 receive 10.1.1.0 255.255.255.255
15 receive 10.1.1.255 255.255.255.255
15 receive 10.10.1.1 255.255.255.255
15 receive 10.10.0.0 255.255.255.255
.
.
.
Console> (enable)

```

Enter the **show mls** command to display a Layer 3 switching summary:

```

Console> (enable) show mls
Total packets switched = 35254
Total bytes switched = 2256256
Total routes = 120569
Total number of Netflow entries = 120000

IP statistics flows aging time = 50 seconds
Long-duration flows aging time = 320 seconds
IP statistics flows fast aging time = 0 seconds, packet threshold = 0

IP Current flow mask is Full-Vlan flow
Netflow Data Export version: 7
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0
Destination Ifindex export is enabled
Source Ifindex export is enabled
Rate limiting is turned off, packets are bridged to router
Load balancing hash is based on source and destination IP addresses and universc
Per-prefix Stats for ALL FIB entries is Enabled
Console> (enable)

```

Understanding the Adjacency Table

For each FIB entry, CEF for PFC2/PFC3A stores Layer 2 information from the designated MSFC2/MSFC3 for adjacent nodes in the adjacency table. Adjacent nodes are nodes that are directly connected at Layer 2. To forward traffic, CEF for PFC2/PFC3A selects a route from a FIB entry, which points to an adjacency entry, and uses the Layer 2 header for the adjacent node in the adjacency table entry to rewrite the packet during Layer 3 switching. CEF for PFC2 supports 256,000 adjacency table entries. CEF for PFC3A supports 1,000,000 adjacency table entries. Only half of the adjacency table entries provide statistics.

Table 13-1 lists the adjacency types.

Table 13-1 Adjacency Types

Adjacency Type	Description
connect	Entry type that contains complete rewrite information
punt	Entry to send traffic to MSFC2/MSFC3
no r/w	Entry to send traffic to MSFC2/MSFC3 when rewrite information is incomplete
frc drp	Entry that is used to drop packets due to ARP throttling
drop, null, loopbk	Entries that are used to drop packets

Enter the **show mls entry cef adjacency** command to display the following:

- FIB information (see the “Understanding the FIB” section on page 13-6)
- Adjacency type (connect, drop, null, loopbk, frc drp, punt, no r/w)
- Next-hop MAC address
- Next-hop VLAN
- Next-hop encapsulation
- Number of packets that are transmitted to this adjacency from the associated FIB entry
- Number of bytes that are transmitted to this adjacency from the associated FIB entry

```

Console> (enable) show mls entry cef adjacency
Mod: 15
Destination-IP: 140.140.1.5      Destination-Mask: 255.255.255.255
FIB-Type: resolved

AdjType  NextHop-IP      NextHop-Mac      Vlan  Encp  Tx-Packets  Tx-Octets
-----
connect  140.140.1.5     00-00-d0-00-00-05  140  ARPA          0          0

Mod: 15
Destination-IP: 150.150.1.5      Destination-Mask: 255.255.255.255
FIB-Type: resolved

AdjType  NextHop-IP      NextHop-Mac      Vlan  Encp  Tx-Packets  Tx-Octets
-----
connect  150.150.1.5     00-00-e0-00-00-05  150  ARPA          0          0

Mod: 15
Destination-IP: 153.153.1.5      Destination-Mask: 255.255.255.255
FIB-Type: resolved

AdjType  NextHop-IP      NextHop-Mac      Vlan  Encp  Tx-Packets  Tx-Octets
-----
connect  153.153.1.5     00-00-e3-00-00-05  153  ARPA          0          0
.
.
.
Console> (enable)

```

Enter the **clear mls entry cef adjacency** command to clear the CEF adjacency information:

