

Multilayer Switching Overview

This chapter provides an overview of Multilayer Switching (MLS).

Note The information in this document, *Cisco IOS Switching Services Configuration Guide*, is a brief summary of the information contained in the *Catalyst 5000 Series Multilayer Switching User Guide*. The commands and configurations described in this guide apply only to the devices that provide routing services. Commands and configurations for Catalyst 5000 series switches are not documented here.

MLS provides high-performance Layer 3 switching for the Catalyst 5000 series LAN switches. MLS switches IP data packets between subnets using advanced application specific integrated circuit (ASIC) switching hardware. Standard routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Routing Information Protocol (RIP), and Intermediate System-to-Intermediate System (IS-IS), are used for route determination.

MLS enables hardware-based Layer 3 switching to offload routers from forwarding unicast IP data packets over shared media networking technologies such as Ethernet. The packet forwarding function is moved onto Layer 3 Catalyst 5000 series 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 still use routers for forwarding packets.

MLS also provides traffic statistics as part of its switching function. These statistics are used for identifying traffic characteristics for administration, planning, and troubleshooting. MLS uses NetFlow Data Export (NDE) to export the flow statistics.

The Route Switch Module (RSM) performs route processing and central configuration and control for the Catalyst 5000 series switch. Routing services can also be provided by an externally attached router.

MLS consists of the following:

- Catalyst 5000 series multilayer LAN switches
- Catalyst RSM, which provides Cisco IOS-based multiprotocol routing and network services

Note Cisco 7500, 7200, 4500, and 4700 series routers also support MLS.

- NetFlow Feature Card (NFFC) which is a modular feature-card upgrade for the Catalyst Supervisor Engine III to provide Layer 3 switching

Note The 10/100BaseTX and 100BaseFX Backbone Fast Ethernet Switching modules have onboard hardware that optimizes MLS performance.

Procedures for configuring MLS and NDE on routers are provided in the following chapter, “Configuring Multilayer Switching.”

Terminology

The following terminology is used:

- Multilayer Switching-Switching Engine (MLS-SE)—A NetFlow Feature Card (NFFC)-equipped Catalyst 5000 series switch.
- Multilayer Switching-Route Processor (MLS-RP)—A Cisco router with MLS enabled.
- Multilayer Switching Protocol (MLSP)—The protocol running between the MLS-SE and MLS-RP to enable MLS.

Key MLS Features

Table 11 lists the key MLS features

Table 11 Summary of Key Features

Feature	Description
Ease of Use	Is autoconfigurable and autonomously sets up its Layer 3 flow cache. Its plug-and-play design eliminates the need for you to learn new IP switching technologies.
Transparency	Requires no end-system changes and no renumbering of subnets. It works with DHCP ¹ and requires no new routing protocols.
Standards Based	Uses IETF ² standard routing protocols such as OSPF and RIP for route determination. You can deploy MLS in a multivendor network.
Investment Protection	Provides a simple feature-card upgrade on the Catalyst 5000 series switches. You can use MLS with your existing chassis and modules. MLS also allows you to use either an integrated RSM or an external router for route processing and Cisco IOS services.
Fast Convergence	Allows you to respond to route failures and routing topology changes by performing hardware-assisted invalidation of flow entries.
Resilience	Provides the benefits of HSRP ³ without additional configuration. This feature enables the switches to transparently switch over to the hot standby backup router when the primary router goes offline, eliminating a single point of failure in the network.
Access Lists	Allows you to set up access lists to filter, or to prevent traffic between members of different subnets. MLS enforces multiple security levels on every packet of the flow at wire speed. It allows you to configure and enforce access control rules on the RSM. Because MLS parses the packet up to the transport layer, it enables access lists to be validated. By providing multiple security levels, MLS enables you to set up rules and control traffic based on IP addresses as well as transport-layer application port numbers.
Accounting and Traffic Management	Allows you to see data flows as they are switched for troubleshooting, traffic management, and accounting purposes. MLS uses NDE to export the flow statistics. Data collection of flow statistics is maintained in hardware with no impact on switching performance. The records for expired and purged flows are grouped together and exported to applications such as NetSys for network planning, RMON ⁴ traffic management and monitoring, and accounting applications.

