



Switch Fabric

This chapter describes the Cisco CRS Series Enhanced 16-Slot Line Card Chassis (LCC) switch fabric. It includes the following sections:

- [Switch Fabric Overview, page 4-1](#)
- [Switch Fabric Operation, page 4-2](#)
- [Switch Fabric Card Description, page 4-5](#)

Switch Fabric Overview

The switch fabric is the core of the Cisco CRS routing system. The Cisco CRS routing system fabric is implemented through multiple redundant switch fabric cards (SFCs) installed in the chassis. The switch fabric uses a cell-switched, buffered, three-stage Benes switch fabric architecture. The switch fabric receives user data from a modular services card (MSC) or Forwarding Processing card (FP) and performs the switching necessary to route the data to the appropriate egress MSC or FP.

The switch fabric is divided into eight planes (plane 0 to plane 7) that are used to evenly distribute traffic across the switch fabric. Each switch fabric plane is independent and not synchronized with one another. Each cell traverses the switch fabric using a single switch fabric plane (cells are not bit-sliced across the switch fabric).

When operating as a single-shelf (standalone) system, the LCC uses one of the following SFCs:

- CRS-16-FC/S (40G)
- CRS-16-FC140/S (140G)
- CRS-16-FC400/S (400G)

Each SFC implements all three stages of the switch fabric. Note the following:

- The CRS-16-FC140/S fabric is able to operate in both 40G mode and 140G mode to allow interconnection between 20G, 40G, or 140G MSCs and FPs.
- The CRS-16-FC400/S fabric is able to operate in 40G, 140G and 400G mode to allow interconnection between the 20G, 40G, 140G, or 400G MSCs and FPs.




Note

The LCC supports either 40G fabric cards (FC/S cards), 140G fabric cards (FC-140/S cards), or 400G fabric cards (FC-400/S cards). A LCC with a mix of 40G, 140G, and 400G fabric cards is not a supported mode of operation. Such a mode is temporarily allowed only during the upgrade process.

When operating as part of a multishelf system for the Cisco CRS-1, Cisco CRS-3, or CRS-X, the LCC accepts either the CRS-16-FC/M (40G), CRS-16-FC140/M (140G), or CRS-16-FC400/M (400G) SFC. In a multishelf system, the SFCs installed in the LCCs perform the S1 stage and S3 stage functions, while the S2 stage of the switch fabric is provided by S2 switch fabric cards in the fabric card chassis (FCC). Like the CRS-16-FC140/S fabric, the CRS-16-FC140/M fabric operates in 40G and 140G mode. With the CRS-16-FC140/M S2 fabric in the FCC, the LCCs in a multishelf system can be a mixture of 40G fabric LCCs and 140G fabric LCCs; in this case, only those LCCs requiring 140G support would need to be upgraded. Also, like the CRS-16-FC400/S fabric, the CRS-16-FC400/M fabric operates in 40G, 140G, and 400G mode. With the CRS-16-FC400/M S2 fabric in the FCC, the LCCs in a multishelf system can be a mixture of 40G fabric LCCs, 140G fabric LCCs, and 400G LCCs; in this case, only those LCCs requiring 400G support would need to be upgraded. See the “[Switch Fabric Operation](#)” section on page 4-2 for details about the stages of the switch fabric.

Figure 4-1 shows the basic path of IP data packets through the switch fabric. Note that the figure shows a single-shelf system, in which all three stages of the switch fabric are provided by SFCs in the LCC. In a multishelf system, Stage 2 of the switch fabric is provided by S2 fabric cards in the FCC.

Figure 4-1 Basic Cisco CRS Series Carrier Routing System Switch Fabric



Ingress data packets are received at a physical interface on a PLIM and transferred to the associated MSC, where the packets are segmented into cells for efficient switching by the switch fabric hardware. Each MSC has multiple connections to each switch fabric plane, which it uses to distribute cells to each fabric plane. On egress, cells are reassembled into data packets before being transmitted by the egress MSC.



Note

The cell structure used in the Cisco CRS switch fabric is a Cisco-specific cell structure and is not related to Asynchronous Transfer Mode (ATM) cells.

