



Understanding Route Processor Redundancy in the Catalyst 8500 MSR

[TAC Notice: What's Changing on TAC Web](#)

Document ID: 19900

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

[Operation of RP Redundancy](#)

[RP Redundancy Requirements](#)

[What Causes a Switchover?](#)

[What Does Not Cause a Switchover?](#)

[RP Redundancy Algorithm](#)

[Arbitration Logic to Determine RP State](#)

[Detection of RP State Changes and Switchover Trigger](#)

[Points](#)

[Configuration Synchronization](#)

[Related Information](#)

Help us help you.

Please rate this document.

- Excellent
- Good
- Average
- Fair
- Poor

This document solved my problem.

- Yes
- No
- Just browsing

Suggestions for improvement:

(256 character limit)

Send

Introduction

Redundancy provides a key building block for high availability by preventing equipment failures from causing service outages, as well as providing a means for hitless maintenance and upgrade activities. Redundancy does not guarantee high availability, which depends more on the levels of redundancy that the equipment is providing. The Cisco Catalyst 8540 Multiservice Switch Router (MSR) provides Route Processor (RP), Switch Processor (SP), and Power Supply (PS) redundancy. However, the Cisco Catalyst 8540 MSR does not provide redundancy at the port, module, or chassis level.

This document describes the following:

- The Cisco Catalyst 8540 one-to-one RP redundancy feature, in detail
- The RP redundancy operation, requirements, and algorithm

- The cause of a switchover from the primary RP to the secondary RP

Refer to the [Route Processor Redundant Operation](#) section of [Configuring the Route Processor](#) for more information.

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](#).

Operation of RP Redundancy

Redundant systems are defined as supporting two RPs. At any given time, one acts as the primary or active RP while the other acts as the secondary or standby RP. Thus, redundant RPs do not distribute the load or cell processing.

The RP redundancy feature provides high availability for a Cisco Catalyst 8500 by switching over to the secondary RP when one of the following conditions occur:

- Cisco IOS® Software failure
- Catastrophic RP hardware failure
- Software upgrade
- Maintenance procedure

The primary and secondary RPs communicate via interprocess communication (IPC) messages. Intelligent processors use IPC messages to exchange messages related to configuration commands and events that need to be reported. IPC messages are sent over the internal switch fabric via an inter-RP permanent virtual circuit (PVC). In addition, a relatively slow serial connection allows the secondary RP to initially request the primary RP to set up the inter-RP PVC. The slow serial connection is also used to communicate short messages, such as sync and switchcard status, that need to succeed even if the inter-RP PVC is down.

RP Redundancy Requirements

RP redundancy requires that the following criteria be met:

- Two RPs are installed.
- Both RPs have identical hardware configurations, including DRAM size, and the presence or absence of a network clock module, and so forth.
- Both RPs have the same functional image. Refer to [Maintaining Functional Images \(Catalyst 8540 MSR\)](#).
- Both RPs are running the same Cisco IOS image.
- Both RPs are set to autoboot (default).

If these requirements are met, the Cisco Catalyst 8540 runs in redundant mode by default. The tasks described in the following sections are optional and used only to change to nondefault values.



Caution: Cisco does not recommend changing the default behavior of the RP redundancy, as the default operation has been designed to be optimal in most circumstances.

What Causes a Switchover?

Switchover is defined as the secondary RP on the switch taking over as the primary RP. The following conditions can lead to an RP switchover:

- Both RPs are running Cisco IOS Software, and the primary RP crashes due to a software failure, or a catastrophic hardware failure. The table below contains detailed information on the type of software or catastrophic hardware failures that can cause an RP switchover.

