



Configuring MLS

This chapter describes how to configure Multilayer Switching (MLS) for the Catalyst 6000 family switches. MLS provides IP and Internetwork Packet Exchange (IPX) unicast Layer 3 switching and IP multicast Layer 3 switching with Supervisor Engine 1, the Policy Feature Card (PFC), and the Multilayer Switch Feature Card (MSFC) or MSFC2.



Note

For complete information on the syntax and usage information for the supervisor engine commands used in this chapter, refer to the *Catalyst 6000 Family Command Reference* publication.

This chapter consists of these sections:

- [Understanding How Layer 3 Switching Works, page 14-1](#)
- [Default MLS Configuration, page 14-10](#)
- [Configuration Guidelines and Restrictions, page 14-11](#)
- [Configuring MLS, page 14-14](#)



Note

Supervisor Engine 2, PFC2, and MSFC2 provide Layer 3 switching with Cisco Express Forwarding for PFC2 (CEF for PFC2). See [Chapter 13, “Configuring CEF for PFC2,”](#) for more information.

Understanding How Layer 3 Switching Works

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



Note

Layer 3 switching supports the routing protocols configured on the MSFC. Layer 3 switching does not replace the routing protocols configured on the MSFC. Layer 3 switching uses IP Protocol Independent Multicast (IP PIM) for multicast route determination.

Layer 3 switching on Catalyst 6000 family switches provides traffic statistics that you can use to identify traffic characteristics for administration, planning, and troubleshooting. Layer 3 switching uses NetFlow Data Export (NDE) to export flow statistics (for more information about NDE, see [Chapter 15, “Configuring NDE”](#)).

These sections describe Layer 3 switching and MLS on the Catalyst 6000 family switches:

- [Understanding Layer 3-Switched Packet Rewrite, page 14-2](#)
- [Understanding MLS, page 14-4](#)

Understanding Layer 3-Switched Packet Rewrite

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 based on information learned from the MSFC so that the packets appear to have been routed by the MSFC.



Note

Rather than just forwarding multicast packets, the switch replicates them as necessary on the appropriate VLANs.

Packet rewrite alters 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)

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

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 the MSFC. The Layer 3 addresses remain the same.

In IP unicast and IP multicast traffic, the switch decrements the Layer 3 Time to Live (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 14-2](#)
- [Understanding IPX Unicast Rewrite, page 14-3](#)
- [Understanding IP Multicast Rewrite, page 14-3](#)

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>MSFC 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>MSFC 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>MSFC 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>MSFC 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:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Group G1 MAC</i>	<i>MSFC MAC</i>	<i>Group G1 IP</i>	<i>Source A IP</i>	<i>n-1</i>	<i>calculation2</i>		

Understanding MLS

**Note**

Supervisor Engine 1, PFC, and MSFC or MSFC2 can only do MLS internally with the MSFC or MSFC2 in the same chassis; an external MLS-RP cannot be used in place of the internal MLS-RP.

Supervisor Engine 1, PFC, and MSFC or MSFC2 provide Layer 3 switching with MLS. Layer 3 switching with MLS identifies flows on the switch after the first packet has been routed by the MSFC and transfers the process of forwarding the remaining traffic in the flow to the switch, which reduces the load on the MSFC.

These sections describe MLS:

- [Understanding MLS Flows, page 14-4](#)
- [Understanding the MLS Cache, page 14-5](#)
- [Understanding Flow Masks, page 14-6](#)
- [Partially and Completely Switched Multicast Flows, page 14-8](#)
- [MLS Examples, page 14-8](#)

Understanding MLS Flows

Layer 3 protocols, such as IP and IPX, are connectionless—they deliver every packet independently of every other packet. However, actual network traffic consists of many end-to-end conversations, or flows, between users or applications.

MLS supports unicast and multicast flows:

- A unicast flow can be any of the following:
 - All traffic to a particular destination
 - All traffic from a particular source to a particular destination
 - All traffic from a particular source to a particular 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.

For example, communication from a client to a server and from the server to the client are separate flows. Telnet traffic transferred from a particular source to a particular destination comprises a separate flow from File Transfer Protocol (FTP) packets between the same source and destination.

**Note**

The PFC uses the Layer 2 multicast forwarding table to identify the ports to which Layer 2 multicast traffic should be forwarded (if any). The multicast forwarding table entries are populated by whichever multicast constraint feature is enabled on the switch (IGMP snooping or Generic Attribute Registration Protocol [GARP] Multicast Registration Protocol [GMRP]). These entries map the destination multicast MAC address to the outgoing switch ports for a given VLAN.

Understanding the MLS Cache

These sections describe the MLS cache:

- [MLS Cache, page 14-5](#)
- [Unicast Traffic, page 14-5](#)
- [Multicast Traffic, page 14-5](#)
- [MLS Cache Aging, page 14-5](#)
- [MLS Cache Size, page 14-6](#)

MLS Cache

The PFC maintains a Layer 3 switching table called the MLS cache for Layer 3-switched flows. The cache also includes entries for traffic statistics that are updated in tandem with the switching of packets. After the PFC creates an MLS cache entry, packets identified as belonging to an existing flow can be Layer 3 switched based on the cached information. The MLS cache maintains flow information for all active flows.

Unicast Traffic

For unicast traffic, the PFC creates an MLS cache entry for the initial routed packet of each unicast flow. Upon receipt of a routed packet that does not match any unicast flow currently in the MLS cache, the PFC creates a new MLS entry.

Multicast Traffic

For multicast traffic, the PFC populates the MLS cache using information learned from the MSFC. Whenever the MSFC receives traffic for a new multicast flow, it updates its multicast routing table and forwards the new information to the PFC. In addition, if an entry in the multicast routing table ages out, the MSFC deletes the entry and forwards the updated information to the PFC.

For each multicast flow cache entry, the PFC maintains a list of outgoing interfaces for the destination IP multicast group. The PFC uses this list to identify the VLANs on which traffic to a given multicast flow should be replicated.

These MSFC IOS commands affect the multicast MLS cache entries on the switch:

- Using the **clear ip mroute** command to clear the multicast routing table on the MSFC clears all multicast MLS cache entries on the PFC.
- Using the **no ip multicast-routing** command to disable IP multicast routing on the MSFC purges all multicast MLS cache entries on the PFC.

MLS Cache 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 MLS entries kept in the MLS cache. 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.

MLS Cache Size

The maximum MLS cache size is 128K entries. The MLS cache is shared by all MLS processes on the switch (IP MLS, IP MMLS, and IPX MLS). An MLS cache larger than 32K entries increases the probability that a flow will not be Layer 3 switched, but will instead be forwarded to the MSFC.

Understanding Flow Masks

The PFC uses flow masks to determine how MLS entries are created.

These sections describe the flow mask modes:

- [Flow Mask Modes, page 14-6](#)
- [Flow Mask Mode and show mls entry Command Output, page 14-7](#)

Flow Mask Modes

The PFC supports only one flow mask (the most specific one) for all MSFCs that are Layer 3 switched by that PFC. If the PFC 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.

