To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 9 March 2003
Subject: T10/03-097r0 SAM-3 Incorrect LUNs in task management requests

Revision History
Revision 0 (9 March 2003) first revision

Related Documents
sam3r05 - SCSI Architecture Model - 3 revision 5

Overview

A Quantum SAS letter ballot comment points out that SAM-3 does not mention how to handle incorrect logical unit numbers for task management functions. Section 5.9.3 describes how this is handled for commands - the target port usually returns CHECK CONDITION status with a sense key of ILLEGAL REQUEST and an additional sense code of LOGICAL UNIT NOT SUPPORTED. Unfortunately, CHECK CONDITION cannot be used for a task management request, so a different rule is needed.

SAM-3 defines the task manager as an object inside a logical unit (many times). According to 4.7.5, task management functions which only refer to an I_T Nexus (TARGET RESET and WAKEUP) are sent to *all* the task managers accessible to the task router. However, 4.7.5 doesn't explain how they coordinate their replies. These functions still must result in a single service result like FUNCTION COMPLETE or FUNCTION REJECTED.

And, as the letter ballot comment notes, nowhere is it mentioned what happens when a logical unit number is specified (for a task management function referring to an I_T_L or I_T_L_Q like ABORT TASK SET) but such a logical unit does not exist.

I suggest this approach:

- define a special task manager inside the target device or target port object (probably the target port)
- have the task router route I_T_L and I_T_L_Q requests with a known L to the corresponding logical unit.
- have the task router route I_T_L and I_T_L_Q based task management functions to the target port task manager if the L is unknown. Define it as returning FUNCTION REJECTED for these lost functions.
- have the task router route I_T based task management functions to the target port task manager for processing. It may communicate with the task managers inside each logical unit to help implement certain functions (e.g. it converts a TARGET RESET into individual LOGICAL UNIT RESET requests).

Current text has the SCSI target device returning CHECK CONDITION for commands with incorrect LUNs; the target port's task manager can be given that job as well.

Suggested Changes

3.1.100 SCSI target port: A SCSI target device object that contains a task router and a task manager and acts as the connection between device servers and task managers and the service delivery subsystem through which indications and responses are routed. When this term is used it refers to a SCSI target port or a SCSI target/initiator port operating as a SCSI target port.

3.1.128 task manager: A server within a logical unit that controls the sequencing of one or more tasks and processes task management functions, or a server within a SCSI target port that processes commands and task management functions specifying unknown logical units and task management functions not specifying logical units.

3.1.129 task router: An object in a SCSI target port that routes commands and task management functions between the service delivery subsystem (see 3.1.114) and the appropriate ~~logical unit's~~ task manager (see 3.1.128).

4.3 The SCSI client-server model

As shown in figure 6, each SCSI target device contains one or more logical units and provides services performed by device servers and task management functions performed by task managers. A logical unit is an object that implements one of the device functional models described in the SCSI command standards and processes SCSI commands such as reading from or writing to the media. Each pending SCSI command or series of linked commands defines a unit of work to be performed by the logical unit. Each unit of work is represented within the SCSI target device by a task that may be externally referenced and controlled through requests issued to the task manager.

[Figure 6 - SCSI client-server model - add target port and its task manager]

...

An application client may request processing of a task management function through a request directed to the task manager within the ~~logical unit~~ target device. The interactions between the task manager and application client when a task management request is processed are shown in 7.11.

...

4.7.2 SCSI target device

A SCSI target device (see figure 12) contains:

a) Zero or more target device names;

b) One or more SCSI target ports each containing a task router, task manager, SCSI target port identifier, and an optional target port name; and

c) One or more logical units.

[Figure 12 - SCSI target device router - add a Task Manager object under SCSI Target Port]

...

A task router routes commands and task management functions between the service delivery subsystem and the appropriate ~~logical unit's~~ task manager (see 4.7.5).

...

4.7.3 SCSI target/initiator device

[apply 4.7.2 changes here too]

4.7.5 SCSI task router

~~The task router routes tasks and task management functions to the selected logical unit. Any task that is sent to a logical unit that is not known to the task router is handled as described in 5.9.4.~~

~~Any task management function that is not sent to a specific logical unit shall be broadcast to all logical units known to the task router.~~

The task router shall send a command to the SCSI target port's task manager for handling as described in 5.9.4 if:

a) the SCSI target device does not support the logical unit;
b) the SCSI target device supports the logical unit, but the peripheral device is not currently attached to the SCSI target device;
c) the SCSI target device supports the logical unit and the peripheral device is attached, but not operational; or

d) the SCSI target device supports the logical unit but is incapable of determining if the peripheral device is attached or is not operational when it is not ready.

Otherwise, the task router shall send a command to the specified logical unit's task manager.

The task router shall send a task management function with a Nexus argument containing an I_T_L Nexus or an I_T_L_Q Nexus specifying a known logical unit to that logical unit's task manager.

The task router shall send a task management function to the SCSI target port's task manager as described in 7.12 if:

a) the task management function has a Nexus argument containing an I_T Nexus (e.g., TARGET RESET); or
b) the task management function has a Nexus argument containing an I_T_L Nexus or an I_T_L_Q Nexus but the logical unit number is unknown.

[Note: do the 4 cases for commands apply here too?]

[Note: some description of the ACCESS CONTROL LUN mapping behavior might belong here]

4.13.3 Multiple port target SCSI device structure

Figure 17 shows the structure of a SCSI target device with multiple SCSI target ports. Each SCSI target port consists of a task router ~~that is shared by a collection of~~, a task manager, and one or more logical units. Each logical unit contains a single task manager and a device server.

