

Date: 3/7/2004

T10: T10 Technical Committee (SCSI)

From: Charles Binford, Sun Microsystems, Inc., (charles.binford@sun.com)

Subject: Target Multiport Model

Revision History

Revision 0 (March, 2004), first revision

Related Documents

SAM3r10 – SCSI Architectural Model – 3, revision 11

Dicsussion

The SCSI Ports section of SAM-3 gives several diagrams of various configurations. For example, figures 20 - 22 shown below.

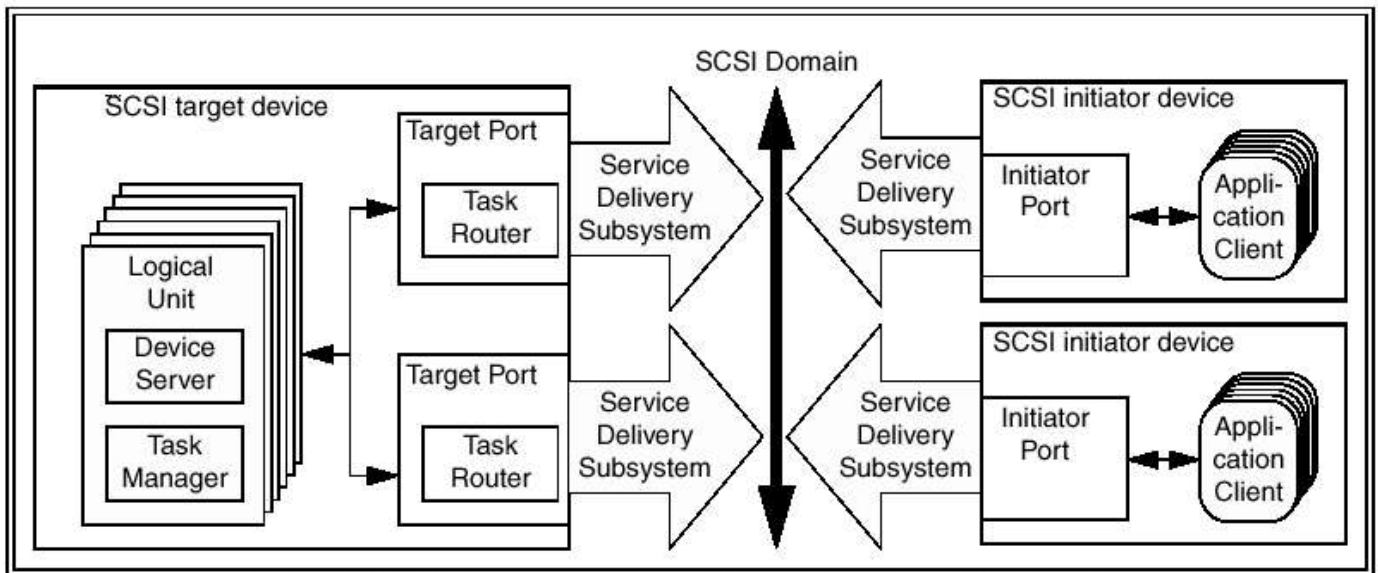


Figure 20 — SCSI target device configured in a single SCSI domain