When the PFC flow mask changes, the entire MLS cache is purged. When the PFC 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 MLS flow masks are as follows:

- **destination-ip**—The least-specific flow mask. The PFC maintains one MLS entry for each Layer 3 destination address. All flows to a given Layer 3 destination address use this MLS entry.
- **destination-ipx**—The only flow mask mode for IPX MLS is destination mode. The PFC maintains one IPX MLS entry for each destination IPX address (network and node). All flows to a given destination IPX address use this IPX MLS entry.
- **source-destination-ip**—The PFC maintains one MLS entry for each source and destination IP address pair. All flows between a given source and destination use this MLS entry regardless of the IP protocol ports.
- **source-destination-vlan**—For IP MMLS. The PFC maintains one MMLS cache entry for each {source IP, destination group IP, source VLAN}. The multicast source-destination-vlan flow mask differs from the IP unicast MLS source-destination-ip flow mask in that, for IP MMLS, the source VLAN is included as part of the entry. The source VLAN is the multicast reverse path forwarding (RPF) interface for the multicast flow.
- **full flow**—The most-specific flow mask. The PFC creates and maintains a separate MLS cache entry for each IP flow. A full flow entry includes the source IP address, destination IP address, protocol, and protocol ports.

Flow Mask Mode and show mls entry Command Output

With the destination-ip flow mask, the source IP, protocol, and source and destination port fields show the details of the last packet that was Layer 3 switched using the MLS cache entry.

This example shows how the **show mls entry** command output appears in destination-ip mode:

```

Console> (enable) show mls entry ip short
Destination-IP Source-IP      Prot  DstPrt SrcPrt Destination-Mac  Vlan
-----
ESrc EDst SPort DPort Stat-Pkts Stat-Byte  Uptime  Age
-----
171.69.200.234 -          -      -      -      00-60-70-6c-fc-22 4
  ARPA SNAP 5/8  11/1  3152    347854    09:01:19 09:08:20
171.69.1.133   -          -      -      -      00-60-70-6c-fc-23 2
  SNAP ARPA 5/8  1/1   2345    123456    09:03:32 09:08:12

Total Entries: 2
* indicates TCP flow has ended
Console> (enable)

```



Note

The **short** keyword exists for some **show** commands and displays the output by wrapping the text after 80 characters. The default is **long** (no text wrap).

With the source-destination-ip flow mask, the protocol, source port, and destination port fields display the details of the last packet that was Layer 3 switched using the MLS cache entry.

This example shows how the **show mls entry** command output appears in source-destination-ip mode:

```

Console> (enable) show mls entry ip short
Destination-IP Source-IP      Prot  DstPrt SrcPrt Destination-Mac  Vlan
-----
ESrc EDst SPort DPort Stat-Pkts Stat-Byte  Uptime  Age
-----
171.69.200.234 171.69.192.41 -      -      -      00-60-70-6c-fc-22 4
  ARPA SNAP 5/8  11/1  3152    347854    09:01:19 09:08:20
171.69.1.133   171.69.192.42 -      -      -      00-60-70-6c-fc-23 2
  SNAP ARPA 5/8  1/1   2345    123456    09:03:32 09:08:12

Total Entries: 2
* indicates TCP flow has ended
Console> (enable)

```

With the full-flow flow mask, because a separate MLS entry is created for every ip flow, details are shown for each flow.

