

Troubleshooting IP Multilayer Switching

Document ID: 10554

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Description of MLS

Troubleshoot IP MLS Technology

- Troubleshoot Flowchart and Steps
- Commands or Screen Captures
- Before You Contact Cisco Technical Support

Related Information

Introduction

This document outlines basic steps to troubleshoot Multilayer Switching (MLS) for IP. This feature has become a highly desired method to accelerate routing performance through the use of dedicated application-specific integrated circuits (ASICs). Traditional routing occurs through a central CPU and software. MLS offloads a significant portion of routing (packet rewrite) to hardware, which is why MLS also bears the term "switching". MLS and Layer 3 switching are equivalent terms. The NetFlow feature of Cisco IOS® Software is distinct; this document does not cover NetFlow. MLS also includes support for Internetwork Packet Exchange (IPX) MLS (IPX MLS) and multicast MLS (MMLS). However, this document exclusively concentrates on basic MLS IP troubleshoot procedures.

For customers with Cisco Catalyst 6500/6000 series switches running Cisco IOS Software, refer to the MLS documentation for your Supervisor Engine:

- Configuring IP Unicast Layer 3 Switching on Supervisor Engine 1
- Configuring IP Unicast Layer 3 Switching on Supervisor Engine 2

Note: This document is not valid for the Catalyst 6500/6000 Supervisor Engine 2 or Supervisor Engine 720, as these Supervisor Engines do not use MLS. The Supervisor Engine 2 and Supervisor Engine 720 use Cisco Express Forwarding (CEF) as a hardware-based forward mechanism. For more information, refer to the document Troubleshoot Unicast IP Routing Involving CEF on Catalyst 6500/6000 Series Switches with a Supervisor Engine 2 and Running CatOS System Software.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Description of MLS

As networks face greater demands, the need for greater performance increases. More and more PCs connect to LANs, WANs, and the Internet. The users require fast access to databases, files and web pages, applications through networks, other PCs, and video stream. To keep connections quick and reliable, networks must be able to rapidly adjust to changes and failures to find the best path. The networks must also remain as invisible as possible to end users. To determine the best path is the primary function of routing protocols, and this can be a CPU-intensive process. Thus, there is a significant performance increase with the offload of a portion of this function to switching hardware. This performance increase is the goal of the MLS feature.

Two of the three major components of MLS are the MLS route processor (MLS-RP) and the MLS switching engine (MLS-SE). The MLS-RP is the MLS-enabled router, which performs the traditional function of routing between subnets/VLANs. The MLS-SE is a MLS-enabled switch, which normally requires a router to route between subnets/VLANs. However, with special hardware and software, MLS-SE can handle the rewrite of the packet. When a packet transverses a routed interface, the change (rewrite) of non-data portions of the packet occurs as the packet heads to the destination, hop by hop. Confusion can arise here because a Layer 2 device appears to take on a Layer 3 task. Actually, the switch only rewrites Layer 3 information and "switches" between subnets/VLANs. The router is still responsible for standards-based route calculations and best-path determination. You can avoid much of this confusion if you mentally keep the routing and switching functions separate, especially when they are within the same chassis (as with an internal MLS-RP). Think of MLS as a much more advanced form of route cache, with a separation of the cache from the router on a switch. MLS requires both the MLS-RP and the MLS-SE, along with respective hardware and software minimums.

The MLS-RP can be internal (installation in a switch chassis) or external (connection via a cable to a trunk port on the switch). Examples of internal MLS-RPs are the Route Switch Module (RSM) and the Route Switch Feature Card (RSFC). You install the RSM or RSFC in a slot or Supervisor Engine of a Catalyst 5500/5000 series switch, respectively. The same applies to the Multilayer Switch Feature Card (MSFC) for the Catalyst 6500/6000 series. Examples of external MLS-RPs include any member of the Cisco 7500, 7200, 4700, 4500 or 3600 series routers. In general, to support the MLS IP feature, all MLS-RPs require a minimum Cisco IOS Software release in the 11.3WA or 12.0WA trains. Refer to Cisco IOS Software release documentation for specifics. Also, you must enable MLS for a router to be a MLS-RP.

