Switch Fabric

This chapter describes the Cisco CRS 8-Slot Line Card Chassis Enhanced router switch fabric.

Switch Fabric

This chapter describes the Cisco CRS 8-Slot Line Card Chassis Enhanced router switch fabric. It includes the following topics:

Switch Fabric Overview

The switch fabric is the core of the Cisco CRS routing system. The switch fabric is implemented through switch fabric cards 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) and performs the switching necessary to route the data to the appropriate egress MSC.

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.) And unlike the Cisco CRS 16-Slot Line Card Chassis Enhanced router the Cisco CRS 8-Slot Line Card Chassis Enhanced router has only 4 physical switch fabric cards

There are four types of switch fabric cards used in the Cisco CRS 8-Slot Line Card Chassis Enhanced router: CRS-8-FC/S, CRS-8-FC-140/S, CRS-8-FC-400/S and the CRS-8-FC-400/M fabric card. Each fabric card implements all three stages of the switch fabric.

The following figure shows the basic path of IP data packets through the Cisco CRS routing system switch fabric. Note that the figure shows a single-shelf system, in which all three stages of the switch fabric are provided by switch fabric cards in the Cisco CRS 8-Slot Line Card Chassis Enhanced router. In a multishelf system, Stage 2 of the switch fabric is provided by S2 fabric cards in the fabric card chassis.

Figure 1: Basic Cisco CRS Carrier Routing System Switch Fabric

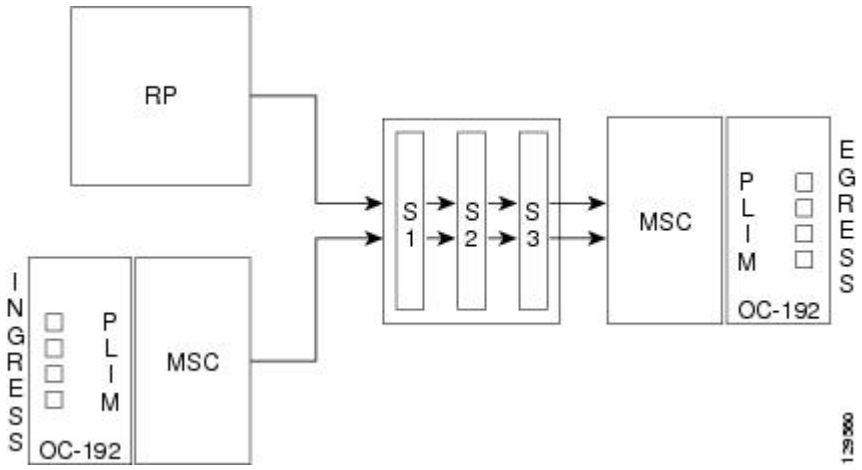

Ingress data packets are received at a physical interface on a PLIM and transferred to the associated MSC, where the packets are segmented into cells for efficient switching by the switch fabric hardware. Each MSC has multiple connections to each switch fabric plane, which it uses to distribute cells 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 switch fabric card 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 Cisco CRS 8-Slot Line Card Chassis Enhanced router (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.

Speed-up 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.

A line card chassis can contain up to 8 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 Card Description

Depending on the card, the fabric card can be used in a CRS single-chassis system or a CRS back-to-back system.