This example shows how the **show mls entry** command output appears in full flow mode:

```

Console> (enable) show mls entry ip short
Destination-IP Source-IP      Prot  DstPrt SrcPrt Destination-Mac  Vlan
-----
ESrc EDst SPort DPort Stat-Pkts Stat-Byte  Uptime  Age
-----
171.69.200.234 171.69.192.41 TCP*  6000  59181 00-60-70-6c-fc-22 4
  ARPA SNAP 5/8  11/1  3152    347854    09:01:19 09:08:20
171.69.1.133   171.69.192.42 UDP   2049  41636 00-60-70-6c-fc-23 2
  SNAP ARPA 5/8  1/1   2345    123456    09:03:32 09:08:12

Total Entries: 2
* indicates TCP flow has ended
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:

- The MSFC 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.
- The MSFC is the first-hop router to the source in PIM sparse mode (in this case, the MSFC must send PIM-register messages to the rendezvous point).
- The multicast TTL threshold is configured on an egress interface for the flow.
- The extended access list deny condition on the RPF interface specifies anything other than the Layer 3 source, Layer 3 destination, or IP protocol (an example is Layer 4 port numbers).
- 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.

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

The PFC prevents multicast traffic in flows that are completely Layer 3 switched from reaching the MSFC, reducing the load on the MSFC. 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 the MSFC, the multicast packet is dropped, because it is switched by the PFC).

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

MLS Examples

[Figure 14-1](#) shows a simple IP MLS 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, an MLS entry for this flow is created (this entry is the second item in the MLS cache shown in [Figure 14-1](#)). The PFC stores the MAC addresses of the MSFC and Host C in the MLS entry when the MSFC forwards the first packet from Host A through the switch to Host C. The PFC uses this information to rewrite subsequent packets from Host A to Host C.

Figure 14-1 IP MLS Example Topology

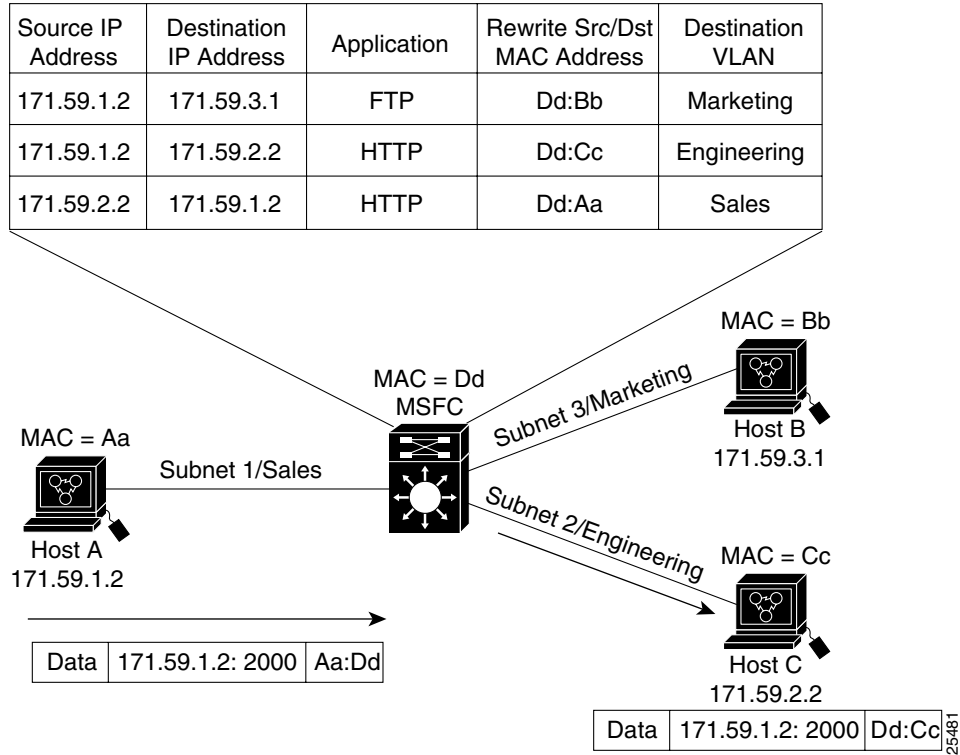
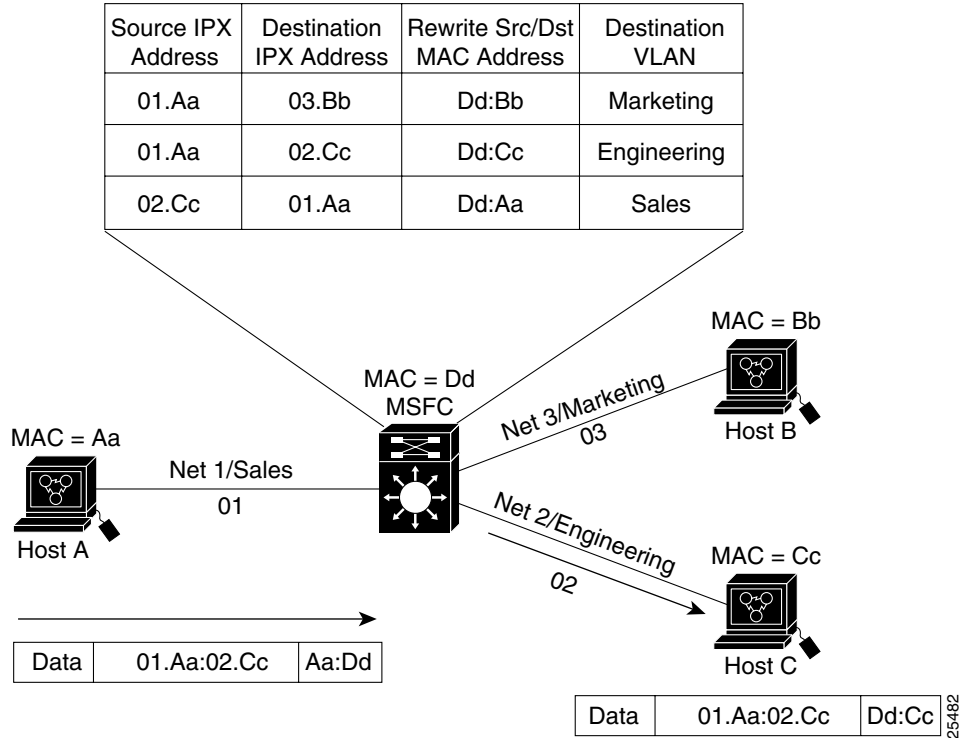


Figure 14-2 shows a simple IPX MLS network topology. In this example, Host A is on the Sales VLAN (IPX address 01.Aa), Host B is on the Marketing VLAN (IPX address 03.Bb), and Host C is on the Engineering VLAN (IPX address 02.Cc).

When Host A initiates a file transfer to Host B, an IPX MLS entry for this flow is created (this entry is the first item in the table shown in Figure 14-1). The PFC stores the MAC addresses of the MSFC and Host B in the IPX MLS entry when the MSFC forwards the first packet from Host A through the switch to Host B. The PFC uses this information to rewrite subsequent packets from Host A to Host B.

Similarly, a separate IPX MLS entry is created in the MLS cache for the traffic from Host A to Host C, and for the traffic from Host C to Host A. The destination VLAN is stored as part of each IPX MLS entry so that the correct VLAN identifier is used when encapsulating traffic on trunk links.

Figure 14-2 IPX MLS Example Topology



Default MLS Configuration

Table 14-1 shows the default IP MLS configuration.

Table 14-1 Default IP MLS Configuration

Feature	Default Value
IP MLS enable state	Enabled
IP MLS aging time	256 seconds
IP MLS fast aging time	0 seconds (no fast aging)
IP MLS fast aging-time packet threshold	0 packets

Table 14-2 shows the default IP MMLS switch configuration.

Table 14-2 Default IP MMLS Supervisor Engine Configuration

Feature	Default Value
Multicast services (IGMP snooping or GMRP)	Disabled
IP MMLS	Enabled

Table 14-3 shows the default IP MMLS MSFC configuration.

Table 14-3 Default IP MMLS MSFC Configuration

Feature	Default Value
Multicast routing	Disabled globally
IP PIM routing	Disabled on all interfaces
IP MMLS Threshold	Unconfigured—no default value
IP MMLS	Enabled when multicast routing is enabled and IP PIM is enabled on the interface

Table 14-4 shows the default IPX MLS configuration.

Table 14-4 Default IPX MLS Configuration

Feature	Default Value
IPX MLS enable state	Enabled
IPX MLS aging time	256 seconds

Configuration Guidelines and Restrictions

These sections describe configuration guidelines and restrictions:

- [IP MLS, page 14-11](#)
- [IP MMLS, page 14-12](#)
- [IPX MLS, page 14-13](#)

IP MLS

These sections describe IP MLS configuration guidelines:

- [Maximum Transmission Unit Size, page 14-11](#)
- [Restrictions on Using IP Routing Commands with IP MLS Enabled, page 14-12](#)

Maximum Transmission Unit Size

The default maximum transmission unit (MTU) for IP MLS is 1500. To change the MTU on an IP MLS-enabled interface, enter the **ip mtu mtu** command.

Restrictions on Using IP Routing Commands with IP MLS Enabled

Enabling certain IP processes on an interface will affect IP MLS on the interface. [Table 14-5](#) shows the affected commands and the resulting behavior.

Table 14-5 IP Routing Command Restrictions

Command	Behavior
clear ip route	Clears all MLS cache entries for all switches performing Layer 3 switching for this MSFC.
ip routing	The no form purges all MLS cache entries and disables IP MLS on this MSFC.
ip security (all forms of this command)	Disables IP MLS on the interface.
ip tcp compression-connections	Disables IP MLS on the interface.
ip tcp header-compression	Disables IP MLS on the interface.

IP MMLS

These sections describe IP MMLS configuration guidelines:

- [IP MMLS Supervisor Engine Guidelines and Restrictions, page 14-12](#)
- [IP MMLS MSFC Configuration Restrictions, page 14-13](#)
- [Unsupported IP MMLS Features, page 14-13](#)

IP MMLS Supervisor Engine Guidelines and Restrictions

These guidelines and restrictions apply when configuring Supervisor Engine 1 for IP MMLS:

- Only ARPA rewrites are supported for IP multicast packets.
- Subnetwork Address Protocol (SNAP) rewrites are not supported.
- You must enable one of the multicast services (IGMP snooping or GMRP) on the switch in order to use IP MMLS.
- IP multicast flows are not multilayer switched if there is no entry in the Layer 2 multicast forwarding table (for example, if no Layer 2 multicast services are enabled or the forwarding table is full). Enter the **show multicast group** command to check for a Layer 2 entry for a particular IP multicast destination.
- If a Layer 2 entry is cleared, the corresponding Layer 3 flow information is purged.
- When using two MSFCs that have one or more interfaces in the same VLAN, the switch uses two reserved VLANs (VLANs 1012 and 1013) internally to forward multicast flows properly.
- The MSFC will not act as an external router for a Catalyst 5000 family switch that has Layer 3 switching hardware.

IP MMLS MSFC Configuration Restrictions

IP MMLS does not perform multilayer switching for an IP multicast flow in the following situations:

- For IP multicast groups that fall into these ranges (where * is in the range 0–255):

224.0.0.* through 239.0.0.*

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



Note In systems with redundant MSFCs, the IP PIM interface configuration must be the same on both the active and redundant MSFCs.

- For flows that are forwarded on the multicast-shared tree (that is, {*,G,*} forwarding) when the interface or group is running IP PIM sparse mode.
- If the shortest-path tree (SPT) bit for the flow is cleared when running IP 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 received on tunnel interfaces (such as MBONE traffic).
- For any RPF interface with multicast tag switching enabled.

Unsupported IP MMLS Features

If you enable IP MMLS, IP accounting for the interface will not reflect accurate values.

IPX MLS

These sections describe configuration guidelines that apply when configuring IPX MLS:

- [IPX MLS Interaction with Other Features, page 14-13](#)
- [IPX MLS and Maximum Transmission Unit Size, page 14-14](#)

IPX MLS Interaction with Other Features

Other IOS software features affect IPX MLS as follows:

- IPX accounting—IPX accounting cannot be enabled on an IPX MLS-enabled interface.
- IPX EIGRP—To support MLS on EIGRP interfaces you must set the Transport Control (TC) maximum to a value greater than the default (16). Enter the **ipx maximum-hop tc_value** global configuration command on the MSFC, with the *tc_value* greater than 16.

IPX MLS and Maximum Transmission Unit Size

In IPX, the two end points of communication negotiate the maximum transmission unit (MTU) to be used. The MTU size is limited by the media type.

Configuring MLS

These sections describe how to configure MLS:

- [Configuring Unicast MLS on the MSFC, page 14-14](#)
- [Configuring MLS on Supervisor Engine 1, page 14-17](#)
- [Configuring IP MMLS, page 14-28](#)

Configuring Unicast MLS on the MSFC

These sections describe how to configure MLS on the MSFC:

- [Disabling and Enabling Unicast MLS on an MSFC Interface, page 14-14](#)
- [Displaying MLS Information on the MSFC, page 14-15](#)
- [Using Debug Commands on the MSFC, page 14-16](#)
- [Using Debug Commands on the SCP, page 14-16](#)

For information on configuring routing on the MSFC, see [Chapter 12, “Configuring InterVLAN Routing.”](#) For information on configuring unicast Layer 3 switching on Supervisor Engine 1, see the [“Configuring MLS on Supervisor Engine 1”](#) section on page 14-17.



Note

The MSFC can be specified as the MLS route processor (MLS-RP) for Catalyst 5000 family switches using MLS. Refer to the *Layer 3 Switching Configuration Guide—Catalyst 5000 Family, 2926G Series, 2926 Series Switches*, for MLS configuration procedures.

Disabling and Enabling Unicast MLS on an MSFC Interface

Unicast MLS for IP and IPX is enabled globally by default, but can be disabled and enabled on a specified interface.

To disable unicast IP or IPX MLS on a specific MSFC interface, perform this task:

Task	Command
Specify an MSFC interface.	Router(config)# interface <i>vlan-id</i>
Disable IP MLS on an MSFC interface.	Router(config-if)# no mls ip
Disable IPX MLS on an MSFC interface.	Router(config-if)# no mls ipx

This example shows how to disable IP MLS on an MSFC interface:

```
Router(config)# interface vlan 100
Router(config-if)# no mls ip
Router(config-if)#
```

This example shows how to disable IPX MLS on an MSFC interface:

```
Router(config)# interface vlan 100
Router(config-if)# no mls ipx
Router(config-if)#
```



Note

Unicast MLS is enabled by default; you only need to enable (or reenable) it if you have previously disabled it.

To enable unicast IP or IPX MLS on a specific MSFC interface, perform this task:

Task	Command
Specify an MSFC interface.	Router(config)# interface <i>vlan-id</i>
Enable IP MLS on an MSFC interface.	Router(config-if)# mls ip
Enable IPX MLS on an MSFC interface.	Router(config-if)# mls ipx

This example shows how to enable IP MLS on an MSFC interface:

```
Router(config)# interface vlan 100
Router(config-if)# mls ip
Router(config-if)#
```

This example shows how to enable IPX MLS on an MSFC interface:

```
Router(config)# interface vlan 100
Router(config-if)# mls ipx
Router(config-if)#
```

Displaying MLS Information on the MSFC

The **show mls status** command displays MLS details.

To display MLS information on the MSFC, perform this task:

Task	Command
Display MLS status.	show mls status

This example shows how to display MLS status on the MSFC:

```
Router# show mls status
MLS global configuration status:

global mls ip:                enabled
global mls ipx:               enabled
global mls ip multicast:      disabled
current ip flowmask for unicast: destination only
current ipx flowmask for unicast: destination only
Router#
```

Using Debug Commands on the MSFC

Table 14-6 describes MLS-related debug commands that you can use to troubleshoot MLS problems on the MSFC.

Table 14-6 *MLS Debug Commands*

Command	Description
[no] debug l3-mgr events	Displays Layer 3 manager-related events.
[no] debug l3-mgr packets	Displays Layer 3 manager packets.
[no] debug l3-mgr global	Displays bugtrace of ip global purge events.
[no] debug l3-mgr all	Turns on all Layer 3 manager debugging messages.

Table 14-7 describes MLS-related debug commands that you can use to troubleshoot MLS problems when using the MSFC as an external router for a Catalyst 5000 family switch.

Table 14-7 *MLS Debug Commands—External Router Function*

Command	Description
[no] debug mls ip	Turns on IP-related events for MLS, including route purging and changes of access lists and flow masks.
[no] debug mls ipx	Turns on IPX-related events for MLS, including route purging and changes of access lists and flow masks.
[no] debug mls rp	Turns on route processor-related events.
[no] debug mls locator	Identifies which switch is switching a particular flow by using MLS explorer packets.
[no] debug mls all	Turns on all MLS debugging events.

Using Debug Commands on the SCP

Table 14-8 describes the Serial Control Protocol (SCP)-related debug commands to troubleshoot the SCP that runs over the Ethernet out-of-band channel (EOBC).

Table 14-8 *SCP Debug Commands*

Command	Description
[no] debug scp async	Displays trace for asynchronous data in and out of the SCP system.
[no] debug scp data	Displays packet data trace.
[no] debug scp errors	Displays errors and warnings in the SCP.
[no] debug scp packets	Displays 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.

Configuring MLS on Supervisor Engine 1

MLS is enabled by default on Catalyst 6000 family switches. You only need to configure Supervisor Engine 1 in these circumstances:

- You want to change the MLS aging time
- You want to enable NDE

These sections describe how to configure MLS on Supervisor Engine 1:

- [Specifying MLS Aging-Time Value, page 14-17](#)
- [Specifying IP MLS Fast Aging Time and Packet Threshold Values, page 14-18](#)
- [Setting the Minimum IP MLS Flow Mask, page 14-19](#)
- [Displaying CAM Entries on the Supervisor Engine, page 14-20](#)
- [Displaying MLS Information, page 14-21](#)
- [Displaying IP MLS Cache Entries, page 14-22](#)
- [Clearing MLS Cache Entries, page 14-26](#)
- [Clearing IPX MLS Cache Entries, page 14-26](#)
- [Displaying IP MLS Statistics, page 14-26](#)
- [Clearing MLS Statistics, page 14-28](#)
- [Displaying MLS Debug Information, page 14-28](#)

For information on configuring VLANs on the switch, see [Chapter 11, “Configuring VLANs.”](#) For information on configuring MLS on the MSFC, see the [“Configuring Unicast MLS on the MSFC”](#) section on [page 14-14](#).

**Note**

When you disable IP or IPX MLS on the MSFC, IP or IPX MLS is automatically disabled on Supervisor Engine 1. All existing protocol-specific MLS cache entries are purged. To disable MLS on the MSFC, see the [“Disabling and Enabling Unicast MLS on an MSFC Interface”](#) section on [page 14-14](#).

**Note**

If NDE is enabled and you disable MLS, you will lose the statistics for existing cache entries—they are not exported.

Specifying MLS Aging-Time Value

The MLS aging time for each protocol (IP and IPX) applies to all protocol-specific MLS cache entries. Any MLS entry that has not been used for *agingtime* seconds is aged out. The default is 256 seconds.

You can configure the aging time in the range of 8 to 2032 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.

**Note**

We recommend that you keep the size of the MLS cache below 32K entries. If the number of MLS entries exceeds 32K, some flows are sent to the MSFC. To help keep the size of the MLS cache down, for IP, enable IP MLS fast aging, as described in the [“Specifying IP MLS Fast Aging Time and Packet Threshold Values”](#) section on page 14-18.

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

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

This example shows how to specify the MLS aging time:

```
Console> (enable) set mls agingtime 512
Multilayer switching agingtime IP and IPX set to 512
Console> (enable)
```

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

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

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

```
Console> (enable) set mls agingtime ip 512
Multilayer switching aging time IP set to 512
Console> (enable)
```

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

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

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

```
Console> (enable) set mls agingtime ipx 512
Multilayer switching aging time IPX set to 512
Console> (enable)
```

Specifying IP MLS Fast Aging Time and Packet Threshold Values

**Note**

IPX MLS does not use fast aging. IPX MLS only operates in destination-source and destination flow modes; therefore, the number of IPX MLS entries in the MLS table is low relative to IP MLS entries in full-flow mode.

To keep the MLS cache size below 32K entries, enable IP MLS fast aging time. The IP MLS fast aging time applies to MLS entries that have no more than *pkt_threshold* packets switched within *fastagingtime* seconds after they are created. A typical cache 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 MLS cache for other data traffic.

The default *fastagingtime* value is 0 (no fast aging). You can configure the *fastagingtime* value to 32, 64, 96, or 128 seconds. Any *fastagingtime* value that is not configured exactly as the indicated values is adjusted to the closest one. You can configure the *pkt_threshold* value to 0, 1, 3, 7, 15, 31, or 63 packets.

If you need to enable IP MLS fast aging time, initially set the value to 128 seconds. If the size of the MLS cache continues to grow over 32K entries, decrease the setting until the cache size stays below 32K. If the cache continues to grow over 32K entries, decrease the normal IP MLS aging time.

Typical values for *fastagingtime* and *pkt_threshold* are 32 seconds and 0 packets (no packets switched within 32 seconds after the entry is created).

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

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

This example shows how to set the IP MLS fast aging time to 32 seconds with a packet threshold of 0 packets:

```
Console> (enable) set mls agingtime fast 32 0
Multilayer switching fast aging time set to 32 seconds for entries with no more than 0
packets switched.
Console> (enable)
```

Setting the Minimum IP MLS Flow Mask

You can set the minimum granularity of the flow mask for the MLS cache on the PFC. The actual flow mask used will be at least of the granularity specified by this command. For information on how the different flow masks work, see the [“Understanding Flow Masks” section on page 14-6](#).

For example, if you do not configure access lists on any MSFC, then the IP MLS flow mask on the PFC is destination-ip by default. However, you can force the PFC to use the source-destination-ip flow mask by setting the minimum IP MLS flow mask using the **set mls flow destination-source** command.



Caution

The **set mls flow destination-source** command purges all existing shortcuts in the MLS cache and affects the number of active shortcuts on the PFC. Exercise care when using this command.

To set the minimum IP MLS flow mask, perform this task in privileged mode:

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

This example shows how to set the minimum IP MLS 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)
```

Displaying CAM Entries on the Supervisor Engine

The **show cam** command displays the content-addressable memory (CAM) entries associated with a specific MAC address. If the MAC address belongs to an MSFC, an “R” is appended to the MAC address.

If you specify a VLAN number, only those CAM entries corresponding to that VLAN number are displayed. If a VLAN is not specified, entries for all VLANs are displayed.

To display CAM entries, perform this task:

Task	Command
Display CAM entries by MAC address.	show cam msfc [vlan]

This example shows how to display the CAM entries:

```
Console> show cam msfc
VLAN  Destination MAC          Destination-Ports or VCs      Xtag  Status
----  -
194   00-e0-f9-d1-2c-00R        7/1                          2     H
193   00-00-0c-07-ac-c1R        7/1                          2     H
193   00-00-0c-07-ac-5dR        7/1                          2     H
202   00-00-0c-07-ac-caR        7/1                          2     H
204   00-e0-f9-d1-2c-00R        7/1                          2     H
195   00-e0-f9-d1-2c-00R        7/1                          2     H
192   00-00-0c-07-ac-c0R        7/1                          2     H
192   00-e0-f9-d1-2c-00R        7/1                          2     H
204   00-00-0c-07-ac-ccR        7/1                          2     H
202   00-e0-f9-d1-2c-00R        7/1                          2     H
194   00-00-0c-07-ac-5eR        7/1                          2     H
196   00-e0-f9-d1-2c-00R        7/1                          2     H
194   00-00-0c-07-ac-c2R        7/1                          2     H
193   00-e0-f9-d1-2c-00R        7/1                          2     H
Total Matching CAM Entries Displayed = 14
Console>
```

This example shows how to display the CAM entries for a specified VLAN:

```
Console> show cam msfc 192
VLAN  Destination MAC          Destination-Ports or VCs      Xtag  Status
----  -
192   00-00-0c-07-ac-c0R        7/1                          2     H
192   00-e0-f9-d1-2c-00R        7/1                          2     H
Console>
```

Displaying MLS Information

The **show mls** command displays protocol-specific MLS information and MSFC-specific information.

To display protocol-specific MLS information and MSFC-specific information, perform this task:

Task	Command
Display general IP or IPX MLS information and MSFC-specific information for all MSFCs.	show mls {ip ipx} [mod¹]

1. The **mod** keyword specifies the module number of the MSFC; either 15 (if the MSFC is installed on Supervisor Engine 1 in slot 1) or 16 (if the MSFC is installed on Supervisor Engine 1 in slot 2).

