



## CHAPTER 3

# Information About Topology

---

Revised: June 13, 2011, OL-24143-01

## Introduction

This chapter describes the possible deployment topologies of the SCE 1000. The Cisco SCE solution offers several basic topology options that permit the user to tailor the SCE Platform to fit the needs of a particular installation. An understanding of the various issues and options is crucial to designing, deploying, and configuring the topology that best meets the requirements of the individual system.

The Cisco SCE solution offers several basic topology options that permit the user to tailor the SCE Platform to fit the needs of a particular installation. An understanding of the various issues and options is crucial to designing, deploying, and configuring the topology that best meets the requirements of the individual system.

- [Information About Topology Considerations, page 3-2](#)
- [Functionality, page 3-4](#)
- [Information About Physical Installation, page 3-5](#)
- [Information About Redundancy, page 3-7](#)
- [Information About Failure and Recovery, page 3-9](#)
- [Information About Topology-Related Parameters, page 3-10](#)

# Information About Topology Considerations

There are several issues that must be considered to arrive at the optimum configuration of the topology-related parameters:

- **Functionality**—Will the system be used solely to monitor traffic flow, with report functionality only, or will it be used for traffic flow control, with enforcement as well as report functionality?
- **Physical installation configuration**—Will the SCE Platform be installed as inline? Or will the SCE Platform use an optical splitter?
- **Redundancy**—Must the system be designed to guarantee uninterrupted service? If so, there must be a backup SCE Platform to assume operation in case of failure of the primary data link.
- **Link failure and recovery**—How should the SCE Platform respond to platform failure and subsequent recovery? Should traffic flow continue even though the unit is not operating, or be halted until the platform is repaired/replaced? Should the unit actually resume operation when it is again operational?

## SCE Platform Configuration

There are four topology-related parameters:

- **Connection mode**—Can be Inline or Receive-only, depending on the physical installation of the SCE 1000:  
May be configured via either the setup command or the connection-mode command.
- **Bypass mode when the SCE 1000 is not operational (on-failure)**—This parameter determines whether the system cuts the traffic or bypasses it when the SCE 1000 has failed.  
May be configured via either the setup command or the connection-mode command.
- **Status after reboot caused by fatal error or abnormal shutdown**—This parameter determines whether the SCE 1000 returns to normal operational state after a failure.  
May be configured via either the setup command or the failure-recovery operation-mode command.
- **Link failure reflection**—This parameter determines the behavior of the system when there is a link problem. In some topologies it is required that link failure on one port be reflected to the other port, to allow the higher layer redundancy protocol in the network to function correctly.  
May be configured via the link failure-reflection command only.

## Failure Detection Mechanism

**Boot time diagnostics failure.** When there is a failure in diagnostics testing at boot time the system will remain in failure status.

- **Boot time diagnostics failure.** When there is a failure in diagnostics testing at boot time the system will remain in failure status.

- Watchdog mechanism. There are two types of watchdogs:
  - HW watchdog. A hardware mechanism that detects control entity failure.
  - SW watchdog. A software mechanism that periodically checks for software failures in the SCE 1000. If a failure is detected, an error message is sent and the SCE 1000 reboots.
- Run time hardware tests. The system periodically tests the hardware components for error. If a hardware component is malfunctioning, it will be discovered by the system within seconds.

## Bypass Mechanism

The SCE 1000 includes a Network Interface Card with a bypass mechanism that is enabled upon SCE 1000 failure. In addition, when connected in-line it can also be enabled in normal operation to simultaneously bypass traffic flow to the other side and direct it internally for analysis. In this case it maintains “receive-only”-like monitoring functions, when control functionality is not required.

The bypass card supports the following four modes:

- Bypass — The bypass mechanism preserves the network link, but traffic is not processed for monitoring or for control.
- Forwarding — This is the normal operational mode, in which the SCE 1000 processes the traffic for monitoring and control purposes.
- Sniffing — The bypass mechanism preserves the network link, while in parallel allowing the SCE 1000 to process the traffic for monitoring only.
- Cutoff — There is no forwarding of traffic, and the physical link is forced down (cutoff functionality at layer 1).

