



Configuring IP Unicast Layer 3 Switching on Supervisor Engine 2

This chapter describes how to configure IP unicast Layer 3 switching for Policy Feature Card 2 (PFC2), Distributed Forwarding Cards (DFCs), and Multilayer Switch Feature Card 2 (MSFC2).



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 and the publications at this URL:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/tsd_products_support_eol_series_home.html

This chapter consists of these sections:

- [Understanding How Layer 3 Switching Works, page 17-1](#)
- [Default Hardware Layer 3 Switching Configuration, page 17-4](#)
- [Layer 3 Switching Configuration Guidelines and Restrictions, page 17-4](#)
- [Configuring Hardware Layer 3 Switching, page 17-5](#)
- [Displaying Hardware Layer 3 Switching Statistics, page 17-6](#)



Note

- Supervisor Engine 2, PFC2, and MSFC2 support IPX with fast switching on the MSFC2. For more information, refer to this URL:

http://www.cisco.com/en/US/docs/ios/12_1/switch/configuration/guide/xcdipsp.html

- For information about IP multicast Layer 3 switching, see [Chapter 18, “Configuring IP Multicast Layer 3 Switching.”](#)

Understanding How Layer 3 Switching Works

These sections describe Layer 3 switching with PFC2 and DFCs:

- [Understanding Hardware Layer 3 Switching on PFC2 and DFCs, page 17-2](#)
- [Understanding Layer 3-Switched Packet Rewrite, page 17-2](#)

Understanding Hardware Layer 3 Switching on PFC2 and DFCs

Hardware Layer 3 switching allows the PFC2 and DFCs, instead of the MSFC2, to forward IP unicast traffic between subnets. Hardware Layer 3 switching provides wire-speed forwarding on the PFC2 and DFCs, instead of in software on the MSFC2. Hardware Layer 3 switching requires minimal support from the MSFC2. The MSFC2 routes any traffic that cannot be hardware Layer 3 switched.

Hardware Layer 3 switching supports the routing protocols configured on the MSFC2. Hardware Layer 3 switching does not replace the routing protocols configured on the MSFC2.

Hardware Layer 3 switching, which runs equally on the PFC2 and DFCs to provide IP unicast Layer 3 switching locally on each module, consists of the following functions:

- Hardware access control list (ACL) switching—For policy-based routing (PBR)
- Hardware NetFlow switching—For TCP intercept, reflexive ACL forwarding decisions, Web Cache Communication Protocol (WCCP), and server load balancing (SLB)
- Hardware Cisco Express Forwarding (CEF) switching—For all other IP unicast traffic

Hardware Layer 3 switching on the PFC2 supports modules that do not have a DFC. The MSFC2 forwards traffic that cannot be Layer 3 switched.

Traffic is hardware Layer 3 switched after being processed by access lists and quality of service (QoS).

Hardware Layer 3 switching makes a forwarding decision locally on the ingress-port module for each packet and sends the rewrite information for each packet to the egress port, where the rewrite occurs when the packet is transmitted from the Catalyst 6500 series switch.

Hardware Layer 3 switching generates flow statistics for Layer 3-switched traffic. Hardware Layer 3 flow statistics can be used for NetFlow Data Export (NDE). (See [Chapter 33, “Configuring NDE”](#).)

Understanding Layer 3-Switched Packet Rewrite

When a packet is Layer 3 switched from a source in one subnet to a destination in another subnet, the Catalyst 6500 series switch performs a packet rewrite at the egress port based on information learned from the MSFC2 so that the packets appear to have been routed by the MSFC2.

Packet rewrite alters five fields:

- Layer 2 (MAC) destination address
- Layer 2 (MAC) source address
- Layer 3 IP Time to Live (TTL)
- Layer 3 checksum
- Layer 2 (MAC) checksum (also called the frame checksum or FCS)



Note

Packets are rewritten with the encapsulation appropriate for the next-hop subnet.

If Source A and Destination B are in different subnets and Source A sends a packet to the MSFC2 to be routed to Destination B, the switch recognizes that the packet was sent to the Layer 2 (MAC) address of the MSFC2.

To perform Layer 3 switching, the switch rewrites the Layer 2 frame header, changing the Layer 2 destination address to the Layer 2 address of Destination B and the Layer 2 source address to the Layer 2 address of the MSFC2. The Layer 3 addresses remain the same.