This example shows how to display IP MLS information and MSFC-specific information:

```
Console> (enable) show mls ip
Total Active MLS entries = 0
Total packets switched = 0
IP Multilayer switching enabled
IP Multilayer switching aging time = 256 seconds
IP Multilayer switching fast aging time = 0 seconds, packet threshold = 0
IP Flow mask: Full Flow
Configured flow mask is Destination flow
Active IP MLS entries = 0
Netflow Data Export version: 8
Netflow Data Export disabled
Netflow Data Export port/host is not configured
Total packets exported = 0
```

MSFC ID	Module	XTAG	MAC	Vlans
52.0.03	15	1	01-10-29-8a-0c-00	1,10,123,434,121 222,666,959

```
Console> (enable)
```

This example shows how to display IPX MLS information:

```
Console> (enable) show mls ipx
IPX Multilayer switching aging time = 256 seconds
IPX flow mask is Destination flow
IPX max hop is 15
Active IPX MLS entries = 356
```

IPX MSFC ID	Module	XTAG	MAC	Vlans
22.1.0.56	15	1	00-10-07-38-29-18	2,3,4,5,6, 7,8,9,10,11, 12,13,14,15,16, 17,18,19,20,66, 77
			00-d0-d3-9c-e3-f4	25
			00-10-07-38-29-18	26,111
			00-d0-d3-9c-e3-f4	112

```

22.1.1.0.58      16      2      00-10-07-38-22-22 2,3,4,5,6,
                                     7,8,9,10,11,
                                     12,13,14,15,16,
                                     17,18,19,20
00-d0-d3-33-17-8c 25
00-10-07-38-22-22 26,66,77,88,99,
                                     111
00-d0-d3-33-17-8c 112

Console> (enable)

```

Displaying IP MLS Cache Entries

These sections describe how to display MLS cache entries on Supervisor Engine 1:

- [Displaying All MLS Entries, page 14-22](#)
- [Displaying MLS Entries for a Specific IP Destination Address, page 14-23](#)
- [Displaying IPX MLS Entries for a Specific IPX Destination Address, page 14-23](#)
- [Displaying Entries for a Specific IP Source Address, page 14-24](#)
- [Displaying Entries for a Specific IP Flow, page 14-24](#)
- [Displaying IPX MLS Entries for a Specific MSFC, page 14-25](#)



Note

For a description of how the flow mask mode affects the screen displays when showing MLS entries, see the [“Flow Mask Mode and show mls entry Command Output” section on page 14-7](#).

Displaying All MLS Entries

To display all MLS entries (IP and IPX), perform this task in privileged mode:

Task	Command
Display all MLS entries.	show mls entry [short long]

This example shows how to display all MLS entries (IP and IPX):

```

Console> (enable) show mls entry short
Destination-IP  Source-IP      Prot  DstPrt  SrcPrt  Destination-Mac  Vlan
-----
ESrc EDst SPort DPort  Stat-Pkts  Stat-Bytes  Created  LastUsed
-----
171.69.200.234 171.69.192.41  TCP*  6000    59181    00-60-70-6c-fc-22  4
  ARPA SNAP 5/8  11/1  3152      347854      09:01:19 09:08:20
171.69.1.133   171.69.192.42  UDP   2049    41636    00-60-70-6c-fc-23  2
  SNAP ARPA 5/8  1/1    2345      1234567     09:03:32 09:08:12
171.69.1.133   171.69.192.42  UDP   2049    41636    00-60-70-6c-fc-23  2
  SNAP ARPA 5/8  1/1    2345      1234567     09:03:32 09:08:12
171.69.1.133   171.69.192.42  UDP   2049    41636    00-60-70-6c-fc-23  2
  SNAP ARPA 5/8  1/1    2345      1234567     09:03:32 09:08:12
171.69.1.133   171.69.192.42  UDP   2049    41636    00-60-70-6c-fc-23  2
  SNAP ARPA 5/8  1/1    2345      1234567     09:03:32 09:08:12

Total IP entries: 5
* indicates TCP flow has ended.

```

```

Destination-IPX          Source-IPX-net Destination-Mac  Vlan Port
  Stat-Pkts Stat-Bytes
-----
BABE.0000.0000.0001      -                00-a0-c9-0a-89-1d 211 13/37
  30230      1510775
201.00A0.2451.7423      -                00-a0-24-51-74-23 201 14/33
  30256      31795084
501.0000.3100.0501      -                31-00-05-01-00-00 501 9/37
  12121      323232
401.0000.0000.0401      -                00-00-04-01-00-00 401 3/1
  4633       38676

Total IPX entries: 4
Console>

```

Displaying MLS Entries for a Specific IP Destination Address

To display MLS entries for a specific destination IP address, perform this task in privileged mode:

Task	Command
Display MLS entries for the specified destination IP address.	show mls entry ip destination [<i>ip_addr</i>]

This example shows how to display MLS entries for a specific destination IP address:

```

Console> (enable) show mls entry ip destination 172.20.22.14/24
Destination-IP Source-IP      Prot  DstPrt SrcPrt Destination-Mac  Vlan
  EDst ESrc DPort  SPort  Stat-Pkts Stat-Bytes  Uptime  Age
-----
MSFC 172.20.25.1 (Module 15):
172.20.22.14    -          -      -      -          00-60-70-6c-fc-22 4
  ARPA ARPA 5/39   5/40   115        5290        00:12:20 00:00:04
MSFC 172.20.27.1 (Module 16):

Total entries:1
Console> (enable)

```

Displaying IPX MLS Entries for a Specific IPX Destination Address

To display IPX MLS entries for a specific destination IPX address, perform this task in privileged mode:

Task	Command
Display IPX MLS entries for a specific destination IPX address (net_address.node_address).	show mls entry ipx destination <i>ipx_addr</i>

This example shows how to display IPX MLS entries for a specific destination IPX address:

```

Console> (enable) show mls entry ipx destination 3E.0010.298a.0c00
Destination IPX          Source IPX net Destination Mac  Vlan Port
-----
MSFC 22.1.0.56 (Module 15):
3E.0010.298a.0c00          13 00-00-00-00-00-09 26 4/7

Console> (enable)

```

Displaying Entries for a Specific IP Source Address

To display MLS entries for a specific source IP address, perform this task in privileged mode:

Task	Command
Display MLS entries for the specified source IP address.	show mls entry ip source [<i>ip_addr</i>]

This example shows how to display MLS entries for a specific source IP address:

```

Console> (enable) show mls entry ip source 10.0.2.15
Destination-IP Source-IP      Prot  DstPrt SrcPrt Destination-Mac  Vlan
  EDst  ESrc  DPort  SPort  Stat-Pkts Stat-Bytes  Uptime   Age
-----
MSFC 172.20.25.1 (Module 15):
172.20.22.14   10.0.2.15      TCP   Telnet 37819   00-e0-4f-15-49-ff 51
  ARPA ARPA 5/39   5/40   115     5290         00:12:20 00:00:04
MSFC 172.20.27.1 (Module 16):

Total entries:1
Console> (enable)

```

Displaying Entries for a Specific IP Flow

The **show mls entry ip flow** command displays MLS entries for a specific IP flow. The *protocol* argument can be **tcp**, **udp**, **icmp**, or a decimal number for other protocol families. The *src_port* and *dst_port* arguments specify the protocol ports if the protocol is TCP or User Datagram Protocol (UDP). A value of zero (0) for *src_port*, *dst_port*, or *protocol* is treated as a wildcard and all entries are displayed (unspecified options are treated as wildcards). If the protocol selected is not TCP or UDP, set the *src_port* and *dst_prt* to 0 or no flows will display.

