

## Chapter 7

# Managing High Availability

This chapter describes how to manage Juniper Networks high availability (HA) software features for the E-series router. Use this chapter with *Chapter 6, Managing Modules* to fully manage the SRP and high availability features.

This chapter contains the following sections:

- Understanding High Availability on page 377
- Platform Considerations on page 378
- Redundancy Modes of Operation on page 379
- Understanding SRP State Behavior on page 381
- Application Support on page 384
- Before Activating High Availability on page 389
- Activating High Availability on page 389
- Deactivating High Availability on page 390
- Upgrading Software on page 391
- Monitoring High Availability on page 391

## Understanding High Availability

---

High availability is the idea of reducing or eliminating single points of failure. When applied to the E-series router, high availability provides both hardware-specific and software-specific methods to ensure minimal downtime and ultimately improve the performance of your network.

For hardware components, Juniper Networks provides redundancy solutions to ensure that the router continues to operate in the event of a hardware fault. This redundancy can exist on various router models in the form of multiple power supplies, cooling fans, switching planes, routing engines and, in some cases, interfaces. Redundancy also allows for hot-swapping various components within your Juniper Networks router.



**NOTE:** For information about E-series hardware redundancy features, see the *ERX Hardware Guide* or the *E120 and E320 Hardware Guide*.

## Platform Considerations

High availability is supported on all E-series routers except for the ERX-310 router.

For information about the modules supported on E-series routers:

- See the *ERX Module Guide* for modules supported on ERX-7xx models and ERX-14xx models.
- See the *E120 and E320 Module Guide* for modules supported on the E120 router and the E320 router.

## Module Requirements

The following table lists which SRPs support or do not support the high availability mode (stateful SRP switchover) feature.

| SRP Model    | Supported |
|--------------|-----------|
| SRP-5G       | No        |
| SRP-5G +     | Yes       |
| SRP-10G      | Yes       |
| SRP-40G      | No        |
| SRP-40G PLUS | Yes       |
| SRP-100      | Yes       |



**NOTE:** High availability requires two SRP modules with 1 GB of memory or more.

## Redundancy Modes of Operation

---

The switch route processor (SRP) modules can operate in one of two redundancy modes—file system synchronization and high availability.

### File System Synchronization Mode

File system synchronization is the default behavior mode for E-series routers that contain redundant SRPs. Available only to SRP modules, this mode has been available since the 2.x release. In this mode:

- Files and data (for example, configuration files and releases) in nonvolatile storage (NVS) remain synchronized between the primary and standby SRP modules.
- SRP modules will reload all line modules and restart from saved configuration files.
- If the active SRP module switches over to the standby SRP, the router cold-restarts as follows:
  - All line modules are reloaded.
  - User connections are lost, and forwarding through the chassis stops until the router SRP module recovers.
  - The standby SRP module boots from the last known good configuration from NVS.

For additional information about the default SRP functionality, see *Chapter 6, Managing Modules*.

### High Availability Mode

Currently applicable to the SRP module, the Juniper Networks high availability mode uses an initial bulk file transfer and subsequent, transaction-based mirroring to ensure rapid SRP module recovery after a switchover. This process is referred to in this chapter as *stateful SRP switchover*.

In addition to keeping the contents of NVS, high availability mode keeps state and dynamic configuration data from the SRP memory synchronized between the primary and standby SRP modules.

When stateful SRP switchover is enabled, an SRP switchover keeps line modules up and forwarding data, and the newly active SRP module continues from the point of switchover.

By using transaction-based mirroring instead of file synchronization, high availability mode keeps the standby SRP module synchronized with the active SRP module. Mirroring occurs from memory on the active SRP module to memory on the standby SRP module by way of transactions. When a transaction is committed on the active SRP module, the data associated with the transaction is sent to the standby SRP module.

In high availability mode:

- The contents of the NVS in the primary and standby SRP modules remain synchronized.
- If a switchover occurs:
  - The standby SRP module warm-restarts using the mirrored data to restore itself to the state of the system before the switchover.
  - During the warm restart:
    - User connections remain active, and forwarding continues through the chassis.
    - New user connection attempts during switchover are denied until switchover is complete.
    - New configuration changes are prevented until switchover is complete (or after 5 minutes).