In IP unicast and IP multicast traffic, the switch decrements the Layer 3 TTL value by 1 and recomputes the Layer 3 packet checksum. The switch recomputes the Layer 2 frame checksum and forwards (or, for multicast packets, replicates as necessary) the rewritten packet to Destination B's subnet.

A received IP unicast packet is formatted (conceptually) as follows:

Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>MSFC2 MAC</i>	<i>Source A MAC</i>	<i>Destination B IP</i>	<i>Source A IP</i>	<i>n</i>	<i>calculation1</i>		

After the switch rewrites an IP unicast packet, it is formatted (conceptually) as follows:

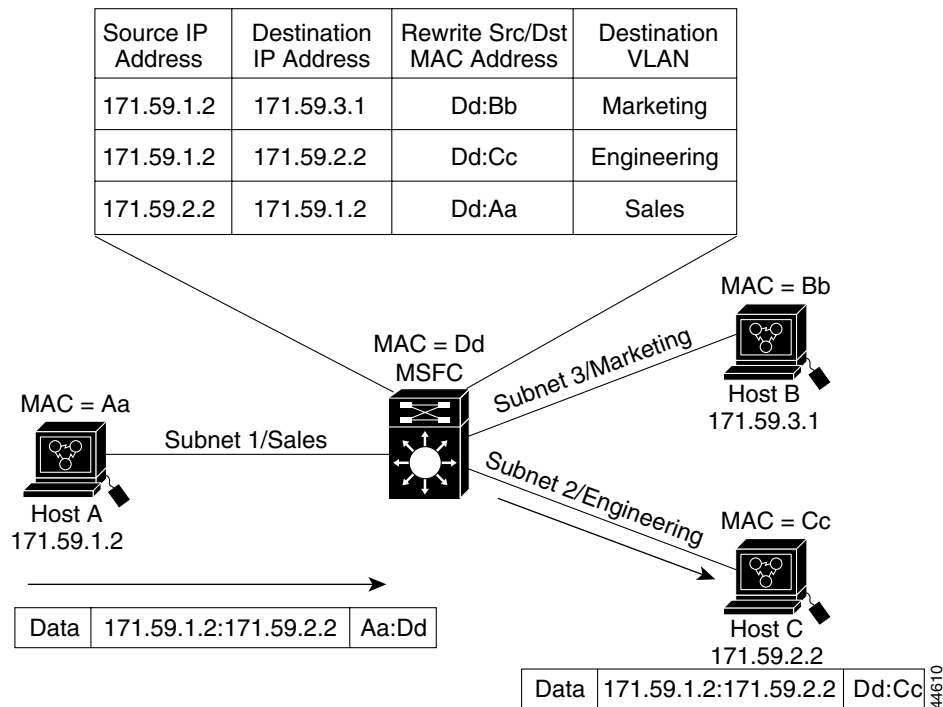
Layer 2 Frame Header		Layer 3 IP Header				Data	FCS
Destination	Source	Destination	Source	TTL	Checksum		
<i>Destination B MAC</i>	<i>MSFC2 MAC</i>	<i>Destination B IP</i>	<i>Source A IP</i>	<i>n-1</i>	<i>calculation2</i>		

Hardware Layer 3 Switching Examples

Figure 17-1 on page 17-3 shows a simple network topology. In this example, Host A is on the Sales VLAN (IP subnet 171.59.1.0), Host B is on the Marketing VLAN (IP subnet 171.59.3.0), and Host C is on the Engineering VLAN (IP subnet 171.59.2.0).

When Host A initiates an HTTP file transfer to Host C, Hardware Layer 3 switching uses the information in the local forwarding information base (FIB) and adjacency table to forward packets from Host A to Host C.

Figure 17-1 Hardware Layer 3 Switching Example Topology



Default Hardware Layer 3 Switching Configuration

Table 17-1 shows the default hardware Layer 3 switching configuration.

