CHAPTER 10

Redundancy and Fail-Over

Redundancy and Fail-Over

This chapter presents the fail-over and redundancy capabilities of the Cisco SCE8000 platform. It first defines relevant terminology, as well as pertinent theoretical aspects of the redundancy and fail-over solution. It then explains specific recovery procedures for both single and dual link topologies. It also explains specific update procedures to be used in a cascaded SCE platform deployments. When fail-over is required in a deployment, a topology with two cascaded Cisco SCE8000 platforms is used. This cascaded solution provides both network link fail-over, and fail-over of the functionality of the SCE platform, including updated subscriber state.

Terminology and Definitions

Following is a list of definitions of terms used in the chapter as they apply to the Cisco fail-over solution, which is based on cascaded SCE platforms.

- Fail-over — A situation in which the SCE platform experiences a problem that makes it impossible for it to provide its normal functionality, and a second SCE platform device immediately takes over for the failed SCE platform.

- Hot standby — When two SCE platforms are deployed in a fail-over topology, one SCE platform is active, while the second SCE platform is in standby, receiving from the active SCE platform all subscriber state updates and keep alive messages.
- Primary/Secondary — The terms Primary and Secondary refer to the default status of a particular SCE platform. The Primary SCE Platform is active by default, while the Secondary device is the default standby.

  Note These defaults apply only when both devices are started together. However, if the primary SCE platform fails and then recovers, it will not revert to active status, but remains in standby status, while the secondary device remains active.

- Subscriber state fail-over — A fail-over solution in which subscriber state is saved.

Redundant Topologies

The Cisco SCE 8000 includes SPA Interface Processor card with an internal electrical bypass module, which provides the capability of preserving the network link in case of failure. However, preserving the SCE platform functionality in case of a failure, requires a redundant SCE platform. Cisco provides a unique solution for this scenario, through deploying two cascaded SCE platforms.

The cascading is implemented by connecting the two SCE platforms using two of the data links. In each SCE platform, two of the four data interfaces are connected to each of the network links, while the other two data interfaces are used for cascading between the SCE platforms. (See the *Cisco SCE 8000 Installation and Configuration Guide* for specific cabling procedures for redundant topologies.) The cascade ports are used for transferring network traffic, keep-alive messages and subscriber state updates.

External Bypass

The Cisco SCE8000 platform can control an external bypass device, which bypasses the traffic during a power failure and also under specific control command from the SCE8000. The SCE8000 automatically activates the external bypass device during reload for the short period (less than 5 seconds) in which the SPA Interface Processor card does not forward traffic between traffic ports. In addition, the Cisco SCE8000 can be configured to activate the external bypass device in the following cases:

- After executing the **external-bypass** command , until the **no external-bypass** command is executed
- When the SCE8000 is in failure state.

Note In a cascaded configuration, an external bypass device should be connected only for the traffic ports. The cascade ports should be directly connected between the two Cisco SCE8000 platforms (see Figure 10-1).

In-line Dual Link Redundant Topology

This topology serves inline deployments where the SCE platform functionality should be preserved in case of a failure, in addition to preserving the network link.

Figure 10-1 In-line Dual Link Redundant Topology

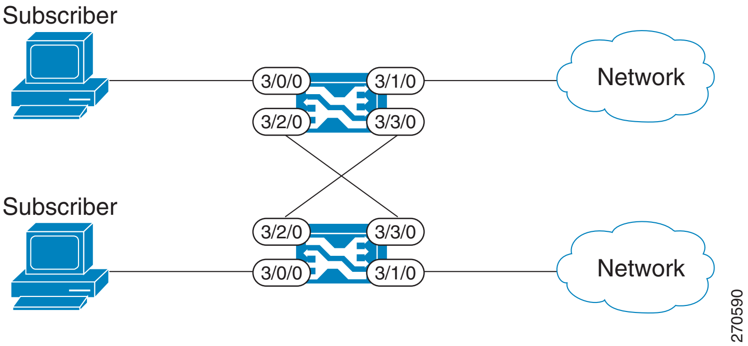

Failure Detection

The Cisco SCE 8000 platform has several types of mechanisms for detecting failures:

- Internal failure detection — The SCE platform monitors for hardware and software conditions such as overheating and fatal software errors.
- Inter-device failure detection — The SCE platform sends periodic keep-alive messages via the cascade ports
- SCE platform-Subscriber Manager (SM) communication failure detection — A failure to communicate with the SM may be regarded as a cause for fail-over. However, this communication failure is not necessarily a problem in the SCE platform. If the connection to the SM of the active SCE platform has failed, while the connection to the SM of the standby SCE platform is alive, a fail-over process will be initiated to allow the SCE platform proper exchange of information between the SCE platforms and the SM.
- Link failure — The system monitors all three types of links for failures:
  - Traffic port link failure — Traffic cannot flow through the SCE platform.
  - Cascade port link failure — Traffic cannot flow between the SCE platforms through the cascade ports.
  - Management port link failure — This is not a failure that interrupts traffic on the link in and of itself. However, when SM is used, management port link failure will cause an SM connection failure and this, in turn, will be declared as a failure of the SCE platform.

  This type of failure, in most cases, does not require reboot of the SCE platform. When the connection with the SM is re-established the SCE platform is again ready for hot standby. If both SCE platforms lose their connections with the SM, it is assumed that it is the SM which has failed, thus, no action will be taken in the SCE platform.

Link Failure Reflection

The SCE platforms are transparent at Layers 2 and 3. The Cisco SCE8000 platform operates in promiscuous mode, and the network elements on both sides of the SCE platform, are using the MAC address of the other network element when forwarding traffic.

To assist the network elements on both sides of the SCE platform to identify the link failures as quickly as possible, the SCE platform supports a functionality of reflecting to the other side of the SCE platforms events of link failure. When the link on one side of the SCE platform fails, the corresponding link on the other side is forced down, to reflect the failure. Link failure reflection is done on the traffic ports. When operating in deployments of single SCE platform with two data links, link failure is reflected between the two ports of each link.

When working with two cascaded SCE platforms, link failure is reflected in two cases:

- Reflection between the traffic ports of each SCE platform.
- If there is a failure in the cascade port link, the two SCE platforms can no longer support proper processing of the two links, since the traffic flowing on the standby SCE platform's link must be forwarded to the active SCE platform for processing. In this case the link failure is reflected from the cascade ports to the traffic ports of the standby SCE platform, in order to force the network to switch all the traffic only through the link of the active SCE platform.

Hot Standby and Fail-over

The fail-over solution requires two SCE platforms connected in a cascade manner.

- Hot Standby, page 10-4
- Fail-over, page 10-5
- Hardware Crash Mode, page 10-6
- Failure in the Cascade Connection, page 10-8
- Installing a Cascaded System, page 10-8

Hot Standby

In fail-over solution, one of the SCE platforms is used as the active SCE platform and the other is used as the standby. Although traffic enters both the active and the standby SCE platforms, all traffic processing takes place in the SCE platform which is currently the active one. The active SCE platform processes the traffic coming on both links, its own link and the link connected to the standby SCE platform, as follows

- All traffic entering the active SCE platform through its traffic ports is processed in that SCE platform and then forwarded to the line.
- All traffic entering the standby SCE platform through its traffic ports is forwarded through the cascade ports to the active SCE platform where it is processed, and then returned to the standby SCE platform through the cascade ports to be forwarded to the original line from which it came.

Since only one SCE platform processes all traffic at any given time, split flows, which are caused by asymmetrical routing, that exist in the two data links are handled correctly.

To support subscriber-state fail-over, both SCE platforms hold subscriber states for all parties, and subscriber state updates are exchanged between the active SCE platform and the standby. This way, if the active SCE platform fails, the standby SCE platform is able to start serving the line immediately with a minimum loss of subscriber-state.

The two SCE platforms also use the cascade channel for exchanging periodic keep-alive messages.

Fail-over

In fail-over solution, the two SCE platforms exchange keep alive messages via the cascade ports. This keep alive mechanism enables fast detection of failures between the SCE platforms and fast fail-over to the standby SCE platform when required.

If the active SCE platform fails, the standby SCE platform then assumes the role of the active SCE platform.