Figure 2: CRS-8-FC400/S Switch Fabric Card

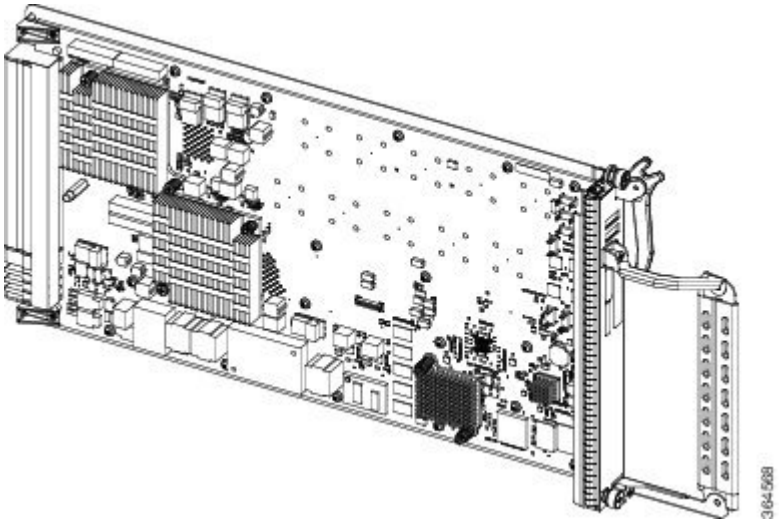

Figure 3: CRS-8-FC400/M Back-to-Back Switch Fabric Card

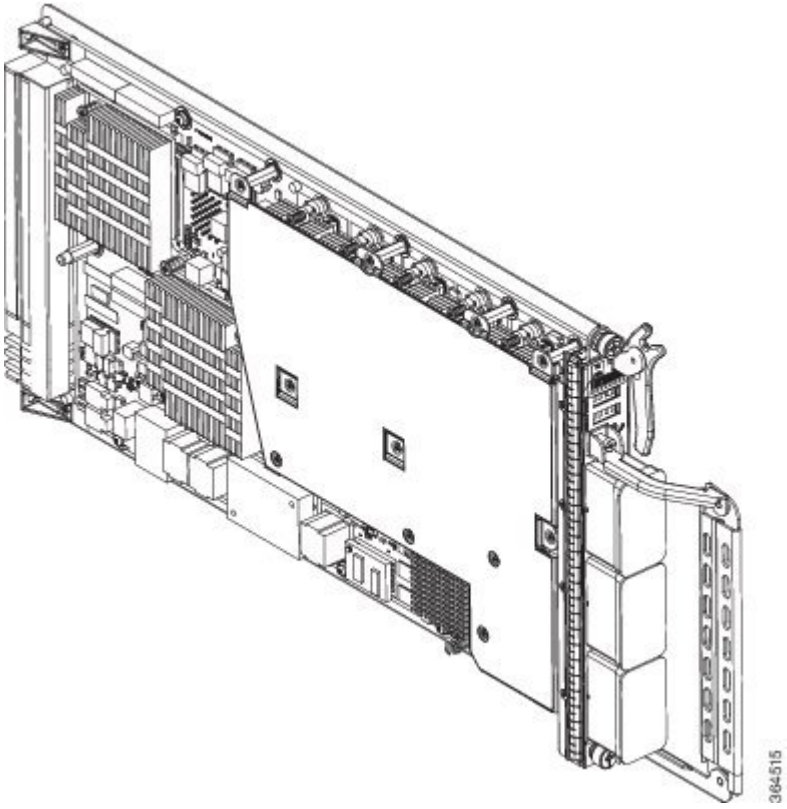

The following figure is a block diagram of the switch fabric card.

Figure 4: \Block Diagram of the Cisco CRS 8-Slot Line Card Chassis Switch Fabric Card