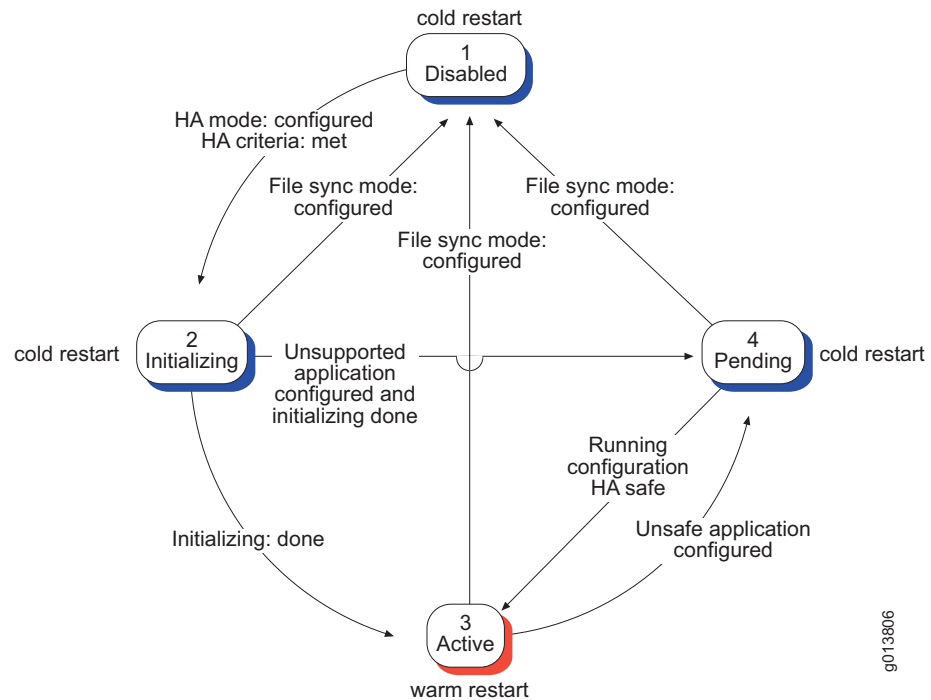
**NOTE:** If the switchover does not finish within 5 minutes, the SRP module cancels the operation and reenables CLI configuration.

---

## Understanding SRP State Behavior

The SRP progresses through various high availability states. These states are illustrated in Figure 27.

**Figure 27: High Availability States**



### Disabled State

The initial, default state for high availability mode is disabled. While in this state, the router continues to use file system synchronization. If a switchover occurs while the router is in this state, the standby SRP module performs a cold restart.

The router enters this state when you power up the router or when the router warm-restarts from an SRP switchover.

Once you enable high availability, the system must meet the following criteria before it can enter the initializing state:

- High availability mode is configured.
- Active SRP hardware supports high availability.
- Network core dump feature is disabled.
- Running configuration will allow high availability to operate (that is, no unsupported applications are configured).
- Standby SRP hardware supports high availability.

- Standby SRP module is online and capable of mirroring.
- Standby SRP module is running the same release.

During the disabled state:

- If any one criterion is not met, the system remains in the disabled state, until the criterion is met.
- If a switchover occurs while the system is in the disabled state, the system cold-restarts.

While in the disabled state, the system operates as if it were configured for file system synchronization (for example, NVS is synchronized every 5 minutes, if autosynchronization is enabled).

If all criteria are met, high availability mode transitions to the initialization state.

## Initializing State

After the SRP module transitions into the initializing state, bulk synchronization of the memory and NVS occurs. This includes the following:

- File synchronization of the primary NVS with the standby NVS
- Mirroring of appropriate state and dynamic configuration information from the active SRP (memory) to the standby SRP (memory)



**NOTE:** Depending on the size of the configuration, this process can take several minutes.

---

During the initializing state:

- If an unsupported application is configured during initialization, the system completes initializing and enters the pending state.
- If any other criterion becomes false (or is no longer met), the system enters the disabled state.
- If a switchover occurs while the system is in this state, the system cold-restarts.

Once initialization is completed, the system enters the active state.

## Active State

During the active state, the data that was synchronized from the active SRP module to the standby SRP module during initialization remains synchronized through mirroring updates.

Mirroring updates occur as follows:

1. When making changes or updates, applications create individual transactions, perform the updates on the active SRP module, and post the transactions.
2. Following the updates, the active SRP module sends the changes to the standby SRP module.
3. The standby SRP module replays the updates (in the order in which they were committed on the active SRP module) and makes the appropriate changes for each changed application.
4. Updates that need to be stored in NVS (that is, for static configurations) are updated in NVS.




---

**NOTE:** While in the active and pending states, the CLI **synchronize** command does not update configuration files; these files are updated by the mirroring process.

---

During the active state:

- If a switchover occurs while the router is in the active state, the standby SRP module performs a warm restart (that is, stateful SRP switchover is in effect); the standby SRP module uses the configuration located in NVS.
- If an unsupported application is configured, the system transitions to the pending state.
- If any other criterion changes (is no longer met), the system transitions to the disabled state.




---

**NOTE:** Changes made in manual commit mode are maintained, uncommitted, in the standby SRP memory until a trigger to commit occurs; if a switchover occurs while in this mode, the standby SRP module uses the configuration in memory.

---

## Pending State

The system transitions to the pending state if an unsupported application is configured. When a transition to the pending state occurs, the system generates SNMP traps and log messages.