The failed SCE platform uses its electrical bypass mechanism, which is a hardware entity that is separate from the main board and processors, to forward traffic to the other SCE platform, and to forward processed traffic back to the link. The previously standby SCE platform now processes all the traffic of this other link that is forwarded to it by the previously active SCE platform in addition to the traffic of its own link.

When the failed SCE platform recovers, it will remain in standby, while the previously standby SCE platform remains active. Switching the SCE platforms back to their original roles may be performed manually, if required, after the failed SCE platform has either recovered or been replaced.

If the failure is in the standby SCE platform, it will continue to forward traffic to the active SCE platform and back to its link, while the active SCE platform continues to provide its normal processing functionality to the traffic of the two links.

Note For information regarding the synchronization of subscriber information between cascaded SCE platforms and the effect of fail-over on the subscriber databases, see Anonymous Groups and Subscriber Templates, page 9-5.

There are three user-configurable options that are relevant in a situation when an SCE platform fails:

- **Bypass** — Maintain the link in bypass mode (continue sending traffic to the other SCE platform, and then continue forwarding the processed traffic back to the link). The incoming traffic in the failed SCE platform is forwarded to the working SCE platform, where it is processed and then sent back to the original SCE platform and back to the link. This is the default configuration.
  - Effect on the network link — negligible.
  - Effect on the SCE platform functionality — The effect on the SCE platform functionality is dependent on the failed SCE platform.
  - If the failure is in the standby SCE platform — the active SCE platform continues providing its normal functionality, processing the traffic of the two links.
  - If the failure is in the active SCE platform — the standby SCE platform takes over processing the traffic, and becomes the active SCE platform.

- **Cutoff** — Change the link of the failed SCE platform to cutoff (layer 1) forcing the network to switch all traffic through the line of the working SCE platform. This will, of course, decrease the network capacity by 50%, but may be useful in some cases.

  This option is available for use in special cases, and requires special configuration.

  - Effect on the network link — The network loses 50% of its capacity (until the failed SCE platform has recovered).
  - Effect on the SCE platform functionality — The effect on the SCE platform functionality is dependent on the failed SCE platform.
  - If the failure is in the standby SCE platform — the active SCE platform continues providing its normal functionality, processing the traffic of the two links.
  - If the failure is in the active SCE platform — the standby SCE platform takes over processing the traffic, and becomes the active SCE platform.
- **External-bypass** — Activate the external bypass device connected to the failed SCE platform, passing all traffic through the line without being serviced by the failed SCE platform. Although this may cause the traffic passing through the other link (that of the non-failed SCE platform) to get service in split-flow conditions if asymmetric routing is present, it may be useful in some cases.

  This option is available for use in special cases, and requires special configuration.

  - Effect on the network link — negligible
  - Effect on the SCE platform functionality — Since the active SCE platform only sees the traffic of a single link, split-flow effects might occur. The link connected to the failed SCE platform gets no service.

Hardware Crash Mode

There are three hardware components that operate together to produce the desired behavior upon failure of the Cisco SCE 8000 platform:

- external optical bypass: The external optical bypass module should always be connected in a cascade setup. The optical bypass may be either activated or deactivated by the hardware during a failure.

  The external optical bypasses protect against a second Cisco SCE 8000 platform failure. In the case of a second failure, the bypass module that is connected to the last Cisco SCE 8000 to fail will be enabled. This preserves one of the network links, assuming the *on-failure* configuration is is **bypass** . If the *on-failure* configuration is **external-bypass** , the external optical bypass is activated even during a single failure by the failed SCE platform.
- internal electrical bypass: The Cisco SCE 8000 contains internal electrical bypasses connecting SPA modules 0 and 2 and modules 1 and 3. These bypasses transmit the traffic to and from the cascade connections between the platforms.
- SPA modules: Under some conditions (such as if the *on-failure* configuration is **cutoff**), the SPA modules are disabled when failure occurs.

The collective behavior of these three components is known as the hardware crash mode and is dependent on the configuration of the *on-failure* parameter of the **connection-mode** command, as well as whether the platform is the active or standby platform.

In the standby platform, hardware crash mode behavior is as follows for:

For **on-failure bypass**:

- The external optical bypass is deactivated (traffic is sent to the platform).
- The electrical bypass is enabled (cascade ports transmit traffic to the active platform for processing)
- The SPA modules are enabled (all ports and links are functioning)

This means that whether the standby platform is operational or has failed, it transmits traffic to the active platform for processing via the electrial bypasses.