Switch Fabric Operation

Several switch element components on each SFC perform the functions to implement each of the three stages (S1, S2, and S3) of the switch fabric. Each stage performs a different function:

- Stage 1 (S1)—Distributes traffic to Stage 2 of the fabric plane. Stage 1 elements receive cells from the ingress MSC and PLIM (or RP) and distribute the cells to Stage 2 (S2) of the fabric plane. Cells are distributed to S2 elements in round-robin fashion; one cell goes to the first S2 element, the next cell goes to the next S2 element, the next cell goes to the third S2 element, and so on, and then back to the first S2 in sequence.

- Stage 2 (S2)—Performs switching, provides 2 times (2x) speedup of cells, and performs the first stage of the multicast function. Stage 2 elements receive cells from Stage 1 and route them toward the appropriate:
 - Egress MSC and PLIM (single-shelf system)
 - Egress line card chassis (multishelf system)
- Stage 3 (S3)—Performs switching, provides 2 times (2x) speedup of cells, and performs a second level of the multicast function. Stage 3 elements receive cells from Stage 2 and perform the switching necessary to route each cell to the appropriate egress MSC and PLIM.

Speedup Function



Note

The speed up function only applies to the Cisco CRS-1 and Cisco CRS-3. It is not applicable for the Cisco CRS-X.

An LCC can contain up to 16 MSCs, each with up to 140 Gbps of bandwidth. To provide 140 Gbps of switching capacity for each MSC, the switch fabric must actually provide additional bandwidth to accommodate cell overhead, buffering, and congestion-avoidance mechanisms.

Congestion can occur in the switch fabric if multiple input data cells are being switched to the same destination egress MSC. Typically, little congestion exists between the S1 and S2 stages because there is little or no contention for individual links between the switch components. However, as multiple cells are switched from the S2 and S3 stages to the same egress MSC, cells might contend for the same output link.

To reduce the possibility of data cells being delayed during periods of congestion, the switch fabric uses 2 times (2x) speedup to reduce contention for S2 and S3 output links. The switch fabric achieves 2x speedup by providing two output links for every input link at the S2 and S3 stages.

S2 and S3 Buffering

Buffering is also used at the S2 and S3 stages of the switch fabric to alleviate any additional congestion that the switch fabric speedup does not accommodate. To ensure that this buffering does not cause cells to arrive out of sequence, the MSC resequences the cells before reassembling them into packets. To limit the amount of buffering required, a back-pressure mechanism is used for flow control (which slows the transmission of data cells to a congested destination). Back-pressure messages are carried in fabric cell headers.

Failure Operation

The routing system can withstand the loss of a single plane of the switch fabric with no impact on the system. The loss of multiple planes results in linear and graceful degradation of performance, but does not cause the routing system to fail.

Note the following:

- For the Cisco CRS-1 and Cisco CRS-3 routing systems, at least two planes of the switch fabric (an even plane and an odd plane) must be active at all times. Otherwise, the switch fabric fails, causing a system failure.
- For the Cisco CRS-X routing system, at least three planes of the switch fabric (an even plane, an odd plane, and one plane in either slot 6 or 7) must be active at all times for the router to operate. Otherwise, the switch fabric fails, causing a system failure.

Switch Fabric Upgrade

You can upgrade the switch fabric to a higher capacity SFC. To avoid traffic loss, you must upgrade the switch fabric one plane at a time, which allows the system to continue operating with seven fabric planes. During the upgrade, some fabric planes may run in one configuration while others run in a different configuration.

To upgrade a fabric plane, first shut down the fabric plane and then remove the fabric card that implements that plane. Then, replace the original fabric card with the new fabric card and restore service to that fabric plane before upgrading the next plane.

For information about how to perform online insertion and removal (OIR) of SFCs in a single-chassis system, see [Cisco CRS Carrier Routing System Getting Started Guide](#). For information about how to upgrade from a single-shelf system to a multishelf system, see [Cisco CRS-1 Carrier Routing System Single-Shelf to Multishelf Upgrade Guide](#).


Switch Fabric Card Description

The single-chassis supports the CRS-16-FC/S, CRS-16-FC140/S, or CRS-16-FC400/S switch fabric card.