```

Console> (enable) clear mls entry cef adjacency
Adjacency statistics has been cleared.
Console> (enable)

```

Partially and Completely Switched Multicast Flows

Some flows might be partially Layer 3 switched instead of completely Layer 3 switched in these situations:

- MSFC2/MSFC3 is configured as a member of the IP multicast group (using the **ip igmp join-group** command) on the RPF interface of the multicast source.
- MSFC2/MSFC3 is the first-hop router to the source in PIM sparse mode (in this case, MSFC2/MSFC3 must send PIM-register messages to the rendezvous point).
- The multicast TTL threshold is configured on an egress interface for the flow.
- The multicast helper is configured on the RPF interface for the flow, and multicast to broadcast translation is required.
- Multicast tag switching is configured on an egress interface.
- Network Address Translation (NAT) is configured on an interface, and source address translation is required for the outgoing interface.

**Note**

CEF for PFC2/PFC3A provides Layer 3 switching when the extended access list deny condition on the RPF interface specifies something other than the Layer 3 source, Layer 3 destination, or IP protocol (an example is the Layer 4 port numbers).

For partially switched flows, all multicast traffic belonging to the flow reaches MSFC2/MSFC3 and is software switched for any interface that is not Layer 3 switched.

**Note**

All (*,G) flows are always partially Layer 3 switched.

PFC2/PFC3A prevents multicast traffic in flows that are completely Layer 3 switched from reaching MSFC2/MSFC3, reducing the load on MSFC2/MSFC3. The **show ip mroute** and **show mls ip multicast** commands identify completely Layer 3-switched flows with the text string “RPF-MFD.” Multicast Fast Drop (MFD) indicates that from the perspective of MSFC2/MSFC3, the multicast packet is dropped, because it is switched by the PFC2/PFC3A.

For all completely Layer 3-switched flows, PFC2/PFC3A periodically sends multicast packet and byte count statistics to MSFC2/MSFC3, because MSFC2/MSFC3 cannot record multicast statistics for completely switched flows, which it never sees. MSFC2/MSFC3 uses the statistics to update the corresponding multicast routing table entries and reset the appropriate expiration timers.

CEF for PFC2/PFC3A Examples

Figure 13-1 shows a simple IP CEF network topology. In this example, Host A is on the Sales VLAN (IP subnet 171.59.1.0), Host B is on the Marketing VLAN (IP subnet 171.59.3.0), and Host C is on the Engineering VLAN (IP subnet 171.59.2.0).

When Host A initiates an HTTP file transfer to Host C, PFC2/PFC3A uses the information in the FIB and adjacency table to forward packets from Host A to Host C.

Figure 13-1 IP CEF Example Topology

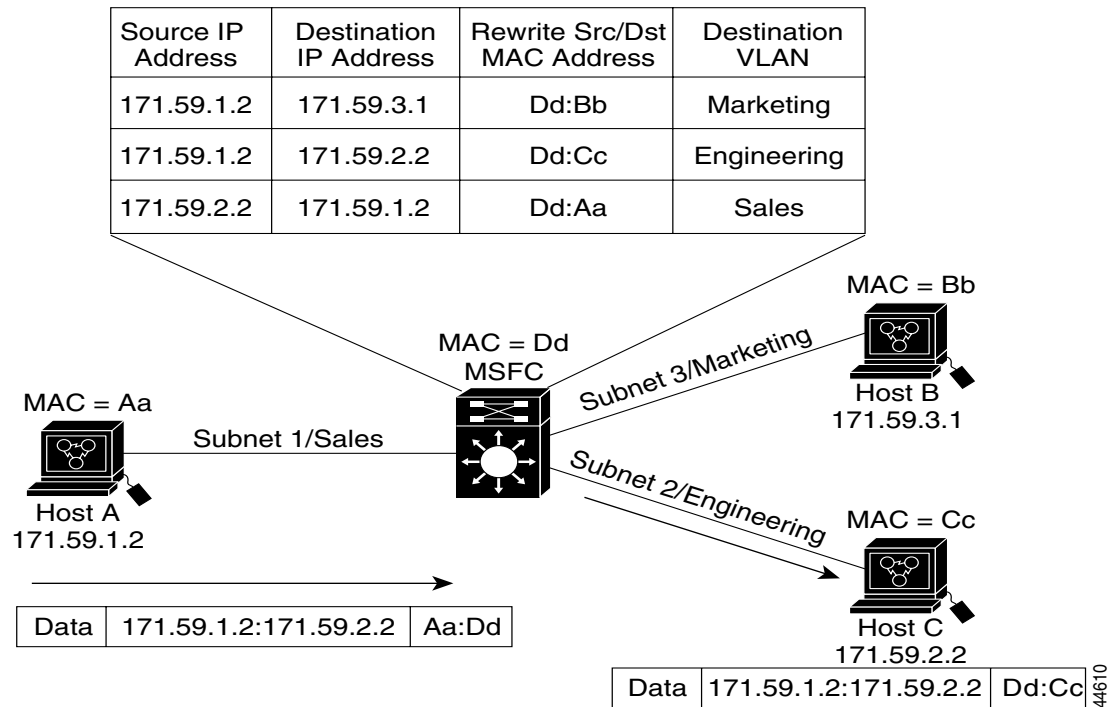
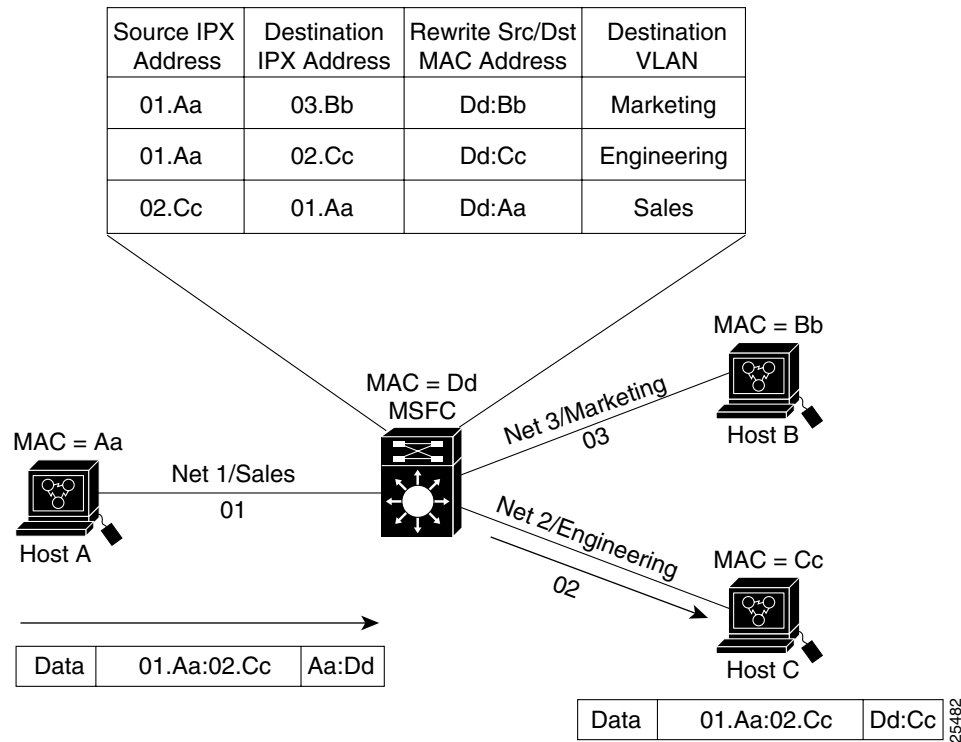


Figure 13-2 shows a simple IPX CEF network topology. In this example, Host A is on the Sales VLAN (IPX address 01.Aa), Host B is on the Marketing VLAN (IPX address 03.Bb), and Host C is on the Engineering VLAN (IPX address 02.Cc).

When Host A initiates a file transfer to Host C, PFC2 uses the information in the FIB and adjacency table to forward packets from Host A to Host C.

Figure 13-2 IPX CEF Example Topology



Understanding the NetFlow Statistics



Note

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.

These sections describe NetFlow statistics:

- [NetFlow Statistics Overview](#), page 13-11
- [NetFlow Table Entry Aging](#), page 13-12
- [Flow Masks](#), page 13-12

NetFlow Statistics Overview

CEF for PFC2/PFC3A generates flow statistics for Layer 3-switched traffic, which are stored in the NetFlow table. NetFlow statistics can be displayed with **show** commands and are also available to NetFlow Data Export (NDE).



Note

A NetFlow table with more than 32,000 entries increases the probability that there will be insufficient room to store statistics. To reduce the number of entries in the NetFlow table, you can exclude specified IP protocols from the statistics or use the least granular flow mask (see the [“Excluding the IP Protocol Entries from the NetFlow Table”](#) section on page 13-31).

NetFlow statistics support unicast and multicast flows as follows:

- A unicast flow can be any of the following:
 - Destination only: All traffic to a particular IP destination
 - Destination-source: All traffic from a particular IP source to a particular IP destination
 - Full-flow: All traffic from a particular IP source to a particular IP destination that shares the same protocol and transport-layer information
- A multicast flow is all traffic with the same protocol and transport-layer information from a particular source to the members of a particular destination multicast group.

NetFlow Table Entry Aging

The state and identity of flows are maintained while packet traffic is active; when traffic for a flow ceases, the entry ages out. You can configure the aging time for the NetFlow table entries that are kept in the NetFlow table. If an entry is not used for the specified period of time, the entry ages out and statistics for that flow can be exported to a flow collector application.

Flow Masks

Flow masks determine how the NetFlow table entries are created. CEF for PFC2 supports only one flow mask (the most specific one) for all statistics. If NetFlow for PFC2 detects different flow masks from different MSFCs for which it is performing Layer 3 switching, it changes its flow mask to the most specific flow mask detected (this applies to the PFC2/MSFC2 only).

When the flow mask changes, the entire NetFlow table is purged. When CEF for PFC2/PFC3A exports cached entries, flow records are created based on the current flow mask. Depending on the current flow mask, some fields in the flow record might not have values. Unsupported fields are filled with a zero (0).

The statistics flow masks are as follows:

- destination-ip—The least-specific flow mask for IP
- destination-ipx—The only flow mask for IPX
- source-destination-ip—For IP
- source-destination-vlan—For IP multicast
- full flow—The most-specific flow mask
- full vlan—The same fields as in full flow plus the source VLAN

Enter the **show mls statistics entry** command to display the contents of the NetFlow table and the current flow mask. Use the keyword options to display information for specific traffic (refer to the *Catalyst 6500 Series Switch Command Reference* publication for more information).

Default CEF for PFC2/PFC3A Configuration

Table 13-2 shows the default CEF for PFC2/PFC3A configuration.

Table 13-2 Default CEF for PFC2/PFC3A Configuration

Feature	Default Value
CEF for PFC2 enable state	Enabled (cannot be disabled)
CEF enable state on MSFC2/MSFC3	Enabled (cannot be disabled)
Multicast services (IGMP snooping)	Enabled
Multicast services (GMRP)	Disabled
Multicast routing on MSFC2/MSFC3	Disabled globally
PIM routing on MSFC2/MSFC3	Disabled on all interfaces
IP MMLS Threshold	Unconfigured—no default value
IP MMLS	Enabled when multicast routing is enabled and IGMP snooping is enabled

CEF for PFC2/PFC3A Configuration Guidelines and Restrictions

**Note**

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.

This section describes the guidelines and restrictions for configuring CEF for PFC2/PFC3A:

- PFC2 supports a maximum of 16 unique Hot Standby Router Protocol (HSRP) group numbers. You can use the same HSRP group numbers in different VLANs. If you configure more than 16 HSRP groups, this restriction prevents use of the VLAN number as the HSRP group number.

**Note**

Identically numbered HSRP groups use the same virtual MAC address, which might cause errors if you configure bridging on the MSFC.

- Because of the restriction to 16 unique HSRP group numbers, CEF for PFC2 cannot support the **standby use-bia** HSRP command.
- PFC3A supports 256 HSRP groups.
- CEF for PFC2 supports the following ingress and egress encapsulations:

**Note**

CEF for PFC3A supports Ethernet V2.0 (ARPA) only.

- For IP unicast:
 - Ethernet V2.0 (ARPA)
 - 802.3 with 802.2 with 1 byte control (SAP1)
 - 802.3 with 802.2 and SNAP
- For IPX:
 - Ethernet V2.0 (ARPA)
 - 802.3 (raw)
 - 802.2 with 1 byte control (SAP1)
 - SNAP



Note When the ingress encapsulation for IPX traffic is SAP1, CEF for PFC2 provides Layer 3 switching only when the egress encapsulation is also SAP1. MSFC2 routes IPX SAP1 traffic that requires an encapsulation change.

- For IP multicast—Ethernet V2.0 (ARPA)

CEF for PFC2/PFC3A does not provide Layer 3 switching for an IP multicast flow in the following cases:

- For IP multicast groups that fall into the range 224.0.0.* (where * is in the range 0–255), which is used by routing protocols. CEF for PFC2/PFC3A supports 225.0.0.* through 239.0.0.* and 224.128.0.* through 239.128.0.*.



Note Groups in the 224.0.0.* range are reserved for routing control packets and must be flooded to all forwarding ports of the VLAN. These addresses map to the multicast MAC address range 01-00-5E-00-00-xx, where xx is in the range 0–0xFF.

- For PIM auto-RP multicast groups (IP multicast group addresses 224.0.1.39 and 224.0.1.40).



Note In systems with redundant MSFC2s/MSFC3s, the PIM interface configuration must be the same on both the active and the redundant MSFC2/MSFC3.

- If the shortest-path tree (SPT) bit for the flow is cleared when running PIM sparse mode for the interface or group.
- For fragmented IP packets and packets with IP options. However, packets in the flow that are not fragmented or that do not specify IP options are multilayer switched.
- For source traffic that is received on tunnel interfaces (such as MBONE traffic).
- For any RPF interface with multicast tag switching enabled.

Configuring CEF for PFC2/PFC3A on the Switch

These sections describe how to configure CEF for PFC2/PFC3A:

- [Displaying the Layer 3-Switching Entries on the Supervisor Engine, page 13-15](#)
- [Configuring CEF on MSFC2/MSFC3, page 13-16](#)
- [Specifying CEF Maximum Routes, page 13-16](#)
- [Configuring IP Multicast on MSFC2/MSFC3, page 13-18](#)
- [Displaying IP Multicast Information, page 13-20](#)



Note For information on configuring routing on MSFC2/MSFC3, see [Chapter 12, “Configuring InterVLAN Routing.”](#)

Displaying the Layer 3-Switching Entries on the Supervisor Engine

CEF for PFC2/PFC3A is permanently enabled on Supervisor Engine 2 with PFC2 and MSFC2 and on Supervisor Engine 720 with PFC3A and MSFC3. No configuration is required.

To display all the Layer 3-switching entries on the supervisor engine, perform this task:

Task	Command
Display Layer 3-switching information.	<code>show mls entry [pbr-route] [cef] [netflow-route] [qos]</code>

This example shows how to display the Layer 3-switching entries:

```

Console> (enable) show mls entry
Mod FIB-Type Destination-IP Destination-Mask NextHop-IP Weight
-----
 15 receive 0.0.0.0 255.255.255.255
 15 receive 255.255.255.255 255.255.255.255
 15 receive 127.0.0.12 255.255.255.255
 16 receive 127.0.0.0 255.255.255.255
 16 receive 127.255.255.255 255.255.255.255
 15 resolved 127.0.0.11 255.255.255.255 127.0.0.11 1
 15 receive 21.2.0.4 255.255.255.255
 16 receive 21.0.0.0 255.255.255.255
 16 receive 21.255.255.255 255.255.255.255
 15 receive 44.0.0.1 255.255.255.255
 16 receive 44.0.0.0 255.255.255.255
 16 receive 44.255.255.255 255.255.255.255
 15 receive 42.0.0.1 255.255.255.255
 16 receive 42.0.0.0 255.255.255.255
 16 receive 42.255.255.255 255.255.255.255
 15 receive 43.0.0.99 255.255.255.255
 15 receive 43.0.0.0 255.255.255.255
 15 receive 43.255.255.255 255.255.255.255
 15 receive 192.20.20.20 255.255.255.255
 16 receive 21.2.0.5 255.255.255.255
 16 receive 42.0.0.20 255.255.255.255
 15 connected 43.0.0.0 255.0.0.0
 15 drop 224.0.0.0 240.0.0.0
 15 wildcard 0.0.0.0 0.0.0.0

Mod FIB-Type Dest-IPX-net NextHop-IPX Weight
-----
 15 connected 21
 15 connected 44
 15 connected 42
 15 resolved 450 42.0050.3EA9.ABFD 1
 15 resolved 480 42.0050.3EA9.ABFD 1
 15 wildcard 0

```

```

Destination-IP  Source-IP      Prot  DstPrnt SrcPrnt Destination-Mac  Vlan EDst Stat-Pkts
Stat-Bytes      Uptime      Age      TcpDltSeq TcpDltAck
-----
0.0.0.5         0.0.0.5       5       204     104     cc-cc-cc-cc-cc-cc 5   ARPA 0
0                01:03:18 01:00:51 cccccccc cccccccc
0.0.0.2         0.0.0.2       2       201     101     cc-cc-cc-cc-cc-cc 2   ARPA 0
0                01:03:21 01:00:51 cccccccc cccccccc
0.0.0.4         0.0.0.4       4       203     X       cc-cc-cc-cc-cc-cc 4   ARPA 0
0                01:03:19 01:00:51 cccccccc cccccccc
0.0.0.1         0.0.0.1       ICMP    200     100     cc-cc-cc-cc-cc-cc 1   ARPA 0
0                01:03:25 01:00:52 cccccccc cccccccc
0.0.0.3         0.0.0.3       3       202     102     cc-cc-cc-cc-cc-cc 3   ARPA 0
0                01:03:20 01:00:52 cccccccc cccccccc
0.0.0.6         0.0.0.6       TCP     205     105     cc-cc-cc-cc-cc-cc 6   ARPA 0
0                01:03:18 01:00:52 cccccccc cccccccc
Console> (enable)

```

Enter the **show mls entry cef** command to display only the FIB entries. Enter the **show mls entry netflow-route** command to display only the entries from the TCP intercept feature and reflexive access control lists (ACLs). Enter the **show mls entry pbr-route** command to display only the PBR entries. Enter the **show mls entry qos** command to display only the QoS entries.

Configuring CEF on MSFC2/MSFC3

CEF is permanently enabled on MSFC2/MSFC3. No configuration is required to support CEF for PFC2/PFC3A.



Note

The **ip load-sharing per-packet**, **ip cef accounting per-prefix**, and **ip cef accounting non-recursive** Cisco IOS CEF commands on MSFC2/MSFC3 apply only to traffic that is switched by CEF on MSFC/MSFC3. The commands do not affect traffic that is switched by CEF for PFC2/PFC3A on the supervisor engine.

Specifying CEF Maximum Routes



Note

This feature is only available with Supervisor Engine 720.

To specify the maximum number of routes that can be programmed in the FIB TCAM for a protocol, use the **set mls cef maximum-routes {ip | ip-multicast} routes** command. The syntax is as follows:

- **ip**—Specifies IP MLS.
- **ip-multicast**—Specifies IP multicasting MLS.
- **routes**—Specifies the number of routes that can be programmed in the FIB TCAM.

Follow these guidelines when specifying the maximum number of routes that can be programmed in the FIB TCAM:

- Routes that exceed the specified number of routes are not installed in the hardware. Packets that take those routes are switched by the MSFC. The routes argument is a unit of 1,000 entries. Setting the routes argument to 0 returns the system to a system-determined default value.
- When no protocols are set, an initial default value is assigned for each protocol. When at least one protocol is set, the default value for other unassigned protocols might change as the system tries to assign the remaining space to the unassigned protocols.

This command has the following characteristics:

- Changing the setting takes effect only after rebooting the active supervisor engine. The change does not take effect after a switchover.
- The setting on the standby supervisor engine is synchronized with the active supervisor engine. If the standby supervisor engine is inserted, both the bootup setting and new setting, if existing, on the active supervisor engine are synchronized with the standby supervisor engine. The standby supervisor engine uses the bootup setting to configure the FIB TCAM. The standby supervisor engine might need to be reset if its original bootup setting is different from the bootup setting of the active supervisor engine. An informational message (FIB_MAXROUTES_RESET) is printed on the active supervisor engine console if this situation occurs.
- To maximize the TCAM utilization, we recommend that you set the maximum routes for IP unicast as a multiple of 16,000 and set the maximum routes for IP multicast as a multiple of 8,000. The internal allocation scheme uses 16,000 as the allocation unit for unicast and 8,000 as the allocation unit for multicast. For example, if IP unicast is set to 1,000, 16,000 entries are reserved, but only 1,000 is allowed.
- When the maximum routes are exceeded or the allocated TCAM space for a protocol is full, a system message (FIB_ALLOC_TCAM_FULL) displays. Because of the internal software allocation scheme, the allocated TCAM space might be full before the maximum routes are exceeded.

**Note**

The sum of the number of maximum routes for all protocols cannot exceed 256,000.

**Note**

If the routes values for all protocols are set to 0, the bootup default is used. When you set the routes value for one protocol to a non-zero value, the default value for the other protocol changes to the remaining size.

**Note**

If the maximum number of routes is not set for an MLS protocol, a system-determined default value is shown. The default value for a protocol might not be fixed, as the system tries to assign the remaining space to the unassigned protocols. If the maximum-routes configuration is changed after bootup, the **show mls cef maximum-routes** command displays two kinds of information: one for the current (bootup) configuration and the other for the new configuration that takes effect after reboot.

To specify the maximum number of routes that can be programmed in the FIB TCAM for a protocol, perform these tasks in privileged mode:

Task	Command
Specify the maximum number of routes that can be programmed in the FIB TCAM for a protocol.	set mls cef maximum-routes {ip ip-multicast} routes
Display the maximum number of routes that are configured for each MLS protocol.	show mls cef maximum-routes

This example shows how to specify the maximum number of routes for IP unicast:

```

Console> (enable) set mls cef maximum-routes ip 220
Configuration change will take effect after next reboot.
Console> (enable) show mls cef maximum-routes
Current:
  IPv4          :192k (default)
  IPv4 multicast : 32k (default)
User configured:(effective after reboot)
  IPv4          :220k
  IPv4 multicast : 16k (adjusted default)
Console> (enable)

```

Configuring IP Multicast on MSFC2/MSFC3

These sections describe how to configure MSFC2/MSFC3 for IP multicast:

- [Enabling IP Multicast Routing Globally, page 13-18](#)
- [Enabling IP PIM on an MSFC2/MSFC3 Interface, page 13-19](#)
- [Configuring the IP MMLS Global Threshold, page 13-19](#)
- [Enabling IP MMLS on MSFC2/MSFC3 Interfaces, page 13-20](#)



Note

This section describes how to enable IP multicast routing on MSFC2/MSFC3. For more detailed IP multicast configuration information, refer to the “IP Multicast” section of the *Cisco IOS IP and IP Routing Configuration Guide* at this URL:

http://www.cisco.com/en/US/docs/ios/12_1/iproute/configuration/guide/ip_c.html

Enabling IP Multicast Routing Globally

You must enable IP multicast routing globally on MSFC2/MSFC3 before you can enable PIM on MSFC2/MSFC3 interfaces.

To enable IP multicast routing globally on MSFC2/MSFC3, perform this task in global configuration mode:

Task	Command
Enable IP multicast routing globally.	Router(config)# ip multicast-routing

This example shows how to enable IP multicast routing globally:

```
Router(config)# ip multicast-routing
Router(config)#
```

Enabling IP PIM on an MSFC2/MSFC3 Interface

You must enable PIM on MSFC2/MSFC3 interfaces before IP multicast will function on those interfaces.

To enable IP PIM on an MSFC2/MSFC3 interface, perform this task in interface configuration mode:

Task	Command
Enable IP PIM on an MSFC2/MSFC3 interface.	Router(config-if)# ip pim { dense-mode sparse-mode sparse-dense-mode }

This example shows how to enable PIM on an MSFC2/MSFC3 interface using the default mode (**sparse-dense-mode**):

```
Router(config-if)# ip pim
Router(config-if)#
```

This example shows how to enable PIM sparse mode on an MSFC2/MSFC3 interface:

```
Router(config-if)# ip pim sparse-mode
Router(config-if)#
```

Configuring the IP MMLS Global Threshold

You can configure a global multicast rate threshold, specified in packets per second, below which all multicast traffic is routed by MSFC2/MSFC3. This prevents creation of MLS entries for short-lived multicast flows, such as join requests.



Note

This command does not affect flows that are already being routed. To apply the threshold to existing routes, clear the route and let it reestablish.

To configure the IP MMLS threshold, perform this task:

Task	Command
Configure the IP MMLS threshold.	Router(config)# [no] mls ip multicast threshold <i>ppsec</i>

This example shows how to configure the IP MMLS threshold to 10 packets per second:

```
Router(config)# mls ip multicast threshold 10
Router(config)#
```

Use the **no** keyword to deconfigure the threshold.

Enabling IP MMLS on MSFC2/MSFC3 Interfaces

IP MMLS is enabled by default on the MSFC2/MSFC3 interface when you enable IP PIM on the interface. Perform this task only if you disabled IP MMLS on the interface and you want to reenabling it.



Note

You must enable IP PIM on all participating MSFC2/MSFC3 interfaces before IP MMLS will function. For information on configuring IP PIM on MSFC2/MSFC3 interfaces, see the [“Enabling IP PIM on an MSFC2/MSFC3 Interface”](#) section on page 13-19.

To enable IP MMLS on an MSFC2/MSFC3 interface, perform this task:

Task	Command
Enable IP MMLS on an MSFC2/MSFC3 interface.	Router(config-if)# [no] mls ip multicast

This example shows how to enable IP MMLS on an MSFC2/MSFC3 interface:

```
Router(config-if)# mls ip multicast
Router(config-if)#
```

Use the **no** keyword to disable IP MMLS on an MSFC2/MSFC3 interface.

Displaying IP Multicast Information

These sections describe how to display IP multicast information:

- [Displaying IP Multicast Information on MSFC2/MSFC3](#), page 13-21
- [Displaying the IP Multicast Information on the Supervisor Engine](#), page 13-24

Displaying IP Multicast Information on MSFC2/MSFC3

These sections describe displaying IP multicast information on MSFC2/MSFC3:

- [Displaying IP MMLS Interface Information, page 13-21](#)
- [Displaying the IP Multicast Routing Table, page 13-21](#)
- [Displaying IP Multicast Details, page 13-22](#)
- [Using the Debug Commands, page 13-24](#)
- [Using the Debug Commands on the SCP, page 13-24](#)

Displaying IP MMLS Interface Information

The **show ip pim interface count** command displays the IP MMLS enable state on MSFC2/MSFC3 IP PIM interfaces and the number of packets that are received and sent on the interface. The output lists the PIM interfaces that are fast switched and process switched, and the packet counts for these interfaces. An “H” is displayed on interfaces where IP MMLS is enabled.

The **show ip interface** command displays the IP MMLS enable state on an MSFC2/MSFC3 interface. To display IP MMLS information for an IP PIM MSFC2/MSFC3 interface, perform one of these tasks:

Task	Command
Display IP MMLS interface information.	Router# show ip pim interface [<i>type number</i>] count
Display the IP MMLS interface enable state.	Router# show ip interface

This example shows how to display information about the IP MMLS interfaces:

```
Router# show ip pim interface count
States: FS - Fast Switched, H - Hardware Switched

Address          Interface      FS  Mpackets In/Out
-----
192.168.10.2     Vlan10        * H 40886/0
192.168.11.2     Vlan11        * H 0/40554
192.168.12.2     Vlan12        * H 0/40554
192.168.23.2     Vlan23        *   0/0
192.168.24.2     Vlan24        *   0/0

