CHAPTER 20

Configuring IPX Unicast Layer 3 Switching on Supervisor Engine 1

Note The features described in this chapter are supported only on Supervisor Engine 1, PFC, and MSFC or MSFC2. For information about Supervisor Engine 2, PFC2, and MSFC2 see Chapter 17, "Configuring IP Unicast Layer 3 Switching on Supervisor Engine 2."

Supervisor Engine 1 with PFC and MSFC or MSFC2 provide Layer 3 switching with Multilayer Switching (MLS). This chapter describes how to configure Internetwork Packet Exchange (IPX) Layer 3 switching on the Catalyst 6500 series switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 6500 Series Switch Cisco IOS Command Reference* publication.

This chapter consists of these sections:

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, 4000 Family, 2926G Series, 2926 Series, and 2948G* for MLS configuration procedures.

Understanding How IPX MLS Works

These sections provide an overview of MLS and describe how MLS works:

IPX MLS Overview

IPX MLS provides high-performance hardware-based Layer 3 switching for Catalyst 6500 series switches. IPX MLS switches unicast IPX data packet flows between networks using advanced application-specific integrated circuit (ASIC) switching hardware, offloading processor-intensive packet routing from network routers.

The packet forwarding function is moved onto Layer 3 switches whenever a partial or complete switched path exists between two hosts. Packets that do not have a partial or complete switched path to reach their destinations are still forwarded by routers. Standard routing protocols, such as Routing Information Protocol (RIP), Enhanced Interior Gateway Protocol (EIGRP), and NetWare Link Services Protocol (NLSP), are used for route determination.

IPX 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.

A flow is a unidirectional sequence of packets between a particular source and destination that share the same protocol and network-layer information. Communication from a client to a server and from the server to the client are separate flows.

Flows are based only on Layer 3 addresses, which allow IPX traffic from multiple users or applications to a particular destination to be carried on a single flow if only the destination IPX address is used to identify a flow.

Layer 3 MLS Cache

The Policy Feature Card (PFC) maintains a Layer 3 switching table (MLS cache) for the Layer 3-switched flows. The cache includes entries for traffic statistics that are updated as packets are switched. After the MLS cache is created, 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.

An IPX MLS cache entry is created for the initial packet of each flow. Upon receipt of a packet that does not match any flow currently in the MLS cache, a new IPX MLS entry is created.

The state and identity of the flow are maintained while packet traffic is active; when traffic for a flow ceases, the entry ages out. You can configure the aging time for IPX 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.

The maximum MLS cache size is 128K entries. However, an MLS cache larger than 32K entries increases the probability that a flow will not be switched by the PFC and will get forwarded to the MSFC.

Flow Masks

The PFC uses flow mask modes to determine how IPX MLS entries are created. The flow mask mode is based on the access lists configured on the IPX MLS router interfaces.

These sections describe how the flow mask modes work:

- Flow Mask Modes, page 20-3
- Flow Mask Mode and show mls entry Command Output, page 20-3

Flow Mask Modes

The PFC supports only one flow mask (the most specific one). 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 mode. Depending on the current mode, some fields in the flow record might not have values. Unsupported fields are filled with a dash (-).

The flow mask modes for IPX MLS are as follows:

- **destination mode**—The least-specific flow mask 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. This mode is used if there are no access lists configured based on source IPX addresses on any of the IPX MLS router interfaces.
- **destination-source mode**—The PFC maintains one MLS entry for each destination (network and node) and source (network only) IPX address pair. All flows between a given source and destination use this MLS entry regardless of the IPX sockets. This mode is used if there is an access list on any of the IPX MLS interfaces that filters on source network.

Flow Mask Mode and show mls entry Command Output

The flow mask mode impacts the screen output of the **show mls ipx** command. In destination mode, the destination IPX address of the switched flows are displayed, along with the packet rewrite information: rewritten destination MAC, rewritten VLAN, and egress interface.

Layer 3-Switched Packet Rewrite

When a packet is Layer 3 switched from a source host to a destination host, the PFC performs a packet rewrite, based on information learned from the MSFC and stored in the MLS cache.

