

Configuring and Troubleshooting IP MLS on Catalyst 6500/6000 Switches with an MSFC

[TAC Notice: What's Changing on TAC Web](#)

Contents

[Introduction](#)

[Before You Begin](#)

[Conventions](#)

[Prerequisites](#)

[Components Used](#)

[Overview of MLS on the MSFC](#)

[Example 1: MSFC Routing Between Two VLANs](#)

[Reading show mls Entries and Basic Checking](#)

[Reading show mls Output Tables](#)

[Troubleshooting a Specific Connection](#)

[Example 2: Two MSFCs in the Same Chassis Configured for HSRP](#)

[Issuing show mls Commands](#)

[Issuing show mls entry Commands](#)

[Troubleshooting Tips](#)

[Creation of Flows](#)

[IP Accounting excludes MLS traffic](#)

[Unable to configure interface-full Flowmask](#)

[Related Information](#)

Help us help you.

Please rate this document.

Excellent

Good

Average

Fair

Poor

This document solved my problem.

Yes

No

Just browsing

Suggestions for improvement:

(256 character limit)

Introduction

This document is a guide on verifying and how to read Multilayer Switching (MLS) commands on the Catalyst 6500/6000. The document contains a very brief review of what MLS is, and also gives examples of how to use MLS. Based on these examples, this document shows how to verify the MLS operation, and provides brief troubleshooting tips for configuring MLS.

This document applies only to the Catalyst 6500/6000 series switch equipped with the following hardware:

- Supervisor Engine 1A running Catalyst OS (CatOS) Software
- Policy Feature Card (PFC)

- Multilayer Switch Feature Card (MSFC)

Note: This document is not valid when using any other hardware configuration such as Supervisor Engine 2 or Multilayer Switch Module (MSM). It is also not valid when running Cisco IOS® Software on both the Supervisor Engine 1A and MSFC.

For similar information regarding troubleshooting unicast routing on a Catalyst 6500/6000 series switch with Supervisor Engine 2 and running CatOS software, refer to [Troubleshoot Unicast IP Routing Involving CEF on Catalyst 6500/6000 Series Switches with a Supervisor Engine 2 and Running CatOS System Software](#).

For a more complete description of MLS terminology and operation, see the [Related Information](#) section.

Before You Begin

Conventions

For more information on document conventions, see the [Cisco Technical Tips Conventions](#).

Prerequisites

There are no specific prerequisites for this document.

Components Used

The information in this document is based on the software and hardware versions below.