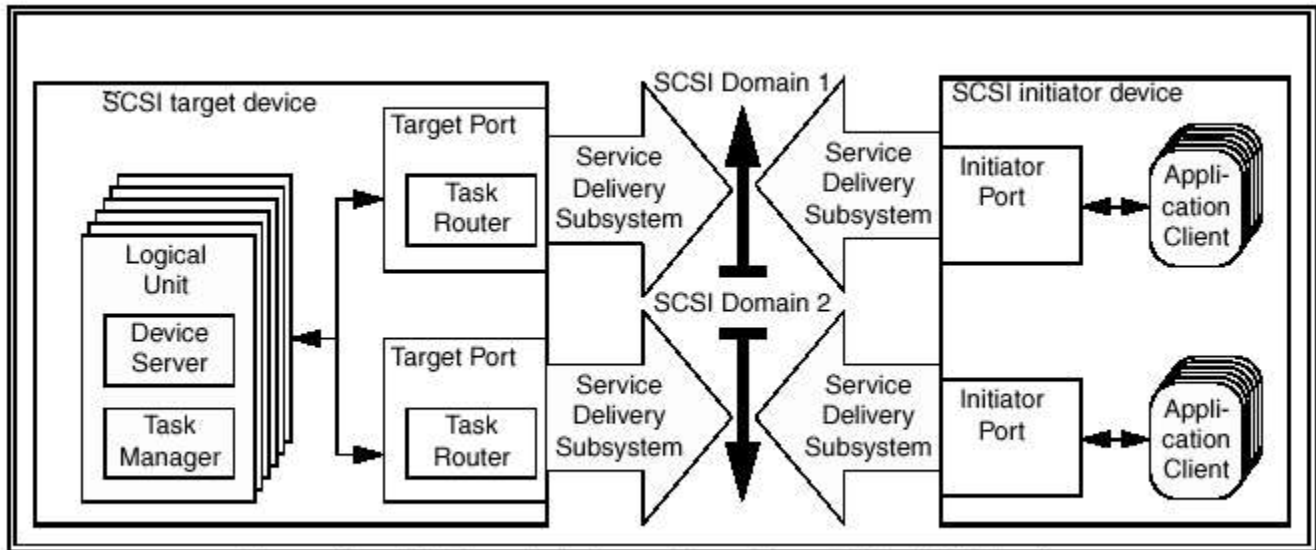


Figure 21 — SCSI target device configured in multiple SCSI domains

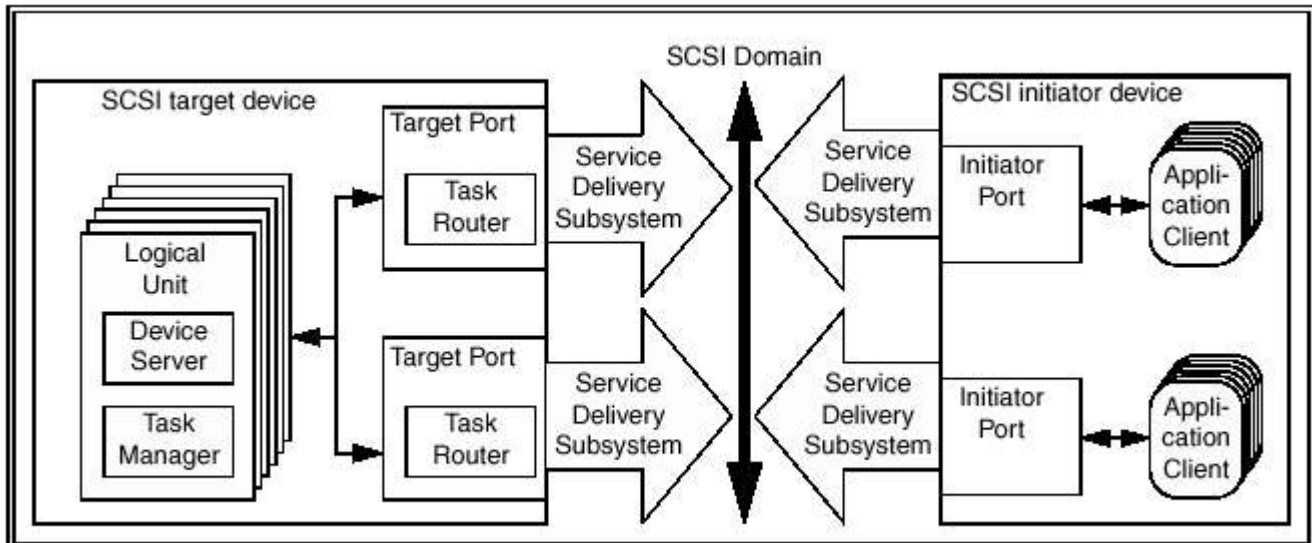


Figure 22 — SCSI target device and SCSI initiator device configured in a single SCSI domain

Based on discussion from last CAP meeting, I believe we also need a diagram that shows the following configuration:

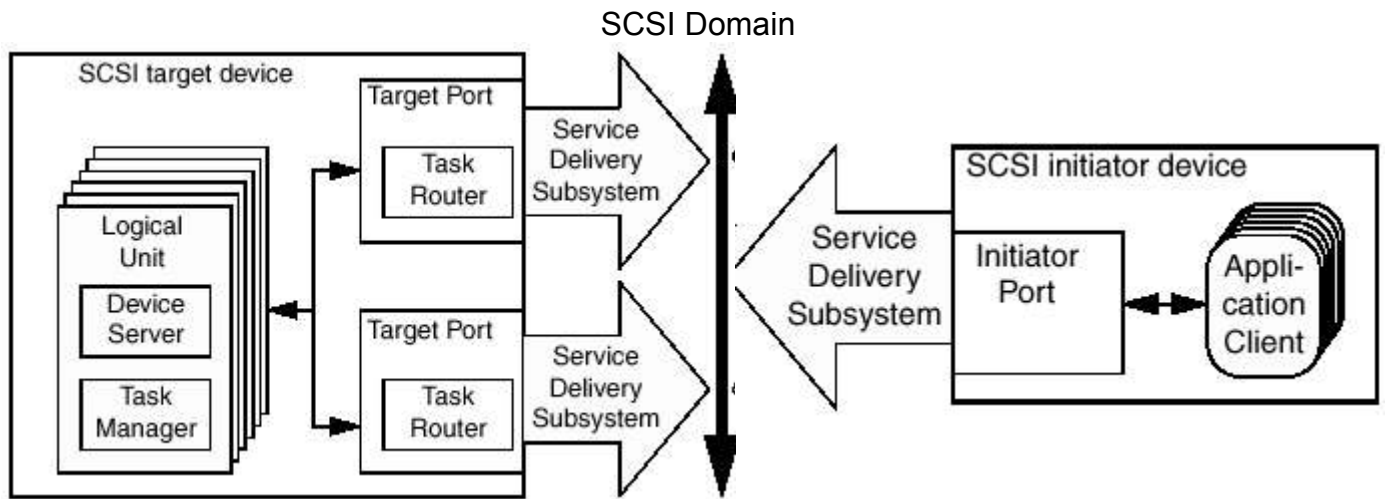


Figure X

The key question this figure will answer is how many “initiators” does the target see? Should the target differentiate this from figure 21 above?