If Host A and Host B are on different VLANs and Host A sends a packet to the MSFC to be routed to Host B, the PFC recognizes that the packet was sent to the MAC address of the MSFC. The PFC checks the MLS cache and finds the entry matching the flow in question.

Received IPX packets are formatted (conceptually) 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		
MSFC MAC	*Source A MAC*	*n*	*Destination B IPX*	*Source A IPX*		

The PFC rewrites the Layer 2 frame header, changing the destination MAC address to the MAC address of Host B and the source MAC address to the MAC address of the MSFC (these MAC addresses are stored in the IPX MLS cache entry for this flow). The Layer 3 IPX addresses remain the same. The PFC rewrites the switched Layer 3 packets so that they appear to have been routed by a router.

The PFC forwards the rewritten packet to Host B's VLAN (the destination VLAN is saved in the IPX MLS cache entry) and Host B receives the packet.

After the switch rewrites an IPX packet, it is formatted (conceptually) 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		
Destination B MAC	*MSFC MAC*	*n+1*	*Destination B IPX*	*Source A IPX*		

IPX MLS Operation

Figure 20-1 displays a conceptual 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 C, an IPX MLS entry for this flow is created (this entry is the second item in the table shown in Figure 20-1). The PFC stores the MAC addresses of the MSFC and Host C in the IPX 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 Station A to Station C.

Similarly, a separate IPX MLS entry is created in the MLS cache for the traffic from Host A to Host B, and for the traffic from Host B 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 20-1 IPX MLS Example Topology

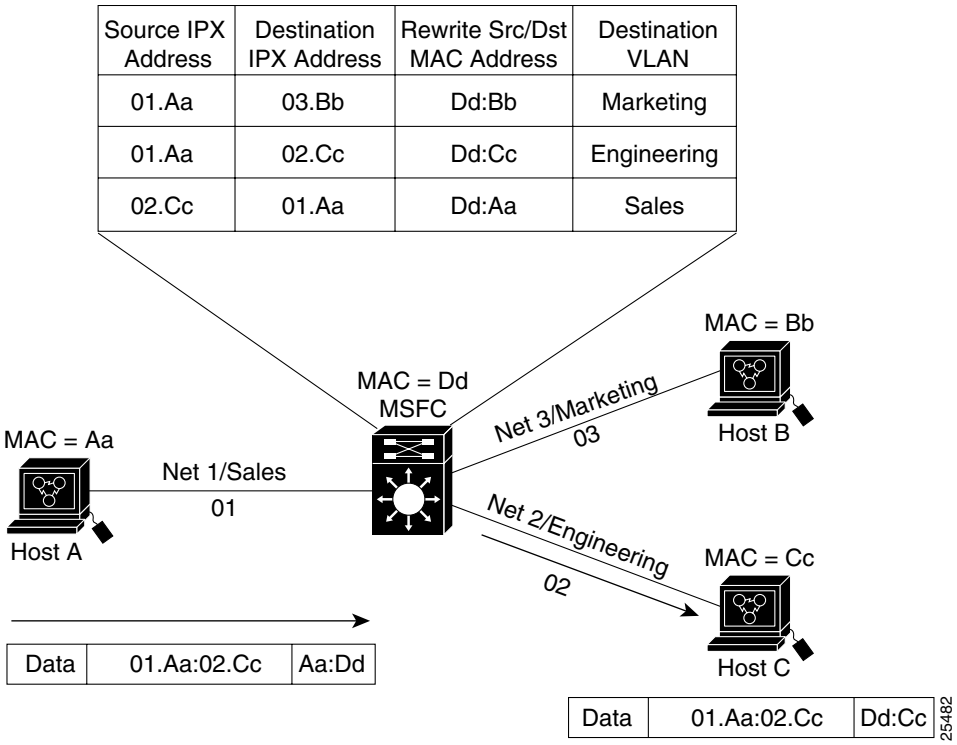

Default IPX MLS Configuration

Table 20-1 displays the default IPX MLS configuration.