- Catalyst 6500/6000 with an MSFC

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

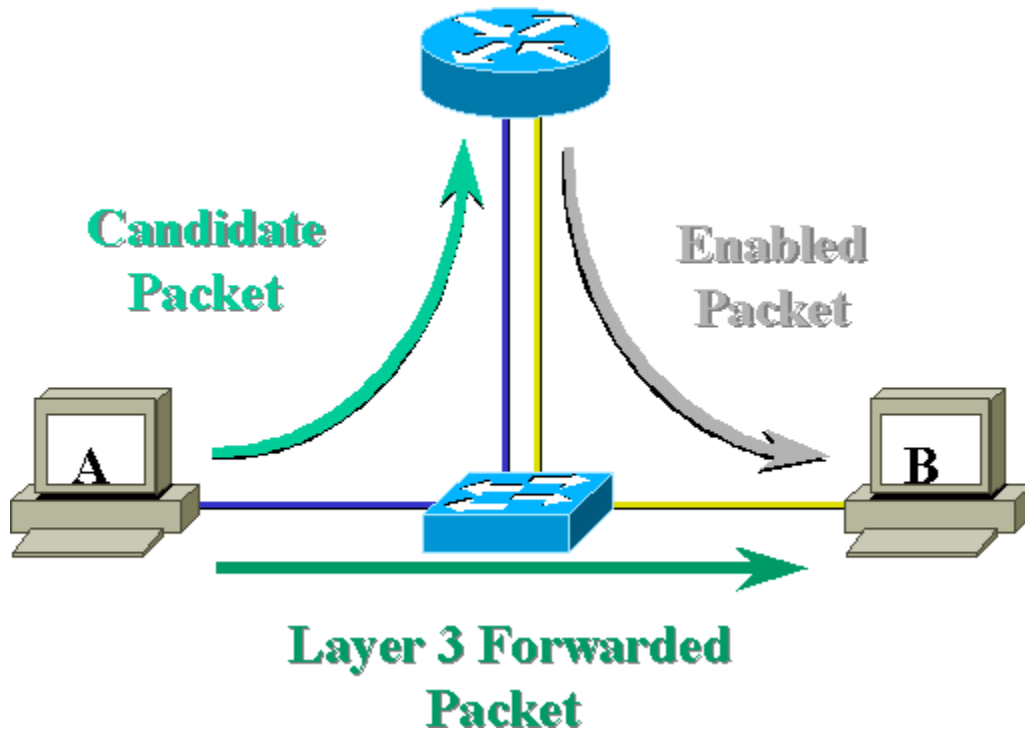
Overview of MLS on the MSFC

The MSFC is the second generation routing engine for the Catalyst 6500/6000 series switch that can route 15 million packets per second. The MSFC only works with Supervisor Engines that have the PFC. The MSFC performs an internal MLS with the PFC, which acts similar to the NetFlow Feature Card (NFFC) on the Catalyst 5000. This internal MLS is not visible and is purely limited to the switch: you have nothing to configure to make it work and it supports hardware shortcuts for IP, IPX, and IP multicast. The configuration of the MSFC is similar to the configuration of an RSM or an RSFC using VLAN interfaces. You can access it using session 15 (for MSFC on the Supervisor Engine in slot 1) or session 16 (for MSFC on the Supervisor Engine in slot 2).

The principle is similar to Multilayer Switching Protocol (MLSP) on the Catalyst 5000. The first packet is routed by the MSFC, and the PFC creates a shortcut that is used by all subsequent packets of the same flow. Unlike MLSP on the Catalyst 5000, which requires IP communication between the MLS-SE and

the MLS-RP, MLS on the Catalyst 6500/6000 works by communication between the MSFC and PFC over a serial channel (SCP).

The PFC cannot be the MLS-SE for a Catalyst 5000 MLS environment; however, the MSFC can be MLS-RP for other Catalyst 5000s in the network. In that case, you should configure the MSFC using the same **mls rp ip** command as you would for any Cisco IOS router used as the MLS RP.



Example 1: MSFC Routing Between Two VLANs

The MLS on the Catalyst 6500/6000 for unicast IP is plug and play. You do not need to configure it. Below is a sample configuration where **tamer** is a Catalyst 6500/6000, which has an MSFC called **tamer-msfc**. Routing between VLAN 11 and 12 is configured on the MSFC without a single command related to MLS. Keep in mind that the Supervisor Engine will not have any MLS-specific configuration.

```
tamer-msfc#wr t
Building configuration...

Current configuration:
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname tamer-msfc
!
boot system flash bootflash:c6msfc-ds-mz.121-1.E2.bin
!
ip subnet-zero
ip cef
```

```

!
interface Vlan11
 ip address 11.1.1.2 255.255.255.0
!
interface Vlan12
 ip address 12.1.1.2 255.255.255.0
!
router eigrp 1
 network 11.0.0.0
 network 12.0.0.0
 no auto-summary
!
ip classless
no ip http server
!
!
line con 0
 transport input none
line vty 0 4
 login
!
end

```

Reading show mls Entries and Basic Checking

Below is the output that you should get when issuing the **show mls** command on the Supervisor Engine.

```

tamer (enable) show mls
Total packets switched = 8

!--- Number of packet shortcutcd by MLS.

Total Active MLS entries = 0

!--- Number of flows currently existing.

MSFC 11.1.1.2 (Module 15) entries = 0
IP Multilayer switching aging time = 256 seconds
IP Multilayer switching fast aging time = 0 seconds, packet threshold = 0
IP Current flow mask is Destination flow
Active IP MLS entries = 0
Netflow Data Export version: 7
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0

IP MSFC ID      Module XTAG MAC                               Vlans
-----
11.1.1.2        15      1    00-d0-d3-9c-9e-3c 12,11

!--- MSFC recognized by the switch.

```

The **show mls** command output should always have an IP MSFC ID. If you do not see the IP MSFC ID when you issue the **show mls** command, verify the following:

- MSFC is up and running (not stuck in ROMmon mode, for example).

- MLS is still enabled on the MSFC.

You can verify this by issuing the following commands on the MSFC:

```
tamer-msfc#show mls status
MLS global configuration status:
global mls ip:                          enabled

!--- Should be enabled for unicast IP.

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

By issuing the **show mls status** command, you can determine whether MLS is enabled for IP, IPX, and IP multicast. These features should always be enabled by default; however, they can be disabled by issuing the following command in configuration mode:

```
no mls ip
```

The **no mls ip** command should only be used for debugging purposes. The command is also available as a hidden command in global configuration mode. You can also disable MLS on a per-VLAN interface basis by issuing the **no mls ip** command in interface configuration mode

Note: Do not issue the **show mls rp** command on the MSFC. This command output indicates that MLS is disabled. However, the **show mls** command output issued on the Supervisor Engine above indicated that MLS was working correctly. This discrepancy occurs because the **show mls rp** command should be used when doing MLS-rp in conjunction with a Catalyst 5000 switch.