The MLS-SE is a switch with special hardware. For a Catalyst 5500/5000 series switch, MLS requires the installation of a NetFlow Feature Card (NFFC) on the Supervisor Engine. The Supervisor Engine IIG and IIIG have a NFFC by default. In addition, a bare minimum of Catalyst OS (CatOS) 4.1.1 software is also a requirement.

Note: The CatOS 4.x train is now in General Deployment (GD). The software passed rigorous end-user criteria and field-experience targets for stability. Refer to Cisco.com for the latest releases.

The Catalyst 6500/6000 hardware and software with the MSFC/Policy Feature Card (PFC) supports and automatically enables IP MLS. (The default for MLS is disabled on other routers.)

Note: IPX MLS and MMLS may have different hardware and software (Cisco IOS Software and CatOS) requirements. More Cisco platforms support the MLS feature. Also, you must enable MLS for a switch to be a MLS-SE.

The third major component of MLS is the Multilayer Switching Protocol (MLSP). You must understand the

basics of MLSP to get at the heart of MLS and perform effective MLS troubleshoot procedures. MLS–RP and MLS–SE use MLSP to communicate with one another. Tasks include:

- The enable of MLS.
- Installation of MLS flows (cache information).
- Update or deletion of flows.
- Management and export of flow statistics.

Note: Other documents cover NetFlow Data Export.


MLSP also allows the MLS–SE to:

- Learn the Layer 2 MAC addresses of the MLS–enabled router interfaces.
- Check the flowmask of the MLS–RP.

Note: The Troubleshoot IP MLS Technology section of this document covers this procedure.

- Confirm that the MLS–RP is operational.

The MLS–RP sends out multicast "hello" packets every 15 seconds with use of MLSP. If the MLS–SE misses three of these intervals, the MLS–SE recognizes that the MLS–RP has failed or that connectivity to the MLS–RP is lost.



This diagram illustrates three essentials that you must complete (with use of MLSP) to create a shortcut: the candidate, enable, and cache steps. The MLS–SE checks for the cache MLS entry. If the MLS cache entry and packet information match (a "hit"), the packet header rewrite occurs locally on the switch. This rewrite is a shortcut or bypass of the router. The packet does not forward to the router as normally occurs. Packets that do not match are forwarded to the MLS–RP as candidate packets. A local switch may occur for these packets. After the pass of the candidate packet through the MLS flowmask (which Step 7 of the section Troubleshoot IP MLS Technology explains) and the rewrite of the information in the packet header (without contact with the data portion), the router sends the packet toward the next hop along the destination path. The packet is now an enabler packet. If the packet returns to the same MLS–SE from which the packet left, a MLS shortcut is created and placed into the MLS cache. Now, instead of the router software, the switch hardware locally rewrites that packet and all similar packets that follow (a "flow").

The same MLS–SE must see both the candidate and enabler packets for a particular flow for the creation of a MLS shortcut. (This requirement is why network topology is important to MLS.) Remember, the purpose of MLS is to allow the communication path between two devices in different VLANs, with connection off the

same switch, to bypass the router. This action enhances network performance.

With use of the flowmask, which is essentially an access list, the administrator can adjust the degree of similarity of these packets. The administrator can adjust the scope of these flows:

- Destination address.
- Destination and source addresses.
- Destination, source, and Layer 4 information.

Note: The first packet of a flow always passes through the router. From then on, the flow is locally switched. Each flow is unidirectional. Communication between PCs, for example, requires the setup and use of two shortcuts. The main purpose of MLSP is to set up, create, and maintain these shortcuts.

These three components (the MLS–RP, the MLS–SE, and MLSP) free up vital router resources through the allowance of other network components to take on some of the router functions. For certain topologies and configurations, MLS provides a simple and highly effective method to increase network performance in the LAN.

