



CHAPTER 4

Configuring a Route Switch Processor 720

This chapter describes how to configure a route switch processor 720 (RSP720). The RSP720 is the newest type of supervisor engine available for Cisco 7600 series routers. The RSP720 consists of a full-size board and two integrated daughter cards: the MSFC4 and a PFC3C or PFC3CXL. The RSP720 has an integrated switch fabric that interconnects all of the line cards in the Cisco 7600 router with point-to-point 20-Gbps full-duplex serial channels.

See [Appendix C, “Cisco IOS Release 12.2SRB Software Images,”](#) for information about the Cisco IOS software images available for the RSP720, Sup720, and Sup32.

This chapter contains these sections:

- [RSP720 PFC Compatibility Matrix, page 4-2](#)
- [RSP720 Features, page 4-2](#)
- [Accessing Flash Memory on the RSP720, page 4-6](#)
- [Configuring route switch processor 720 Ports, page 4-6](#)
- [Configuring and Monitoring the Switch Fabric Functionality, page 4-7](#)

For complete syntax and usage information for the commands in this chapter, see the Cisco 7600 series routers command references at this URL:

http://www.cisco.com/en/US/products/hw/routers/ps368/prod_command_reference_list.html



Note

- The RSP720 is supported in all Cisco 7600 chassis except the Cisco 7603 and Cisco OSR-7609.
- With a 4-slot chassis, install the RSP720 in slot 1 or 2.
- With a 6-slot or a 9-slot chassis (including enhanced [-S] chassis), install the RSP720 in slot 5 or 6.
- With a 13-slot chassis, install the RSP720 in slot 7 or 8.

RSP720 PFC Compatibility Matrix

The route switch processor 720 (RSP720) is configured with one of two types of Policy Feature Card: a PFC3C or a PFC3CXL. The PFC on the RSP720 can interoperate with cards that have a lower version number, such as a Distributed Forwarding Card (DFC) at version 3B (a DFC3B). A card's version number indicates its operating mode: 3B, 3BXL, 3C, or 3CXL.

For the RSP720 to interoperate with lower-version cards, the system determines the lowest operating mode of all installed cards (3B, 3BXL, 3C, or 3CXL) and applies this mode to all of the cards. The cards provide the features supported in the selected operating mode (even if a card has a higher version number). For example, a PFC3C operating in 3B mode offers only those features supported by a PFC3B.

Here are some examples of mode setting among different version cards:

- A system with an RSP720-3CXL, a DFC3BXL, and an Ethernet Services Module (ES20-3C) operates in 3BXL mode (which is the lowest operating mode among all cards).
- A system with an RSP720-3CXL and an ES20-3C operates in 3C mode.
- A system with an RSP720-3C and a DFC3B operates in 3B mode.
- A system with an RSP720-3CXL and an ES20-3CXL operates in 3CXL mode.


Note

Use the **show platform hardware pfc mode** command to display the PFC operating mode.


Note

OSMs are not supported on RSP 720 Supervisor card.

RSP720 Features

The RSP720 is the newest version of supervisor engine for the Cisco 7600 series router. Along with its two new integrated daughter cards (a PFC3C or PFC3CXL and an MSFC4), the RSP720 provides many enhancements and new features over previous supervisor engines. These enhancements and features are described in the sections that follow.


Note

The RSP720 supports all of the features as the Supervisor Engine 720 (with PFC3B or PFC3BXL). Unless otherwise noted, the configuration and operation of the features described in later chapters of this document is the same for both types of processors (RSP720 and Sup720).

Hardware

- Two new integrated daughter cards: a PFC3C or PFC3CXL and an MSFC4
- Faster CPUs and more default memory on the route processor (RP) and switch processor (SP)
 - RSP720-3C-GE: 1-GB DRAM on RP and SP
 - RSP720-3CXL-GE: 2-GB DRAM (RP) and 1-GB DRAM (SP)
- Additional memory provides a larger MAC address table
- Layer 2 and Layer 3 functions have been integrated on a single ASIC
- ASIC (hardware) forwarding of IP and MPLS traffic

For information about hardware support for the RSP720, see the “Route Switch Processor 720” section in Chapter 2 of the *Cisco 7600 Series Router Supervisor Engine and Route Switch Processor Guide*.