Table 20-1 Default IPX MLS Configuration

Feature	Default Value
IPX MLS enable state	Enabled
IPX MLS aging time	256 seconds
IPX MLS fast aging time	32 seconds
IPX MLS fast aging time packet threshold	100 packets
IPX MLS long aging time	900 seconds

Configuration Guidelines and Restrictions

- These Cisco IOS software features and commands affect IPX MLS:
  - IPX accounting—IPX accounting cannot be enabled on an IPX MLS-enabled interface.
  - IPX EIGRP—MLS is supported for EIGRP interfaces if the Transport Control (TC) maximum is set to a value greater than the default (16).
  - The **clear ipx route** command clears all IPX MLS cache entries.

  - The **no ipx routing** command purges all IPX MLS cache entries and disables IPX MLS.
  - The **ipx security** interface command disables IPX MLS on the interface.
- In IPX, the two end points of communication negotiate the maximum transmission unit (MTU) to be used. MTU size is limited by media type.

Configuring IPX MLS

These sections describe how to configure IPX MLS:

Note For information on configuring VLANs on the switch, see Chapter 7, "Configuring LAN Ports for Layer 2 Switching."

Note With Release 12.1(11b)E and later, when you are in configuration mode you can enter EXEC mode-level commands by entering the **do** keyword before the EXEC mode-level command.

Enabling IPX MLS Globally

IPX MLS is enabled globally and cannot be disabled.

Enabling IPX MLS on a Layer 3 Interface

To enable IPX MLS on an interface, perform this task:

	Command	Purpose
Step 1	`Router(config)# interface {{vlan vlan_ID} | {type[1] slot/port} | {port-channel number}}`	Selects the interface to configure.
Step 2	`Router(config-if)# mls ipx`	Enables IPX MLS.
	`Router(config-if)# no mls ipx`	Disables IPX MLS.

1. *type* = **ethernet**, **fastethernet**, **gigabitethernet**, or **tengigabitethernet**

This example shows how to enable IPX MLS for Fast Ethernet interface 5/5:

```
Router(config)# interface fastethernet 5/5
Router(config-if)# mls ipx
Router(config-if)#
```

To display the IPX MLS interface configuration, perform this task:

Command	Purpose
Router# show [**ipx** [**interface** {{**vlan** *vlan_ID*} | {*type*[1] *slot/port*} | {**port-channel** *number*}}] | **nde**]	Displays MLS details for an interface.

1. *type* = **ethernet**, **fastethernet**, **gigabitethernet**, or **tengigabitethernet**

This example shows how to display interface IPX MLS information for interface VLAN 200:

```
Router# show ipx interface vlan 200
Vlan200 is up, line protocol is up
  IPX address is 2.0050.3e8d.6400, NOVELL-ETHER [up]
  Delay of this IPX network, in ticks is 1 throughput 0 link delay 0
  IPXWAN processing not enabled on this interface.
  IPX SAP update interval is 60 seconds
  IPX type 20 propagation packet forwarding is disabled
  Incoming access list is not set
  Outgoing access list is not set
  IPX helper access list is not set
  SAP GGS output filter list is not set
  SAP GNS processing enabled, delay 0 ms, output filter list is not set
  SAP Input filter list is not set
  SAP Output filter list is not set
  SAP Router filter list is not set
  Input filter list is not set
  Output filter list is not set
  Router filter list is not set
  Netbios Input host access list is not set
  Netbios Input bytes access list is not set
  Netbios Output host access list is not set
  Netbios Output bytes access list is not set
  Updates each 60 seconds aging multiples RIP: 3 SAP: 3
  SAP interpacket delay is 55 ms, maximum size is 480 bytes
  RIP interpacket delay is 55 ms, maximum size is 432 bytes
  RIP response delay is not set
  IPX accounting is disabled
  IPX fast switching is configured (enabled)
  RIP packets received 0, RIP packets sent 1, 0 Throttled
  RIP specific requests received 0, RIP specific replies sent 0
  RIP general requests received 0, 0 ignored, RIP general replies sent 0
  SAP packets received 0, SAP packets sent 1, 0 Throttled
  SAP GNS packets received 0, SAP GNS replies sent 0
  SAP GGS packets received 0, 0 ignored, SAP GGS replies sent 0
  IPX mls switching is enabled
Router#
```

Configuring the MLS Aging Time

The MLS aging time applies to all MLS cache entries. See the "Configuring the MLS Aging Time" section on page 33-10.