Background

This subject came up at January's CAP meeting and it quickly became clear that there were different opinions/interpretations of how this configuration is viewed per SAM. Some felt that given the current transports defined for SCSI, that the target can and should figure out that the SCSI Initiator Port Name is the same at treat IOs delivered via either target port as if they were from the same SCSI initiator port.

Others (at least me! :-)) feel that should SAM require the target to treat this configuration as if there were two separate initiators.

With old parallel SCSI there was no question. No mechanism existed that would allow a target device to “figure it out”, so it was required to **assume** multiple initiators. While some may argue that parallel SCSI is not applicable here, I say we would be wise to not change the model from what it has been in the past without good reason.

With FCP, for example, it is possible for a device server to figure out whether or not a SCSI Initiator port visible from target port A is the same as the SCSI Initiator port visible from target port B. But, just because it is possible doesn't mean it is a good idea and should be required.

Relevant Text

The following two paragraphs talk about the case where a single host has two separate ports. In this scenario it is specifically prohibited from treating the two ports as a single initiator.

SAM-3 r11 page 45

4.13.5 Multiple port target/initiator SCSI device structure

....

This standard does not specify or require the definition of any mechanisms by which a SCSI target device would have the ability to discover that it is communicating with multiple

SCSI initiator ports on a single SCSI target/initiator device. **In those SCSI transport protocols where such mechanisms are defined, they shall not have any effect on how commands are processed** (e.g., reservations shall be handled as if no such mechanisms exist).

SAM-3 r11, page 48

4.13.7 SCSI target device view of a multiple port SCSI initiator device

This standard does not require a SCSI target device to have the ability to detect the presence of a SCSI initiator device with multiple SCSI initiator ports. Therefore, a SCSI target device handles a SCSI initiator device with multiple SCSI initiator ports exactly as it would handle multiple separate SCSI initiator devices. For example, a SCSI target device handles the configurations shown in figure 21 and figure 22 in exactly the same way it handles the configuration show in figure 20.

NOTE 5 - The implications of this view of a SCSI initiator device are more far reaching than are immediately apparent. For example, if a SCSI initiator device makes an exclusive access reservation via one SCSI initiator port, then access will be denied to the other SCSI initiator port(s) on that same SCSI initiator device.

No where can I find language that directly applies to Figure X. However, the following text addresses the Figure X from the Initiator's point of view.

SAM-3 r11, page 47, immediately following Figure 22

This model for application client determination of multiple SCSI target port configurations relies on information that is available only to the application clients via SCSI commands. The SCSI initiator ports in the SCSI **initiator devices** (figure 20) or SCSI initiator device (figure 21 and figure 22) **are unable to distinguish the multiple SCSI target ports from individual SCSI target ports in two separate SCSI target devices.**

I maintain that the target and initiator **MUST** have the same view of the world. Thus, if the initiator device can't tell if Figure X is a single target device with multiple ports or two separate target devices, then to preserve the semantics of task management and reservations the target must treat Figure X configurations as two separate initiators.

Port Definition

3.1.96 SCSI port: A SCSI device resident object that connects the application client, device server or task manager to the service delivery subsystem through which requests and responses are routed. SCSI port is synonymous with port. A SCSI port is either a SCSI initiator port (see 3.1.95) or a SCSI target port (see 3.1.100).

3.1.97 SCSI port identifier: A value by which a SCSI port is referenced within a domain. The SCSI port identifier is either an initiator port identifier (see 3.1.53) or a target port identifier (see 3.1.121).

3.1.98 SCSI port name: A name (see 3.1.67) of a SCSI port that is world wide unique within the SCSI transport protocol of the SCSI domain of that SCSI port (see 4.7.7). The name may be made available to other SCSI devices or SCSI ports in that SCSI domain in SCSI transport protocol specific ways.

In FCP, the port identifier is NOT world wide unique. Thus a target can not tell from it whether or not there is one or multiple SCSI domains. (see table A.2 and A.3 in SAM-3)

Why Does It Matter?

The decision taken on this subject has an impact on the queuing model when TST =001b as well as how Qerr=11b behaves. For example, at the last CAP meeting we changed the following table:

SPC-3 r16

Table 222 — Task set type

| Value | Description |
|-------------|---|
| 000b | The logical unit maintains one task set for all initiator ports |
| 001b | The logical unit maintains separate task sets for each initiator port |
| 010b - 111b | Reserved |

SPC-3 r17

Table 224 — Task set type

| Value | Description |
|-------------|---|
| 000b | The logical unit maintains one task set for all initiator ports |
| 001b | The logical unit maintains separate task sets for each initiator port regardless of target port |
| 010b - 111b | Reserved |

My port model says this was an incorrect change.

For a redundant array controller configuration, having a device server view Figure X as two separate initiator ports allows for a much looser coupling of the task sets when TST=001b or Qerr=11b.