Error	Description
Software watchdog timeout	<p>Occurs when a Cisco IOS process runs for more than 2 minutes without suspending and therefore completely utilizes CPU resources without yielding to any other processes. If Cisco IOS Software encounters such a condition, it reloads the system and reports a software watchdog timeout as the reload reason.</p> <p>Note: Cisco IOS Software reports a CPUHOG error message only when the condition persists for 2 seconds.</p>
Software forced crashes	<p>Occurs when Cisco IOS Software accesses a null or low memory address (0x0), or encounters corrupted data structures, due to memory corruption problems. Refer to Understanding Software-forced Crashes for further information.</p>

Bus error	<p>Occurs when Cisco IOS Software tries to access memory at a nonexistent address. This is typically caused by memory corruption, which occur when:</p> <ul style="list-style-type: none"> • A Cisco IOS process allocates a buffer of a certain size and writes data larger than the allocated buffer into memory. • A Cisco IOS process frees a memory buffer, but retains ownership of the buffer by holding the pointer to the buffer in memory. The Cisco IOS process eventually tries to write to the freed location using the pointer. It is possible that the same memory location has been allocated by another code and, when the original piece of code writes to this location, it corrupts the data being used by the current code that owns the buffer location. <p>Refer to Troubleshooting Bus Error Crashes for further information.</p>
Catastrophic hardware failure	<p>Defined as a hardware failure which prevents the normal execution of Cisco IOS code on the RP. Examples of serious hardware failures are listed below:</p> <ul style="list-style-type: none"> • Failure of the oscillator submodule, which causes failure of the system clock. • Failure of the address or data bus connected to the RP. • Failure of the CPU submodule. • Any hardware failure that leads to a software failure, such as a software-forced crash or exceptions.

- Executing the **reload** command on the primary RP effectively implements a controlled form of RP switchover.
- Manually forcing a switchover with the **redundancy force-failover main-cpu** command. This command can be executed on the primary RP only. When this command is executed, it implements a controlled switchover.

- Removing the primary RP. Although not recommended, hot swap or hot insertion of RPs and SPs can be done. Removing the primary RP during normal operation causes the secondary RP to take over. The **redundancy prepare-for-cpu-removal** command must be executed before the RP is hot swapped from the chassis. The reason for this is that the **redundancy prepare-for-CPU-removal** command does a last synchronization of running-config and dynamic soft-vc information, which minimizes the traffic interruption of a switchover. However, this command will cause the CPU to remain in ROM monitor (ROMMON) rather than autobooting. If the active CPU is removed without any prior command, a switchover will also occur, but some of the configuration may be lost and a much longer traffic interruption may occur.

What Does Not Cause a Switchover?

The following conditions will not lead to an RP switchover:

- Failure of a hardware component that does not prevent normal execution of code on the primary RP. Examples of such failures are listed below:
 - Failure of one of the two installed power supplies.
 - Inability for the RP to read the contents of flash memory.
 - Inability to read to or write from NVRAM, which stores the configuration file on the RP.
 - A failing line card that causes incessant interruptions to the primary RP, which then hangs. If the secondary RP were to take over, the primary RP would again hang due to the persistent interruptions. The only way to recover from this failure is to identify the faulty line card and remove it from the system.
- The break sequence is sent to the primary RP.
- The primary RP crashes while the secondary RP is booting, or if there are failures during booting. With this sequence of events, the secondary RP takes over as the primary RP with the startup configuration in its NVRAM.
- The primary RP cannot find a valid Cisco IOS image.
- A failure occurs when the primary RP is running the ROMMON software and has a yellow status LED.
- A failure occurs when the secondary RP is running the ROMMON software and has a yellow status LED.
- A failure occurs when both the primary and the secondary RPs are in ROMMON.
- A failure occurs before both RPs complete the initial configuration synchronization. With this sequence of events, the secondary RP takes over as primary RP with the startup configuration in its NVRAM.
- The system experiences a *double fault*, such as a crash during a switchover.

The Cisco Catalyst 8540 does not allow a secondary RP to become primary when the primary RP is in a valid state. The secondary RP can become primary only if the current primary RP changes to the *not-primary* state. If a primary RP hangs before changing its state, the secondary RP cannot take over on its own accord. This protection mechanism is designed to ensure that there is never more than one primary RP at a time in the system. If both RPs are in the primary state, the line cards do not know which RP is controlling the chassis. This condition can lead to damaged line card hardware since both RPs may attempt to read/write from line cards simultaneously, thus overdriving the current to the line cards.