Note IPX MLS does not use fast aging.

Configuring the Minimum IPX MLS Flow Mask

You can configure 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 "Flow Masks" section on page 20-3.

Caution This command purges all existing shortcuts in the MLS cache and affects the number of active shortcuts on the PFC. Be careful when using this command.

To configure the minimum IPX MLS flow mask, perform this task:

Command	Purpose
`Router(config)# mls flow ipx {destination | destination-source}`	Configures the minimum IPX MLS flow mask.
`Router(config)# no mls flow ipx`	Reverts to the default IPX MLS flow mask.

This example displays how to configure the minimum IPX MLS flow mask to destination:

```
Router(config)# mls flow ipx destination
Router(config)#
```

To display the IPX MLS flow mask configuration, perform this task:

Command	Purpose
`Router# show mls flowmask`	Displays the flow mask configuration.

This example displays the MLS flow mask configuration:

```
Router# show mls flowmask
current ip flowmask for unicast: destination only
current ipx flowmask for unicast: destination only
```

Displaying IPX MLS Information

This section describes the commands used to display IPX MLS configuration and statistics for the switch and the various interfaces and is separated into the following:

- Displaying IPX MLS Cache Entries, page 20-9
- Displaying the IPX MLS Contention Table, page 20-11
- Displaying IPX MLS VLAN Statistics, page 20-12

Displaying IPX MLS Cache Entries

The **show mls ipx** command displays IPX shortcut cache entries. You can specify the following parameters to focus the information displayed:

- **source** and **destination** parameters display the source and or destination IPX network addresses associated with those entries.
- **interface** arguments display only entries associated with a specific interface number.
- **slot** displays only the cache entries associated with a specific slot number.
- **macs** and **macd** arguments and the IPX network address display the source and destination MAC addresses associated with those entries.

To display all IPX MLS entries on the switch, perform this task:

Command	Purpose
`Router# show mls ipx [destination ipx_network_address | interface {{vlan vlan_ID} | {type[1] slot/port} | {port-channel number}} | macd destination_mac_address | macs source_mac_address | source ipx_network_address]]`	Displays various IPX MLS cache entries.

1. *type* = **ethernet**, **fastethernet**, **gigabitethernet**, or **tengigabitethernet**

These sections provide examples of how to display specific IPX MLS cache entries on the switch:

- Displaying All IPX MLS Cache Entries, page 20-9
- Displaying IPX MLS Cache Entries for a Specific Destination Address, page 20-10
- Displaying IPX MLS Cache Entries for a Specific Source Address, page 20-10
- Displaying IPX MLS Cache Entries for a Specific Interface, page 20-11
- Displaying IPX MLS Cache Entries for a Specific MAC Destination or Source Address, page 20-11

Displaying All IPX MLS Cache Entries

To display all IPX MLS cache entries on the switch, perform this task:

Command	Purpose
`Router# show mls ipx`	Displays all IPX MLS entries.

This example shows how to display all IPX MLS entries on the switch:

```
Router# show mls ipx
DstNet-DstNode          SrcNet   DstVlan-DstMac      Pkts          Bytes
------------------------------------------------------------------------
SrcDstPorts   SrcDstEncap Age   LastSeen
---------------------------------------

 Number of Entries Found = 0

Router#
```

Displaying IPX MLS Cache Entries for a Specific Destination Address

To display IPX MLS cache entries for a specific destination IPX address, perform this task:

Command	Purpose
Router# **show mls ipx** [**destination** *ipx_addr*]	Displays IPX MLS entries for a specific destination address (*net_address.node_address*).

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

```
Router# show mls ipx destination 1.2.2.2
DstNet-DstNode          SrcNet   DstVlan-DstMac      Pkts          Bytes
------------------------------------------------------------------------
SrcDstPorts   SrcDstEncap Age   LastSeen
---------------------------------------

 Number of Entries Found = 0

Router#
```

Displaying IPX MLS Cache Entries for a Specific Source Address

To display IPX MLS cache entries for a specific source network address, perform this task:

Command	Purpose
Router# **show mls ipx source** *ipx_address*	Displays IPX MLS entries for a specific source network address (net_address).