```
tamer-msfc#show mls rp
ip multilayer switching is globally disabled
ipx multilayer switching is globally disabled
ipx mls inbound acl override is globally disabled
mls id is 00d0.d39c.9e04
mls ip address 0.0.0.0
mls ip flow mask is unknown
mls ipx flow mask is unknown
number of domains configured for mls 0
```

A candidate packet is a packet that can potentially initiate the setup of a MLS shortcut. Its destination MAC address is equal to the MAC address of a router running MLS. In this case, the MAC address of the MSFC is 00-d0-d3-9c-9e-3c (seen by issuing the **show mls** command). To verify that the switch knows that this MAC address is a router MAC address, issue the **show cam mac_address** command, as shown below.

```
tamer (enable) show cam 00-d0-d3-9c-9e-3c
* = Static Entry. + = Permanent Entry. # = System Entry. R = Router Entry.
X = Port Security Entry
VLAN  Dest MAC/Route Des      [CoS]  Destination Ports or VCs / [Protocol Type]
-----
11    00-d0-d3-9c-9e-3c R#          15/1
12    00-d0-d3-9c-9e-3c R#          15/1
Total Matching CAM Entries Displayed = 2
```

This output confirms that the switch knows that this MAC address is a router entry linked to port 15/1 (MSFC port in slot 1).

If you still do not see the MSFC when you issue the command **show mls** on the switch Supervisor Engine, issue the following command:

```
tamer (enable) show mls rlog 12
SWLOG at 815d0c50: magic 1008, size 51200, cur 815d4170, end 815dd460
Current time is: 08/08/00,17:13:25

118 08/08/00,17:13:16:(RouterConfig)Router_cfg: router_add_MAC_to_earl 00-d0-d3
9c-9e-3c added for mod 15/1 Vlan 12 Earl AL =0
117 08/08/00,17:13:16:(RouterConfig)Router_Cfg: Process add mls entry for mod 1
/1 vlan 12, i/f 1, proto 0, LC 3
116 08/08/00,17:13:16:(RouterConfig)Router_cfg: router_add_MAC_to_earl 00-d0-d3
9c-9e-3c added for mod 15/1 Vlan 11 Earl AL =0
115 08/08/00,17:13:16:(RouterConfig)Router_Cfg: Process add mls entry for mod 1
/1 vlan 11, i/f 1, proto 0, LC 3
```

This command shows you the messages the switch is receiving from the MSFC, and that the router entries are added.

Reading show mls Output Tables

Issue the **show mls entry** command to view the full MLS table with all the shortcuts. The output below shows all of the flows received.

```
tamer (enable) show mls entry
Destination-IP  Source-IP      Prot  DstPrt  SrcPrt  Destination-MAC  Vlan  EDst
-----
MSFC 11.1.1.2 (Module 15):
10.68.5.1      -              -     -       -       00-d0-00-3f-a3-ff 8      ARPA
12.1.1.1      -              -     -       -       00-00-0c-8c-70-88 12     ARPA
11.1.1.1      -              -     -       -       00-00-0c-09-50-66 11     ARPA
ESrc DPort  SPort  Stat-Pkts  Stat-Bytes  Uptime  Age
-----
ARPA 1/3    7/3    4          400         00:00:02 00:00:02
ARPA 7/4    7/3    4          400         00:00:08 00:00:08
ARPA 7/3    7/4    9          900         00:00:08 00:00:08

Destination-IPX      Destination-MAC  Vlan  EDst  ESrc  Port  Stat-Pkts
-----
Stat-Bytes  Uptime  Age
-----

MSFC 11.1.1.2 (Module 15):

Total entries displayed: 2
tamer (enable)
```

Note: One flow is created per destination. A ping from 12.1.1.1 to 11.1.1.1 will create two flows (one for each direction), as indicated by the last two lines in the output shown above.

The following are some descriptions of the information that is found in the table:

- Destination IP, source IP, Prot, DstPrt, and SrcPrt are the fields used to create shortcuts. In this case, destination only flow is used. Only the destination IP address of a flow is cached. This can be changed by modifying the flowmask, which is described later in this document.
- Destination MAC is the MAC address that will be used to rewrite the destination MAC of the packet. The source MAC address is rewritten with the MAC address of the MSFC.
- VLAN indicates the destination VLAN needed to reach that IP address. The destination VLAN is important, for example, if the packet needs to be sent on a trunk.
- DPort and Sport are the destination and source port of the flow.
- Stat-Pkts and Stat-Bytes give you statistics on the number of packets that have used this shortcut since the creation of the flow.
- Uptime is the time since the flow has been created.
- Age is the elapsed time since the flow was last used.

Change the flow to destination-source. The **show mls entry** command output displays both the source IP address and destination IP address are caching. A different flow for each source IP address communicating to the same destination IP address is now created, as shown below.