How the router behaves depends on which HA state the application is in when it shifts to a pending state:

- From disabled state—The router remains in the disabled state.
- From initializing state—The router completes the initializing state and transitions to the pending state after initialization is complete.
- Active State—The router transitions to the pending state.

The system remains in the pending state until the configuration of the unsupported application is removed. However, even though it is in the pending state, the system continues mirroring updates from the primary SRP module to the standby SRP module.



**NOTE:** You can use the **show redundancy srp** command to display the name of any unsupported applications that are configured.

---

If a switchover occurs while the system is in the pending state, the system cold-restarts.

## Application Support

---

Applications are either supported or unsupported by high availability.

- Supported—You can configure supported applications without having any adverse impact to high availability. When a switchover occurs, supported applications can react to switchovers in one of two different ways:
  - Gracefully recover using mirrored static and dynamic information (for example, IP, PPP, and PPPoE)
  - Recover using static configuration only; that is, no runtime state is restored after a switchover. Dynamic configuration and state information are lost. (For example, CLI sessions are restarted, telnet sessions are dropped, multicast routes must be rebuilt, and so on.)
- Unsupported—We recommend that you not configure unsupported applications on a chassis running in high availability mode. Although configured unsupported applications suspend high availability or prevent high availability from becoming active, they do not cause any problems with the function of the router.

Table 45 indicates which applications support or do not support stateful SRP switchover.



**Table 45: Application Support for Stateful SRP Switchover**

| Application                                       | Supported | Unsupported | Notes   |
|---|-----------|-------------|---|
| <b>Physical Layer Protocols</b>                   |           |             |   |
| DS1   | a         | –           | –   |
| DS3   | a         | –           | –   |
| HDLC  | a         | –           | –   |
| SONET/SDH   | a         | –           | –   |
| SONET/SDH VT                                      | a         | –           | –   |
| <b>Link-Layer Protocols</b>                       |           |             |   |
| ATM   | a         | –           | Static and dynamic interfaces, with the exception of ATM subscribers, are supported.<br><br>In this case, <i>ATM subscribers</i> refers to a technology on the E-series router where the ATM layer does authentication (that is, not PPP or IP subscriber manager). |
| ATM 1483 bulk configuration of dynamic interfaces | a         | –           | –   |
| Bridged Ethernet                                  | a         | –           | –   |
| Cisco HDLC  | a         | –           | –   |
| Ethernet (with and without VLANs)                 | a         | –           | –   |
| Frame Relay                                       | a         | –           | –   |
| PPP   | a         | –           | –   |
| PPPoE   | a         | –           | –   |
| Transparent bridging                              | a         | –           | –   |
| <b>Unicast Routing</b>                            |           |             |   |
| Access Routes                                     | a         | –           | –   |
| BGP   | a         | –           | Supported only when the graceful restart extension is enabled.  |
| FTP   | a         | –           | Static recovery support only.   |
| IP  | a         | –           | –   |
| IPv6  | –         | a           | –   |
| IPSec Transport                                   | –         | a           | –   |
| IPSec Tunnels                                     | a         | –           | Completed IKE phase 1 and phase 2 negotiations supported only.  |
| IS-IS   | a         | –           | Supported only when the graceful restart extension is enabled.  |

**Table 45: Application Support for Stateful SRP Switchover (continued)**

| Application                   | Supported | Unsupported | Notes  |
|-------------------------------|-----------|-------------|--|
| OSPF                          | a         | –           | Supported only when the graceful restart extension is enabled.   |
| RIP                           | a         | –           | Static recovery support only.  |
| Static Routes                 | a         | –           | –  |
| Telnet                        | a         | –           | Static recovery support only.  |
| <b>IPv4 Multicast Routing</b> |           |             |  |
| Multicast Routing             | a         | –           | Static recovery support only. During switchover, the system mirrors the multicast queue so that IP can use the same queue without needing to recreate a different connection.  |
| DVMRP                         | a         | –           | Static recovery support only. DVMRP gives the restart complete indication to the IP routing table after getting a peer update (60-second time-out).  |
| IGMP                          | a         | –           | IC IGMP deletes its interface and membership state on SRP failover (controller down). As part of SRP warm start, IGMP interfaces are reconfigured from NVS and dynamic IGMP interfaces are reconfigured from mirrored storage. IGMP hosts are queried as IP interfaces come back up, the join state is re-established, and SC IGMP state is created. After the maximum query response time (across all interfaces) expires to allow hosts to re-establish join state, IGMP notifies MGMTM that graceful restart is complete.   |
| PIM                           | a         | –           | Static recovery support only. For warm start, PIM interfaces are reconfigured from NVS and a Hello message with a new Generation ID is issued as IP interfaces come up. A neighbor that receives this Hello determines that the upstream neighbor has lost state and needs to be refreshed. A VR-global configurable graceful restart timer is required for PIM to time out the re-establishment of the join state for sparse-mode interfaces. After this timer expires, PIM notifies MGMTM that graceful restart is complete. |
| <b>IPv6 Multicast Routing</b> |           |             | No multicast routing state information remains following a switchover. Incremental support for multicast routing is planned for future releases.   |
| Multicast Routing             | –         | a           | –  |
| MLD                           | –         | a           | –  |
| PIM                           | –         | a           | –  |