Note This task should be performed with IPX flow in destination-source mode. For more information, see the "Flow Mask Modes" section on page 20-3.

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

```
Router# show mls ipx source 1.2.2.2
DstNet-DstNode          SrcNet   DstVlan-DstMac      Pkts          Bytes
------------------------------------------------------------------------
SrcDstPorts   SrcDstEncap Age   LastSeen
---------------------------------------

 Number of Entries Found = 0

Router#
```

Displaying IPX MLS Cache Entries for a Specific Interface

To display IPX MLS entries for a specific interface, perform this task:

Command	Purpose
Router# **show mls ipx interface** {{**vlan** *vlan_ID*} | {*type*[1] *slot/port*} | {**port-channel** *number*}}	Displays IPX MLS cache entries for a specific interface.

1. *type* = **ethernet**, **fastethernet**, **gigabitethernet**, or **tengigabitethernet**

This example shows how to display IPX MLS entries for IPX Fast Ethernet interface 5/4:

```
Router# show mls ipx interface fastethernet 5/4
DstNet-DstNode          SrcNet   DstVlan-DstMac      Pkts          Bytes
------------------------------------------------------------------------
SrcDstPorts   SrcDstEncap Age   LastSeen
---------------------------------------

 Number of Entries Found = 0

Router#
```

Displaying IPX MLS Cache Entries for a Specific MAC Destination or Source Address

To display IPX MLS entries for a specific MAC destination or source address, perform this task:

Command	Purpose
Router# **show mls ipx** [**macd** *destination_address* | **macs** *source_address*]	Displays IPX MLS cache entries for a specific destination or source MAC address.

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

```
Router# show mls ipx macd aaaa.bbbb.bbbb
DstNet-DstNode          SrcNet   DstVlan-DstMac      Pkts          Bytes
------------------------------------------------------------------------
SrcDstPorts   SrcDstEncap Age   LastSeen
---------------------------------------

 Number of Entries Found = 0

Router#
```

Displaying the IPX MLS Contention Table

The **show mls table-contention** command displays the flow contention level for the switch. The table contention level (TCL) is indicated with a number ranging from 0 (normal) to 3 (maximum). When reaching levels 1 through 3, accelerated aging starts, which begins to age out entries at a rate suitable to reduce the current contention rate. The detailed option displays the breakdown of contention between different flows.

To show the MLS contention table and VLAN statistics for the switch, perform this task:

Command	Purpose
Router# **show mls table-contention** [**detailed** | **summary**]	Displays the IPX MLS contention table.

This example displays the IPX MLS contention table for the switch:

```
Router# show mls table-contention detailed
Detailed Table Contention Level Information
===========================================
Layer 3
-------
L3 Contention Level:     0
Page Hits Requiring 1 Lookup    =        10
Page Hits Requiring 2 Lookups   =        0
Page Hits Requiring 3 Lookups   =        0
Page Hits Requiring 4 Lookups   =        0
Page Hits Requiring 5 Lookups   =        0
Page Hits Requiring 6 Lookups   =        0
Page Hits Requiring 7 Lookups   =        0
Page Hits Requiring 8 Lookups   =        0
Page Misses                     =        0

Router#
```

Displaying IPX MLS VLAN Statistics

The **show mls vlan-statistics** command displays VLAN-based statistics for IPX MLS cache entries. Specifying a VLAN identifier results in a display with only the shortcuts for that VLAN. If you specify a slot, only the information about that slot is shown; otherwise, all entries are displayed.

To display the IPX MLS VLAN statistics for the switch, perform this task:

Command	Purpose
Router# **show mls vlan-statistics** *1-1024*	Displays the IPX MLS VLAN statistics.