```
tamer (enable) set mls flow destination-source
Configured IP flowmask is set to destination-source flow.
Warning: Configuring more specific flow mask may increase the number of MLS
entries dramatically.
tamer (enable) 2000 Aug 09 17:05:12 %MLS-5-FLOWMASKCHANGE:IP flowmask changed
from DEST to DEST-SRC

tamer (enable) show mls entry
Destination-IP  Source-IP          Prot  DstPrt  SrcPrt  Destination-MAC  Vlan  EDst
-----
MSFC 11.1.1.2 (Module 15):
11.1.1.1      12.1.1.1           -    -       -       00-00-0c-09-50-66 11    ARPA
11.1.1.1      10.68.5.1          -    -       -       00-00-0c-09-50-66 11    ARPA
10.68.5.1     11.1.1.1           -    -       -       00-d0-00-3f-a3-ff 8     ARPA
12.1.1.1     11.1.1.1           -    -       -       00-00-0c-8c-70-88 12    ARPA
MSFC 0.0.0.0 (Module 16):

ESrc  DPort  Sport  Stat-Pkts  Stat-Bytes  Uptime  Age
-----
ARPA  7/3    7/4    4           400         00:00:02 00:00:02
ARPA  7/3    1/3    4           400         00:00:32 00:00:32
ARPA  1/3    7/3    4           400         00:00:32 00:00:32
ARPA  7/4    7/3    4           400         00:00:02 00:00:02

Destination-IPX          Destination-MAC  Vlan  EDst  ESrc  Port
-----
Stat-Pkts  Stat-Bytes  Uptime  Age
-----

MSFC 11.1.1.2 (Module 15):
MSFC 0.0.0.0 (Module 16):
```

```
Total entries displayed: 4
tamer (enable)
```

The third option is set MLS to full flow. Perform a few pings and Telnet sessions to see how the different flows are created for each TCP port. Below is how the MLS table should look after making a few pings and Telnet sessions. Using full flow, the number of flows created increases very rapidly. The TCP port information is cached, and appears in the MLS table.

```
tamer (enable) set mls flow full
Configured IP flowmask is set to full flow.
Warning: Configuring more specific flow mask may increase the number of MLS
entries dramatically.
Tamer (enable) 2000 Aug 09 17:30:01 %MLS-5-FLOWMASKCHANGE:IP flowmask changed
from DEST to FULL
```

```
tamer (enable)
tamer (enable) show mls entry
```

Destination-IP	Source-IP	Prot	DstPrt	SrcPrt	Destination-MAC	Vlan	EDst
MSFC 11.1.1.2 (Module 15):							
12.1.1.1	11.1.1.1	ICMP	-	-	00-00-0c-8c-70-88	12	ARPA
11.1.1.1	12.1.1.1	TCP	11001	Telnet	00-00-0c-09-50-66	11	ARPA
12.1.1.1	11.1.1.1	TCP*	Telnet	11001	00-00-0c-8c-70-88	12	ARPA
11.1.1.1	10.68.5.1	TCP	11002	Telnet	00-00-0c-09-50-66	11	ARPA
10.68.5.1	11.1.1.1	ICMP	-	-	00-d0-00-3f-a3-ff	8	ARPA
10.68.5.1	11.1.1.1	TCP*	Telnet	11002	00-d0-00-3f-a3-ff	8	ARPA
11.1.1.1	10.68.5.1	ICMP	-	-	00-00-0c-09-50-66	11	ARPA
11.1.1.1	12.1.1.1	ICMP	-	-	00-00-0c-09-50-66	11	ARPA

ESrc	DPort	Sport	Stat-Pkts	Stat-Bytes	Uptime	Age
ARPA	7/4	7/3	4	400	00:00:30	00:00:30
ARPA	7/3	7/4	16	688	00:00:26	00:00:24
ARPA	7/4	7/3	18	757	00:00:26	00:00:24
ARPA	7/3	1/3	61	4968	00:00:16	00:00:06
ARPA	1/3	7/3	4	400	00:00:33	00:00:33
ARPA	1/3	7/3	69	2845	00:00:17	00:00:06
ARPA	7/3	1/3	4	400	00:00:33	00:00:33
ARPA	7/3	7/4	4	400	00:00:32	00:00:31

Destination-IPX	Destination-MAC	Vlan	EDst	ESrc	Port	Stat-Pkts
Stat-Bytes	Uptime	Age				

```
MSFC 11.1.1.2 (Module 15):
Total entries displayed: 8
```

Notes

- In a live network, the number of flows created can be up to several thousands. Issue the **show mls entry ip [destination/source]** command to display specific flow instead of displaying the full flow table, as shown below.