**Table 45: Application Support for Stateful SRP Switchover (continued)**

| Application  | Supported | Unsupported | Notes  |
|--|-----------|-------------|--|
| <b>Multiprotocol Label Switching</b>                       |           |             |  |
| MPLS   | a         | –           | <p>MPLS is HA-unsafe during a graceful restart. It is HA-unsafe until all the configured MPLS signaling protocols have completed their graceful restart procedures and any stale forwarding elements have been flushed from the line modules.</p> <p>If you force an SRP switchover while MPLS is HA-unsafe, the SRP module switches but the SRP module and the line modules undergo a cold restart.</p> <p>If the primary SRP module resets while MPLS is HA-unsafe, the router undergoes a cold restart.</p> |
| BGP signaling  | a         | –           |  |
| LDP signaling  | a         | –           | <p>To provide uninterrupted service during an SRP switchover in a scaled configuration, such as one with 32,000 Martini circuits, set the LDP graceful restart reconnect time to the maximum 300 seconds and set the LDP graceful restart recovery timer to the maximum 600 seconds. This requirement is true for all SRP switchovers, including those in the context of a unified in-service software upgrade.</p>  |
| RSVP signaling   | a         | –           |  |
| Local cross-connects between layer 2 interfaces using MPLS | a         | –           | –  |
| <b>Policies and QoS</b>                                    |           |             |  |
| Policies   | a         | –           | –  |
| QoS  | a         | –           | Static recovery support only.  |
| <b>Remote Access</b>                                       |           |             |  |
| AAA  | a         | –           | –  |
| DHCP External Server and Packet Trigger                    | a         | –           | <p>Following a switchover, the DHCP lease (that is, time remaining) is recalculated based on when the lease started. When the release timer for a client expires, the client is deleted and the access route is removed, along with the dynamic subscriber interface if it was created. If the client requests a new lease, DHCP external server resynchronizes with the new lease time.</p>   |
| DHCP Packet Capture  | a         | –           | –  |

**Table 45: Application Support for Stateful SRP Switchover (continued)**

| Application                   | Supported | Unsupported | Notes  |
|-------------------------------|-----------|-------------|--|
| DHCP Proxy Client             | –         | a           | –  |
| DHCP Relay Proxy              | –         | a           | –  |
| DHCP Relay Server             | a         | –           | <p>Before HA support, clients identified by the DHCP relay server were maintained on a switchover (their state was stored to NVS); DHCP relay server always had some level of HA support.</p> <p>Currently, following a switchover, the DHCP lease (that is, time remaining) is reset. When the release timer for a client expires, the client requests a new lease. The E-series router DHCP relay server then synchronizes with the new state.</p> |
| DHCPv4 Local Server           | a         | –           | –  |
| DHCPv6 Local Server           | –         | a           | HA does not support IPv6.  |
| L2TP                          | a         | –           | –  |
| L2TP Dialout                  | –         | a           | –  |
| Local Address Pools           | a         | –           | The internal local address server state supports only static recovery. However, the AAA application reallocates active addresses on a switchover. The resulting effect is the local address server having full HA support.   |
| RADIUS Client                 | a         | –           | Similar to local address server, AAA recovers disrupted RADIUS communication on a switchover. The resulting effect is the RADIUS client having full HA support.  |
| RADIUS Dynamic-Request Server | a         | –           | Static recovery support only.  |
| RADIUS Initiated Disconnect   | a         | –           | –  |
| RADIUS Relay Server           | a         | –           | –  |
| RADIUS Route-Download Server  | a         | –           | –  |
| SRC Client                    | a         | –           | –  |
| TACACS+                       | a         | –           | Static recovery support only.  |
| <b>Miscellaneous</b>          |           |             |  |
| Firewall                      | a         | –           | –  |
| J-Flow (IP flow statistics)   | a         | –           | –  |
| Line Module Redundancy        | a         | –           | –  |

**Table 45: Application Support for Stateful SRP Switchover (continued)**

| Application                 | Supported | Unsupported | Notes                         |
|-----------------------------|-----------|-------------|-------------------------------|
| Network Address Translation | a         | –           | –                             |
| NTP                         | a         | –           | –                             |
| Resource Threshold Monitor  | a         | –           | –                             |
| Response Time Reporter      | a         | –           | –                             |
| Route Policy                | a         | –           | Static recovery support only. |
| Subscriber Interfaces       | a         | –           | –                             |
| Tunnels (GRE and DVMRP)     | a         | –           | –                             |
| VRRP                        | a         | –           | Static recovery support only. |



**CAUTION:** When IP tunnels are configured on an HA-enabled router and the Service Module (SM) carrying these tunnels is reloaded, HA transitions to the pending state. HA remains in the pending state for 5 minutes after the successful reloading of the SM. This amount of time allows for IP tunnel relocation and for the tunnels to become operational again on the SM. If an SRP switchover occurs while HA is in the pending state, the router performs a cold restart.