Table 11 Summary of Key Features (Continued)

Feature	Description
Network Design Simplification	Enables you to speed up your network while retaining the existing subnet structure. It makes the number of Layer 3 hops irrelevant in campus design, enabling you to cope with increases in any-to-any traffic.
Media Speed Access to Server Farms	You do not have to centralize servers in multiple VLANs to get direct connections. By providing security on a per-flow basis, you can control access to the servers and filter traffic based on subnet numbers and transport-layer application ports without compromising Layer 3 switching performance.
Faster Interworkgroup Connectivity	Addresses the need for higher-performance interworkgroup connectivity by intranet and multimedia applications. By deploying MLS, you gain the benefits of both switching and routing on the same platform.

- 1 DHCP = Dynamic Host Configuration Protocol
- 2 IETF = Internet Engineering Task Force
- 3 HSRP = Hot Standby Router Protocol
- 4 RMON2 = Remote Monitoring 2

Introduction to Multilayer Switching

Layer 3 protocols, such as IP and Internetwork Packet Exchange (IPX), are connectionless—they deliver each packet independently of each other. 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 transport-layer information. Communication from a client to a server and from the server to the client are separate flows. For example, Hypertext Transfer Protocol (HTTP) Web packets from a particular source to a particular destination are a separate flow from File Transfer Protocol (FTP) file transfer packets between the same pair of hosts.

Flows can be based on only Layer 3 addresses. This feature allows IP traffic from multiple users or applications to a particular destination to be carried on a single flow if only the destination IP address is used to identify a flow.

The NFFC maintains a Layer 3 switching table (MLS cache) for the Layer 3-switched flows. The cache also includes entries for traffic statistics that are updated in tandem with the switching of packets. 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. When the Layer 3-switching entry for a flow ages out, the flow statistics can be exported to a flow collector application.


Multilayer Switching Implementation

This section provides a step-by-step description of MLS implementation.

Note The MLS-RPs shown in the figures represent either a Route Switch Module (RSM) or an externally attached Cisco router.

Step 1 The MLSP informs the Catalyst 5000 series switch of the MLS-RP MAC addresses used on different VLANs and the MLS-RP's routing and access-list changes. Through this protocol, the MLS-RP multicasts its MAC and VLAN information to all MLS-SEs. When the MLS-SE hears the MLSP *hello* message indicating an MLS initialization, the MLS-SE is programmed with the MLS-RP MAC address and its associated VLAN number (see Figure 13).

Figure 13 MLS Implementation: Step 1




Step 2 In Figure 14, host A and host B are located on different VLANs. Host A initiates a data transfer to host B. When host A sends the first packet to the MLS-RP, the MLS-SE recognizes this packet as a *candidate packet* for Layer 3 switching because the MLS-SE has learned the MLS-RP's destination MAC address and VLAN through MLSP. The MLS-SE learns the Layer 3 flow information (such as the destination address, source address, and protocol port numbers), and forwards the first packet to the MLS-RP. A partial MLS entry for this Layer 3 flow is created in the MLS cache.

The MLS-RP receives the packet, looks at its route table to determine how to forward the packet, and applies services such as access control lists and class of service (COS) policy.

The MLS-RP rewrites the MAC header adding a new destination MAC address (host B's) and its own MAC address as the source.

Figure 14 **MLS Implementation: Step 2**


Because the Catalyst switch has learned the MAC and VLAN information of the MLS-RP, the switch starts the MLS process for the Layer 3 flow contained in this packet, the *candidate packet*



Step 3 The MLS-RP routes the packet to host B. When the packet appears back on the Catalyst 5000 series switch backplane, the MLS-SE recognizes the source MAC address as that of the MLS-RP, and that the packet’s flow information matches the flow for which it set up a candidate entry. The MLS-SE considers this packet an *enabler packet* and completes the MLS entry (established by the candidate packet) in the MLS cache (see Figure 15).