```
tamer (enable) show mls entry ip destination 12.1.1.1
```

Destination-IP	Source-IP	Prot	DstPrt	SrcPrt	Destination-MAC	Vlan
MSFC 11.1.1.2 (Module 15):						
12.1.1.1	-	-	-	-	00-00-0c-8c-70-88	12

ESrc	DPort	Sport	Stat-Pkts	Stat-Bytes	Uptime	Age
ARPA	7/4	7/3	4	400	00:00:30	00:00:30

- You can verify the statistics for the flow by issuing the command **show mls statistics**, as shown below.

```
tamer (enable) show mls statistics entry ip 15
```

Destination IP	Source IP	Prot	DstPrt	SrcPrt	Stat-Pkts	Stat-Bytes
MSFC 11.1.1.2 (Module 15):						
12.1.1.1	11.1.1.1	ICMP	-	-	9	900
11.1.1.1	10.68.5.1	TCP	11005	Telnet	20	913
11.1.1.1	10.68.5.1	TCP	11004	Telnet	0	0
10.68.5.1	11.1.1.1	ICMP	-	-	4	400
10.68.5.1	12.1.1.1	ICMP	-	-	9	900
12.1.1.1	10.68.5.1	ICMP	-	-	4	400
11.1.1.1	10.68.5.1	ICMP	-	-	4	400
11.1.1.1	12.1.1.1	ICMP	-	-	9	900

Troubleshooting a Specific Connection

If you are having a connectivity problem to a specific IP address or between two specific hosts, try the following to troubleshoot:

- Issue the **show mls entry ip [destination|source]** command to see if the flow has been created.
- Issue the **show mls statistics entry [source|destination]** command several times in a row to see if the counters of stat-pakts for that shortcut are increasing.
- Verify the relevant flow.

For example, for an FTP session of a big file between TFTP server 12.1.1.1 and TFTP client 11.1.1.1, you need to verify the following two flows:

- One with destination 12.1.1.1 that should be hit only by the TFTP acknowledgment (small packet) (source of the flow 12.1.1.1 if destination-source flow used).
- One with destination 11.1.1.1 that should be hit by a lot of big packets (the actual file transfer) (source of the flow 11.1.1.1 if destination-source flow is used).

This is an example of TFTP between 12.1.1.1. and 11.1.1.1 of a file of about 7.6 MB. Following is the MLS stat table before the start of the TFTP:

```
tamer (enable) show mls statistics entry
                Last      Used
Destination IP  Source IP      Prot DstPrt SrcPrt Stat-Pkts  Stat-Bytes
-----
MSFC 11.1.1.2 (Module 15):
12.1.1.1       11.1.1.1      ICMP -      -      4      400
11.1.1.1       12.1.1.1      ICMP -      -      4      400
12.1.1.1       11.1.1.1      TCP  11000  Telnet 20      894
```

The TFTP has just started. The two additional flows created for TFTP traffic (UDP port 69) are shown below.

```
tamer (enable) show mls statistics entry
                Last      Used
Destination IP  Source IP      Prot DstPrt SrcPrt Stat-Pkts  Stat-Bytes
-----
MSFC 11.1.1.2 (Module 15):
12.1.1.1       11.1.1.1      ICMP -      -      4      400
12.1.1.1       11.1.1.1      UDP  69      50532 343      10997
11.1.1.1       12.1.1.1      ICMP -      -      4      400
11.1.1.1       12.1.1.1      UDP  50532  69      343      186592
12.1.1.1       11.1.1.1      TCP  11000  Telnet 20      894
```

The TFTP transfer just ended. Roughly 8.1 MB are transferred from server to client in 14,903 packets, which creates an average size of 544 bytes per packet. In the other direction, the same amount of packets are received with an average size of 476,949 divided by 14,904, which makes 33 bytes.