Performance

- Faster software bootup
- Faster protocol convergence (BGP, OSPF) and ARP learning
- Improved IGMP snooping times
- Faster speeds for establishing DHCP servers, Label Distribution Protocol (LDP) sessions, IP sessions, and Traffic Engineering (TE)
- Faster processing for Bidirectional Forwarding Detection (BFD), Resource Reservation Setup Protocol (RSVP), and other control-plane functions
- Improved speeds for accessing and copying local files

Scalability

- 30 million packets-per-second (Mpps) forwarding rates for Layer 2 and Layer 3 traffic. The RSP720 uses hardware-based Cisco Express Forwarding (CEF). Forwarding rates are:
 - IP forwarding rates—30 Mpps
 - MPLS forwarding rates—20 Mpps
- Support for larger customer configurations and more interfaces:
 - 32000 IP subscriber sessions
 - 1 million routes
 - 96000 MAC addresses maximum (80000 in real life), up from 64000
 - 32000 VLANs
 - 128000 Address Resolution Protocol (ARP) entries

High-Availability Features

- Online insertion and removal (OIR)
- Route processor redundancy (RPR and RPR+)
- Nonstop forwarding with stateful switchover (NSF/SSO)
- Fast-fabric switchover
- In Service Software Upgrade (ISSU) and enhanced Fast Software Upgrade (eFSU)
(Cisco IOS Release 12.2SRB1 and later)

IPv6 ACL Enhancements (Security)

Support for 2K access control list (ACL) labels and 16K access control entries (ACEs), up from 1K masks and 8K ACEs

Rate-Limiting of Unknown Unicast Packets

Allows you to limit the number of unknown unicast packets that the router processes and thus keep the packets from flooding the network. If the number of unknown packets received by the router exceeds the specified rate, excess packets are not forwarded. See the next section (“[Configuration Guidelines for Unknown Unicast Packet Rate-Limiting](#)”) for configuration guidelines.

The following new commands are provided to configure and verify this feature:

- Use the following command to configure rate-limiting, where *pps* is the maximum number of unknown unicast packets to allow per second (from 10 to 1000000) and *packets-in-burst* is an optional packet burst rate (from 1 to 255, with a default value of 10). The **no** form of the command turns off rate-limiting for unknown packets.

```
Router(config)# mls rate-limit layer2 unknown pps [packets-in-burst]
Router(config)# no mls rate-limit layer2 unknown
```



Note If any physical ports on the router are configured for routing, issue the **mac-address-table learning interface** *interface* command (in global configuration mode) on each of those ports. Otherwise, the rate-limiting counts might not be accurate.

- Use the **show mls rate-limit** command to verify that rate-limiting of unknown unicast packets is enabled. If rate-limiting for unknown unicast packets is enabled, the output will include the following rate-limiter type:

```
Router(config)# show mls rate-limit

Sharing Codes: S - static, D - dynamic
Codes dynamic sharing: H - owner (head) of the group, g - guest of the group

Rate Limiter Type      Status    Packets/s    Burst    Sharing
-----  -----  -----  -----  -----
UCAST IP TINY FRAG     On        100000       100      Not sharing
```

Configuration Guidelines for Unknown Unicast Packet Rate-Limiting

Observe these guidelines when configuring unknown unicast packet rate-limiting:

- This feature is available only with the PFC3C and PFC3CXL (RSP720). It is not available with the PFC3B or PFC3BXL.
- If you run the Remote Switched Port Analyzer (RSPAN) with unknown unicast rate-limiting configured, be aware that traffic amounts might differ between RSPAN source and destination ports. This difference occurs if the traffic being monitored contains unknown unicast packets. In this case, the unknown unicast traffic is rate-limited before being sent to the RSPAN destination port, resulting in a mismatch between the amount of traffic at the RSPAN source and destination ports.

Packet Fragmentation over GRE Tunnels

Support for packet fragmentation over Generic Routed Encapsulation (GRE) tunnels. With the PFC3C and PFC3CXL, you can use the [**no**] **mls cef tunnelfragment** command to set the don't fragment (DF) bit to zero, which allows the PFC3C or PFC3CXL to reassemble fragmented GRE traffic. The **no** form of the command turns off tunnel fragmentation and causes fragmented GRE traffic to be dropped.



Note To use this feature, the router at the other end of the tunnel must support tunnel fragmentation.

