To: T10 T10/98-135r0
From: Bill Ham, Simon Fok
Subject: Section of EPI describing SCSI switches
Date: March 17, 1998

This document proposes additions and modifications to the sections of EPI to include switching expanders.

Bridging expanders

Bridging expanders perform the electrical isolation functions of simple expanders but may not be bound by all the same configuration rules. In their most general form, bridging expanders appear as multiple internally interconnected SCSI devices (each with its own independent ID and external [to the expander] connector) that are capable of passing payload data within the expander between the SCSI ports. This architecture is shown in Figure 1.

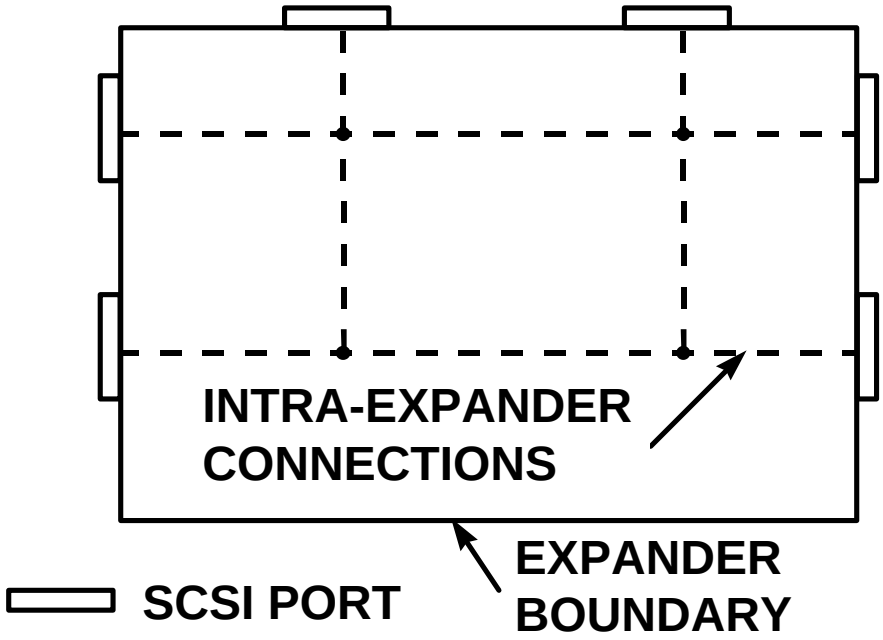

Figure 1 - General bridging expander

It is not the intent of this document to describe general bridging expanders since the number of possibilities are quite diverse. For example, the boundary in Figure 1 could be a server system with three dual port SCSI-to-PCI adapters.

There are three special cases of bridging expanders that will be discussed, however: (1) two port caching expanders that allow some provision for data buffering, (2) address enhancing expanders that use the LUN field presently defined in SCSI as the mechanism for the address

expansion and (3) multiport SCSI switches that route information to specific switch ports depending on the final destination.

Two port caching expanders

The least complex caching expander has two ports and some means for buffering data between the ports. This structure could be considered for use in cases where one desires to allow different data phase speed operation in the same domain thereby necessitating buffering of data in the expander. Another use could be to accept or deliver data from/to out of band (i.e. non SCSI) sources. Caching expanders are prone to deadlock problems if the buffer or cache ever becomes full and it requires management beyond SCSI to use this type of expander effectively. Therefore, we will not consider two port caching expanders further in this document. They could be integrated into a single chip but tend to be only useful for very specialized applications.

Multiport (more than two ports) expanders with caching are considered in section xx (Switching expanders).

Logical unit bridges

A type of non-caching bridging expander termed “LUN Bridge” are considered in this section.

[insert present epi material here]

Switching expanders

A form of caching bridging expander that acts as a multiple data path switch (as opposed to a mechanical switch) is considered in this section. This architecture requires considerable non-SCSI logic to control the data paths and also requires caching functionality.

The switch is different from a case 1 LUN bridge because a more general set of ports is allowed and the requirements for device placement are relaxed. Case 1 LUN bridges require a single primary port and may not have any initiators on the secondary busses (except for the LUN bridge itself). With a switch, all ports are initiators and there are no restrictions on the number of initiators on any ports. There is no distinction between primary and secondary busses. A full range of ID/LUN’s is available to use for every switch port. Every switch port has a SCSI ID for the port and that ID is part of the ID set for the domain attached to that port. LUN(0) on the port ID is reserved for the switch port.

A data path can be constructed within the switch between any two switch ports by using source and destination information contained within the SCSI information phases of the received data. By reading this information the SCSI switch can determine the final destination for the information in terms of the SCSI ports on the switch and the ID/LUN of devices on the switch port.

A SCSI switch enables an expansion of the concept of a SCSI domain to include switch ports in addition to the ID/LUN's on the switch ports. [SCSI domains are defined as the set of SCSI devices that are addressable from an initiator or target.]

SCSI switches constitute a separate SCSI device type.

SCSI switches can sustain simultaneous data paths between any two ports. This capability allows much better utilization of system resources when compared to other types of expander because the busses attached to switch ports can be used simultaneously for different tasks.