```
Tamer (enable) show mls statistics entry
                Last      Used
Destination IP  Source IP      Prot DstPrt SrcPrt Stat-Pkts  Stat-Bytes
-----
MSFC 11.1.1.2 (Module 15):
12.1.1.1       11.1.1.1      ICMP -      -      4      400
12.1.1.1       11.1.1.1      UDP  69      50532 14904    476949
11.1.1.1       12.1.1.1      ICMP -      -      4      400
11.1.1.1       12.1.1.1      UDP  50532  69      14903    8107224
12.1.1.1       11.1.1.1      TCP  11000  Telnet 20      894
```

These tables should give you an idea of what your traffic pattern should look like.

Example 2: Two MSFCs in the Same Chassis Configured for HSRP

Below is the running configuration of the two MSFCs configured for HSRP and the output of the **show standby** command. The MSFC in slot 15 is active for VLAN 12, and the MSFC in slot 16 is active for VLAN 11.

Slot 15	Slot 16
<pre>tamer-msfc#wr t Building configuration... Current configuration: !</pre>	<pre>tamer-msfc-2#wr t Building configuration... Current configuration: !</pre>

```

version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname tamer-msfc
!
boot system flash bootflash:
  c6msfc-ds-mz.121-1.E2.bin
!
ip subnet-zero
ip cef
!
!
!
!
interface Vlan1
  ip address 10.200.11.120 255.255.252.0
!
interface Vlan8
  ip address 10.68.5.2 255.255.252.0
!
interface Vlan11
  ip address 11.1.1.2 255.255.255.0
  no ip redirects
  standby 11 preempt
  standby 11 ip 11.1.1.3
!
interface Vlan12
  ip address 12.1.1.2 255.255.255.0
  no ip redirects
  standby 12 priority 105 preempt
  standby 12 ip 12.1.1.3
!
router eigrp 1
  network 10.0.0.0
  network 11.0.0.0
  network 12.0.0.0
  no auto-summary
!
ip classless
!
line con 0
  transport input none
line vty 0 4
  login
!
end

```

```

tamer-msfc>show standby
Vlan11 - Group 11
  Local state is Standby, priority 100,
    may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:00.814
  Hot standby IP address is 11.1.1.3 configured
  Active router is 11.1.1.4 expires in 00:00:09

```

```

version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname tamer-msfc-2
!
boot system flash
  bootflash:c6msfc-jsv-mz.121-2.E
!
ip subnet-zero
!
!
!
!
interface Vlan1
  ip address 10.200.11.121 255.255
!
interface Vlan8
  ip address 10.68.5.4 255.255.252
!
interface Vlan11
  ip address 11.1.1.4 255.255.255.
  no ip redirects
  standby 11 priority 105 preempt
  standby 11 ip 11.1.1.3
!
interface Vlan12
  ip address 12.1.1.4 255.255.255.
  no ip redirects
  standby 12 preempt
  standby 12 ip 12.1.1.3
!
router eigrp 1
  network 10.0.0.0
  network 11.0.0.0
  network 12.0.0.0
  no auto-summary
!
ip classless!
!
line con 0
  transport input none
line vty 0 4
  login
!
end

```

```

tamer-msfc-2#show standby
Vlan11 - Group 11
  Local state is Active, priority
    may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:02.846
  Hot standby IP address is 11.1.
  Active router is local

```

```

Standby router is local
Standby virtual MAC address is 0000.0c07.ac0b
4 state changes, last state change 00:06:36
Vlan12 - Group 12
Local state is Active, priority 105,
  may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:02.380
Hot standby IP address is 12.1.1.3 configured
Active router is local
Standby router is 12.1.1.4 expires in 00:00:09
Standby virtual MAC address is 0000.0c07.ac0c
2 state changes, last state change 00:12:22

```

```

Standby router is 11.1.1.2 expires in 00:00:02.518
Standby virtual MAC address is 0000.0c07.ac0a
2 state changes, last state change 00:06:36
Vlan12 - Group 12
Local state is Standby, priority 105,
  may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:02.518
Hot standby IP address is 12.1.1.3 configured
Active router is 12.1.1.2 expires in 00:00:09
  priority 105
Standby router is local
Standby virtual MAC address is 0000.0c07.ac0b
4 state changes, last state change 00:06:36

```

All the information in the previous example is still valid. To verify what has changed after configuring HSRP, view the output of the MLS commands below.

Issuing show mls Commands

```

tamer (enable) show mls
Total packets switched = 29894
Total Active MLS entries = 0
  MSFC 11.1.1.2 (Module 15) entries = 0
  MSFC 10.200.11.121 (Module 16) entries = 0
IP Multilayer switching aging time = 256 seconds
IP Multilayer switching fast aging time = 0 seconds, packet threshold = 0
IP Current flow mask is Full flow
Active IP MLS entries = 0
Netflow Data Export version: 7
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0

```

IP MSFC ID	Module	XTAG	MAC	Vlans
11.1.1.2	15	1	00-d0-d3-9c-9e-3c 00-00-0c-07-ac-0c	12,11,8,1
10.200.11.121	16	2	00-d0-bc-f0-07-b0 00-00-0c-07-ac-0b	1,8,11,12