Router#
```

Displaying the IP Multicast Routing Table

The **show ip mroute** command displays the IP multicast routing table on MSFC2/MSFC3.

To display the IP multicast routing table, perform this task:

Task	Command
Display the IP multicast routing table.	Router# show ip mroute [<i>group[source]</i>] [summary] [count] [active kbps]

This example shows how to display the IP multicast routing table:

```
Router# show ip mroute 239.252.1.1
IP Multicast Routing Table
Flags:D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
      R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT
      M - MSDP created entry, X - Proxy Join Timer Running
      A - Advertised via MSDP
Outgoing interface flags:H - Hardware switched
Timers:Uptime/Expires
Interface state:Interface, Next-Hop or VCD, State/Mode

(*, 239.252.1.1), 04:04:59/00:02:59, RP 80.0.0.2, flags:SJ
  Incoming interface:Vlan800, RPF nbr 80.0.0.2
  Outgoing interface list:
    Vlan10, Forward/Dense, 01:29:57/00:00:00, H

(22.0.0.10, 239.252.1.1), 00:00:19/00:02:41, flags:JT
  Incoming interface:Vlan800, RPF nbr 80.0.0.2, RPF-MFD
  Outgoing interface list:
    Vlan10, Forward/Dense, 00:00:19/00:00:00, H
```

Displaying IP Multicast Details

The **show mls ip multicast** command displays detailed information about IP MMLS.

To display detailed MMLS information on MSFC2/MSFC3, perform one of these tasks:

Task	Command
Display IP MMLS group information.	Router# show mls ip multicast group <i>group-address</i> [interface type number statistics]
Display IP MMLS details for all interfaces.	Router# show mls ip multicast interface <i>type number</i> [statistics summary]
Display a summary of IP MMLS information.	Router# show mls ip multicast summary
Display IP MMLS statistics.	Router# show mls ip multicast statistics
Display IP MMLS source information.	Router# show mls ip multicast source <i>ip-address</i> [interface type number statistics]

This example shows how to display IP MMLS statistics on MSFC2/MSFC3:

```
Router# show mls ip multicast statistics
MLS Multicast configuration and state:
  Router Mac:0050.0f2d.9bfd, Router IP:1.12.123.234
  MLS multicast operating state:ACTIVE
  Maximum number of allowed outstanding messages:1
  Maximum size reached from feQ:1
  Feature Notification sent:5
  Feature Notification Ack received:4
  Unsolicited Feature Notification received:0
  MSM sent:33
  MSM ACK received:33
  Delete notifications received:1
  Flow Statistics messages received:248
```

```

MLS Multicast statistics:
  Flow install Ack:9
  Flow install Nack:0
  Flow update Ack:2
  Flow update Nack:0
  Flow delete Ack:0
  Complete flow install Ack:10
  Complete flow install Nack:0
  Complete flow delete Ack:1
  Input VLAN delete Ack:4
  Output VLAN delete Ack:0
  Group delete sent:0
  Group delete Ack:0
  Global delete sent:7
  Global delete Ack:7

  L2 entry not found error:0
  Generic error :3
  LTL entry not found error:0
  MET entry not found error:0
  L3 entry exists error :0
  Hash collision error :0
  L3 entry not found error:0
  Complete flow exists error :0

```

This example shows how to display information on a specific IP MMLS entry on MSFC2/MSFC3:

```

Router# show mls ip multicast 224.1.1.1
Multicast hardware switched flows:
(1.1.13.1, 224.1.1.1) Incoming interface: Vlan13, Packets switched: 61590
Hardware switched outgoing interfaces: Vlan20 Vlan9
RFD-MFD installed: Vlan13

(1.1.9.3, 224.1.1.1) Incoming interface: Vlan9, Packets switched: 0
Hardware switched outgoing interfaces: Vlan20
RFD-MFD installed: Vlan9

(1.1.12.1, 224.1.1.1) Incoming interface: Vlan12, Packets switched: 62010
Hardware switched outgoing interfaces: Vlan20 Vlan9
RFD-MFD installed: Vlan12

(1.1.12.3, 224.1.1.1) Incoming interface: Vlan12, Packets switched: 61980
Hardware switched outgoing interfaces: Vlan20 Vlan9
RFD-MFD installed: Vlan12

(1.1.11.1, 224.1.1.1) Incoming interface: Vlan11, Packets switched: 62430
Hardware switched outgoing interfaces: Vlan20 Vlan9
RFD-MFD installed: Vlan11

(1.1.11.3, 224.1.1.1) Incoming interface: Vlan11, Packets switched: 62430
Hardware switched outgoing interfaces: Vlan20 Vlan9
RFD-MFD installed: Vlan11

Total hardware switched installed: 6
Router#

```

This example shows how to display a summary of IP MMLS information on MSFC2/MSFC3:

```

Router# show mls ip multicast summary
7 MMLS entries using 560 bytes of memory
Number of partial hardware-switched flows:2
Number of complete hardware-switched flows:5
Router#

```

Using the Debug Commands

Table 13-3 describes the IP MMLS-related debug troubleshooting commands.

Table 13-3 IP MMLS Debug Commands

Command	Description
[no] debug mls ip multicast group <i>group_id group_mask</i>	Configures filtering that applies to all other multicast debugging commands.
[no] debug mls ip multicast events	Displays the IP MMLS events.
[no] debug mls ip multicast errors	Turns on the debug messages for multicast MLS-related errors.
[no] debug mls ip multicast messages	Displays the IP MMLS messages from/to the hardware switching engine.
[no] debug mls ip multicast all	Turns on all the IP MMLS messages.
[no] debug mdss error	Turns on the Multicast Distributed Switching Services (MDSS) error messages.
[no] debug mdss events	Turns on the MDSS-related events.
[no] debug mdss all	Turns on all the MDSS messages.

Using the Debug Commands on the SCP

Table 13-4 describes the Serial Control Protocol (SCP)-related debug commands to troubleshoot the SCP that runs over the Ethernet out-of-band channel (EOBC).

Table 13-4 SCP Debug Commands

Command	Description
[no] debug scp async	Displays the trace for asynchronous data in and out of the SCP system.
[no] debug scp data	Shows the packet data trace.
[no] debug scp errors	Displays the errors and warnings in the SCP.
[no] debug scp packets	Displays the packet data in and out of the SCP system.
[no] debug scp timeouts	Reports timeouts.
[no] debug scp all	Turns on all SCP debugging messages.

Displaying the IP Multicast Information on the Supervisor Engine

These sections describe how to display the IP multicast information:

- [Displaying the IP Multicast Statistics, page 13-25](#)
- [Clearing the IP Multicast Statistics, page 13-26](#)
- [Displaying the IP Multicast Entries, page 13-26](#)

Displaying the IP Multicast Statistics

The **show mls multicast statistics** command displays the IP multicast statistics.

To display the IP multicast statistics, perform this task:

Task	Command
Display the IP multicast statistics.	show mls multicast statistics [<i>ip_addr</i>]

This example shows how to display the IP multicast statistics for MSFC2/MSFC3:

```

Console (enable) show mls multicast statistics
Router IP          Router Name      Router MAC
-----
1.1.9.254          ?                00-50-0f-06-3c-a0

Transmit:
  Delete Notifications:          23
  Acknowledgements:            92
  Flow Statistics:               56

Receive:
  Open Connection Requests:      1
  Keep Alive Messages:          72
  Shortcut Messages:            19
    Shortcut Install TLV:        8
    Selective Delete TLV:       4
  Group Delete TLV:             0
  Update TLV:                   3
  Input VLAN Delete TLV:        0
  Output VLAN Delete TLV:       0
  Global Delete TLV:            0
  MFD Install TLV:              7
  MFD Delete TLV:              0
Router IP          Router Name      Router MAC
-----
1.1.5.252          ?                00-10-29-8d-88-01

Transmit:
  Delete Notifications:          22
  Acknowledgements:            75
  Flow Statistics:               22

