

Date: 8/24/2004

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

Revision 1 (August, 2004), SAM-3 wording changes in separate "fdf" file.

Revision 2 (August, 2004), Re did 'fdf' comments per conference call feedback.

Revision 3 (August, 2004), One minor change I missed from con-call notes.

Related Documents

SAM3r13 – SCSI Architectural Model – 3, revision 13

R1 Update:

I've uploaded an Acrobat comment extract file (.FDF) that contains all proposed changes to SAM-3 with regard to this issue.

R2 / 3 Update:

Updates based on 8/10 conference call were made to the PDF comments and a new FDF file was created and uploaded.

Note, I will focus on SBC-2 next as because of letter ballot timing. SPC-3 comments will follow that.

Remainder of this doc unchanged....

Discussion

The current T10 docs (e.g. SAM-3 and SPC-3) often refer to "SCSI initiator port" when a more accurate reference would be to the I_T nexus. For example, from SAM-3 r13 5.7.1:

An action transmitted via a SCSI initiator port may abort task(s) created via the **SCSI initiator port** itself, task(s) created via another **SCSI initiator port**, or both its own tasks and tasks created via another **SCSI initiator port**.

SCSI Domain

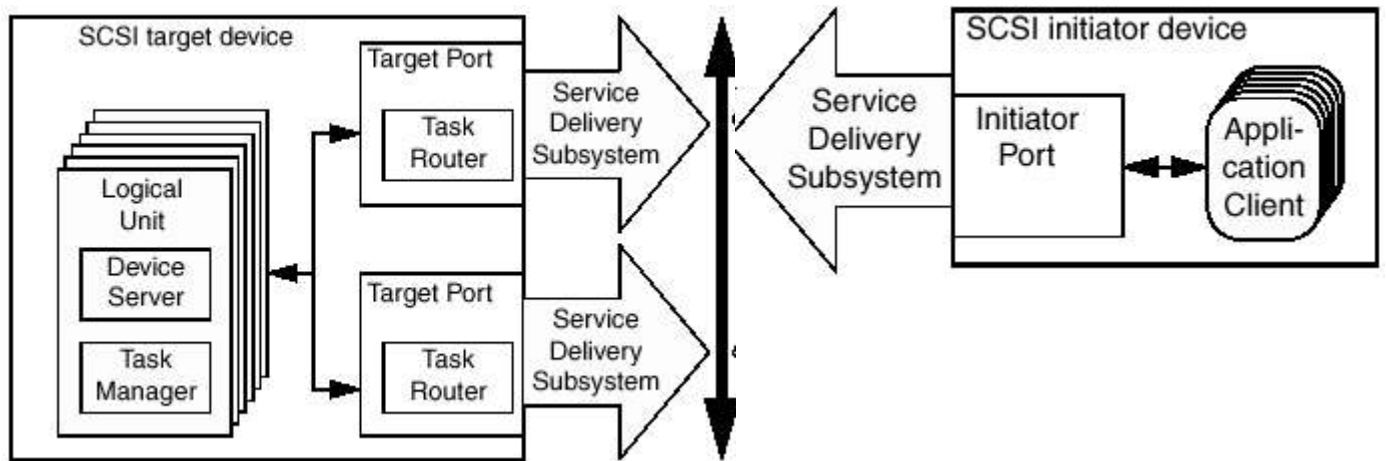


Figure X

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.

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.