```

IPX Multilayer switching aging time = 256 seconds
IPX flow mask is Destination flow
IPX max hop is 15
Active IPX MLS entries = 0

```

IPX MSFC ID	Module	XTAG	MAC	Vlans
11.1.1.2	15	1	-	-
10.200.11.121	16	2	-	-

- There are now two MLS routers seen by the PFC.
- For each router seen, the MAC address used by the HSRP group is 00-00-0c-07-ac-xx. These MAC addresses are the virtual MAC addresses used by HSRP. You are only seeing the MAC

address of group 11 linked to the router which is active for that group (slot 15 for VLAN 12 and slot 16 for VLAN 11). That means that in addition to packets with destination MAC address being the MSFC MAC address, candidate packets are also considered, which are packets with destination MAC being the HSRP address.

- As stated in the first example, you also need to see these HSRP address in the Layer 2 CAM table pointing to the MSFC.

```
tamer (enable) show cam 00-00-0c-07-ac-0c
```

```
* = Static Entry. + = Permanent Entry. # = System Entry. R = Router Entry.
X = Port Security Entry
```

```
VLAN  Dest MAC/Route Des      [COs]  Destination Ports or VCS / [Protocol Type]
-----
12     00-00-0c-07-ac-0c R#          15/1
Total Matching CAM Entries Displayed = 1
tamer (enable)
```

```
tamer (enable) show cam 00-00-0c-07-ac-0b
```

```
* = Static Entry. + = Permanent Entry. # = System Entry. R = Router Entry.
X = Port Security Entry
```

```
VLAN  Dest MAC/Route Des      [COs]  Destination Ports or VCS / [Protocol Type]
-----
11     00-00-0c-07-ac-0b R#          16/1
Total Matching CAM Entries Displayed = 1
tamer (enable)
```

Issuing show mls entry Commands

```
tamer (enable) show mls entry
```

```
Destination-IP  Source-IP      Prot  DstPrt  SrcPrt  Destination-MAC  Vlan  EDst
-----
MSFC 11.1.1.2 (Module 15):
11.1.1.1        12.1.1.1      ICMP  -        -        00-00-0c-09-50-66 11    ARPA
MSFC 10.200.11.121 (Module 16):
12.1.1.1        11.1.1.1      ICMP  -        -        00-10-7b-3b-af-3b 12    ARPA
```

```
ESrc DPort  Sport  Stat-Pkts  Stat-Bytes  Uptime  Age
-----
ARPA 7/3    7/4    4           400         00:00:03 00:00:03
ARPA 7/4    7/3    4           400         00:00:04 00:00:03
```

```
Destination-IPX      Destination-MAC  Vlan  EDst  ESrc  Port  Stat-Pkts
-----
Stat-Bytes  Uptime  Age
-----
```