For **on-failure cutoff**:

- The external optical bypass is deactivated (traffic is sent to the platform).
- The electrical bypass is disabled.
- The SPA modules are powered off.

This means that if the standby platform fails, the link connected to it SCE is severed.

For **on-failure external-bypass**:

- The external optical bypass is activated (traffic is bypassed).
- The electrical bypass is enabled between ports 0 and 1. (This is a special configuration, done just in case the optical bypass device does not function, which has a very low probability.)
- The SPA modules are enabled (all ports and links are functioning)

This means that when the standby SCE platfirm is failed, the external optical bypass is used to ensure link continuity at the expense of not servicing the traffic on that link.

In the active platform, the hardware crash mode behavior is exactly the same as on a standby platform. The active platform assumes that when it fails, the standby platform will take over and process the traffic.

If the standby platform has failed, failure of the active platform means that the entire system has failed (this state is also called a 'second failure'). The hardware crash mode behavior in the active platform in this case depends on the configuration of the *on-failure* parameter, which determines whether traffic is bypassed via the external bypass, maintaining traffic flow through one link, although with no processing, or whether traffic is completely cut off.

If the standby platform has failed and the *on-failure* configuration is **bypass** or **external-bypass**, hardware crash mode behavior in the active platform is as follows:

- The external optical bypass is activated (traffic is bypassed via the external bypass).
- The electrical bypass is enabled between ports 0 and 1. (This is a special configuration, done just in case the optical bypass device does not function, which has a very low probability)
- The SPA modules are enabled (all ports and links are functioning)

Since neither platform is operational at this point, there is no processing taking place and traffic is simply bypassed via the external optical bypass.

If the standby platform has failed and the *on-failure* configuration is **cutoff**, hardware crash mode behavior in the active platform is as follows:

- The external optical bypass is deactivated (traffic is sent to the platform).
- The electrical bypass is disabled (cascade ports do not transmit traffic to the standby platform since it is also not operational)
- The SPA modules are disabled (ports and links are not functioning)

Since neither platform is operational at this point, there is no processing taking place. In addition, although traffic reaches the platform, since the internal bypasses as well as all the ports are disabled, this results in a 'dead end' , cutting off all traffic on both links.

Failure in the Cascade Connection

The effect of a failure in the cascade connection between the two SCE platforms depends on whether one or both connections fail:

- Only one cascade connection is down — In this case, both SCE platforms can still communicate, so each still knows the status of the peer.

  As long as one cascade connection remains up, the standby will cut off its traffic links so that all traffic is routed via the active SCE platform. Therefore, split flow is avoided, but at the expense of half line capacity.
- Both cascade links are down — In this case, neither SCE platform knows anything about the status of the peer. Each platform then works in standalone mode, which means that each SCE platform processes on its own traffic only. This results in split flows.

Installing a Cascaded System

This section outlines the installation procedures for a redundant solution with two cascaded SCE platforms.

Refer to the *Cisco SCE8000 Installation and Configuration Guide* for information on topologies and connections.

Refer to the *Cisco SCE8000 CLI Command Reference* for details of the CLI commands.

Note When working with two SCE platforms with split-flow and redundancy, it is extremely important to follow this installation procedure.

Step 1 Install both SCE platforms, power them up, and perform the initial system configuration.

Step 2 Make sure that the external optical bypass modules are connected correctly and are operational. Use the **show external-bypass** command.

Note Each Cisco SCE8000 platform must be connected to an external optical bypass module.

Step 3 Connect both SCE platforms to the management station.

Step 4 Connect the cascade ports.

The cascade ports may be connected either directly in Layer 1 (dark fibers) or through a switch. When connecting the cascade ports through a switch, it is important to assign each cascade link to a different VLAN, otherwise the traffic will be forwarded incorrectly (between different links) by the switch.

Step 5 Set topology configurations for each SCE platform via the connection-mode options. (See Topology-Related Parameters for Redundant Topologies, page 10-11)

Step 6 Make sure that the SCE platforms have synchronized and active SCE platform was selected. Use the **show interface linecard 0 connection-mode** command.