To display MLS entries for a specific IP flow (when the flow mask mode is full flow), perform this task in privileged mode:

Task	Command
Display entries for a specific IP flow (when the flow mask mode is full flow).	show mls entry ip flow [<i>protocol src_port dst_port</i>]

This example shows how to display MLS entries for a specific IP flow:

```

Console> (enable) show mls entry ip flow tcp 23 37819
Destination IP Source IP      Port DstPrt SrcPrt Destination Mac  Vlan Port
-----
MSFC 51.0.0.3:
10.0.2.15     51.0.0.2      TCP  37819   Telnet 08-00-20-7a-07-75 10  3/1
Console> (enable)

```

Displaying IPX MLS Entries for a Specific MSFC

To display IPX MLS entries for a specific MSFC, perform this task in privileged mode:

Task	Command
Display IPX MLS entries for a specific MSFC.	show mls entry ipx mod¹

1. The **mod** keyword specifies the module number of the MSFC; either 15 (if the MSFC is installed on Supervisor Engine 1 in slot 1) or 16 (if the MSFC is installed on Supervisor Engine 1 in slot 2).

This example shows how to display IPX MLS entries for a specific MSFC:

```

Console> (enable) show mls entry ipx 15
Destination-IPX      Destination-Mac      Vlan EDst ESrc  Port  Stat-Pkts
Stat-Bytes  Uptime  Age
-----
MSFC 22.1.0.56 (Module 15):
11.0000.0000.2B10      00-00-00-00-2b-10  11  ARPA ARPA  -    7869
361974      00:15:52  00:00:00
11.0000.0000.A810      00-00-00-00-a8-10  11  ARPA ARPA  -    3934
180964      00:15:52  00:00:00
11.0000.0000.3210      00-00-00-00-32-10  11  ARPA ARPA  -    7871
362066      00:15:52  00:00:00
11.0000.0000.B110      00-00-00-00-b1-10  11  ARPA ARPA  -    3935
181010      00:15:52  00:00:00
11.0000.0000.1910      00-00-00-00-19-10  11  ARPA ARPA  -    7873
362158      00:15:52  00:00:00
11.0000.0000.9A10      00-00-00-00-9a-10  11  ARPA ARPA  -    3936
181056      00:15:52  00:00:00
11.0000.0000.0010      00-00-00-00-00-10  11  ARPA ARPA  3/11  7875
362250      00:15:52  00:00:00
11.0000.0000.8310      00-00-00-00-83-10  11  ARPA ARPA  -    3937
181102      00:15:52  00:00:00
10.0000.0000.0109      00-00-00-00-01-09  10  ARPA ARPA  3/10  96364
4432744     00:15:52  00:00:00
11.0000.0000.4F10      00-00-00-00-4f-10  11  ARPA ARPA  -    7877
362342      00:15:53  00:00:00
11.0000.0000.CC10      00-00-00-00-cc-10  11  ARPA ARPA  -    3938
181148      00:15:53  00:00:00
11.0000.0000.5610      00-00-00-00-56-10  11  ARPA ARPA  -    7879
362434      00:15:53  00:00:00
11.0000.0000.D510      00-00-00-00-d5-10  11  ARPA ARPA  -    3939
181194      00:15:53  00:00:00
11.0000.0000.7D10      00-00-00-00-7d-10  11  ARPA ARPA  -    3940
181240      00:15:53  00:00:00
11.0000.0000.FE10      00-00-00-00-fe-10  11  ARPA ARPA  -    3941
181286      00:15:53  00:00:00
11.0000.0000.6410      00-00-00-00-64-10  11  ARPA ARPA  -    7883
362618      00:15:53  00:00:00
11.0000.0000.E710      00-00-00-00-e7-10  11  ARPA ARPA  -    3941
181286      00:15:53  00:00:00
11.0000.0000.6010      00-00-00-00-60-10  11  ARPA ARPA  -    7885
362710      00:15:53  00:00:00
11.0000.0000.E310      00-00-00-00-e3-10  11  ARPA ARPA  -    3942
181332      00:15:53  00:00:00
11.0000.0000.7910      00-00-00-00-79-10  11  ARPA ARPA  -    3943
181378      00:15:54  00:00:00

Console> (enable)

```

Clearing MLS Cache Entries

The **clear mls entry** command removes specific MLS cache entries. The **all** keyword clears all MLS entries. 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, and entries for all protocols are cleared (unspecified options are treated as wildcards).
- TCP or UDP source and destination port numbers (*src_port* and *dst_port*)—If the protocol 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, and entries for all source or destination ports are cleared (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 an MLS entry, perform this task in privileged mode:

Task	Command
Clear an MLS entry.	clear mls 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 MLS entries with destination IP address 172.20.26.22:

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

This example shows how to clear MLS entries with destination IP address 172.20.22.113, TCP source port 1652, and TCP destination port 23:

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

Clearing IPX MLS Cache Entries

The **clear mls entry ipx** command removes specific IPX MLS cache entries. The **destination** and **source** keywords specify the source and destination IPX addresses. The **all** keyword clears all MLS entries.

Displaying IP MLS Statistics

These sections describe how to display a variety of IP MLS statistics:

- [Displaying IP MLS Statistics by Protocol, page 14-27](#)
- [Displaying Statistics for MLS Cache Entries, page 14-27](#)

Displaying IP MLS Statistics by Protocol

The **show mls statistics protocol** command displays IP MLS statistics by protocol (such as Telnet, FTP, and WWW). The **protocol** keyword functions only if the flow mask mode is full flow. Enter the **show mls** command to see the current flow mask.

To display IP MLS statistics by protocol, perform this task in privileged mode:

Task	Command
Display IP MLS statistics by protocol (only if IP MLS is in full flow mode).	show mls statistics protocol

This example shows how to display IP MLS statistics by protocol:

```

Console> (enable) show mls statistics protocol
Protocol  TotalFlows  TotalPackets  Total Bytes
-----  -
Telnet    900          630           4298
FTP       688          2190          3105
WWW       389          42679         623686
SMTP      802          4966          92873
X         142          2487          36870
DNS       1580         52            1046
Others    82           1             73
Total     6583         53005         801951
Console> (enable)

```

Displaying Statistics for MLS Cache Entries

The **show mls statistics entry** command displays IP MLS statistics for MLS cache entries. Specify the destination IP address, source IP address, protocol, and source and destination ports to see specific MLS cache entries.