```
MSFC 11.1.1.2 (Module 15):
MSFC 10.200.11.121 (Module 16):
```

- There are now two shortcut tables: one for flows created by the first MSFC, and one for flows created by the second MSFC.

- By pinging between 11.1.1.1 and 12.1.1.1, which are two PCs configured with the HSRP address as their default gateway, packets from 12.1.1.1 to 11.1.1.1 coming in VLAN 12 on the switch were shortcutted by the MSFC in slot 15 (as it is the active HSRP router for VLAN 12), and a shortcut for destination 11.1.1.1 is created by the MSFC in slot 15. The shortcuts for packets from 11.1.1.1 to 12.1.1.1 have been created on the other end by the MSFC in slot 16.

Troubleshooting Tips

Creation of Flows

If a flow has not been created, use the following troubleshooting tips:

- Does the Supervisor Engine see the MSFC with all expected MAC addresses when issuing the **show mls** command?
 - If yes, go to the next step.
 - If no, make sure that the MSFC is not stuck in ROMmon mode. Verify that MLS is enabled on the MSFC by issuing the command **show mls status**.
- Is the MAC address of the MSFC present when issuing the **show cam** command? Does it show up as a router cam entry (R#)?
- Verify that you have not configured a feature on the MSFC that disables MLS.
 - Check the features that may impact MLS by reviewing the release notes for the software version you are running.
 - Examples of restrictions are shown in the table below.

IP Router 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.
----------------------------------	-----------------------------------

- Verify that you have enabled an access list that requests software handling instead of being processed by hardware shortcuts. Refer to [Hardware and Software Handling of IOS ACLs](#) for more information.

If you have gone through all the tips listed above, and are still having problems, verify that the MSFC is still being hit by a lot of packets.

- Something may be causing the entries to be purged continuously. Some possible causes of purging the flow table may include the following:
 - Route flapping or any Layer 3 instability.
 - ARP cache changes on the MSFC.
 - Flowmask changed on the Supervisor Engine.
 - Destination VLAN is deleted.
 - VLAN interface is shutdown on the MSFC.
- Some other reasons for software forwarding (packet hitting the MSFC) instead of hardware shortcuts may include the following:
 - Packet with IP option set.
 - Packet with TTL less than or equal to 1.
 - Packet that needs to be fragmented.
- You can have up to 128 K of flow, however, a hashing algorithm is used. If you exceed 32 K flows, you may begin to have hash collision that will cause packets being routed by software.
 - A way to avoid having too many flows is to configure a fast aging for MLS flow.
- Remember that you can only have MLS for IP, IPX, and IP Multicast. If you have other types of traffic (for example, AppleTalk), they will be software routed, and can cause CPU peaks on the MSFC or excessive packet hitting on the MSFC.
- As specified, IP MLS and IPX MLS are enabled by default, however, IP multicast MLS is not enabled by default. If you are using IP multicast, make sure to enable MLS for multicast as specified in the configuration guide.

Note: A Spanning Tree Topology Change Notification (TCN) or flapping ports on a Catalyst 6500/6000 series switch will not cause the MLS flow table to be cleared, as was the case for MLS on Catalyst 5000 switches.

IP Accounting excludes MLS traffic

In the Cisco Catalyst 6500 series, Multiple Layer Switching (MLS) is deployed in such a way that once a flow is established, traffic is directly switched at PFC (hardware-switched) and is not processed by the MSFC, hence the lack of continuous accounting. Only new or process-switched flows (software-switched) is recorded by IP accounting when enabled, and even then only until the entry is entered into the database. Thus, the previous warning message is normally displayed when you enable IP accounting on such a platform.

```
6500(config)#int fa8/40
6500(config-if)#ip accounting
Accounting will exclude mls traffic when mls is enabled.
```

NetFlow accounting is the preferred method. Refer to [Configuring NetFlow](#) for more information regarding to NetFlow.

Unable to configure interface-full Flowmask

```
C6500#mls flow ip interface-full
% Unable to configure flow mask for ip protocol: interface-full. Reset
to the default flow mask type: none
```

Issue the **show fm fie flowmask detail** command and verify if NAT is enabled and uses *Intf Full Flow* mask.

```
C6500#show fm fie flowmask detail

!--- Part of the output not shown

Primary Flowmasks registered by Features
-----+-----+-----
Feature                               Flowmask                               Flowmask Status
-----+-----+-----
IP_ACCESS_INGRESS                     Intf Full Flow                         Disabled/Unused
IP_ACCESS_EGRESS                       Intf Full Flow                         Disabled/Unused
NAT_INGRESS                             Intf Full Flow                         Enabled
NAT_EGRESS                              Intf Full Flow                         Enabled

!--- Remaining part of the output not shown
```

If NAT uses the *Intf Full Flow* mask and when you try to configure the *interface-full* for Netflow, there is an issue, as there is a flowmask conflict. If you want to use Netflow stat, you can try to use *interface-destination-source* (**mls flow ip interface-destination-source** command), which does not conflict with the Netflow mask usage.

NDE assumes that all flows are created with the same flow mask. Due to this restriction, NDE cannot be enabled with certain features that require conflicting flow masks. One specific case is hardware-accelerated NAT. NDE and hardware-accelerated NAT are mutually exclusive.

NDE fails if any of these events occur:

- Hardware-accelerated NAT is enabled.
- Two or more features with conflicting flow masks have been configured on the switch.

Conversely, once NDE is successfully configured, NAT cannot be configured to work in the hardware and two different features with conflicting flow mask requirements cannot be configured on the switch.

Related Information

- [Troubleshooting IP MultiLayer Switching](#)
 - [Troubleshoot Unicast IP Routing Involving CEF on Catalyst 6500/6000 Series Switches with a Supervisor Engine 2 and Running CatOS System Software](#)
 - [Configuring IP MultiLayer Switching on the Catalyst 5000](#)
 - [Technical Support & Documentation - Cisco Systems](#)
-

Home	How to Buy	Login	Profile	Feedback	Site Map	Help
----------------------	----------------------------	-----------------------	-------------------------	--------------------------	--------------------------	----------------------

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2007 - 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)