Improved Load Balancing on GE Bundles

Load balancing improvements on Gigabit Ethernet (GE) bundles configured as 802.1q trunks:

- The VLAN ID is now included in the bundle hash for multicast traffic.
- Multicast receivers that handle multicast traffic for multiple VLANs can load balance the traffic across the member links in the bundle.
- The router provides more efficient load balancing of fragmented traffic.

QoS Enhancements

- On the PFC3C and PFC3CXL you can configure ingress and egress policers to operate independently of each other (in *serial mode*). Normally, ingress and egress policers operate in parallel mode, where action by one policer causes a corresponding action in the other. For example, if the egress policer drops a packet, the ingress policer does not count the packet either. Note that this change does not affect marking using policers.

To enable serial mode for ingress and egress policers on the PFC3C or PFC3CXL, use the following new command in global configuration mode. The **no** form of the command disables serial mode and resets the policing mode to parallel.

[no] mls qos police serial

- Marking packets after recirculation. Rather than using the trust of the original input interface, the PFC3C and PFC3CXL treat recirculated packets as untrusted. This enhancement allows recirculated packets to be marked by an ingress policy.
- Ingress IP DSCP and MPLS EXP marking at the IP-to-MPLS edge. This PFC3C and PFC3CXL enhancement allows you to mark both the IP DSCP bits (**set ip dscp**) and the MPLS EXP bits (**set mpls exp**) during MPLS label imposition. Note that if you do not issue the **set mpls exp** command, the router copies the IP DSCP bits to EXP.
- Ingress EXP marking does not affect locally routed IP-to-IP traffic. With the PFC3C and PFC3CXL, you can use the **no mls qos rewrite ip dscp** command to turn off the egress QoS rewrite of PFC QoS logic, which keeps locally routed IP-to-IP traffic from being affected by EXP marking.
- Concurrent CoS and DSCP transparency for Layer 2 VPNs. This PFC3C and PFC3CXL enhancement enables customers to deploy a combination of Layer 2 VPNs and Layer 3 VPNs for use in a triple-play network (video, Voice over IP [VoIP], and data access [Internet]). It also supports Quality of Service (QoS) guarantees for traffic. The feature results in the following enhancements:
 - You can preserve the CoS and DSCP settings for VPLS and SVI-based EoMPLS, by using the **platform vfi dot1q-transparency** command in conjunction with the **no mls qos rewrite ip dscp** command.
 - The **no mls qos rewrite ip dscp** command can now be used with MPLS. Note that the router must be in PFC3C or PFC3CXL mode, which means that the router cannot contain any Cisco 7600 SIP-600 or WS-X6xxx cards (with a DFC3B or DFC3BXL).
 - Because the **no mls qos rewrite ip dscp** command is now compatible with MPLS, Layer 3 VPNs can now be terminated on the same provider edge (PE).
- A new cli command **mls qos recirc untrust** is used to prevent QoS data from getting reset during the second pass lookup over internal vlans for the mvpn case.



Note

For complete syntax and usage information for the command **mls qos recirc untrust**, refer to the Cisco 7600 Series Routers Command References at this URL:
http://www.cisco.com/en/US/docs/ios/qos/command/reference/qos_book.html

Other Enhancements

- Support for Ethernet over MPLS (EoMPLS) control word
- ACL counters have been enhanced to include counts of packets coming from the route processor (RP)
- IPv6 packet fragments handled in the hardware

Unsupported Features

The following Sup720 features are not supported on the RSP720:

■ Accessing Flash Memory on the RSP720

- Server load balancing (SLB)

Accessing Flash Memory on the RSP720

Table 4-1 lists the names of the flash memory devices on the RSP720. To access the appropriate flash memory (internal or external), use these keywords when you issue software commands through the command line interface (CLI):

Table 4-1 CLI Keywords for RSP720 Flash Memory Devices

CLI Keyword	Used to access...
bootdisk:	Internal flash memory on the active route processor (RP)
sup-bootdisk:	Internal flash memory on the active switch processor (SP)
slavebootdisk:	Internal flash memory on the redundant route processor (RP)
slavesup-bootdisk:	Internal flash memory on the redundant switch processor (SP)
disk0:	External flash memory on active RSP (Disk 0 on front panel)
disk1:	External flash memory on active RSP (Disk 1 on front panel)
slavedisk0:	External flash memory on redundant RSP (Disk 0 on front panel)
slavedisk1:	External flash memory on redundant RSP (Disk 1 on front panel)

Configuring route switch processor 720 Ports

route switch processor 720 port 1 has a small form-factor pluggable (SFP) connector and has no unique configuration options.

route switch processor 720 port 2 has an RJ-45 connector and an SFP connector (default). To use the RJ-45 connector, you must change the configuration.