Table 17-1 Default Hardware Layer 3 Switching Configuration

Feature	Default Value
Hardware Layer 3 switching enable state	Enabled (cannot be disabled)
Cisco IOS CEF enable state on MSFC2	Enabled (cannot be disabled)
Cisco IOS dCEF ¹ enable state on MSFC2	Enabled (cannot be disabled)
IGMP ² snooping	Enabled
Multicast routing on MSFC2	Disabled globally
PIM ³ routing on MSFC2	Disabled on all Layer 3 interfaces
IP multicast Layer 3 switching threshold	Unconfigured—no default value
IP multicast Layer 3 switching	Enabled when multicast routing is enabled and IP PIM is enabled on the interface

1. dCEF = Distributed Cisco Express Forwarding
2. IGMP = Internet Group Management Protocol
3. PIM = Protocol Independent Multicast

Layer 3 Switching Configuration Guidelines and Restrictions

Follow these guidelines and restrictions when configuring hardware Layer 3 switching:

- The PFC2 supports a maximum of 16 unique Hot Standby Routing Protocol (HSRP) group numbers. You can use the same HSRP group numbers in different VLANs. If you configure more than 16 HSRP groups, this restriction prevents use of the VLAN number as the HSRP group number.



Note Identically numbered HSRP groups use the same virtual MAC address, which might cause errors if you configure bridge groups.

- Hardware Layer 3 switching supports the following ingress and egress encapsulations:
 - Ethernet V2.0 (ARPA)
 - 802.3 with 802.2 with 1 byte control (SAP1)
 - 802.3 with 802.2 and SNAP



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.

Configuring Hardware Layer 3 Switching



Note

For information on configuring unicast routing on the MSFC2, see [Chapter 12, “Configuring Layer 3 Interfaces.”](#)

Hardware Layer 3 switching is permanently enabled on Supervisor Engine 2 with PFC2, MSFC2, and Distributed Feature Card (DFC). No configuration is required.

To display information about Layer 3-switched traffic, perform this task:

Command	Purpose
Router# show interface {{type ¹ slot/port} {port-channel number}} begin L3	Displays a summary of Layer 3-switched traffic.

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

This example shows how to display information about hardware Layer 3-switched traffic on Fast Ethernet port 3/3:

```
Router# show interface fastethernet 3/3 | begin L3
L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 12 pkt, 778 bytes mcast
L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
4046399 packets input, 349370039 bytes, 0 no buffer
Received 3795255 broadcasts, 2 runts, 0 giants, 0 throttles
<...output truncated...>
Router#
```



Note

The Layer 3 switching packet count is updated approximately every five seconds.

Cisco IOS CEF and dCEF are permanently enabled on the MSFC2. No configuration is required to support hardware Layer 3 switching.

The Cisco IOS CEF **ip load-sharing per-packet**, **ip cef accounting per-prefix**, and **ip cef accounting non-recursive** commands on the MSFC2 apply only to traffic that is CEF-switched in software on the MSFC2. The commands do not affect traffic that is hardware Layer 3 switched on the PFC2 or on DFC-equipped switching modules.

For information about Cisco IOS CEF and dCEF on the MSFC2, refer to these publications:

- The “Cisco Express Forwarding” section at this URL:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_c/xcprt2/index.htm
- The *Cisco IOS Switching Services Command Reference* publication at this URL:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_r/index.htm

Displaying Hardware Layer 3 Switching Statistics

Hardware Layer 3 switching statistics are obtained on a per-VLAN basis.

To display hardware Layer 3 switching statistics, perform this task:

Command	Purpose
Router# show interfaces <i>{{type¹ slot/port}</i> <i>{port-channel number}</i>	Displays hardware Layer 3 switching statistics.

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

This example shows how to display hardware Layer 3 switching statistics:

```
Router# show interfaces gigabitethernet 9/5 | include Switched
L2 Switched: ucast: 8199 pkt, 1362060 bytes - mcast: 6980 pkt, 371952 bytes
L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes mcast
L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
```

To display adjacency table information, perform this task:

Command	Purpose
Router# show adjacency <i>[{{type¹ slot/port}</i> <i>{port-channel number}</i>] detail internal summary	Displays adjacency table information. The optional detail keyword displays detailed adjacency information, including Layer 2 information.

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

This example shows how to display adjacency statistics:

```
Router# show adjacency gigabitethernet 9/5 detail
Protocol Interface Address
IP GigabitEthernet9/5 172.20.53.206(11)
504 packets, 6110 bytes
00605C865B82
000164F83FA50800
ARP 03:49:31
```



Note

Adjacency statistics are updated approximately every 60 seconds.