Figure 15 **MLS Implementation: Step 3**

The MLS-RP routes this packet to Host B. Because the MLS-SE has learned both this MLS-RP and the Layer 3 flow in this packet, it completes the MLS entry in the MLS cache. The first routed packet is called the *enabler packet*

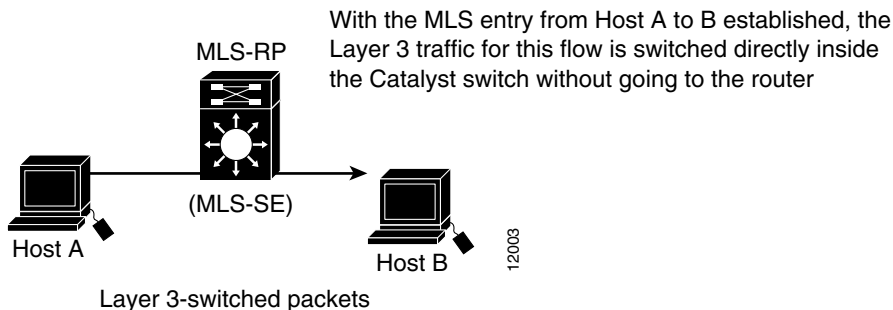


Step 4 After the MLS entry has been completed in Step 3, all Layer 3 packets with the same flow from host A to host B are Layer 3 switched directly inside the switch from host A to host B, bypassing the router (see Figure 16). After the Layer 3-switched path is established, the packet from host A is rewritten by the MLS-SE before it is forwarded to host B. The rewritten information includes the MAC addresses, encapsulations (when applicable), and some Layer 3 information.

The resultant packet format and protocol behavior is identical to that of a packet that is routed by the RSM or external Cisco router.

Note MLS is unidirectional. For host B to talk to host A, another Layer 3-switched path needs to be created from host B to host A.

Figure 16 **MLS Implementation: Step 4**



See the *Catalyst 5000 Series Multilayer Switching User Guide* for additional network implementation examples that include network topologies that do not support MLS.

Standard and Extended Access Lists

Note Router interfaces with input access lists *cannot* participate in MLS. However, any input access list can be translated to an output access list to provide the same effect on the interface. For complete details on how input and output access lists affect MLS, see the following chapter, *Configuring Multilayer Switching*.

MLS allows you to enforce access lists on every packet of the flow without compromising MLS performance. When you enable MLS, standard and extended access lists are handled at wire speed by the MLS-SE. Access lists configured on the MLS-RP take effect automatically on the MLS-SE.

Additionally, route topology changes and the addition of access lists are reflected in the switching path of MLS.

Consider the case where an access list is configured on the MLS-RP to deny access from station A to station B. When station A wants to talk to station B, it sends the first packet to the MLS-RP. The MLS-RP receives this packet and checks to see if this packet flow is permitted. If an access control list is configured for this flow, the packet is discarded. Because the first packet for this flow does not return from the MLS-RP, an MLS cache entry is not established by the MLS-SE.

In another case, access lists are introduced on the MLS-RP while the flow is already being Layer 3 switched within the MLS-SE. The MLS-SE immediately enforces security for the affected flow by purging it.

Similarly, when the MLS-RP detects a routing topology change, the appropriate MLS cache entries are deleted in the MLS-SE. The techniques for handling route and access list changes apply to both the RSM and directly attached external routers.

Restrictions on Using IP Router Commands with MLS Enabled

When you issue some Cisco IOS commands you will affect Multilayer Switching on your router. The commands that will affect MLS are as follows:

- **clear ip-route**—Clears all MLS cache entries for all Catalyst 5000 series switches performing Layer 3 switching for this MLS-RP.

- **ip routing**—The **no** form purges all MLS cache entries and disables MLS on this MLS-RP.
- **ip security** (all forms of this command)—Disables MLS on the interface.
- **ip tcp compression-connections**—Disables MLS on the interface.
- **ip tcp header-compression**—Disables MLS on the interface.

General Guidelines

- When you enable MLS, the RSM or externally attached router continues to handle all non-IP protocols while offloading the switching of IP packets to the MLS-SE.
- Do not confuse MLS with the NetFlow switching supported by Cisco routers. MLS uses both the RSM or directly attached external router and the MLS-SE. With MLS, you *are not* required to use NetFlow switching on the RSM or directly attached external router; any switching path on the RSM or directly attached external router will work (process, fast and so on).