# Functionality

The SCE 1000 can serve one of two general functions:

- **Monitoring and Control**—The SCE 1000 monitors and controls traffic flow. Decisions are enforced by the SCE 1000 depending on the results of the monitoring functions of the SCE 1000 and the configuration of the Service Control Application for Broadband or Mobile solution.

To perform control functions, the SCE 1000 must be physically installed as an inline installation and the connection mode must be “inline”.

- **Monitoring only**—The SCE 1000 monitors traffic flow, but cannot control it.

Either an inline installation or an optical splitter installation may be used for monitoring only. In the latter case connection mode must be “receive-only”.

# Information About Physical Installation

There are two options for the physical installation configuration of the SCE 1000 Platform:

- Inline installation (provides control functionality).
- Out-of-line installation utilizing an external optical splitter

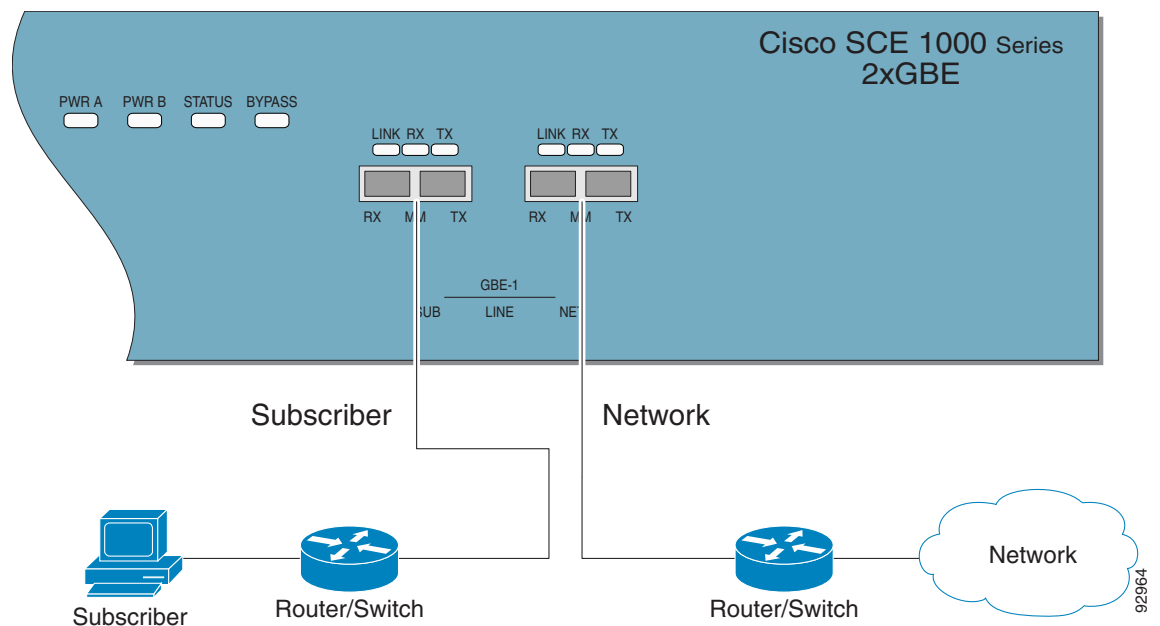
The physical installation determines the connection mode that should be configured.

- [Bump-in-the-Wire \(Inline\) Topology, page 3-5](#)
- [External Splitting \(Receive-only\) Topology, page 3-6](#)

## Bump-in-the-Wire (Inline) Topology

Typically, the SCE 1000 is connected on a full duplex line between two devices (Router, BRAS, and so forth). When the SCE 1000 is installed as a bump-in-the-wire, it physically resides on the data link between the subscriber side and the network side, and can both receive and transmit traffic ([Figure 3-1](#)).