This example displays the VLAN statistics for VLAN 1 for every slot in the switch:

```
Router# show mls vlan-statistics 1
Slot 0
=======
Vlan 1 Statistics Information:
------------------------------
65280 Layer 2 Packets Bridged, 0 Bytes
65280 Layer 3 Packets Input, 0 Bytes
65280 Layer 3 Packets Output, 0 Bytes
Slot 1
=======
Vlan 1 Statistics Information:
------------------------------
65280 Layer 2 Packets Bridged, 0 Bytes
65280 Layer 3 Packets Input, 0 Bytes
65280 Layer 3 Packets Output, 0 Bytes
Slot 2
=======
```

```
Vlan 1 Statistics Information:
------------------------------
65280 Layer 2 Packets Bridged, 0 Bytes
65280 Layer 3 Packets Input, 0 Bytes
65280 Layer 3 Packets Output, 0 Bytes
Slot 3
=======
Vlan 1 Statistics Information:
------------------------------
65280 Layer 2 Packets Bridged, 0 Bytes
65280 Layer 3 Packets Input, 0 Bytes
65280 Layer 3 Packets Output, 0 Bytes
Slot 4
=======
Vlan 1 Statistics Information:
------------------------------
65280 Layer 2 Packets Bridged, 0 Bytes
65280 Layer 3 Packets Input, 0 Bytes
65280 Layer 3 Packets Output, 0 Bytes
Slot 5

(Information Deleted)
```

Clearing IPX MLS Cache Entries

Clear IPX shortcut entries in the Layer 3 table based on the entered criteria. The **clear mls ipx** command clears shortcut entries in the Layer 3 tables matching configured parameters. If none of the following parameters are entered, all IPX Layer 3 entries in the table are cleared:

- **destination** or **source—**Describes the IPX addresses of the origin and termination points being purged.
- **interface** and its arguments must be specified, which limits the purge to entries associated with the specified interface.
- **macd** (MAC **destination**) or **macs** (MAC **source**)—Specifies the source port or destination interface arguments to consider when searching for entries to purge.
- **slot**—Clears only the entries associated with a specific slot number.

To clear the IPX MLS statistics, perform this task:

Command	Purpose
Router# **clear mls ipx** [**exclude protocol** [**all port** *1-96* | **tcp port** *1-96* | **udp port** *1-96*] | [**destination** [*hostname* | *ipx_address*] **interface** {{**vlan** *vlan_ID*} | {*type*[1] *slot/port*} | {**port-channel** *number*}} | **macd** *destination_mac_address* | **macs** *source_mac_address*]]	Clears the IPX MLS statistics.

1. *type* = **ethernet**, **fastethernet**, **gigabitethernet**, or **tengigabitethernet**

This example shows how to clear the MLS cache entries in the Layer 3 table IPX interface, Fast Ethernet interface 5/5:

```
Router# clear mls ipx interface fastethernet 5/5
Router#
```

To display the MLS entries and confirm they have been cleared, see the "Displaying IPX MLS Information" section on page 20-8.

Troubleshooting IPX MLS

Table 20-2 describes debug commands that you can use to troubleshoot IPX MLS problems.

Table 20-2 IPX 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 bug trace of IP global purge events.
[no] debug l3-mgr all	Turns on all Layer 3 manager debugging messages.
[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 locator	Identifies which switch is switching a particular flow by using MLS explorer packets.
[no] debug mls all	Turns on all MLS debugging events.

To configure the debug commands that you can use to troubleshoot IPX MLS problems, perform this task:

Command	Purpose
Router(config)# debug mls {ipx {all | error | events | messages} | rp {all | error| events | ip | ipx | locator | packets | verbose}}	Configures IPX MLS debugging.
Router(config)# {no | undebug} mls {all | ipx {all | error | events | messages} | rp {all | error| events | ipx | locator | packets | verbose}}	Disables MLS debugging.

This example displays how to configure all IPX debugging:

```
Router# debug mls ipx all
mls ip all debugging is on
Router#
```

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

Table 20-3 describes the Serial Control Protocol (SCP)-related debug commands to troubleshoot the SCP that runs over the Ethernet out-of-band channel (EOBC).

Table 20-3 SCP Debug Commands

Command	Description
[no] debug scp async	Displays trace for async 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 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.