Note The 10/100BaseTX and 100BaseFX Backbone Fast Ethernet Switching modules for the Catalyst 5000 have onboard hardware that optimizes MLS performance.

Software and Hardware Requirements

MLS requires these software and hardware versions:

- Catalyst 5000 series supervisor engine software
 - Release 4.1(1) or later
- Cisco IOS router software
 - 11.3(2)WA4(4) or later
- Supervisor Engine III with NetFlow Feature Card (NFFC)
- Route Switch Module (RSM) or Cisco 7500, 7200, 4500, or 4700 series router

Guidelines for External Routers

When using an external router, follow these guidelines:

- We recommend one directly attached external router per Catalyst 5000 series switch to ensure that the MLS-SE caches the appropriate flow information from both sides of the routed flow.
- You can use Cisco high-end routers (Cisco 7500, 7200, 4500, and 4700 series) for MLS when they are externally attached to the Catalyst 5000 series switch. You can make the attachment with multiple Ethernets (one per subnet), by using Fast Ethernet with the Inter-Switch Link (ISL), or with Fast Etherchannel.
- You can connect end hosts through any media (Ethernet, Fast Ethernet, ATM, and Fiber Distributed Data Interface [FDDI]) but the connection between the external router and the Catalyst 5000 series switch must be through standard 10/100 Ethernet interfaces, ISL links, or Fast Etherchannel.

Features that Affect MLS

This section describes how certain features affect MLS.

Access Lists

The following sections describe how access lists affect MLS.

Input Access Lists

Router interfaces with input access lists *cannot* participate in MLS. If you configure an input access list on an interface, all packets for a flow that are destined for that interface go through the router (even if the flow is allowed by the router it is not Layer 3 switched). Existing flows for that interface get purged and no new flows are cached.

Note Any input access list can be translated to an output access list to provide the same effect on the interface.

Output Access Lists

If an output access list is applied to an interface, the MLS cache entries for that interface are purged. Entries associated with other interfaces are not affected; they follow their normal aging or purging procedures.

Applying an output access list to an interface, when the access list is configured using the **log**, **precedence**, **tos**, or **establish** options, prevents the interface from participating in MLS.

Access List Impact on Flow Masks

Access lists impact the flow mask advertised by an MLS-RP. When there is no access list on any MLS-RP interface, the flow mask mode is destination-ip (the least specific). When there is a standard access list on any of the MLS-RP interfaces, the mode is source-destination-ip. When there is an extended access list on any of the MLS-RP interfaces, the mode is ip-flow (the most specific).

Reflexive Access Lists

Router interfaces with reflexive access lists *cannot* participate in Layer 3 switching.

IP Accounting

Enabling IP accounting on an MLS-enabled interface disables the IP accounting functions on that interface.

Note To collect statistics for the Layer 3-switched traffic, enable NetFlow Data Export (NDE).

Data Encryption

MLS is disabled on an interface when the data encryption feature is configured on the interface.

Policy Route-Map

MLS is disabled on an interface when a policy route-map is configured on the interface.

TCP Intercept

With MLS interfaces enabled, the TCP intercept feature (enabled in global configuration mode) might not work properly. When you enable the TCP intercept feature, the following message displays:

```
Command accepted, interfaces with mls might cause inconsistent behavior.
```

Network Address Translation

MLS is disabled on an interface when Network Address Translation (NAT) is configured on the interface.

Committed Access Rate

MLS is disabled on an interface when Committed Access Rate (CAR) is configured on the interface.

Maximum Transmission Unit

The maximum transmission unit (MTU) for an MLS interface must be the default Ethernet MTU, 1500 bytes.

To change the MTU on an MLS-enabled interface, you must first disable MLS on the interface (enter **no mls rp ip** on the interface). If you attempt to change the MTU with MLS enabled, the following message displays:

```
Need to turn off the mls router for this interface first.
```

If you attempt to enable MLS on an interface that has an MTU value other than the default value, the following message will be displayed:

```
mls only supports interfaces with default mtu size
```