Troubleshoot IP MLS Technology

This section includes a flow diagram for basic IP MLS troubleshooting. The diagram derives from the most common types of MLS–IP service requests that customers make with Cisco Technical Support. MLS is a robust feature with which you should have no problems. However, if an issue does arise, this section should help you resolve the problem. To troubleshoot, these items must be true:

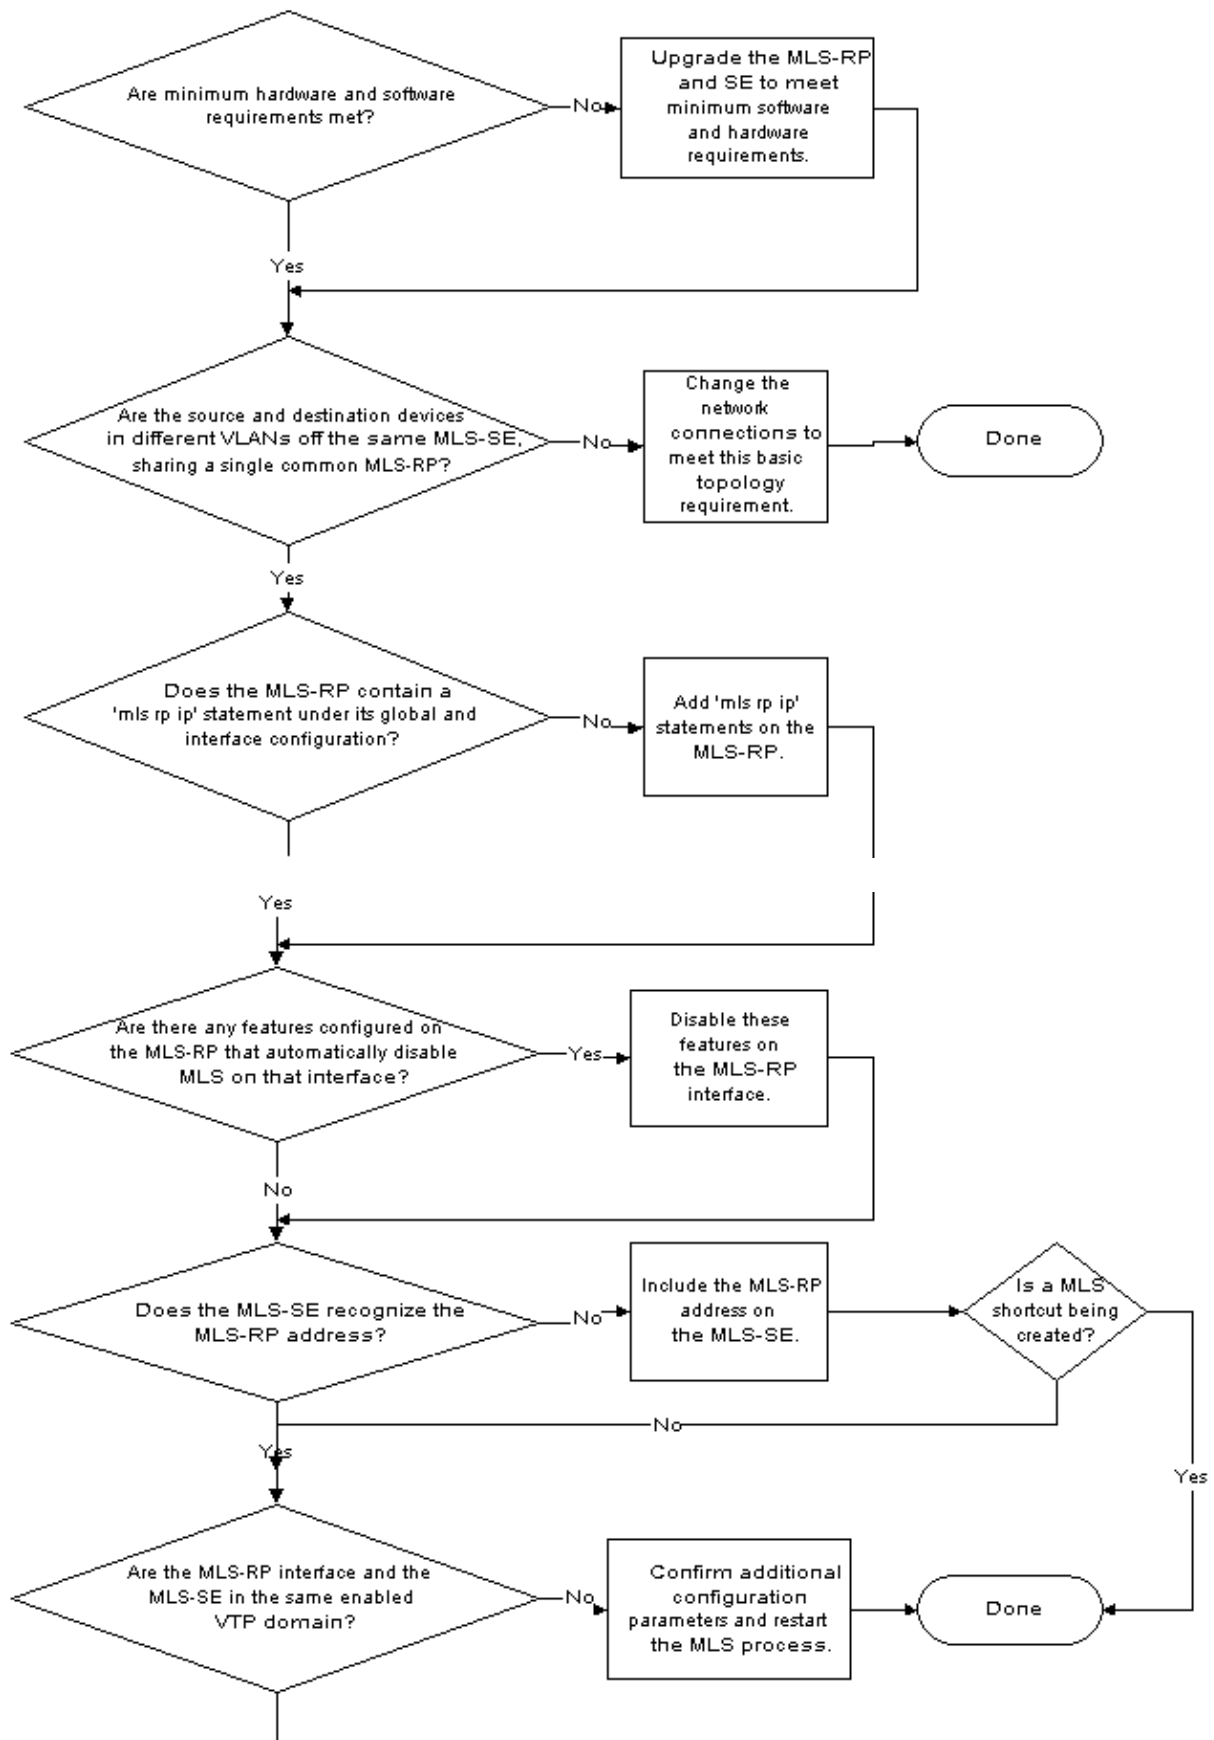
- You are familiar with and have completed the basic configuration steps necessary to enable IP MLS on the router and switches. See the Related Information section of this document for further information.
- You have IP routing enabled on the MLS–RP (default). If the command **no ip routing** appears in the global configuration of a **show run** command, IP routing is off. In this case, IP MLS does not function.
- IP connectivity exists between the MLS–RP and MLS–SE. Ping the IP addresses of the router from the switch. Then, look for the display of exclamation points (bangs) in return.
- MLS–RP interfaces are in an "up/up" state on the router. Issue the **show ip interface brief** command on the router to confirm the state.