Step 7 If you want to start in bypass mode, change the link mode to bypass in both SCE platforms. The bypass mode will be applied only to the active SCE platform. (See About the Link Mode, page 7-4.)

Step 8 Verify the link mode configuration. (See Monitoring the System, page 10-12.) Use the **show interface linecard 0 link mode** command.

Step 9 Connect the traffic port of SCE platform #1. This will cause a momentary down time until the network elements from both sides of the SCE platform auto-negotiate with it and start working (when working inline).

Step 10 Connect the traffic port of SCE platform #2, this will cause a momentary down time until the network elements from both sides of the SCE platform auto-negotiate with it and start working (when working inline).

Step 11 When full control is needed, change the link mode on both SCE platforms on both links to ‘forwarding’. It is recommended to first configure the active SCE platform and then the standby. (See About the Link Mode, page 7-4.)

Step 12 You can now start working with the Subscriber Manager.

Recovery

- Replacing the SCE platform (manual recovery), page 10-9
- Reboot only (fully automatic recovery), page 10-10

This section specifies the procedure for recovery after a failure. The purpose of the recovery procedure is to restore the system to fully functional status. After the recovery procedure, the behavior of the system is the same as after installation.

A failed SCE platform may either recover automatically or be replaced (manual recovery). Whether recovery is automatic or manual depends on the original cause of the failure:

- Power failure — manual or automatic recovery can be implemented.
- Any failure resulting in a reboot — manual or automatic recovery can be implemented (this is configurable).
- 3-consecutive reboots within half an hour — manual recovery only
- Cascade ports link-failure — automatic recovery when link revives.
- Traffic link failure — automatic recovery when link revives.
- Failure in the communications with the SM — automatic by SM decisions after connection is re-established.
- Hardware malfunction — manual recovery, after replacing the malfunctioning SCE platform.

Replacing the SCE platform (manual recovery)

This is done in two stages, first manual installation steps performed by the technician, and then automatic configuration steps performed by the system.

- Manual steps, page 10-10
- Automatic steps (in parallel with the manual steps, requires no user intervention), page 10-10

Manual steps

Step 1 Disconnect the failed SCE platform from the network

Step 2 Connect a new SCE platform to the management link and the cascade links (leave network ports disconnected.)

Step 3 Configure the SCE platform.

Step 4 Basic network configurations done manually (first time).

Step 5 Load application software (*Service Control Application for Broadband*) to the SCE platform. Provide application configuration.

Step 6 Connect the traffic ports to the network links.

Automatic steps (in parallel with the manual steps, requires no user intervention)

Step 1 Establishment of inter-SCE platform communication.

Step 2 Synchronization with the SM.

Step 3 Copying updated subscriber states from the active SCE platform to the standby.

Reboot only (fully automatic recovery)

Step 1 Reboot of the SCE platform.

Step 2 Basic network configurations.

Step 3 Establishment of inter-SCE platform communication.

Step 4 Selection of the active SCE platform.

Step 5 Synchronization of the recovered SCE platform with the SM.

Step 6 Copying updated subscriber states from the active SCE platform to the standby.

CLI Commands for Cascaded Systems

This section presents CLI commands relevant to the configuration and monitoring of a redundant system.

Use the following commands to configure and monitor a redundant system:

- connection-mode
- [no] force failure-condition
- show interface linecard 0 connection-mode
- show interface linecard 0 physically-connected links

Topology-Related Parameters for Redundant Topologies

All four of the topology-related parameters are required when configuring a redundant topology.

- **Connection mode** — Redundancy is achieved by cascading two SCE platforms. Therefore the connection mode for both SCE platforms may be either:
  - Inline-cascade
  - Receive-only-cascade
- **Physically-connected-links** — For each of the cascaded SCE platforms, this parameter defines the number of the link (Link 0 or Link 1) connected to this SCE platform.
- **Priority** — For each of the cascaded SCE platforms, this parameter defines whether it is the primary or secondary device.
- **On-failure** — For each of the cascaded SCE platforms, this parameter determines whether the system cuts the traffic or bypasses it when the SCE platform either has failed or is booting.

Configuring the Connection Mode

Use the following command to configure the connection mode, including the following parameters.