## Before Activating High Availability

Before you activate high availability on the SRP modules, review *Chapter 6, Managing Modules* and any high availability–related changes to SRP management commands.

## Activating High Availability

You activate high availability (stateful SRP switchover) by launching Redundancy Configuration mode and issuing the **mode high-availability** command.

When activating high availability, keep the following in mind:

- In an E-series router that supports stateful SRP switchover, both SRP modules must be running the same software release version in order to activate high availability mode.
- If high availability mode cannot become active because of different releases on the active and standby SRP modules, the system reverts to its default mode (file system synchronization).

- When active or pending, the router configuration files are mirrored from the active SRP module to the standby SRP module. All other files shared between the active and standby SRP modules are automatically synchronized using legacy synchronization methods.

To enable high availability, enter the following:

```
host1(config)#redundancy
host1(config-redundancy)#mode high-availability
```

### **mode**

- Use the redundancy **mode** command to enable high availability.
- The **high-availability** keyword enables high availability mode for stateful SRP switchover. In this mode, the router uses mirroring to keep the configuration and state of the standby SRP module coordinated with the configuration and state of the active SRP module.



**NOTE:** High availability mode is currently available only on ERX-1440, ERX-1400, and ERX-700 routers that support dual SRPs.

- The **file-system-synchronization** keyword reverts the redundancy mode to its default. In this mode, the router uses file synchronization to keep the configuration of the standby SRP module coordinated with the configuration of the active SRP module.
- Example  

```
host1(config-redundancy)#mode high-availability
```
- Use the **no** version to return high availability mode operation to its default (file system synchronization).

### **redundancy**

- Use to enter Redundancy Configuration mode.
- Example  

```
host1(config)#redundancy
host1(config-redundancy)#
```
- There is no **no** version.

## **Deactivating High Availability**

To deactivate high availability support, enter the following:

```
host1(config)#redundancy
host1(config-redundancy)#mode file-system-synchronization
```

or

```
host1(config)#redundancy
host1(config-redundancy)#no mode
```

## Upgrading Software

---

As already mentioned, you cannot activate high availability when a different release of software is running on the standby SRP module. The router determines whether a release is the same by viewing the build date, the release file name, and the internal version number for the software on each SRP module.

The most efficient way to upgrade the software is to ensure that the standby SRP module is armed with the new release and then reload the standby SRP module. This reload occurs automatically after you download and arm a new release onto the active SRP module and the active SRP module subsequently synchronizes with the standby SRP module.

After reloading, and even though high availability mode is configured, the active SRP module reverts to using the file-system-synchronization operational mode for synchronizing updates. To complete the upgrade and place the system back in high-availability operational mode, you must execute the **srp switch** command to force the standby SRP module to take over as the active SRP module.



**NOTE:** Executing the **srp switch** command results in a cold restart of the router.

---

Once the switchover is initiated, the formerly active SRP module reloads the software and starts running the same release as the newly active SRP module. When the formerly active SRP module becomes operational as the standby SRP module, the newly active SRP module detects that the release it is running is the same as that on the standby SRP module and allows the originally active SRP module to resume the high-availability operational mode.

If a fault occurs when the active SRP module is in file-system-synchronization operational mode, the standby SRP module detects the fault and takes over, and the router cold-restarts. For this reason, it is important that you arm the new release only when you can accept the resulting window of vulnerability where high availability is disabled (that is, until the active and standby SRP modules are again running the same release).

## Monitoring High Availability

---

This section shows how to use the **show** commands to view your high availability configuration and how to clear the high availability switchover history for the router.

### High Availability show Commands

You can monitor various aspects of high availability using **show** commands. These aspects include redundancy modes and status, redundancy clients, historical information about redundancy on the router, and specific redundancy information for line modules and SRPs.

**show redundancy**

- Use to display the supported redundancy modes and other status relating to high availability. In particular, the output indicates any conditions that are preventing high availability from operating.
- Field descriptions
  - SRP
    - high-availability state—State of the high availability mode (disabled, active, or pending)
    - current redundancy mode—Redundancy mode currently being used by this router (high-availability or file-system-synchronization)
    - last activation type—Last type of activation that occurred on this router (that is, the method by which the SRP last booted [cold-start or warm-start])
  - Criteria Required for High Availability to be Active—criteria required for high availability to be active.