**Figure 3-1** *Bump-in-the-Wire (Inline) Topology*

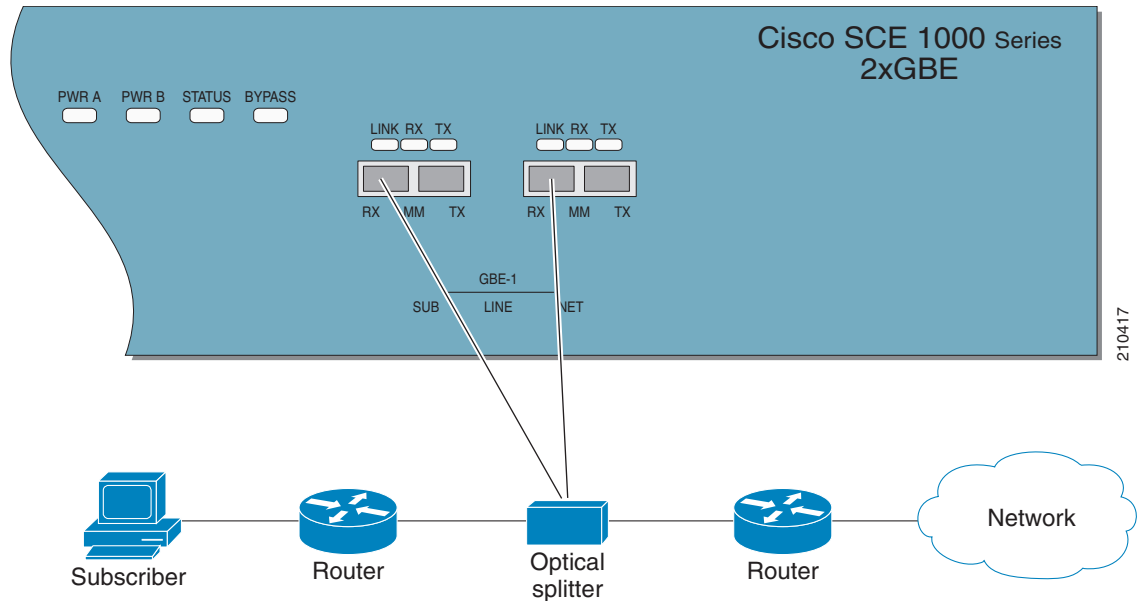


A bump-in-the-wire installation is referred to as inline connection mode.

## External Splitting (Receive-only) Topology

In external splitting, an external optical splitter resides physically on the GBE link between the subscriber side and the network side (Figure 3-2). In this topology, the traffic passes through the external splitter, which splits traffic to the SCE 1000. The external splitter is connected to the SCE 1000 via Rx links only. The SCE 1000, therefore, receives traffic only. It does not transmit.

**Figure 3-2 External Splitting (Receive-only) Topology**



An external splitting installation is referred to as receive-only connection mode.

Note that in an external splitting installation, the SCE 1000 has only traffic monitoring capabilities.



### Note

Receive-only topologies can also be implemented using a switch. Such a switch must support SPAN functionality that includes separation between ingress and egress traffic and multiple SPAN-ports destinations.

# Information About Redundancy

- [Redundancy, page 3-7](#)
- [Two Platforms on Parallel Links in Bump-in-the-Wire Topology, page 3-7](#)

## Redundancy

When a high degree of reliability is desired, a second SCE 1000 Platform should be installed to provide backup operation capabilities. This redundant SCE 1000 guarantees uninterrupted functioning of all SCE 1000 operations.

**Note**

---

Redundancy is possible in inline installations only.

---

A single SCE 1000 platform does not provide redundancy for SCE 1000 control functions. In case of failure of the SCE 1000 unit, the SCE 1000 simply bypasses the traffic; the traffic link is not cut, but no control or monitoring functionality is available.

## Two Platforms on Parallel Links in Bump-in-the-Wire Topology

Using two SCE 1000 platforms on parallel links provides redundancy for all SCE 1000 features. In case of failure in the active platform, the backup SCE 1000 unit takes over.

Using redundant SCE 1000 platforms is applicable as an overlay to a customer's redundant topology, on condition that the entire traffic of a specific subscriber (end station, subnet or VLAN) is flowing through one link only. Both links may be active, providing that the subscriber traffic is mutually exclusive.

This redundancy solution addresses any failure in the SCE 1000 platform itself. It is based on the idea that any fatal hardware or software failure will cause the platform to “cut” the link. A “cut” link will cause the routers/switches on both ends to switch the traffic to the standby link. On the standby link, the traffic is analyzed and policies enforced by the standby SCE 1000, which, after the failure, acts as the active SCE 1000.

