

High Availability Overview

Cisco High Availability (HA) enables network-wide protection by providing fast recovery from faults that may occur in any part of the network. With Cisco High Availability, network hardware and software work together and enable rapid recovery from disruptions to ensure fault transparency to users and network applications.

The unique hardware and software architecture of the router is designed to maximize router uptime during any network event, and thereby provide maximum uptime and resilience within any network scenario.

This chapter covers the aspects of High Availability that are unique to the router. It is not intended as a comprehensive guide to High Availability, nor is it intended to provide information on High Availability features that are available on other Cisco routers that are configured and implemented identically on the router. The Cisco IOS feature documents and guides should be used in conjunction with this chapter to gather information about High Availability-related features that are available on multiple Cisco platforms and work identically on the router.

- Hardware Redundancy Overview, on page 1
- Stateful Switchover, on page 2
- Bidirectional Forwarding Detection, on page 3

Hardware Redundancy Overview

The router supports redundant Route Switch Processors (RSPs) and power supplies. Redundancy is not supported on interface modules.

Note Some interface modules require a reload during a software upgrade, briefly interrupting traffic.

Note Route Processor Redundancy (RPR) is *not* supported on the router. Stateful Switchover (SSO) is supported. See Stateful Switchover, on page 2.

Hardware redundancy provides the following benefits:

- A failover option—If a processor fails, the standby processor immediately becomes the active processor with little or no delay. The failover happens completely within the same router, so a second standby router is not needed.

- No downtime upgrades—Using features like ISSU, a software upgrade can be handled on the standby processor while the active processor continues normal operation.

Table 1: Hardware Redundancy Overview

Hardware	Support for Dual Hardware Configuration	Failover Behavior
Route Switch Processor	Yes	If an active RSP experiences an event that makes it unable to forward tra as a hardware failure, a software failure, an OIR, or a manual switch) and RSP is configured, the standby RSP immediately becomes the active RS **Note** The dual RSP reaches the STANDBY HOT state even if the system imag different on Active and Standby modules, as long as they are ISSU comp This is not applicable on the RSP3 module.
Interface module	No	No standby configurations are available for interface modules. If an interfac fails, it cannot forward traffic. In the event of an interface module shutdown, all other interface module fully operational.

Stateful Switchover

The Stateful Switchover (SSO) feature takes advantage of processor redundancy by establishing one of the processors as the active processor while the other RSP is designated as the standby processor, and then synchronizing critical state information between them. Following an initial synchronization between the two processors, SSO dynamically maintains RSP state information between the dual processors.

Stateful Switchover is particularly useful in conjunction with Nonstop Forwarding. SSO allows the dual processors to maintain state at all times, and Nonstop Forwarding lets a switchover happen seamlessly when a switchover occurs.

It is important to note that in most cases, SSO requires less downtime for switchover and upgrades than RPR. RPR should only be used when there is a compelling reason to not use SSO.

For additional information on NSF/SSO, see the Cisco Nonstop Forwarding document.

SSO-Aware Protocol and Applications

SSO-supported line protocols and applications must be SSO-aware. A feature or protocol is SSO-aware if it maintains, either partially or completely, undisturbed operation through an RSP switchover. State information for SSO-aware protocols and applications is synchronized from active to standby to achieve stateful switchover for those protocols and applications.

The dynamically created state of SSO-unaware protocols and applications is lost on switchover and must be reinitialized and restarted on switchover.

To see which protocols are SSO-aware on your router, use the following commands **show redundancy client** or **show redundancy history**.

Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning is easier, and reconvergence time is consistent and predictable.