**NOTE:** All criteria must be “yes” for high availability to be active.

---

- Line Card
  - automatic reverting—State of automatic reverting (on or off)
  - slot(s)—Slots in which the line modules reside
  - hardware role—Function of the line module: primary or spare
  - lockout config—Status of redundancy on this line module (protected—line module redundancy is enabled; locked out—line module redundancy is disabled)
  - backed up by slot—Slot that contains the line module that is a spare for this primary line module
  - sparing for slot—Slot that contains the primary line module for which this line module is a spare
  - revert at—Time at which you want line module to revert
  - midplane type—Identifier for the type of midplane
  - midplane rev—Hardware revision number of the redundancy midplane
  - fabric slice redundancy—Status of the fabric slice on the SRP modules or SFMs on the E120 and E320 routers
  - slot—Slot in which the fabric slice resides
  - slice state—State of the fabric slice (online, not present)
  - type—Identifier for the type of hardware (SRP module or SFM)



- Example 1—Displays a summary of the redundancy status of an ERX-14xx model

```
host1#show redundancy
```

SRP  
---

```
high-availability state: disabled
current redundancy mode: high-availability
last activation type: cold-switch
```

### Criteria Preventing High Availability from being Active

| ----- criterion -----                           | met |
|---|-----|
| Standby SRP is online and capable of mirroring? | No  |

Line Card

automatic reverting is off

| slot | hardware<br>role | lockout<br>config | backed<br>up<br>by<br>slot | sparing<br>for<br>slot | revert<br>at |
|------|------------------|-------------------|----------------------------|------------------------|--------------|
| 3    | ---              | ---               | ---                        | ---                    | ---          |
| 8    | spare            | ---               | ---                        | ---                    | ---          |
| 12   | primary          | protected         | ---                        | ---                    | ---          |

|        |                  |                 |
|--------|------------------|-----------------|
| slots  | midplane<br>type | midplane<br>rev |
| -----  | -----            | -----           |
| 8 - 13 | 6                | 0               |

- Example 2—Displays the redundancy status of an ERX-14xx model in detail

```
host1#show redundancy detail
```

SRP

```
high-availability state: disabled
current redundancy mode: file-system-synchronization
last activation type: cold-start
```

### Criteria Required for High Availability to be Active

| ----- criterion -----                                | met |
|--|-----|
| Active SRP hardware supports High Availability?      | Yes |
| High Availability mode configured?                   | No  |
| Mirroring Subsystem present?                         | Yes |
| Mirroring activity levels within limits?             | Yes |
| Network Core Dumps disabled?                         | Yes |
| Running configuration is safe for High Availability? | Yes |
| Standby SRP hardware supports High Availability?     | Yes |
| Standby SRP is online and capable of mirroring?      | Yes |
| Standby SRP is running the same release?             | Yes |

```
Line Card
-----

automatic reverting is off

                                backed
                                up
                                by
slot  hardware  lockout  slot  sparing  revert
-----
3      ---      ---      ---      ---      ---
8      spare      ---      ---      ---      ---
12     primary  protected  ---      ---      ---

slots  midplane  midplane
-----  type      rev
8 - 13  6          0
```

**show redundancy clients**

- Use to display high availability clients and their various levels of high availability support.
- Specify an optional client type that you want to view (all, supported, unsafe, unsupported)



**NOTE:** Issuing this command without the optional client type results in showing only unsupported high availability clients (the default).

- Field descriptions
  - client—High availability client
  - mode—Whether the client is supported or unsupported for high availability
  - configuration—Safety level of the configuration based on whether or not the client is supported or unsupported and, in the case of those unsupported, whether or not the client has been configured (for example, if an unsupported client is configured on a router with high availability enabled, the configuration would read “unsafe”)

■ Example 1

host1#**show redundancy clients**

Unsupported High Availability Clients

| client                | configuration |
|-----------------------|---------------|
| DHCP Proxy Client     | safe          |
| Global Ipv6           | safe          |
| IPsec Transport (ITM) | safe          |
| l2tpDialoutGenerator  | safe          |
| DHCPv6 Local Server   | safe          |
| Radius Relay Server   | safe          |

- Example 2

```
host1#show redundancy clients all
```

```
High Availability Client Information
```