Note that when both links are simultaneously independently active and redundant for the other link (as is the case when HSRP with two virtual routers is used), if one link fails, its traffic is directed to the other link. However, the overall supported load in the link that is now carrying all the traffic is only equal to one link, not two.

During setup of this topology, the configuration of the two SCE 1000 platforms is done through multi-box configuration. This ensures that both hold the same configurations and policies. The functional operation of switching from the active to standby, SCE 1000 is contingent upon the fact that the two SCE 1000 platforms are in the same Domain. All configurations performed on this Domain are automatically updated on both SCE Platforms. Both boxes should also be assigned to the same Subscriber Domain. For more information on Domains, see the [Cisco Service Control Management Suite Subscriber Manager User Guide](#).

The common protocols used for redundancy traffic switching between network elements such as routers and switches in networks are Spanning-Tree in layer2, HSRP in layer3 (usually used in data-centers), and other common routing protocols like OSPF or RIP.

**Note**

---

When using routing/switching protocols that perform load balancing as well, the load balancing capabilities should be disabled.

---

The transition to the backup SCE 1000 platform is transparent. Once the routers/switches detected that traffic has been cut, they start sending traffic through the redundant link. After this occurs, the failed SCE 1000 can be fixed/replaced with no downtime, since the box is effectively disconnected from the network. After fixing/replacing the failed SCE 1000, you must copy the configuration of the current active SCE 1000 to the fixed/replaced SCE 1000.

The backup and restore procedures used for copying policies and Service Configurations from one SCE 1000 to the next are detailed in the [Cisco Service Control Application for Broadband User Guide](#).



# Information About Failure and Recovery

It is important to decide how the system should behave in case of the failure of the SCE 1000, both during the time that the unit is down and after recovery. This decision is influenced by several factors:

- Physical installation (connection mode)
- Redundancy
- Relative importance of maintaining connectivity vs. the continuity of the value-added services that the SCE 1000 enables.

## Physical Installation

In a link connection via an external optical splitter, SCE 1000 failure does not affect traffic flow, which continues through the external optical splitter. When the SCE 1000 detects a failure that requires a recover by reboot, it immediately switches to Cutoff mode, stopping all traffic flow over the link until the SCE 1000 unit is restored to operation.

When operation resumes, the defined operational bypass mode is automatically resumed.

The configuration of a bump-in-the-wire installation depends on the remaining two factors.

## Redundancy

Redundancy requires two platforms on parallel links, one active and one standby, in inline topology. When the active SCE 1000 platform detects a failure situation, it will immediately switch to Cutoff mode, causing the routers/switches on both ends to switch the traffic to the standby link and thus activate the standby SCE 1000 platform.

There are two options when the failed SCE 1000 platform is finished reloading:

- It may either actually resume operation in the defined operational bypass mode, returning to its status as the active SCE 1000 platform.
- It may remain inactive in the failure bypass mode.

## Maintaining the Network Links Versus Maintaining SCE 1000 Platform Functionality

When a single SCE 1000 is deployed, the user may decide that in case of a failure, maintaining the network link is more important than providing the SCE 1000 functionality. In this scenario, when the SCE 1000 detects a failure that requires a reboot process for recovering, it immediately switches to Bypass mode, allowing all traffic to bypass the SCE 1000. The SCE 1000 stays in Bypass mode maintaining the network link, albeit without SCE 1000 processing, until the SCE 1000 fully recovers from the failure and is ready to resume normal functioning.

Alternatively, the user may decide that the SCE 1000 functionality is sufficiently crucial to require severing the link if the SCE 1000 platform fails. In this case, when the SCE 1000 detects a failure that requires a reboot process for recovering, it immediately switches to Cutoff mode, stopping all traffic flow. The SCE 1000 stays in Cutoff mode, halting all traffic, until it fully recovers from the failure and is ready to resume normal functioning. In Cutoff the physical interface is blocked, enabling the network device connected to the SCE 1000 to sense that the link is down.

# Information About Topology-Related Parameters

Refer to the following sections to determine the correct values for all topology-related parameters before beginning run the initial setup of the SCE 1000.