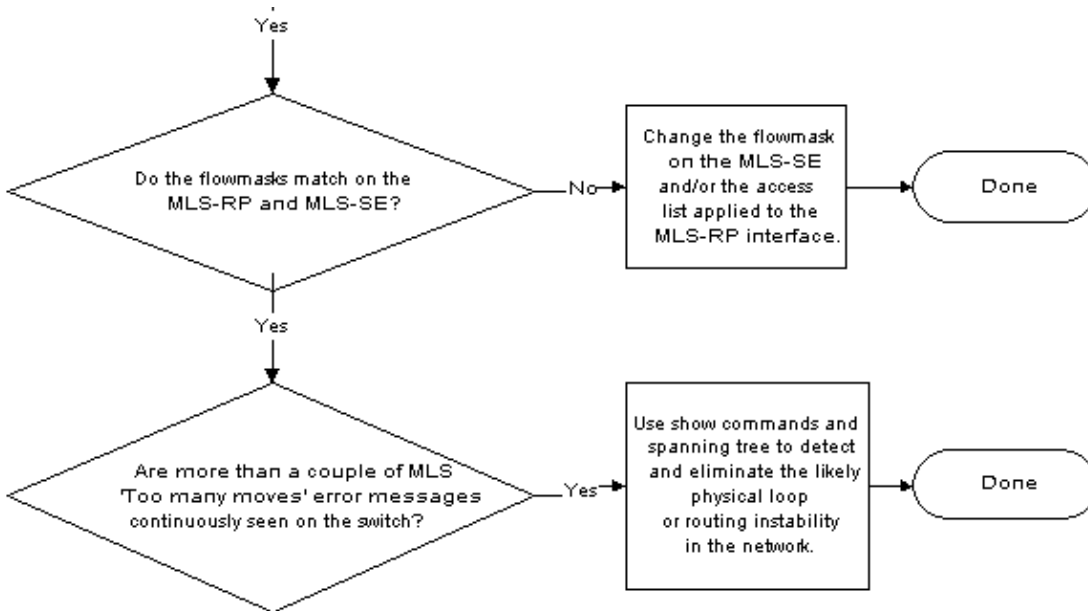


Caution: Whenever you make configuration changes to a router that you intend to be permanent, remember to save those changes with the **copy running–config starting–config** command. Shorter versions of this command include **copy run start** and **write memory**. Any configuration modifications are lost if the router reloads or you reset the router. The RSM, RSFC, and MSFC are routers, not switches. In contrast, the automatic save of changes occurs when the changes are made at the switch prompt of a Catalyst 5500/5000 or 6500/6000 series switch.

Troubleshoot Flowchart and Steps

Note: The procedure that appears below the flowchart provides further detail about each step in the flowchart.





1. Are minimum hardware and software requirements met?

Upgrade the MLS–RP and MLS–SE to meet minimum software and hardware requirements. For the MLS–RP, no additional hardware is necessary. Although you can configure MLS on nontrunked interfaces, the connection to the MLS–SE is generally through VLAN interfaces (as with a RSM) or support trunking. (You can also configure trunking to support MLS on multiple VLANs if you configure Inter–Switch Link Protocol [ISL] or IEEE 802.1Q trunking on the switchport and router interface.) Also, only members of the Cisco 7500, 7200, 4700, 4500 and 3600 series routers support MLS externally. Currently, only these external routers and the routers that fit into the Catalyst 5500/5000 or 6500/6000 switch series can be MLS–RPs. (Examples include the RSM and RSFC for the Catalyst 5500/5000 series and the MSFC or MSFC2 for the Catalyst 6500/6000 series.) The MSFC requires the PFC as well. You must install both on the Catalyst 6500/6000 Supervisor Engine. IP MLS is now a standard feature in Cisco IOS Software Release 12.0 and later. Cisco IOS Software earlier than Cisco IOS Software Release 12.0 generally requires a special train. For such IP MLS support, install the latest images in Cisco IOS Software Release 11.3 that have the letters "WA" in the file names.

For the MLS–SE, a NFFC is necessary for a member of the Catalyst 5500/5000 series. You install this card in the Supervisor Engine module of the Catalyst switch. Newer Catalyst 5500/5000 series Supervisor Engines (since 1999) include the card as standard hardware. Supervisor Engines I and II do not support the NFFC; NFFC is an option on early Supervisor Engine IIIs. Also, you need CatOS 4.1.1, at minimum, for IP MLS. In contrast, for Catalyst 6500/6000 series switches with Supervisor Engine 1 or 1A, there is support for IP MLS from the first CatOS software release, 5.1.1. (In fact, IP MLS is an essential and default ingredient for the high performance of this software.) With the release of new platforms and software that support IP MLS, you need to check the documentation and release notes. Generally, install the latest release in the lowest train that meets your feature requirements. Always check the release notes and consult with your local Cisco sales office for new MLS support and feature developments.

To determine the hardware and software that you have installed, use the **show version** command on the router and the **show module** command on the switch.

Note: The Catalyst 6500/6000 series switches do not support an external MLS–RP. The MLS–RP must be a MSFC.