RP Redundancy Algorithm

Arbitration Logic to Determine RP State

An RP in the Cisco Catalyst 8540 can be in one of the following four states:

- Primary
- Secondary
- Non-participant
- Not present

When the system is powered up, the two RPs follow an arbitration process to determine the initial states of the two RPs. The arbitration rules are listed below:

- The current state of the RPs is saved to NVRAM. If a system is power-cycled, the RPs come up with the same state that they had prior to the power-cycle. The exception is a new installation, when the RPs are powered up for the first time. On a new system powered up for the first time, the RP in slot 4 comes up as the primary, and the RP in slot 8 comes up as secondary.
- If two RPs with the same initial state stored in the NVRAM are inserted into a chassis, the arbitration logic makes the RP in slot 4 as primary and the RP in slot 8 as secondary. This case is used to account for situations where an RP is moved between multiple chassis.
- Once the initial RP state is determined by the arbitration logic, the state is retained until one of the following occurs:
 - Cisco IOS Software prompts a software-forced crash on the primary RP. In this event, the state of the RP is changed within the software to the non-participant state, indicating that this RP has encountered a software failure and is not participating in redundancy. If the RP is reset (automatically occurs if it is configured for autoboot), then it comes out of the non-participant state and moves to either the primary or secondary RP, depending on the state of the other RP.
 - The primary RP encounters a catastrophic hardware failure that prevents it from executing Cisco IOS Software. In the event of such a failure, the RP automatically moves to a non-participant state. A keepalive timer is maintained in the RP hardware to protect against such failures. Once Cisco IOS Software is booted on the RP, it periodically refreshes the keepalive timer. In the event of a hardware failure that prevents Cisco IOS Software from

executing, the keepalive timer is not refreshed and eventually expires. When the keepalive timer expires, the RP automatically goes to the non-participate state, thus allowing the other RP to take over via a switchover.

Detection of RP State Changes and Switchover Trigger Points

Each RP is aware of the other RP's state by monitoring each other's state via traces that run across the backplane of the chassis. If one RP changes its state, due to one of the reasons explained [earlier](#), a state-change interrupt is generated to the other RP. The interrupted handler in the other RP can then determine the new state of the other RP and take any switchover action as needed.

A switchover is triggered if the primary RP changes its state to *non-participant* due to a failure. The secondary RP receives a signal and initiates a switchover by changing its state to primary. The original primary RP should come up in the secondary state after rebooting, if autoboot is configured. The reason for the failure should be investigated.

A switchover is not triggered if a secondary RP changes its state to *non-participant*. The primary RP is aware that the secondary RP has encountered a failure. When the secondary RP boots Cisco IOS Software again, the primary RP resends synchronization information to the secondary RP. At this point, the secondary RP is a proper backup for the primary RP and can take over in the event of a failure on the primary RP.

The following section explains the information that is synchronized between RPs.

Configuration Synchronization

During bootup, the startup and running configurations of the primary RP are synchronized automatically to the secondary RP. As of Cisco IOS Software Release 12.1(7a)EY, the Cisco Catalyst 8540 supports the three configuration synchronization types listed in the table below:

Synchronization Type	Time of Execution
Startup configuration sync	When the write memory command is issued at the command-line interface (CLI).
Running configuration sync	When you exit from the configuration mode using the end command.
Dynamic sync	Designed to preserve switched virtual channels (VCs) and soft VCs, including ATM, circuit emulation service (CES), and Frame Relay. When the structure states of the concerned software module changes, a sync is done. For example, when a switched virtual circuit (SVC) is installed or released, the dynamic-sync mechanism starts. A bulk update is done as soon as the command is parsed, IPC is

up at that time, and the two RPs can communicate.

Related Information

- [Route Processor Redundant Operation](#)
 - [More ATM Information](#)
 - [Tools & Resources - Cisco Systems](#)
 - [Technical Support & Documentation - Cisco Systems](#)
-

Home	How to Buy	Login	Profile	Feedback	Site Map	Help
----------------------	----------------------------	-----------------------	-------------------------	--------------------------	--------------------------	----------------------

All contents are Copyright © 1992-2005 Cisco Systems, Inc. All rights reserved. [Important Notices](#) and [Privacy Statement](#).

Updated: May 05, 2005

Document ID: 19900