- [Connection Mode Parameter, page 3-10](#)
- [On-Failure Mode Parameter, page 3-10](#)
- [Link Failure Reflection Parameter, page 3-11](#)
- [Status of the SCE 1000 After Abnormal Boot, page 3-11](#)

## Connection Mode Parameter

The connection mode parameter refers directly to the physical topology in which the SCE 1000 is installed. Installation is possible in either of the two following modes:

- **Inline** — The SCE 1000 resides on the data link between the subscriber side and the network side, thus both receiving and transmitting packets.
- **Receive-only** — The SCE 1000 does not reside physically on the data link. Data is forwarded to the SCE 1000 via an external optical splitter. The SCE 1000 itself receives only and does not transmit.



### Note

---

Default value = Inline

---

The connection mode parameter is determined by the physical deployment of the SCE 1000 as follows:

- Bump-in-the-wire installation = Inline connection mode.
- External optical splitter installation = Receive-only connection mode.

## On-Failure Mode Parameter

As described in the section *The Bypass Mechanism*, the bypass card supports four different modes. The following two modes are possible when the SCE 1000 is not operational due to platform failure or boot:

- **Bypass** — The optical splitter forwards traffic with no intervention of the control application running in the SCE 1000 platform, but monitoring functions continue uninterrupted.
- **Cutoff** — There is no forwarding of traffic. The link is forced down, resulting in traffic cutoff at Layer1.

The Forwarding mode enables control of traffic flow and is not compatible with the non-operational status.

In a single SCE 1000 topology, the value of this parameter is determined by whether or not the link can be completely cut when the SCE 1000 fails, or whether traffic flow should continue across the link in spite of platform failure.

- Cutoff mode is required for the following:
  - Redundant inline topology.
  - Non-redundant inline topology if value-added services are crucial and are more important than maintaining connectivity.
- Bypass mode is required for the following:
  - Non-redundant inline topology if connectivity is crucial.

## Link Failure Reflection Parameter

The link failure reflection refers to the behavior of the SCE 1000 when one of the data links fails. Some network redundant topologies require a layer 1 cutoff in order for the network element to recognize the link failure and translate it into action (switch to redundant link). In this case, if one of the ports fails, it must be reflected to the other port as well.

- Link failure-reflection — When one data port link fails, the SCE 1000 forces the other port link down as well. The port will be forced down as long as the first port link is down. When the problematic port link goes up, the other port link will also be turned on again.
- No link failure-reflection — Link failure is not reflected to the other port.

**Note**

---

Default value = no link failure-reflection

---

## Status of the SCE 1000 After Abnormal Boot

This parameter determines whether the SCE 1000 returns to normal operational state after a reboot caused by fatal error or abnormal shutdown. In general, it is desirable that the SCE 1000 resume operation, and as promptly as possible. However, in a redundant topology, a recovered SCE 1000 may remain non-operational. In this case the platform that had been the backup and is currently active will remain active.

The two options for this parameter are:

- Operational—The status of the SCE 1000 after abnormal boot is operational. The platform automatically resumes functioning in the defined operational link bypass mode.
- Not Operational—The status of the SCE 1000 after abnormal boot is not operational. The platform remains in the defined failure link bypass mode.

This option is to be used only in a redundant topology where a second, operational platform exists.

**Note**

---

Default value = Operational for all non-redundant systems.

---

**Note**

---

Must be explicitly configured for redundant topologies.

---

**Table 3-1**      **Topology Configuration Summary**

<b>Description</b>	<b>Connection mode</b>	<b>On-failure link bypass mode</b>	<b>Admin status after abnormal boot</b>
Link connection via external switch with port-mirroring	Receive-only	Bypass <sup>1</sup>	Operational
Bump-in-the-wire			
Bump-in-the-wire, monitor and control, not redundant	Inline	Bypass	Operational
Bump-in-the-wire, monitor only, not redundant	Inline	Bypass	Operational
Bump-in-the-wire, monitor and control, redundant	Inline	Cutoff	Operational <sup>2</sup>
Bump-in-the-wire, monitor only, redundant	Inline	Cutoff	Operational

1. Italicized values represent automatically applied defaults that are applied based on previously defined parameters. These values can be changed only via specific CLI commands.
2. In a redundant topology, it is also possible to configure admin status after abnormal boot to be Not operational. In this case, though, the SCE 1000 would have to be manually reloaded to resume full functionality.