2. Are the source and destination devices in different VLANs off the same MLS–SE sharing a single common MLS–RP?

A basic topology requirement of MLS is that the router have a path to each of the VLANs. Remember that the purpose of MLS is to create a shortcut between two VLANs so that the switch can perform the "routing" between the two end devices. Then, the router is free to perform other tasks. The switch does not actually route, but rewrites the frames so that the end devices appear to talk through the router. If the two devices are in the same VLAN, the MLS-SE switches the frame locally without the need to utilize MLS, as switches do in such a transparently bridged environment. Therefore, there is no creation of a MLS shortcut. You can have multiple switches and routers in the network, and even multiple switches along the flow path. However, the path between the two end devices for which you want a MLS shortcut must include a single MLS-RP in that VLAN for that path. In other words, the flow from source to destination must cross a VLAN boundary on the same MLS-RP; also, the same MLS-SE must see a candidate and enabler packet pair for the creation of a MLS shortcut. If the topology does not meet these criteria, the packet routes normally without the use of MLS. See the Related Information section of this document for diagrams and discussions with regard to network topologies with support and without support.

3. Does the MLS-RP contain an **mls rp ip** statement under both its global and interface configuration?

If one is not present, add **mls rp ip** statements appropriately on the MLS-RP. Except for routers that automatically enable IP MLS (such as the Catalyst 6500/6000 MSFC and MSFC2), the configuration requires this step. For most MLS-RPs (routers that you configure for IP MLS), the **mls rp ip** statement must appear both in the global configuration and under the interface configuration.

Note: When you configure the MLS-RP, also remember to issue the **mls rp management-interface** command under one of the IP MLS interfaces of the MLS-RP. This required step tells the MLS-RP on which interface the MLS-RP should send MLSP messages to communicate with the MLS-SE. Again, you need to issue this command under one interface only.

4. Are there any features configured on the MLS-RP that automatically disable MLS on that interface?

There are several configuration options on the router that are not compatible with MLS. These options include IP accounting, encryption, compression, IP security, Network Address Translation (NAT), and committed access rate (CAR). For further information, see the links that relate to IP MLS configuration in the Related Information section of this document. Packets that traverse a router interface that you have configured with any of these features must route normally; the creation of a MLS shortcut does not occur. For MLS to work, you must disable these features on the MLS-RP interface.

Another important feature that affects MLS is access lists, both input and output. Further discussion of this option appears in Step 7 of this section.

5. Does the MLS-SE recognize the MLS-RP address?

For MLS to function, the switch must recognize the router as a MLS-RP. The MLS-SE in which you have installed an internal MLS-RPs automatically recognizes the MLS-RP. (Examples of internal MLS-RPs include the RSM or RSFC in a Catalyst 5500/5000 series switch and the MSFC/MSFC2 in a Catalyst 6500/6000 series switch.) For external MLS-RPs, you must explicitly inform the switch of the router address. This address, which comes from the list of IP addresses on the router interfaces, is not actually an IP address. The address is simply a router ID. For internal MLS-RPs, the MLS-ID is normally not even an IP address on the router. The ID is commonly a loopback address (127.0.0.x) because of the automatic inclusion of internal MLS-RPs. For MLS to function, include on the MLS-SE the MLS-ID found on the MLS-RP.

Use the **show mls rp** command on the router to find the MLS-ID. Then, configure that ID on the switch with the issue of the **set mls include MLS-ID** command. The configuration requires this step when you use external MLS-RPs.



Caution: If you change the IP address of MLS–RP interfaces and then reload the router, the

MLS process on the router may choose a new MLS–ID. This new MLS–ID may differ from the MLS–ID that you manually included on the MLS–SE, which may cause MLS to cease to function. The problem is not a software glitch, but an effect of the switch attempt to communicate with a MLS–ID that is no longer valid. Be sure to include this new MLS–ID on the switch to get MLS to operate again. You may also have to disable/enable IP MLS as well.

Note: When the MLS–SE does not directly connect to the MLS–RP, the address to include on the MLS–SE may appear as the loopback address mentioned in this step: a switch that connects between the MLS–SE and MLS–RP. You must include the MLS–ID even though the MLS–RP is internal. To the second switch, the MLS–RP appears as an *external* router because the MLS–RP and MLS–SE are not in the same chassis.