Figure 2 shows one set of data paths established by the switch for specific tasks. Figure 3 shows the same switch servicing a different set of tasks that need a different set of data paths.

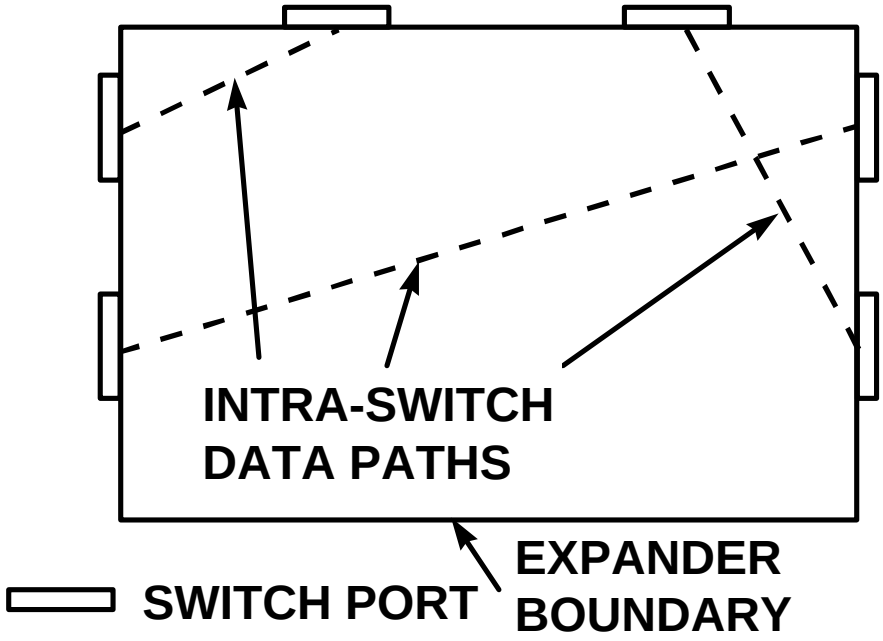

Figure 2 - One set of data paths in a SCSI switch

The data paths can be different for each information transfer requested.

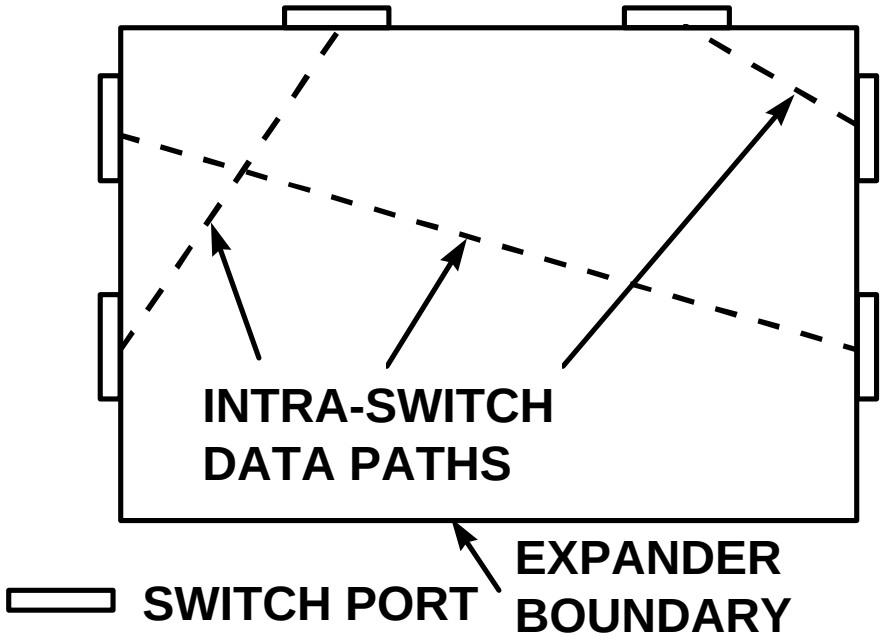

Figure 3 - Another set of data paths in the same SCSI switch

SCSI devices connected to the same switch port communicate with each other as they would if there were no switch present. Transferring information to SCSI devices connected to other switch ports or to the switch itself requires addressing the information to the switch port for the source device under the switch port's ID and LUN with information containing the destination port and the ID/LUN of the destination device. The switch has the responsibility to read this information and direct the data to the destination switch port where the switch will execute the subsequent transfer to the destination device.

All switch ports shall be initiators that support target mode.

The structure of the information directed to the switch and to devices on other switch ports is not defined in this document.

Requirements on switches include:

- TERMPWR shall not be directly connected between switch ports
- DIFFSENS shall not be connected electrically or logically between switch ports
- RESETS shall not be passed from one switch port to another switch port (except if deemed by the switch to be needed and then shall not be implemented by a direct electrical connection)
- Switches may supply TERMPWR to attached busses
- The REPORT LUN’s command shall be supported by every switch port under LUN(0) for reporting the configuration and status of attached devices, switch status, and port status.
- Switches are responsible for avoiding deadlock between ports.