Switch Fabric Card (Single-Chassis System)

In a single chassis, the SFC implements all three stages of the three-stage Benes switch fabric. Each card also implements one plane of the eight-plane switch fabric. [Figure 4-2](#) shows the CRS-16-FC/S switch fabric card. The CRS-16-FC140/S is similar, but the CRS-16-FC400/S is different.

Figure 4-2 CRS-16-FC/S Switch Fabric Card



Switch Fabric Card Components

The SFC contains the following major components:

- S1 switch element—Implements Stage 1 of the switch fabric. Receives cells from the MSC or RP and distributes them to Stage 2. Each S1 switch element is connected to every S2 switch element.
- S2 switch element—Implements Stage 2 of the switch fabric. Receives cells from Stage 1, performs 2x speedup, and routes the cells toward the appropriate egress S3 element. Each S2 switch element is connected to every S3 switch element.

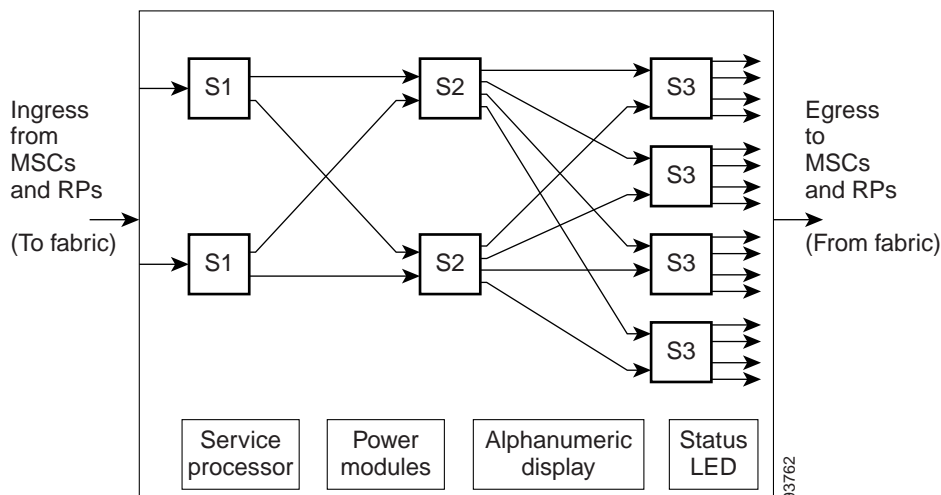


Note The speed up function only applies to the Cisco CRS-1 and Cisco CRS-3. It is not applicable for the Cisco CRS-X.

- S3 switch element—Implements Stage 3 of the switch fabric. Receives data cells from Stage 2 and performs switching and 2x speedup. The S2 and S3 switch elements are pure output-buffered switch elements with a central memory to buffer cells and queuing capabilities to distinguish between high-priority and low-priority traffic.
- Service processor—Controls the operation of the fabric card and provides the interface to the system control plane. The service processor performs card power up and power down, performs link-up and link-down processing, configures switch element components, updates the fabric group ID (FGID) for multicast traffic, and maintains cell configuration.
- Power modules—Take -48 VDC input power from the midplane and convert it to the voltages required by the components on the switch fabric card.

Figure 4-3 shows the major components of the switch fabric card.

Figure 4-3 Block Diagram of the Switch Fabric Card




Note Each stage of the three-stage Benes switch fabric is implemented with the same switch element components. However, during system startup the components are programmed by Cisco IOS XR software to operate in S1, S2, or S3 mode, depending on their functions in the switch fabric. Each switch fabric card contains two S1, two S2, and four S3 components.

Switch Fabric Card Physical Characteristics

Figure 4-4 shows the front panel of the CRS-16-FC/S SFC. The front panel of the CRS-16-FC140/S and CRS-16-FC400/S are similar.

Figure 4-4 CRS-16-FC/S Switch Fabric Card Front Panel



1	Status LED	4	PID/VDN
2	Two Alphanumeric LEDs	5	CLEI
3	SN	6	Ejector

The SFC front panel contains:

- Status LED—Indicates the status of the fabric card.
- Alphanumeric display—Displays switch fabric card messages.
- SN—Serial number
- PID/VDN—Product identification and vector directory number
- CLEI—Common Language Equipment Identifier code label
- Ejector levers