6. Are the MLS–RP interface and the MLS–SE in the same enabled VLAN Trunking Protocol (VTP) domain?

MLS requires MLS components, which include the end stations, to be in the same VTP domain. VTP is a Layer 2 protocol that manages VLANs on several Catalyst switches from a central switch. VTP allows an administrator to create or delete a VLAN on all switches in a domain without the need to do so on every switch in that domain. The MLSP, which the MLS–SE and the MLS–RP use to communicate with one another, does not cross a VTP domain boundary. If you have enabled VTP on the switches, use the **show vtp domain** command on the switch to determine the VTP domain placement of the MLS–SE. (The default for VTP is enabled on Catalyst 5500/5000 and 6500/6000 series switches.)

Complete these steps to add the VTP domain to each of the router MLS interfaces. (The exception to the performance of these steps is with the Catalyst 6500/6000 MSFC and MSFC2, on which MLS is essentially a "plug–and–play" feature.) This procedure permits MLSP multicasts to move between the MLS–RP and MLS–SE and, therefore, allows MLS to function.

- a. Issue the command **no mls rp ip** .

This disables MLS on the affected MLS–RP interface before modification of the VTP domain.

- b. Issue the command **mls rp vtp–domain VTP–domain–name** .

The VTP domain name on each interface for which you have enabled MLS must match the domain name of the switch.

- c. Issue the command **mls rp vlan–id VLAN–ID–number** .

This is only necessary for non–ISL trunking and external MLS–RP interfaces.

- d. Issue the command **mls rp management–interface**.

Issue this command for only one interface on the MLS–RP. This required step tells the MLS–RP to which interface MLS–RP should send MLSP messages.

- e. Issue the command **mls rp ip**.

This command enables MLS on the interface of the MLS–RP.

To change the VTP domain name of the MLS–SE, issue this command at the switch enable prompt:

- ◆ **set vtp domain name VTP–domain–name**

For MLS to work, be sure that you have enabled VTP on the switch with this command:

- ◆ **set vtp enable**

7. Do the flowmasks agree on the MLS–RP and MLS–SE?

A flowmask is a filter that a network administrator configures. MLS uses the filter to determine if the creation of a shortcut is necessary. The process is similar to that of an access list in that, if you set up criteria with great detail, the MLS process must look deep into the packet to verify if the packet meets those criteria. To adjust the scope of shortcuts that the MLS creates, you can make the flowmask more or less specific. The flowmask is essentially a "tuning" device. The three IP MLS modes are:

- ◆ destination–ip
- ◆ source–destination–ip
- ◆ full–flow–ip

When you have not applied an access list to the router interface for which you have enabled MLS, the destination–ip mode (the default) is in use. When you apply a standard access list on MLS–RP, source–destination–ip mode is in use and if an extended access list is in use on MLS–RP, full–flow–ip mode is in effect. The type of access list you apply to the interface implicitly determines the MLS mode on the MLS–RP. In contrast, the MLS mode on the MLS–SE is an explicit configuration. When you choose the appropriate mode, you configure MLS such that *one* of these statements is true:

- ◆ Only the destination address must match for the creation of a MLS shortcut.
- ◆ Both the source and destination information, or even Layer 4 information such as TCP/User Datagram Protocol (UDP) port numbers, must match.

The MLS mode is configurable on both the MLS–RP and the MLS–SE. In general, the modes must match. However, if you deem necessary either the source–destination–ip or full–flow–ip MLS mode, you should configure the mode on the router through the application of the appropriate access list. MLS always chooses the most specific mask. MLS gives precedence to the flowmask on the MLS–RP over the flowmask found on the MLS–SE. Be careful if you change the MLS mode of the switch from the default destination–ip. You should be sure that the MLS mode matches the mode on the router for MLS to work. For source–destination–ip and full–flow–ip modes, remember to apply the access list to the appropriate router interface. If you apply no access list, the mode is simply the default destination–ip, even if you configure the MLS mode otherwise.