| client                | mode        | configuration |
|-----------------------|-------------|---------------|
| atm1483DataService    | supported   | safe          |
| AA83                  | supported   | safe          |
| aaaServer             | supported   | safe          |
| atmAal5               | supported   | safe          |
| AAQS                  | supported   | safe          |
| atm                   | supported   | safe          |
| Bridged Ethernet      | supported   | safe          |
| Transparent Bridging  | supported   | safe          |
| dcm                   | supported   | safe          |
| dhcpExternal          | supported   | safe          |
| DHCP Proxy Client     | unsupported | safe          |
| DS1                   | supported   | safe          |
| DS3                   | supported   | safe          |
| ethernet              | supported   | safe          |
| Flow Inspection       | supported   | safe          |
| frameRelay            | supported   | safe          |
| FT1                   | supported   | safe          |
| Global Ipv6           | unsupported | safe          |
| Global Ip             | supported   | safe          |
| HDLC                  | supported   | safe          |
| IKEP                  | supported   | safe          |
| ipFlowstats           | supported   | safe          |
| IpSubscriberManager   | supported   | safe          |
| IPTU                  | supported   | safe          |
| IPVR                  | supported   | safe          |
| IPsec Transport (ITM) | unsupported | safe          |
| l2tpDialoutGenerator  | unsupported | safe          |
| l2tp                  | supported   | safe          |
| LMGR                  | supported   | safe          |
| DHCPv4 Local Server   | supported   | safe          |
| DHCPv6 Local Server   | unsupported | safe          |
| MPLS                  | supported   | safe          |
| PMGR                  | supported   | safe          |
| pppoe                 | supported   | safe          |
| ppp                   | supported   | safe          |
| qos                   | supported   | safe          |
| Radius Relay Server   | unsupported | safe          |
| RSVP                  | supported   | safe          |
| SCM                   | supported   | safe          |
| slotHelper            | supported   | safe          |
| Cisco HDLC            | supported   | safe          |
| ServiceManager        | supported   | safe          |
| Sonet                 | supported   | safe          |
| SonetPath             | supported   | safe          |
| SonetVT               | supported   | safe          |
| IPsec Tunnel (ST)     | supported   | safe          |

### **show redundancy history**

- Use to display information about dates, times, and the number of occurrences for starts and switchovers.
- Use the **srp** keyword to view SRP module-specific information.
- Use the **detail** keyword to view additional history information.

- Field descriptions

- system up time—Amount of time elapsed since the last cold boot
- last cold start—Date and time the router experienced the last cold start
- last cold switchover—Date and time the router experienced the last cold switchover
- last warm switchover—Date and time the router experienced the last warm switchover
- cold starts—Total number of cold starts the router has experienced
- switchovers—Number of cold, warm, and consecutive warm switchovers the router has experienced

- Example 1

```
host1#show redundancy history
```

```
system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57
```

```
activation statistics:
  cold starts:        92
  switchovers:
    cold:             21
    warm:             147
    consecutive warm:  0
```

- Example 2

```
host1#show redundancy history detail
```

```
system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57
```

```
activation statistics:
  cold starts:        92
  switchovers:
    cold:             21
    warm:             147
    consecutive warm:  0
```

| SRP activation time | type       | slot | system uptime | running release   |
|---------------------|------------|------|---------------|-------------------|
| 2004-09-08 15:10:40 | cold-start | 00   | ---           | erx_6-0-0b1-8.rel |
| 2004-09-08 14:39:10 | cold-start | 00   | ---           | erx_6-0-0b1-1.rel |

**show redundancy line-card**

- Use to display line-module-specific redundancy information.
- Field descriptions
  - automatic reverting—State of automatic reverting (on or off)
  - slot(s)—Slots in which the line modules reside
  - hardware role—Function of the line module: primary or spare
  - lockout config—Status of redundancy on this line module (protected—Line module redundancy is enabled; locked out—Line module redundancy is disabled)
  - backed up by slot—Slot that contains the line module that is a spare for this primary line module
  - sparing for slot—Slot that contains the primary line module for which this line module is a spare
  - revert at—Time at which you want line module to revert
  - midplane type—Identifier for the type of midplane
  - midplane rev—Hardware revision number of the redundancy midplane
- Example

```
host1#show redundancy line-card
```

```
automatic reverting is off
```

| slot | hardware<br>role | lockout<br>config | backed<br>up<br>by<br>slot | sparing<br>for<br>slot | revert<br>at |
|------|------------------|-------------------|----------------------------|------------------------|--------------|
| 3    | ---              | ---               | ---                        | ---                    | ---          |
| 8    | spare            | ---               | ---                        | ---                    | ---          |
| 12   | primary          | protected         | ---                        | ---                    | ---          |

| slots  | midplane<br>type | midplane<br>rev |
|--------|------------------|-----------------|
| 8 - 13 | 6                | 0               |

**show redundancy srp**

- Use to display SRP-module-specific redundancy information.
- Use the brief keyword to display summary information about SRP redundancy.

- Field descriptions
  - SRP
    - high-availability state—State of the high availability mode (disabled, active, or pending)
    - current redundancy mode—Redundancy mode currently being used by this router (high-availability or file-system-synchronization)
    - last activation type—Last type of activation that occurred on this router (that is, the method by which the SRP last booted [cold start or warm start])
  - Criteria Required for High Availability to be Active—Criteria required for high availability to be active.