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

To display statistics for MLS cache entries, perform this task in privileged mode:

Task	Command
Display statistics for MLS cache entries. If you do not specify an MLS cache entry, all statistics are shown.	show mls statistics entry ip [destination ip_addr_spec] [source ip_addr_spec] [flow protocol src_port dst_port]

This example shows how to display statistics for a particular MLS cache 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>

```

Clearing MLS Statistics

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

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

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

Task	Command
Clear IP MLS statistics.	clear mls statistics

This example shows how to clear IP MLS statistics:

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

Displaying MLS Debug Information

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

To display MLS debug information, perform this task:

Task	Command
Display MLS 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 IP MMLS

These sections describe how to configure IP MMLS:

- [Configuring IP MMLS on the MSFC, page 14-28](#)
- [Displaying Global IP MMLS Information on the Supervisor Engine, page 14-34](#)

Configuring IP MMLS on the MSFC

These sections describe how to configure the MSFC for IP MMLS:

- [Enabling IP Multicast Routing Globally, page 14-29](#)
- [Enabling IP PIM on MSFC Interfaces, page 14-29](#)
- [Configuring the IP MMLS Global Threshold, page 14-30](#)
- [Enabling IP MMLS on MSFC Interfaces, page 14-30](#)

- [Displaying IP MMLS Interface Information](#), page 14-31
- [Displaying the IP Multicast Routing Table](#), page 14-31
- [Monitoring IP MMLS on the MSFC](#), page 14-32
- [Using Debug Commands on the IP MMLS MSFC](#), page 14-33
- [Using Debug Commands on the SCP](#), page 14-34

**Note**

For information on configuring routing on the MSFC, see [Chapter 12, “Configuring InterVLAN Routing.”](#)

**Note**

You can specify the MSFC as the MLS route processor (MLS-RP) for Catalyst 5000 family switches using MLS. Refer to the *Layer 3 Switching Configuration Guide—Catalyst 5000 Family, 2926G Series, 2926 Series Switches* for Catalyst 5000 family switch MLS configuration procedures.

**Note**

This section describes how to enable IP multicast routing on the MSFC. For more detailed IP multicast configuration information, refer to the “IP Multicast” section of the Cisco IOS IP and IP Routing Configuration Guide at http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt3/index.htm

Enabling IP Multicast Routing Globally

You must enable IP multicast routing globally on the MSFC before you can enable IP MMLS on MSFC interfaces.

To enable IP multicast routing globally on the MSFC, 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 MSFC Interfaces

You must enable IP PIM on the MSFC interfaces before IP MMLS will function on those interfaces.

To enable IP PIM on an interface, perform this task:

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

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

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

This example shows how to enable IP PIM sparse mode on an 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 (S,G) multicast traffic is routed by the MSFC. 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 MSFC Interfaces

IP MMLS is enabled by default on the MSFC 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 MSFC interfaces before IP MMLS will function. For information on configuring IP PIM on MSFC interfaces, see the [“Enabling IP PIM on MSFC Interfaces” section on page 14-29](#).

To enable IP MMLS on an MSFC interface, perform this task:

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

This example shows how to enable IP MMLS on an MSFC interface:

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

Use the **no** keyword to disable IP MMLS on an MSFC interface.

Displaying IP MMLS Interface Information

The **show ip pim interface count** command displays the IP MMLS enable state on MSFC IP PIM interfaces and the number of packets received and sent on the interface.

The **show ip interface** command displays the IP MMLS enable state on an MSFC interface.

To display IP MMLS information for an IP PIM MSFC 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

Displaying the IP Multicast Routing Table

The **show ip mroute** command displays the IP multicast routing table on the MSFC.

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 for 239.252.1.1:

```
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
```

Monitoring IP MMLS on the MSFC

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

To display detailed IP MMLS information on the MSFC, perform one of these tasks:

Task	Command
Display IP MMLS group information.	Router# show mls ip multicast group <i>group-address</i> [interface <i>type number</i> 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 <i>type number</i> statistics]

This example shows how to display IP MMLS statistics on the MSFC:

```
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 the MSFC:

```
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 the MSFC:

```
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 Debug Commands on the IP MMLS MSFC

Table 14-9 describes IP MMLS-related debug troubleshooting commands.

Table 14-9 IP MMLS Debug Commands

Command	Description
[no] debug mls ip multicast group <i>group_id</i> <i>group_mask</i>	Configures filtering that applies to all other multicast debugging commands.
[no] debug mls ip multicast events	Displays IP MMLS events.
[no] debug mls ip multicast errors	Turns on debug messages for multicast MLS-related errors.
[no] debug mls ip multicast messages	Displays IP MMLS messages from/to the hardware switching engine.
[no] debug mls ip multicast all	Turns on all IP MMLS messages.
[no] debug mdss error	Turns on MDSS ¹ error messages.
[no] debug mdss events	Turns on MDSS-related events.
[no] debug mdss all	Turns on all MDSS messages.

1. MDSS = Multicast Distributed Switching Services

Using Debug Commands on the SCP

Table 14-10 describes the Serial Control Protocol (SCP)-related debug commands to troubleshoot the SCP that runs over the Ethernet out-of-band channel (EOBC).

Table 14-10 SCP Debug Commands

Command	Description
[no] debug scp async	Displays trace for asynchronous data in and out of the SCP system.
[no] debug scp data	Shows packet data trace.
[no] debug scp errors	Displays errors and warnings in the SCP.
[no] debug scp packets	Displays 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 Global IP MMLS Information on the Supervisor Engine

These sections describe how to configure IP MMLS on Supervisor Engine 1:

- [Displaying IP MMLS Configuration Information, page 14-34](#)
- [Displaying IP MMLS Statistics, page 14-35](#)
- [Clearing IP MMLS Statistics, page 14-36](#)
- [Displaying IP MMLS Entries, page 14-36](#)



Note

IP MMLS is permanently enabled on Supervisor Engine 1 and cannot be disabled.



Note

To configure IP MMLS on the MSFC, see the [“Configuring IP MMLS on the MSFC”](#) section on [page 14-28](#).

Displaying IP MMLS Configuration Information

The **show mls multicast** command displays global IP MMLS configuration information and the state of participating MSFCs.

To display global IP MMLS configuration information, perform this task:

Task	Command
Display global IP MMLS configuration information.	show mls multicast

This example shows how to display global IP MMLS configuration information:

```
Console> (enable) show mls multicast
Admin Status: Enabled
Operational Status: Active
Configured flow mask is {Destination-source-vlan flow}
Active Entries = 10
Router include list :
1.1.9.254 (Active)
1.1.5.252 (Active)
Console> (enable)
```

Displaying IP MMLS Statistics

The **show mls multicast statistics** command displays IP MMLS statistics for multicast MSFCs.

To display IP MMLS statistics for multicast MSFCs, perform this task:

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

This example shows how to display IP MMLS statistics for multicast MSFCs:

```
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 IP MMLS Statistics

The **clear mls multicast statistics** command clears IP MMLS statistics for all participating MSFCs.

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

Task	Command
Clear IP MMLS statistics.	clear mls multicast statistics

This example shows how to clear IP MMLS statistics:

```

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

```

Displaying IP MMLS Entries

The **show mls multicast entry** command displays a variety of information about the multicast flows being handled by the PFC. You can display entries based on any combination of the participating MSFC, the VLAN, the multicast group address, or the multicast traffic source.

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

Task	Command
Display information about IP MMLS 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 IP MMLS 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 IP MMLS entries for a specific MSFC:

```

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 IP MMLS 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 IP MMLS entries for a specific MSFC and a specific multicast source address:

```

Console> (enable) show mls multicast entry 15 1.1.5.252 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)

```