Caution: Whenever you change the flowmask, whether on the MLS–RP or MLS–SE, the

purge of all cache MLS flows occurs, and the MLS process restarts. A purge also can occur when you issue the command **clear ip route–cache** on the router. If you issue the global router configuration command **no ip routing**, the command causes a purge and disables MLS. (The **no ip routing** command turns off IP routing and essentially transforms the router into a transparent bridge.) Routing is a prerequisite of MLS. Each of these actions may temporarily, but seriously, affect router performance in a production network. The router experiences a spike in router load until the creation of the new shortcuts because the router handles all the flows that the switch previously processed.

Note: Avoid the very wide use of flowmasks that you have configured with Layer 4 information, especially with a Catalyst 5500/5000 series switch as the MLS–SE. If you force the router to peer deeply into every packet on the interface, you bypass many of the intended benefits of MLS. The wide use of flowmasks is much less an issue when you utilize a Catalyst 6500/6000 series switch as the MLS–SE; with a 6500/6000 as the MLS–SE, the switch ports can recognize Layer 4 information.

Note: Until recently, MLS did not support flowmasks with inbound configuration on a MLS–RP interface, but only with outbound configuration. Now, there is support for an inbound flowmask with use of the **mls rp ip input–acl** command in addition to the normal MLS–RP configuration commands on a router interface.

8. Are more than a couple of MLS "Too many moves" error messages continuously seen on the switch?

As Step 7 notes, if you change a flowmask, clear the route cache, or globally turn off IP routing, the action causes a cache purge. Other circumstances can also cause full purge or many single-entry purges. MLS then indicates "Too many moves". There are several forms of this message, but each contains these three words. Another of the most common causes of this error occurs when the switch learns multiple identical Ethernet MAC addresses within the same VLAN. Ethernet standards do not allow for identical MAC addresses within the same VLAN. If you see the error infrequently, or just a few times consecutively, there is no cause for concern. MLS is a robust feature. Normal network events, such as the move of a PC connection between ports, may cause the message. However, if you see the error continuously for several minutes, the message is likely a symptom of a more serious issue.

When such a situation arises, the common root cause is the presence of two devices with the same MAC address with connection to a VLAN, or a physical loop within the VLAN. (Another possibility is multiple VLANs, if you bridge across these broadcast domains.) Use spanning-tree troubleshooting and the tip below to find the loop and eliminate it. Also, any rapid topology changes can cause temporary network (and MLS) instability. Examples include router interfaces that flap or a bad network interface card (NIC).

Tip: Use the **show mls notification** and **show looktable** commands on the switch to point toward the duplicate MAC address or physical loop. The **show mls notification** command provides a table address (TA) value. The **show looktable TA-value** command returns a possible MAC address that you can trace to the root of the problem.

Commands or Screen Captures

For descriptions and examples in detail of IP MLS router and switch commands, see the Related Information section of this document.

Before You Contact Cisco Technical Support

Before you contact Cisco Technical Support, be sure you have read through this document and completed the actions the document recommends for your system problem.

Additionally, complete these items and document the results for better assistance:

- Capture the output of the **show module** command from all the affected switches.
- Capture the output of the **show vtp domain** command from all the affected switches.
- Capture the output of the **show trunk mod_number/port_number** command from all the affected ports.
- Capture the output of the **show trunk mod_number/port_number capabilities** command from all the affected ports.
- Capture the output of the **show tech-support** command from the MLS-RP.
- Capture the output of the **show mls rp** command on the MLS-RP and both the **show mls** and **show mls include** commands on the MLS-SEs.
- Capture the output of additional commands, as necessary, which depends on the nature of the issue.

A clear network topology and dial-in or Telnet access also help considerably in effective problem resolution.

Related Information

- **IP MLS Supported Network Topologies section and IP MLS Unsupported Network Topologies section of Configuring IP MLS**
- **Layer 3 Switching Software Configuration Guide**

- **Catalyst 5000 Family Release Notes**
 - **Catalyst 6500 Series Release Notes**
 - **Configuring and Troubleshooting IP MLS on Catalyst 6500/6000 Switches with an MSFC**
 - **LAN Product Support Pages**
 - **LAN Switching Support Page**
 - **Technical Support & Documentation – Cisco Systems**
-

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

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Dec 09, 2005

Document ID: 10554