**NOTE:** All criteria must be “yes” for high availability to be active.

- Example 1

```
host1#show redundancy srp
```

```
high-availability state: active
current redundancy mode: high-availability
last activation type: warm-switch
```

- Example 2

```
host1#show redundancy srp detail
```

```
high-availability state: disabled
current redundancy mode: file-system-synchronization
last activation type: cold-start
```

| Criteria Required for High Availability to be Active |     |
|--|-----|
| -----<br>criterion                                   | met |
| Active SRP hardware supports High Availability?      | Yes |
| High Availability mode configured?                   | No  |
| Mirroring Subsystem present?                         | Yes |
| Mirroring activity levels within limits?             | Yes |
| Network Core Dumps disabled?                         | Yes |
| Running configuration is safe for High Availability? | Yes |
| Standby SRP hardware supports High Availability?     | Yes |
| Standby SRP is online and capable of mirroring?      | Yes |
| Standby SRP is running the same release?             | Yes |

**show redundancy switchover-history**

- Use to display the high availability switchover history for the chassis.
- Field descriptions
  - SRP activation time—Amount of time the SRP module has been active
  - type—Type of switchover
  - slot—Slot in which the SRP module resides
  - system uptime—Amount of time the chassis has been operational
  - running release—Release running on the SRP module at the time of the switchover
- Example

```
host1#show redundancy switchover-history
```

| SRP activation time | type        | slot | system uptime | running release       |
|---------------------|-------------|------|---------------|-----------------------|
| -----               | -----       | ---- | -----         | -----                 |
| 2004-07-26 10:44:25 | cold-start  | 07   | ---           | L-07-25-60b1mrg-e.rel |
| 2004-07-25 20:58:57 | warm-switch | 06   | 0 00:15:08    | L-07-25-60b1mrg-e.rel |
| 2004-07-25 20:53:41 | warm-switch | 07   | 0 00:09:51    | L-07-25-60b1mrg-e.rel |
| 2004-07-25 20:44:43 | cold-start  | 06   | ---           | L-07-25-60b1mrg-e.rel |
| 2004-07-25 19:32:01 | cold-start  | 06   | ---           | L-07-25-60b1mrg-d.rel |
| 2004-07-25 18:58:01 | warm-switch | 06   | 0 00:12:01    | L-07-25-60b1mrg-c.rel |
| 2004-07-25 18:51:56 | cold-switch | 07   | 0 00:05:56    | L-07-25-60b1mrg-c.rel |
| 2004-07-25 18:46:54 | cold-start  | 06   | ---           | L-07-25-60b1mrg-c.rel |
| 2004-07-25 17:44:48 | warm-switch | 06   | 0 00:14:32    | L-07-25-60b1mrg-b.rel |
| 2004-07-25 17:31:07 | cold-start  | 07   | ---           | L-07-25-60b1mrg-b.rel |
| 2004-07-25 16:05:08 | cold-start  | 07   | ---           | L-07-25-60b1mrg-a.rel |
| 2004-07-24 23:25:09 | warm-switch | 07   | 0 16:27:03    | L-07-24-60b1mrg-b.rel |
| 2004-07-24 23:18:23 | cold-switch | 06   | 0 16:20:17    | L-07-24-60b1mrg-b.rel |

**Clearing the Redundancy History**

You can use the **clear redundancy history** command to clear the high availability switchover history for the router.

**clear redundancy history**

- Use to clear the detailed high availability switchover history for the router.
- Example

```
host1#show redundancy history detail
```

```
system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57
```

```
activation statistics:
cold starts:         92
switchovers:
cold:                21
warm:                147
consecutive warm:    0
```

| SRP activation time | type        | slot | system uptime | running release       |
|---------------------|-------------|------|---------------|-----------------------|
| 2004-07-26 10:44:25 | cold-start  | 07   | ---           | L-07-25-60b1mrg-e.rel |
| 2004-07-25 20:58:57 | warm-switch | 06   | 0 00:15:08    | L-07-25-60b1mrg-e.rel |
| 2004-07-25 20:53:41 | warm-switch | 07   | 0 00:09:51    | L-07-25-60b1mrg-e.rel |
| 2004-07-25 20:44:43 | cold-start  | 06   | ---           | L-07-25-60b1mrg-e.rel |
| 2004-07-25 19:32:01 | cold-start  | 06   | ---           | L-07-25-60b1mrg-d.rel |
| 2004-07-25 18:58:01 | warm-switch | 06   | 0 00:12:01    | L-07-25-60b1mrg-c.rel |
| 2004-07-25 18:51:56 | cold-switch | 07   | 0 00:05:56    | L-07-25-60b1mrg-c.rel |

host1#clear redundancy history

host1#show redundancy history

```

system up time:      0 00:08:01
last cold start:     2004-07-26 10:44:25
last cold switchover: 2004-07-25 18:51:56
last warm switchover: 2004-07-25 20:58:57

```

```

activation statistics:
  cold starts:      92
  switchovers:
    cold:           21
    warm:           147
  consecutive warm: 0

```

| SRP activation time | type  | slot | system uptime | running release |
|---------------------|-------|------|---------------|-----------------|
| -----               | ----- | ---- | -----         | -----           |

- There is no **no** version.