Receive:
  Open Connection Requests:      1
  Keep Alive Messages:          68
  Shortcut Messages:            6
    Shortcut Install TLV:        4
    Selective Delete TLV:       2
  Group Delete TLV:             0
  Update TLV:                   0
  Input VLAN Delete TLV:        0
  Output VLAN Delete TLV:       0
  Global Delete TLV:            0
  MFD Install TLV:              4
  MFD Delete TLV:              0
Console (enable)

```

Clearing the IP Multicast Statistics

The **clear mls multicast statistics** command clears the IP multicast statistics.

To clear the IP multicast statistics, perform this task in privileged mode:

Task	Command
Clear the IP multicast statistics.	clear mls multicast statistics

This example shows how to clear the IP multicast statistics:

```
Console> (enable) clear mls multicast statistics
All statistics for the MLS routers in include list are cleared.
Console> (enable)
```

Displaying the IP Multicast Entries

The **show mls multicast entry** command displays a variety of information about the multicast flows that are being handled by PFC2/PFC3A. You can display entries that are based on any combination of the participating MSFC2/MSFC3, the VLAN, the multicast group address, or the multicast traffic source.

To display information about the IP multicast entries, perform this task in privileged mode:

Task	Command
Display information about the IP multicast entries.	show mls multicast entry [[[<i>mod</i>] [<i>vlan vlan_id</i>] [<i>group ip_addr</i>] [<i>source ip_addr</i>]] [<i>all</i>]]

This example shows how to display all the IP multicast entries:

```
Console> (enable) show mls multicast entry all
Router IP      Dest IP      Source IP    Pkts      Bytes      InVlan  OutVlans
-----
1.1.5.252      224.1.1.1   1.1.11.1    15870     2761380    20
1.1.9.254      224.1.1.1   1.1.12.3    473220    82340280   12
1.1.5.252      224.1.1.1   1.1.12.3    15759     2742066    20
1.1.9.254      224.1.1.1   1.1.11.1    473670    82418580   11
1.1.5.252      224.1.1.1   1.1.11.3    15810     2750940    20
1.1.9.254      224.1.1.1   1.1.12.1    473220    82340280   12
1.1.5.252      224.1.1.1   1.1.13.1    15840     2756160    20
1.1.9.254      224.1.1.1   1.1.13.1    472770    82261980   13
1.1.5.252      224.1.1.1   1.1.12.1    15840     2756160    20
1.1.9.254      224.1.1.1   1.1.11.3    473667    82418058   11
Total Entries: 10
Console> (enable)
```

This example shows how to display the IP multicast entries for a specific MSFC2/MSFC3:

```
Console> (enable) show mls multicast entry 15
Router IP      Dest IP      Source IP      Pkts      Bytes      InVlan  OutVlans
-----
1.1.5.252     224.1.1.1   1.1.11.1      15870     2761380   20
1.1.5.252     224.1.1.1   1.1.12.3      15759     2742066   20
1.1.5.252     224.1.1.1   1.1.11.3      15810     2750940   20
1.1.5.252     224.1.1.1   1.1.13.1      15840     2756160   20
1.1.5.252     224.1.1.1   1.1.12.1      15840     2756160   20
Total Entries: 5
Console> (enable)
```

This example shows how to display the IP multicast entries for a specific multicast group address:

```
Console> (enable) show mls multicast entry group 226.0.1.3 short
Router IP      Dest IP      Source IP      InVlan Pkts      Bytes      OutVlans
-----
171.69.2.1    226.0.1.3   172.2.3.8     20      171      23512     10,201,22,45
171.69.2.1    226.0.1.3   172.3.4.9     12      25       3120      8,20
Total Entries: 2
Console> (enable)
```

This example shows how to display the IP multicast entries for a specific MSFC2/MSFC3 and a specific multicast source address:

```
Console> (enable) show mls multicast entry 15 source 1.1.11.1 short
Router IP      Dest IP      Source IP      Pkts      Bytes
InVlan  OutVlans
-----
172.20.49.159 224.1.1.6   1.1.40.4      368      57776
  40      23,25
172.20.49.159 224.1.1.71   1.1.22.2      99       65142
  22      30,37
172.20.49.159 224.1.1.8   1.1.22.2      396      235620
  22      13,19
Console> (enable)
```

Configuring the NetFlow Statistics on the Switch



Note

With Supervisor Engine 720 (MSFC3), IPX routing is done through the software.

These sections describe how to configure the NetFlow statistics:

- [Specifying NetFlow Table Entry Creation on a Per-Interface Basis, page 13-28](#)
- [Specifying the NetFlow Table Entry Aging-Time Value, page 13-29](#)
- [Specifying the NetFlow Table IP Entry Fast Aging Time and Packet Threshold Values, page 13-30](#)
- [Setting the Minimum Statistics Flow Mask, page 13-31](#)
- [Excluding the IP Protocol Entries from the NetFlow Table, page 13-31](#)
- [Displaying the NetFlow Statistics, page 13-31](#)
- [Clearing the NetFlow IP and IPX Statistics, page 13-34](#)
- [Displaying the NetFlow Statistics Debug Information, page 13-36](#)

Specifying NetFlow Table Entry Creation on a Per-Interface Basis



Note

This feature requires PFC3B, PFC3BXL or later.

With software release 8.4(1) and later releases, you can create the NetFlow table entries on a per-interface basis. This feature uses the same mechanism as the bridged-flow statistics to create flows. The NetFlow entries are created for both VLANs on which the bridged-flow statistics are enabled and on the VLANs on which NetFlow entry creation is enabled (see the [“Enabling and Disabling Bridged-Flow Statistics on VLANs”](#) section on page 16-12).

For example, if you enable the Layer 3 per-interface entry creation on VLAN 100 and 200 and at the same time you want to enable the bridged-flow statistics on VLAN 150 and 250, the NetFlow entry and the bridged-flow statistics are enabled on all four VLANs. To specify only the bridged-flow statistics for VLAN 150 and 250, you must disable the per-interface entry feature.

In addition, the bridged-flow statistics are automatically enabled when you enable the NetFlow entry creation on a per-interface basis for VLANs. The CLI allows you to disable NetFlow per interface if you do not want this overlap in the Netflow table entry creation.

The status of this feature is displayed as part of the **show mls** command. The VLANs that have entry creation enabled are displayed as part of the VLANs that have the bridged-flow statistics feature enabled.

To enable or disable the NetFlow per-interface table entries, perform this task in privileged mode:

Task	Command
Enable the NetFlow per-interface table entries.	set mls netflow-per-interface [enable disable]

This example shows how to enable the NetFlow per-interface table entries:

```
Console> (enable) set mls netflow-per-interface enable
Console> (enable)
```

You can specify the VLANs for which the NetFlow entries can be enabled or disabled. To control the flow creation on a VLAN basis, perform this task in privileged mode:

Task	Command
Enable the per-VLAN NetFlow table entries.	set mls netflow-entry-create [enable disable] <i>vlan-list</i>

This example shows how to specify the VLANs that are used to create the NetFlow table entries:

```
Console> (enable) set mls netflow-entry-create enable 150, 250
Console> (enable)
```

Specifying the NetFlow Table Entry Aging-Time Value

The entry aging time for each protocol (IP and IPX) applies to all the protocol-specific NetFlow table entries. Any entry that has not been used for *agingtime* seconds is aged out. The default is 16 seconds.

For normal aging time, you can specify the aging time in the range of 1–1092 seconds in 8-second increments. Any aging-time value that is not a multiple of 8 seconds is adjusted to the closest multiple of 8 seconds. For example, a value of 65 is adjusted to 64 and a value of 127 is adjusted to 128.

To specify the entry aging time for both IP and IPX, perform this task in privileged mode:

Task	Command
Specify the aging time for the NetFlow table entries.	set mls agingtime [<i>agingtime</i>]

This example shows how to specify the entry aging time:

```
Console> (enable) set mls agingtime 16
Multilayer switching agingtime IP and IPX set to 16
Console> (enable)
```

To specify the IP entry aging time, perform this task in privileged mode:

Task	Command
Specify the IP entry aging time for the NetFlow table.	set mls agingtime ip [<i>agingtime</i>]

This example shows how to specify the IP entry aging time:

```
Console> (enable) set mls agingtime ip 16
Multilayer switching aging time IP set to 16
Console> (enable)
```

To specify the IPX entry aging time, perform this task in privileged mode:

Task	Command
Specify the IPX entry aging time for the NetFlow table.	set mls agingtime ipx [<i>agingtime</i>]

This example shows how to specify the IPX entry aging time:

```
Console> (enable) set mls agingtime ipx 16
Multilayer switching aging time IPX set to 16
Console> (enable)
```

Specifying the NetFlow Table IP Entry Fast Aging Time and Packet Threshold Values



Note

The IPX entries do not use fast aging.

To increase the utilization of the NetFlow table, enable IP entry fast aging time. The IP entry fast aging time applies to the NetFlow table entries that have no more than *pkt_threshold* packets that are routed within *fastagingtime* seconds after they are created. A typical NetFlow table entry that is removed is the entry for flows to and from a Domain Name Server (DNS) or TFTP server; the entry might never be used again after it is created. Detecting and aging out these entries saves space in the NetFlow table for other data traffic.

The default *fastagingtime* value is 0 (no fast aging). For Supervisor Engine 1 and Supervisor Engine 2, you can configure the *fastagingtime* value from 8–128 seconds in increments of 8 seconds. For Supervisor Engine 720, you can configure the *fastagingtime* value from 0–128 seconds in increments of 1 second. Any *fastagingtime* value that is not configured exactly as the indicated values is adjusted to the closest one. For Supervisor Engine 1 and Supervisor Engine 2, you can configure the *pkt_threshold* value to 0, 1, 3, 7, 15, 31, 63, or 127 packets. For Supervisor Engine 720, you can configure the *pkt_threshold* value from 1–127 packets in increments of 1 packet.

If you need to enable IP entry fast aging time, initially set the value to 128 seconds. If the NetFlow table remains full, decrease the setting. If the NetFlow table continues to remain full, decrease the normal IP entry aging time.

To specify the IP entry fast aging time and packet threshold, perform this task in privileged mode:

Task	Command
Specify the IP entry fast aging time and packet threshold for a NetFlow table entry.	set mls agingtime fast [<i>fastagingtime</i>] [<i>pkt_threshold</i>]

This example shows how to set the IP entry fast aging time to 8 seconds with a packet threshold of 15 packets:

```
Console> (enable) set mls agingtime fast 8 15
Multilayer switching fast aging time set to 8 seconds for entries with no more than 15
packets switched.
Console> (enable)
```

You can force an active flow to age out by entering the **set mls agingtime long-duration** {*longagingtime*} command. You can specify the aging time of the active flow in the range of 64–1920 seconds in increments of 64. The default *longagingtime* is 320.

This example shows how to set the aging time for active flows:

```
Console> (enable) set mls agingtime long-duration 128
Multilayer switching agingtime set to 128 seconds for long duration flows
Console> (enable)
```

Setting the Minimum Statistics Flow Mask

You can set the minimum granularity of the flow mask for the NetFlow table. The actual flow mask will be at least of the granularity that is specified by this command. For information on how the different flow masks work, see the “Flow Masks” section on page 13-12.



Note

Entering the **set mls flow** command purges all the existing entries in the NetFlow table.

To set the minimum NetFlow statistics flow mask, perform this task in privileged mode:

Task	Command
Set the minimum statistics flow mask.	set mls flow { destination destination-source null full }

This example shows how to set the minimum statistics flow mask to destination-source-ip:

```
Console> (enable) set mls flow destination-source
Configured IP flow mask is set to destination-source flow.
Console> (enable)
```

Excluding the IP Protocol Entries from the NetFlow Table

You can configure the NetFlow table to exclude specified IP protocols.

To exclude the IP protocols from the NetFlow table, perform this task in privileged mode:

Task	Command
Exclude the IP protocols from the NetFlow table.	set mls exclude protocol { tcp udp both } <i>port</i>

The *port* parameter can be a port number or a keyword: **dns**, **ftp**, **smtp**, **telnet**, **x** (X-Windows), or **www**.

This example shows how to exclude the Telnet traffic from the NetFlow table:

```
Console> (enable) set mls exclude protocol tcp telnet
NetFlow table will not create entries for TCP packets with protocol port 23.
Note: MLS exclusion only works in full flow mode.
Console> (enable)
```

Displaying the NetFlow Statistics



Note

To display the forwarding decision entries, enter the **show mls entry cef** command (see the “Displaying the Layer 3-Switching Entries on the Supervisor Engine” section on page 13-15).

To display a summary of the NetFlow table entries and statistics, perform this task in privileged mode:

Task	Command
Display all the NetFlow table entries and statistics.	show mls

This example shows how to display all the NetFlow table entries (the display is from a Supervisor Engine 2):

```

Console> (enable) show mls
show mls
=====
Total packets switched = 2
Total bytes switched = 112
Total routes = 48
IP statistics flows aging time = 16 seconds
IP statistics flows fast aging time = 0 seconds, packet threshold = 0
IP Current flow mask is Full flow
Netflow Data Export version:7
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0

IPX statistics flows aging time = 16 seconds
IPX flow mask is Destination flow
IPX max hop is 15

Module 15:Physical MAC-Address 00-50-3e-a9-ab-fc
Vlan Virtual MAC-Address(es)
-----
    42 00-00-0c-07-ac-00
Console>

```

This example shows how to display all the NetFlow table entries (the display is from a Supervisor Engine 720):

```

Console> (enable) show mls
Total packets switched = 35254
Total bytes switched = 2256256
Total routes = 120569
Total number of Netflow entries = 120000

IP statistics flows aging time = 50 seconds
Long-duration flows aging time = 320 seconds
IP statistics flows fast aging time = 0 seconds, packet threshold = 0

IP Current flow mask is Full-Vlan flow
Netflow Data Export version: 7
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0
Destination Ifindex export is enabled
Source Ifindex export is enabled
Rate limiting is turned off, packets are bridged to router
Load balancing hash is based on source and destination IP addresses and universc
Per-prefix Stats for ALL FIB entries is Enabled
Console> (enable)

```

The **show mls statistics entry** command can display all statistics or the statistics for the specific NetFlow table entries. Specify the destination address, source address, and for IP, the protocol, and source and destination ports to see the statistics for a specific NetFlow table entry.

A value of zero (0) for *src_port* or *dst_port* is treated as a wildcard, and all the NetFlow statistics are displayed (unspecified options are treated as wildcards). If the protocol that is specified is not TCP or UDP, set the *src_port* and *dstprt* to 0 or no NetFlow statistics will display.

To display the statistics for the NetFlow table entries, perform this task in privileged mode:

Task	Command
Display the statistics for the NetFlow table entries. If you do not specify a NetFlow table entry, all the NetFlow statistics are shown.	show mls statistics entry [ip ipx uptime] [destination ip_addr_spec] [source ip_addr_spec] [flow protocol src_port dst_port]

This example shows how to display the NetFlow statistics for a particular NetFlow table entry:

```

Console> show mls statistics entry ip destination 172.20.22.14
                Last      Used
Destination IP  Source IP      Prot DstPrt SrcPrt Stat-Pkts Stat-Bytes
-----
MSFC 127.0.0.12:
172.20.22.14   172.20.25.10   6    50648  80    3152    347854
Console>

```

The **show mls statistics entry ip top-talkers** command can display the statistics for the netflows with the maximum amount of network usage. The NetFlow entries are pulled out of the NetFlow table based on the number of packets that each flow has. The results are displayed in descending order with the top talkers being the entries with the largest packet count. You can get the statistics for the network (the top 32 talkers are displayed) or for a specified number of flows such as the top 1 or 2 talkers.

To display the NetFlow top talkers for the NetFlow table entries, perform this task in privileged mode:

Task	Command
Display the NetFlow talkers with the maximum amount of network usage.	show mls statistics entry ip top-talkers

This example shows how to display the NetFlow top talkers for a network:

```

Console> show mls statistics entry ip top-talkers
                Last      Used
Destination IP  Source IP      Prot DstPrt SrcPrt Vlan  Stat-Pkts  Stat-Bytes
-----
12.0.0.5        11.0.0.6       255  N/A    N/A    N/A   387110     17807060
12.0.0.5        11.0.0.7       255  N/A    N/A    N/A   387109     17807014
12.0.0.5        11.0.0.4       TCP   8      7      N/A    20         920
127.0.0.20     127.0.0.19     UDP   67     68    N/A    18         828
12.0.0.5        11.0.0.2       TCP   6      5      N/A    15         690
12.0.0.5        11.0.0.5       TCP   8      7      N/A    15         690
12.0.0.5        11.0.0.3       TCP   6      5      N/A    12         552
Console>

```

This example shows how to display the statistics for a specified number of NetFlows with the maximum network usage:

```

Console> show mls statistics entry ip top-talkers 2
Last      Used
-----
Destination IP   Source IP      Prot  DstPrt  SrcPrt  Vlan  Stat-Pkts  Stat-Bytes
-----
12.0.0.5         11.0.0.6      255   N/A     N/A     N/A   387110     17807060
12.0.0.5         11.0.0.7      255   N/A     N/A     N/A   387109     17807014
Console>

```

Clearing the NetFlow IP and IPX Statistics

These sections describe how to clear the NetFlow statistics:

- [Clearing All the NetFlow Statistics, page 13-34](#)
- [Clearing the NetFlow IP Statistics, page 13-34](#)
- [Clearing the NetFlow IPX Statistics, page 13-35](#)
- [Clearing the NetFlow Statistics Totals, page 13-36](#)



Note

The **clear mls** commands affect only the statistics. None of the **clear mls** commands affect the forwarding entries or the NetFlow table entries that correspond to the forwarding entries.

Clearing All the NetFlow Statistics

To clear all the NetFlow IP and IPX statistics, perform this task in privileged mode:

Task	Command
Clear all the NetFlow statistics.	clear mls statistics entry all

This example shows how to clear all the NetFlow statistics:

```

Console> (enable) clear mls statistics entry all
All MLS IP and IPX entries cleared.
Console> (enable)

```

Clearing the NetFlow IP Statistics

The **clear mls statistics entry ip** command clears the NetFlow IP statistics. Use the **all** keyword to clear all the NetFlow IP statistics. The **destination** and **source** keywords specify the source and destination IP addresses. The destination and source *ip_addr_spec* can be a full IP address or a subnet address in the format *ip_subnet_addr*, *ip_addr/subnet_mask*, or *ip_addr/subnet_mask_bits*.

The **flow** keyword specifies the following additional flow information:

- Protocol family (*protocol*)—Specify **tcp**, **udp**, **icmp**, or a decimal number for other protocol families. A value of zero (0) for *protocol* is treated as a wildcard (unspecified options are treated as wildcards).
- TCP or UDP source and destination port numbers (*src_port* and *dst_port*)—If the protocol that you specify is TCP or UDP, specify the source and destination TCP or UDP port numbers. A value of zero (0) for *src_port* or *dst_port* is treated as a wildcard (unspecified options are treated as wildcards). For other protocols, set the *src_port* and *dst_port* to 0, or no entries will clear.

To clear the statistics for a NetFlow table IP entry, perform this task in privileged mode:

Task	Command
Clear the statistics for a NetFlow table IP entry.	clear mls statistics entry ip [destination <i>ip_addr_spec</i>] [source <i>ip_addr_spec</i>] [flow <i>protocol src_port dst_port</i>] [all]

This example shows how to clear the statistics for NetFlow table entries with destination IP address 172.20.26.22:

```
Console> (enable) clear mls statistics entry ip destination 172.20.26.22
MLS IP entry cleared
Console> (enable)
```

This example shows how to clear the statistics for the NetFlow table entries with destination IP address 172.20.22.113, TCP source port 1652, and TCP destination port 23:

```
Console> (enable) clear mls statistics entry ip destination 172.20.26.22 source 172.20.22.113 flow tcp 1652 23
MLS IP entry cleared
Console> (enable)
```

Clearing the NetFlow IPX Statistics

The **clear mls statistics entry ipx** command clears the NetFlow IPX statistics. Use the **all** keyword to clear all the NetFlow IPX statistics. The **destination** and **source** keywords specify the source and destination IPX addresses.

To clear the statistics for a NetFlow table IPX entry, perform this task in privileged mode:

Task	Command
Clear the statistics for a NetFlow table IPX entry.	clear mls statistics entry ipx [destination <i>ipx_addr_spec</i>] [source <i>ipx_addr_spec</i>] [all]

This example shows how to clear the statistics for the IPX MLS entries with destination IPX address 1.0002.00e0.fefc.6000:

```
Console> (enable) clear mls statistics entry ipx destination 1.0002.00e0.fefc.6000
MLS IPX entry cleared.
Console> (enable)
```

Clearing the NetFlow Statistics Totals

The **clear mls statistics** command clears the following NetFlow statistics:

- Total packets that are switched (IP and IPX)
- Total packets that are exported (for NDE)

To clear the NetFlow statistic totals, perform this task in privileged mode:

Task	Command
Clear the NetFlow statistics totals.	clear mls statistics

This example shows how to clear the NetFlow statistics totals:

```
Console> (enable) clear mls statistics
All mls statistics cleared.
Console> (enable)
```

Displaying the NetFlow Statistics Debug Information

The **show mls debug** command displays the NetFlow statistics debug information that you can send to your technical support representative for analysis if necessary.

To display the NetFlow statistics debug information, perform this task:

Task	Command
Display the NetFlow statistics debug information that you can send to your technical support representative.	show mls debug



Note

The **show tech-support** command displays supervisor engine system information. Use application-specific commands to get more information about particular applications.

Configuring the MLS IP-Directed Broadcasts on the Switch

The IP-directed broadcasts are used primarily for ticker-type (stock quote) devices; however, when the feature is enabled on router interfaces, it provides a means to enable malicious denial-of-service attacks.

An IP-directed broadcast is a datagram that is sent to the broadcast address of a subnet to which the sending machine is not directly connected. The directed broadcast is routed through the network as a unicast packet until it arrives at the target subnet, where it is converted into a link layer broadcast. Due to the IP addressing architecture, only the last router in the chain, the one that is connected directly to the target subnet, can conclusively identify an IP-directed broadcast.

Before supervisor engine software release 7.2(2), the IP-directed broadcast traffic was handled by enabling the IP-directed broadcasts using the **ip directed-broadcast** command on the MSFC. The MSFC handled the traffic at the process level, which caused high CPU utilization.

With software release 7.2(2) and later releases, you can configure MSFC2 to handle the IP-directed broadcasts in the hardware using PFC2.

**Note**

Cisco IOS Release 12.1(11b)E is required on MSFC2.

This example shows how to enable the IP-directed broadcasts:

```
Router(config-if)# mls ip directed-broadcast ?  
  exclude-router  exclude router from recipient list for directed broadcast  
  include-router  include router in recipient list for directed broadcast
```

The **exclude-router** option forwards the IP-directed broadcast packet in the hardware to all the hosts in the VLAN except the router.

The **include-router** option forwards the IP-directed broadcast packet in the hardware to all the hosts in the VLAN including the router. With this option, the router does not forward the IP-directed broadcast packet again.

The **no** form of the command is as follows:

```
Router(config-if)# no mls ip directed-broadcast [exclude-router | include-router]
```

The **no** form returns the interface configuration to the default mode. In the default mode, the IP-directed broadcast packets are not hardware forwarded. They are handled at the process level by MSFC2. The MSFC2 decision to forward or not forward the packet depends on the **ip directed-broadcast** command configuration.

There is no interaction between the **ip directed-broadcast** command and the **mls ip directed-broadcast** command. The **ip directed-broadcast** command involves software forwarding and the **mls ip directed-broadcast** command involves hardware forwarding.