[Figure 17 - Multiple port target SCSI device structure model - add task manager object]

Two-way communications shall be possible between all logical units and all SCSI target ports, however, communications between any logical unit and any SCSI target port may be inactive. Two-way communications shall be available between each task manager and all task routers. Each SCSI target port shall accept commands sent to LUN 0 and the task router shall route them to a device server for processing. The REPORT LUNS commands (see SPC-2) shall be accepted by the logical unit with the logical unit number zero from any SCSI target port and shall return the logical unit inventory available via that SCSI target port. The availability of the same logical unit through multiple SCSI target ports is discovered by matching SCSI port name or identifier values in the INQUIRY command Device Identification VPD page (see SPC-2).

...

4.13.4 Multiple port target/initiator SCSI device structure

[apply 4.13.3 changes here too]

5.9 Command processing considerations and exception conditions

5.9.1 Commands that complete with CHECK CONDITION status [no change]

5.9.2 Auto contingent allegiance (ACA) [no change]

5.9.3 Overlapped commands [no change]

5.9.4 ~~Incorrect~~ Unknown logical unit selection [reword in terms of task manager]

[Note: there is one use of "incorrect logical unit number" in SPC-3 in the Sense key descriptions table that would change if this rename is accepted.]

~~The SCSI target device's response to an incorrect logical unit number is described in this subclause.~~

The SCSI target port's task manager receives commands if the logical unit number is unknown (see 4.7.5). The logical unit number may be incorrect because:

a) The SCSI target device does not support the logical unit (e.g., some SCSI target devices support only one peripheral device).

In response to an INQUIRY command the task manager shall return the INQUIRY data with the peripheral qualifier of 011b (see SPC-2).

In response to a REQUEST SENSE command, the task manager shall return GOOD status and parameter data that contains sense data. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED.

In response to any other command except REQUEST SENSE and INQUIRY, the ~~SCSI target device~~task manager shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED;

b) The SCSI target device supports the logical unit, but the peripheral device is not currently attached to the SCSI target device.

In response to an INQUIRY command the ~~SCSI target device~~ task manager shall return the INQUIRY data with the peripheral qualifier set to ~~the value required in SPC-2~~001b (see SPC-2)..

In response to a REQUEST SENSE command, the ~~SCSI target device~~ task manager shall return GOOD status and parameter data that contains sense data. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED.

In response to any other command ~~except REQUEST SENSE and INQUIRY~~, the ~~SCSI target device~~ task manager shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED;

c) The SCSI target device supports the logical unit and the peripheral device is attached, but not operational.

In response to an INQUIRY command the ~~SCSI target device~~ task manager shall return the INQUIRY data with the peripheral qualifier set to ~~the value required in SPC-2~~000b (see SPC-2)..

In response to REQUEST SENSE, the ~~SCSI target device~~ task manager shall return GOOD status and parameter data that contains sense data appropriate to the condition that is making the logical unit not operational.

The ~~SCSI target device~~ task manager's response to any other command ~~other than INQUIRY and REQUEST SENSE~~ is vendor specific; or

d) The SCSI target device supports the logical unit but is incapable of determining if the peripheral device is attached or is not operational when it is not ready.

In response to an INQUIRY command the ~~SCSI target device~~ task manager shall return the INQUIRY data with the peripheral qualifier set to the value specified in SPC-2.

In response to a REQUEST SENSE command the ~~SCSI target device~~ task manager shall return GOOD status and parameter data that contains sense data with a sense key of NO SENSE.

The ~~SCSI target device~~ task manager's response to any other command is vendor specific.

5.9.5 Task attribute exception conditions [no change]

5.9.6 Sense data [no change]

5.9.7 Unit Attention condition [no change]

...

7 Task management functions

7.1 Introduction [no change]

7.2 ABORT TASK [no change]

7.3 ABORT TASK SET [no change]

7.4 CLEAR ACA [no change]

7.5 CLEAR TASK SET [no change]

7.6 LOGICAL UNIT RESET [no change]

7.7 QUERY TASK [no change]

7.8 TARGET RESET

Function Call:

Service Response = TARGET RESET (IN (I_T Nexus))

Description:

~~Before returning a FUNCTION COMPLETE response, the SCSI target port shall cause logical unit reset functions to be performed as specified in 6.3.3 for every logical unit.~~

The SCSI target port's task manager shall send cause logical unit reset functions to be performed as specified in 6.3.3 for every logical unit. Upon completion of all the logical unit resets, the SCSI target port's task manager shall return a FUNCTION COMPLETE response.

An application client should issue LOGICAL UNIT RESETs only to the logical units it is using rather than issuing a TARGET RESET. This avoids resetting logical units that other SCSI initiator ports, possibly in other SCSI initiator devices, may be using.

NOTE 15 - Previous versions of this standard required TARGET RESET support in all SCSI target devices. SCSI transport protocols may or may not require that TARGET RESET be supported. SCSI transport protocols may require additional actions beyond those specified here.

7.9 WAKEUP [this section is being deleted by another proposal]

7.10 Task management SCSI transport protocol services [no change]

7.11 Task management function example [no change]

7.12 Task management function processing considerations [maybe slip in before 7.11]

The SCSI target port's task manager receives task management functions specifying only an I_T Nexus or if the logical unit number is incorrect (see 4.7.5).

If the task management function species only an I_T Nexus, the task manager shall process the task management function as described in this clause.

If the logical unit number is incorrect, the task manager shall return a response of FUNCTION REJECTED.