IP Unicast Layer 3 Switching

Note
- For complete syntax and usage information for the commands used in this chapter, see these publications:

  http://www.cisco.com/en/US/products/ps9536/prod_command_reference_list.html

- Cisco IOS Release 12.2SY supports only Ethernet interfaces. Cisco IOS Release 12.2SY does not support any WAN features or commands.

- For information about IP multicast Layer 3 switching, see Chapter 40, "IPv4 Multicast Layer 3 Features."

Tip
For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

Participate in the Technical Documentation Ideas forum

Prerequisites for Hardware Layer 3 Switching

None.

Restrictions for Hardware Layer 3 Switching

- 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)

Information About Layer 3 Switching

Hardware Layer 3 Switching

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

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

Hardware Layer 3 switching runs equally on the PF3 and DFCs to provide IP unicast Layer 3 switching locally on each module. Hardware Layer 3 switching provides the following functions:

- Hardware access control list (ACL) switching for policy-based routing (PBR)
- Hardware flow-based switching for TCP intercept and reflexive ACL forwarding decisions
- Hardware Cisco Express Forwarding (CEF) switching for all other IP unicast traffic

Hardware Layer 3 switching on the PFC supports modules that do not have a DFC. The RP 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 switch.

Hardware Layer 3 switching generates flow statistics for Layer 3-switched traffic. Hardware Layer 3 flow statistics can be used for NetFlow. (See Chapter 49, "NetFlow Hardware Support".)

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 switch performs a packet rewrite at the egress port based on information learned from the RP so that the packets appear to have been routed by the RP.

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 RP to be routed to Destination B, the switch recognizes that the packet was sent to the Layer 2 (MAC) address of the RP.

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 RP. 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		
RP MAC	*Source A MAC*	*Destination B IP*	*Source A IP*	*n*	*calculation1*		

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		
Destination B MAC	*RP MAC*	*Destination B IP*	*Source A IP*	*n-1*	*calculation2*		

Hardware Layer 3 Switching Examples

Figure 30-1 on page 30-4 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 30-1   Hardware Layer 3 Switching Example Topology
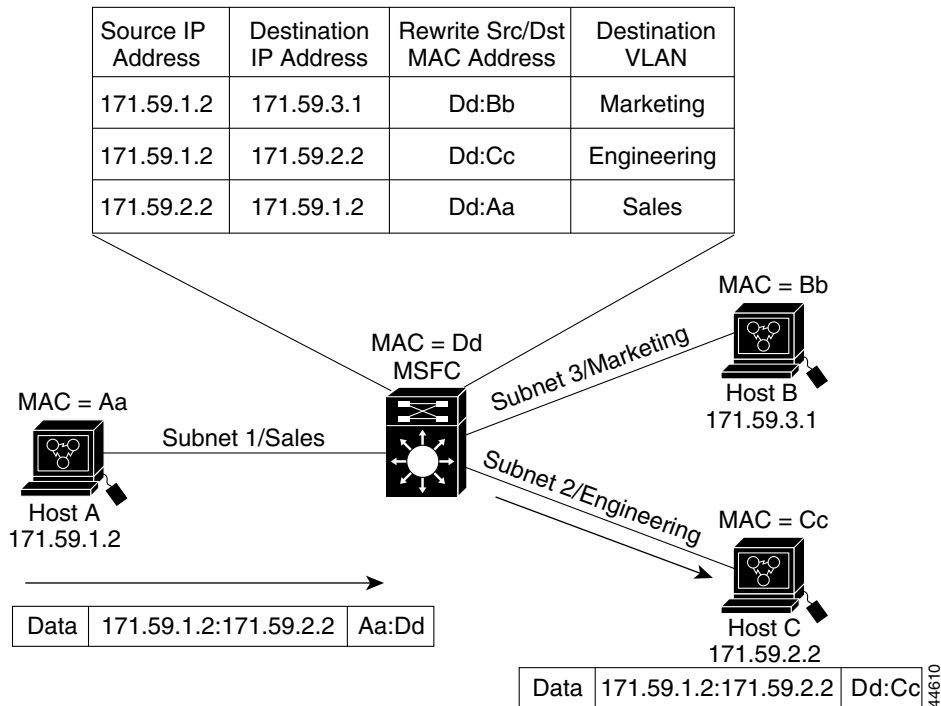
Default Settings for Hardware Layer 3 Switching

Feature	Default Value
Hardware Layer 3 switching enable state	Enabled (cannot be disabled)
Cisco IOS CEF enable state on RP	Enabled (cannot be disabled)
Cisco IOS dCEF enable state on RP	Enabled (cannot be disabled)

How to Configure Hardware Layer 3 Switching

Note   For information on configuring unicast routing on the RP, see Chapter 32, "Layer 3 Interfaces."

Hardware Layer 3 switching is permanently enabled. No configuration is required.

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

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

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

```
Router# show interface gigabitethernet 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. No configuration is required to support hardware Layer 3 switching.

With a PFC (and DFCs, if present), hardware Layer 3 switching uses per-flow load balancing based on IP source and destination addresses. Per-flow load balancing avoids the packet reordering that can be necessary with per-packet load balancing. For any given flow, all PFC- and DFC-equipped switches make exactly the same load-balancing decision, which can result in nonrandom load balancing.

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

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** {{*type slot/port*} | {**port-channel** *number*}}	Displays hardware Layer 3 switching statistics.

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** [{{*type slot/port*} | {**port-channel** *number*}} | **detail** | **internal** | **summary**]	Displays adjacency table information. The optional **detail** keyword displays detailed adjacency information, including Layer 2 information.

This example shows how to display adjacency statistics, which are updated approximately every 60 seconds:

```
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
```

Tip   For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

Participate in the Technical Documentation Ideas forum