- inline/receive only
- physically connected links
- behavior upon failure of the SCE platform
- primary/secondary

To configure the connection mode, use the following command.

Step 1 From the SCE (config if)# prompt, type **connection-mode (inline-cascade|receive-only-cascade) physically-connected-links {link-0|link-1}priority {primary|secondary} on-failure {bypass | external-bypass|cutoff}** and press **Enter**.

Examples

EXAMPLE 1

Use the following command to configure the primary SCE platform in a two-SCE platform inline topology. Link 1 is connected to this SCE platform and the behavior of the SCE platform if a failure occurs is *bypass* , which is the default.

```
SCE(config-if)#connection-mode inline-cascade physically-connected-links link-1 priority
primary
```

EXAMPLE 2

Use the following command to configure the SCE platform that might be cascaded with the SCE platform in Example 1. This SCE platform would have to be the secondary SCE platform, and Link 0 would be connected to this SCE platform, since Link 1 was connected to the primary. The connection mode would be the same as the first, and the behavior of the SCE platform if a failure occurs is external-bypass.

```
SCE(config-if)# connection-mode inline-cascade physically-connected-links link-0 priority
secondary on-failure external-bypass
```

Monitoring the System

Use the following commands to view the current connection mode and link mode parameters.

How to View the Current Connection Mode

Step 1 From the SCE# prompt, type `show interface linecard 0 connection-mode` and press **Enter**.

How to View the Current Link Mode

Step 1 From the SCE# prompt, type `show interface linecard 0 link mode` and press **Enter**.

How to View Current Link Mappings

Step 1 From the SCE 2000# prompt, type `show interface linecard 0 physically-connected-links` and press **Enter**.

Configuring Forced Failure

Use the following commands to force a virtual failure condition, and to exit from the failure condition when performing an application upgrade.

How to Force a Virtual Failure Condition

Step 1 From the SCE(config if)# prompt, type `force failure-condition` and press **Enter**.

Forces the SCE platform into a virtual failure state.

How to Exit a Virtual Failure Condition

Step 1 From the SCE(config if)# prompt, type `no force failure-condition` and press **Enter**.

Exits from the virtual failure state.

System Upgrades

- Firmware Upgrade (package installation), page 10-13
- Application Upgrade, page 10-14
- Simultaneous Upgrade of Firmware and Application, page 10-14

In a redundant solution, it is important that firmware and/or application upgrades be performed in such a way that line and service are preserved.

Refer to the following sections for instructions on how to perform these procedures on two cascaded SCE platforms:

- Upgrade the firmware only
- Upgrade the application only
- Upgrade both the firmware and the application at the same time

Note When upgrading only one component (either firmware only or application only), always verify that the upgraded component is compatible with the component that was not upgraded.

Firmware Upgrade (package installation)

Step 1 Install package on both SCE platforms (open the package and copy configuration).

Step 2 Reload the standby SCE platform.

Step 3 Wait until the standby finishes synchronizing and is ready to work.

Step 4 Make sure that the connection mode configurations are correct.

Step 5 Reload the active SCE platform.

Step 6 After the former active SCE platform reboots and is ready to work manually, it may be left as standby or we can manually switch the SCE platforms to their original state.

Application Upgrade

Step 1 Unload the application in the standby SCE platform.

Step 2 Load new application to the standby SCE platform.

Step 3 Make sure that the connection mode configurations are correct.

Step 4 Wait until the standby SCE platform finishes synchronizing and is ready to work.

Step 5 Force failure condition in the active SCE platform.

Step 6 Upgrade the application in the former active SCE platform.

Step 7 Remove the force failure condition in that platform.

Step 8 After the former active SCE platform recovers and is ready to work, it may remain the standby or can be manually switched back to active.

Simultaneous Upgrade of Firmware and Application

Step 1 In the standby SCE platform:

a. Uninstall the application.

b. Upgrade the firmware (this includes a reboot).

c. Install the new application.

Step 2 Force-failure in the active SCE platform.

This makes the updated SCE platform the active one, and it begins to give the NEW service.

Step 3 Repeat step 1 for the (now) standby SCE platform.

Since this includes a reboot, it is not necessary to undo the force failure command.