To configure port 2 on a route switch processor 720 to use either the RJ-45 connector or the SFP connector, perform this task:

Step	Command	Purpose
Step 1	Router(config)# interface gigabitethernet slot/2	Selects the Ethernet port to be configured.
Step 2	Router(config-if)# media-type {rj45 sfp} or Router(config-if)# no media-type	Selects the connector to use. Reverts to the default configuration (SFP module).

This example shows how to configure port 2 of an RSP720 in slot 5 to use the RJ-45 connector:

```
Router(config)# interface gigabitethernet 5/2
Router(config-if)# media-type rj45
```

Configuring and Monitoring the Switch Fabric Functionality

These sections describe how to configure the switching mode and monitor the switch fabric functionality that is included on a route switch processor 720:

- [Understanding How the Switch Fabric Functionality Works, page 4-7](#)
- [Configuring the Switch Fabric Functionality, page 4-8](#)
- [Monitoring the Switch Fabric Functionality, page 4-9](#)

Understanding How the Switch Fabric Functionality Works

These sections describe how the switch fabric functionality works:

- [Switch Fabric Functionality Overview, page 4-7](#)
- [Forwarding Decisions for Layer 3-Switched Traffic, page 4-7](#)
- [Switching Modes, page 4-8](#)

Switch Fabric Functionality Overview

The switch fabric functionality is built into the route switch processor 720 and creates a dedicated connection between fabric-enabled modules and provides uninterrupted transmission of frames between these modules. In addition to the direct connection between fabric-enabled modules provided by the switch fabric functionality, fabric-enabled modules also have a direct connection to the 32-Gbps forwarding bus.

Forwarding Decisions for Layer 3-Switched Traffic

Either a PFC3 or a Distributed Forwarding Card 3 (DFC3) makes the forwarding decision for Layer 3-switched traffic, as follows:

- A PFC3 makes all forwarding decisions for each packet that enters the router through a module without a DFC3.
- A DFC3 makes all forwarding decisions for each packet that enters the router on a DFC3-enabled module in these situations:
 - If the egress port is on the same module as the ingress port, the DFC3 forwards the packet locally (the packet never leaves the module).
 - If the egress port is on a different fabric-enabled module, the DFC3 sends the packet to the egress module, which sends it out the egress port.
 - If the egress port is on a different nonfabric-enabled module, the DFC3 sends the packet to the route switch processor 720. The route switch processor 720 fabric interface transfers the packet to the 32-Gbps switching bus, where it is received by the egress module and is sent out the egress port.

Switching Modes

With a route switch processor 720, traffic is forwarded to and from modules in one of the following modes:

- Compact mode—The router uses this mode for all traffic when only fabric-enabled modules are installed. In this mode, a compact version of the DBus header (32 bytes) is forwarded over the switch fabric channel, which provides the best possible performance.
- Truncated mode—The router uses this mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the router sends a truncated version of the traffic (the first 64 bytes of the frame) over the switch fabric channel.
- Bus mode—The router uses this mode for traffic between nonfabric-enabled modules and for traffic between a nonfabric-enabled module and a fabric-enabled module. In this mode, all traffic passes between the local bus and the supervisor engine or RSP bus.

Table 4-2 shows the switching modes used with fabric-enabled and nonfabric-enabled modules installed.

Table 4-2 Switch Fabric Functionality Switching Modes

Modules	Switching Modes
Between fabric-enabled modules (when no nonfabric-enabled modules are installed)	Compact ¹
Between fabric-enabled modules (when nonfabric-enabled modules are also installed)	Truncated ²
Between fabric-enabled and nonfabric-enabled modules	Bus
Between nonfabric-enabled modules	Bus

1. In **show** commands, displayed as dcef mode for fabric-enabled modules with DFC3 installed; displayed as fabric mode for other fabric-enabled modules.
2. Displayed as fabric mode in **show** commands.