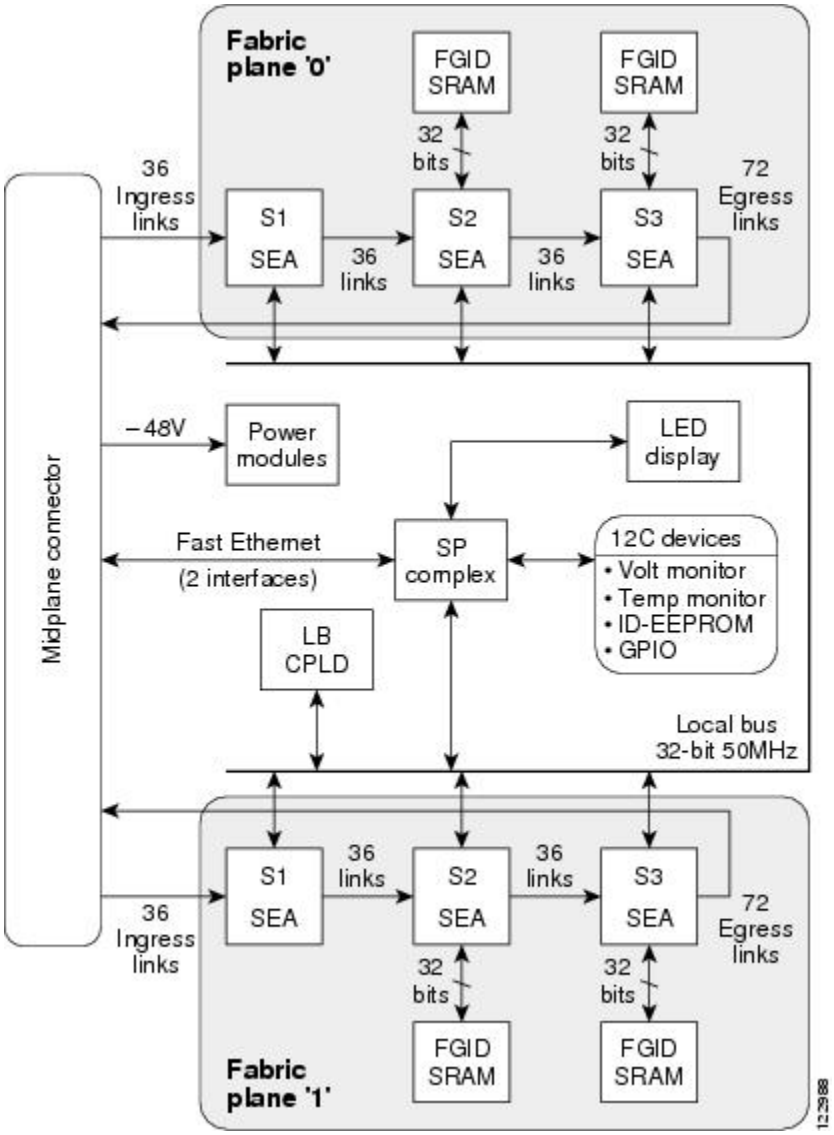

The major functional blocks of the switch fabric card are as follows:

- S1 switch element: Receives data cells from the MSC (or RP) and distributes them to the S2 stage. The S1 switch element is connected to its corresponding S2 switch element within the same fabric plane.
- S2 switch element: Receives data cells from the S1 stage. The S2 switch element is connected to its corresponding S1 and S3 switch elements within the same fabric plane. S2 has 36 inputs and 36 outputs.
- S3 switch element: Receives data cells from the S2 stage and performs switching and fabric speed-up. S3 has 36 inputs and 72 outputs.
- Service processor: Provides the interface to the Cisco CRS-1 control plane. The service processor does the following:
  - Controls power up and power down processes of the switch fabric card.
  - Configures the components in the various switch elements.
  - Updates the FGID (Fabric Group ID), for multicast traffic.

- Maintains cell configuration.
- Controls link-up and link-down processing and status.
- Collects and processes statistics for the switch fabric card.

- Power modules: Takes –48 VDC input power from the midplane and converts it to the voltages required by the components on the switch fabric card.
- Alphanumeric display: Displays switch fabric card messages. Messages are explained in Alphanumeric Displays, page 1-13 .
- Status LED: Indicates status of the switch fabric card.

The following figure shows the CRS-8-FC/S switch fabric card front panel. The front panel of the CRS-8-FC140/S, CRS-8-FC400/S, and CRS-8-FC400/M fabric cards are similar.

Figure 5: CRS-8-FC/S Switch Fabric Card Front Panel

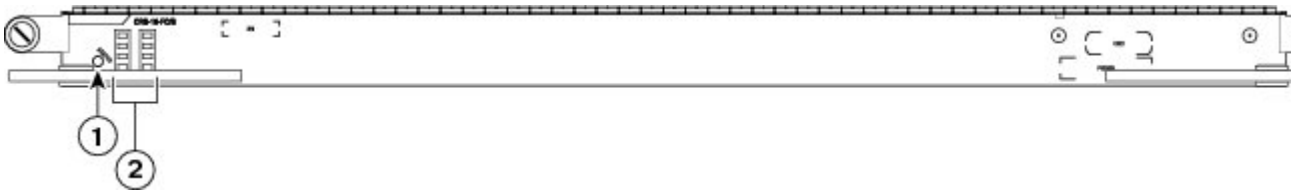

1	Status LED	2	Alphanumeric LEDs

The switch fabric card front panel contains the following elements:

- Status LED
- An alphanumeric display

See Appendix A, “Technical Specifications,” for additional switch fabric card specifications.