Configuring the Switch Fabric Functionality

To configure the switching mode, perform this task:

Command	Purpose
Router(config)# [no] fabric switching-mode allow {bus-mode {truncated [{threshold [number]}]}}	Configures the switching mode.

When configuring the switching mode, note the following information:

- To allow the use of nonfabric-enabled modules or to allow fabric-enabled modules to use bus mode, enter the **fabric switching-mode allow bus-mode** command.
- To prevent the use of nonfabric-enabled modules or to prevent fabric-enabled modules from using bus mode, enter the **no fabric switching-mode allow bus-mode** command.



Caution When you enter the **no fabric switching-mode allow bus-mode** command, power is removed from any nonfabric-enabled modules installed in the router.

- To allow fabric-enabled modules to use truncated mode, enter the **fabric switching-mode allow truncated** command.
- To prevent fabric-enabled modules from using truncated mode, enter the **no fabric switching-mode allow truncated** command.
- To configure how many fabric-enabled modules must be installed before they use truncated mode instead of bus mode, enter the **fabric switching-mode allow truncated threshold number** command.
- To return to the default truncated-mode threshold, enter the **no fabric switching-mode allow truncated threshold** command.

Monitoring the Switch Fabric Functionality

The switch fabric functionality supports a number of **show** commands for monitoring purposes. A fully automated startup sequence brings the module online and runs the connectivity diagnostics on the ports.

These sections describe how to monitor the switch fabric functionality:

- [Displaying the Switch Fabric Redundancy Status, page 4-9](#)
- [Displaying Fabric Channel Switching Modes, page 4-10](#)
- [Displaying the Fabric Status, page 4-10](#)
- [Displaying the Fabric Utilization, page 4-10](#)
- [Displaying Fabric Errors, page 4-11](#)

Displaying the Switch Fabric Redundancy Status

To display the switch fabric redundancy status, perform this task:

Command	Purpose
Router# show fabric active	Displays switch fabric redundancy status.

This example shows how to display the redundancy status of the switch fabric:

```
Router# show fabric active
Active fabric card in slot 5
No backup fabric card in the system
Router#
```

Displaying Fabric Channel Switching Modes

To display the fabric channel switching mode of one or all modules, perform this task:

Command	Purpose
Router# show fabric switching-mode [module slot_number all]	Displays fabric channel switching mode of one or all modules.

This example shows how to display the fabric channel switching mode of all modules:

```
Router# show fabric switching-mode all
%Truncated mode is allowed
%System is allowed to operate in legacy mode

Module Slot      Switching Mode      Bus Mode
      5            DCEF             Compact
      9            Crossbar         Compact
```

Displaying the Fabric Status

To display the fabric status of one or all switching modules, perform this task:

Command	Purpose
Router# show fabric status [slot_number all]	Displays fabric status.

This example shows how to display the fabric status of all modules:

```
Router# show fabric status all
      slot    channel      speed           module      fabric
                  status           status
      1        0          8G           OK          OK
      5        0          8G           OK          Up- Timeout
      6        0         20G           OK          Up- BufError
      8        0          8G           OK          OK
      8        1          8G           OK          OK
      9        0          8G           Down- DDRsync   OK
```

Displaying the Fabric Utilization

To display the fabric utilization of one or all modules, perform this task:

Command	Purpose
Router# show fabric utilization [slot_number all]	Displays fabric utilization.

This example shows how to display the fabric utilization of all modules:

```
Router# show fabric utilization all
Lo% Percentage of Low-priority traffic.
Hi% Percentage of High-priority traffic.
```

slot	channel	speed	Ingress Lo%	Egress Lo%	Ingress Hi%	Egress Hi%
5	0	20G	0	0	0	0
9	0	8G	0	0	0	0

Displaying Fabric Errors

To display fabric errors of one or all modules, perform this task:

Command	Purpose
Router# show fabric errors [slot_number all]	Displays fabric errors.

This example shows how to display fabric errors on all modules:

```
Router# show fabric errors all
```

Module errors:

slot	channel	crc	hbeat	sync	DDR	sync
1	0	0	0	0	0	0
8	0	0	0	0	0	0
8	1	0	0	0	0	0
9	0	0	0	0	0	0

Fabric errors:

slot	channel	sync	buffer	timeout
1	0	0	0	0
8	0	0	0	0
8	1	0	0	0
9	0	0	0	0

Configuring and Monitoring the Switch Fabric Functionality