

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 25 January 2005
Subject: 04-372r2 SAM-4 SPC-4 SAS-2 I_T NEXUS RESET task management function

Revision history

Revision 0 (5 November 2004) First revision

Revision 1 (13 November 2004) Incorporated comments from November CAP WG.

Revision 2 (25 January 2005) Incorporated comments from January SAS and CAP WGs - target proposal at SAS-2 so SAS-1.1 need not depend on SAM-4/SPC-4, and added "(see SAM-4)" in the SPC-4 text.

Related documents

sam3r14 - SCSI Architecture Model - 3 revision 14

spc3r21 - SCSI Primary Commands - 3 revision 21

sas1r06 - Serial Attached SCSI 1.1 revision 6

03-097r0 Incorrect LUNs in task management requests (Rob Elliott, HP) not incorporated into SAM-3

03-173r0 Task Management Incorrect Logical Unit Selection (Ralph Weber, ENDL) incorporated into sam3r07

Overview

Initiator device software may have a target-cognizant level (e.g., an initiator port driver) below the LU-cognizant level(s). The target-cognizant level is the one that knows about I_T nexus losses; it may not know anything about LUs being used. If it experiences an I_T nexus loss, it trashes the appropriate context and tells the LU-cognizant levels about it; they assume all the tasks are aborted, etc. as defined in SAM-3.

When communication resumes on that I_T, the target-cognizant level would like to ensure the target made the same decision before letting the LU-cognizant levels have access. In SAS, I_T nexus loss is determined by timers, which might not be perfectly synchronized. The target-cognizant level could send ABORT TASK SET to each LUN, but this is impractical since it doesn't know about the LUNs in use. It could ask the LU-cognizant levels to each send ABORT TASK SETs, but a) this would be confusing since they were previously told all their activity on those LUs was cleaned up, b) this takes longer, and c) ABORT TASK SET has fewer consequences than I_T nexus loss (e.g. I_T nexus loss clears ACA conditions).

In Fibre Channel and iSCSI, login and logout serve as I_T nexus loss events, and the new login ensures both sides are in agreement that an I_T nexus loss must have occurred. Since SAS doesn't support login/logout, a task management function called I_T NEXUS RESET is proposed to help mimic that behavior.

Additionally, a protocol bridge (see BCC) from a protocol that supports this new function (e.g. SAS) to one that does not (e.g. FCP) can use it to ease the mapping of logins/logouts:

- a) FCP login or logout -> SAS I_T NEXUS RESET task management function;
- b) SAS I_T Nexus Loss timer expiration -> FC logout;
- c) SAS communication resumes -> FC login.

This would be defined in SAM-4 so any protocol could support it. FCP-3 and iSCSI have no need for it; SAS-1.1 (or SAS-2) would define support. It cannot be made mandatory at this late date. SPC-4's REPORT SUPPORTED TASK MANAGEMENT FUNCTION command adds a bit to indicate support for it.

In response to a question in the November CAP WG, editorial changes are proposed throughout SAM-4 to clarify how an I_T based task management function is handled - the task router sends it to all the task managers. Some clarifications about incorrect LUN handling are made as well.

Suggested changes to SAM-4

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. [See 4.8.](#)

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 3.1.113) and the appropriate ~~logical unit's~~ task manager(s) (see 3.1.128). [See 4.7.6.](#)

4.7.2 SCSI target device

...

The task router routes commands and task management functions between the service delivery subsystem and the appropriate ~~logical unit's~~ task manager(s) (see 4.7.6).

...

4.7.3 SCSI target/initiator device

...

When the SCSI target/initiator device is operating as a SCSI target device the task router routes the commands and task management functions between the service delivery subsystem and the appropriate logical unit(s) (see 4.7.6).

...

4.7.6 SCSI task router

~~The task router routes commands 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.~~

The task router routes commands and task management functions as follows:

- a) commands addressed to a valid logical unit number are routed to the task manager in the specified logical unit;
- b) commands addressed to an incorrect logical unit number are handled as described in 5.9.4;
- a) task management functions with I_T_L nexus scope (e.g., ABORT TASK SET, CLEAR TASK SET, CLEAR ACA, and LOGICAL UNIT RESET) or I_T_L_Q nexus scope (e.g., ABORT TASK and QUERY TASK) addressed to a valid logical unit number are routed to the task manager in the specified logical unit;
- b) task management functions with an I_T nexus scope (e.g., I_T NEXUS RESET) are routed to the task manager in each logical unit about which the task router knows; and
- c) task management functions with I_T_L nexus scope or I_T_L_Q nexus scope addressed to an incorrect logical unit number are handled as described in 7.8.

4.9 Logical unit numbers

4.9.1 Introduction

Subclause 4.9 defines the construction of logical unit numbers to be used by SCSI target devices. Application clients should use only those logical unit numbers returned by a REPORT LUNS command. The ~~device-server~~ task router shall respond to logical unit numbers other than those returned by a REPORT LUNS command (i.e., incorrect logical unit numbers) as specified in 5.9.4 and 7.8.

4.10 Well known logical units

...

If a SCSI target ~~device~~port receives a W-LUN and the well known logical unit specified by the W-LUN does not exist, ~~a task manager~~ the task router shall follow the rules for selection of incorrect logical units described in 5.9.4 and 7.8.

...

5.5 Task and command lifetimes

The application client maintains an application client task to interact with the task from the time the Send SCSI Command SCSI transport protocol service request is invoked until it receives one of the following SCSI target device responses:

- d) A service response of TASK COMPLETE for that task;
- e) Notification of a unit attention condition with one of the following additional sense codes;
 - A) COMMANDS CLEARED BY ANOTHER INITIATOR, if in reference to the task set containing the task;
 - B) Any additional sense code whose ADDITIONAL SENSE CODE field contains 29h (e.g., POWER ON, RESET, OR BUS DEVICE RESET OCCURRED; POWER ON OCCURRED; SCSI BUS

RESET OCCURRED; BUS DEVICE RESET FUNCTION OCCURRED; DEVICE INTERNAL RESET; or I_TNEXUS LOSS OCCURRED);

- f) Notification that the task manager has detected the use of a duplicate I_T_L_Q nexus (see 5.9.3);
- g) A service response of FUNCTION COMPLETE following an ABORT TASK task management function directed to the specified task;
- h) A service response of FUNCTION COMPLETE following an ABORT TASK SET or a CLEAR TASK SET task management function directed to the task set containing ~~the specified that~~ task; ~~or~~
- i) [A service response of FUNCTION COMPLETE following an I_T NEXUS RESET task management function delivered on the I_T nexus used to deliver that task; or](#)
- j) A service response of FUNCTION COMPLETE in response to a LOGICAL UNIT RESET task management function directed to the logical unit.

If a service response of SERVICE DELIVERY OR TARGET FAILURE is received for a command ([e.g., when an I_T nexus loss is detected by the SCSI initiator port](#)), the application client shall maintain an application client task to interact with the task until the application client has determined (~~e.g., by completion of an ABORT TASK task management function~~) that the task is no longer known to the device server. [An application client may determine that a task is no longer known to the device server by detecting:](#)

- a) [Completion of an ABORT TASK task management function specifying that task;](#)
- b) [Completion of an ABORT TASK SET or an I_T NEXUS RESET task management function on the I_T nexus used to deliver that task; or](#)
- c) [Completion of a CLEAR TASK SET or LOGICAL UNIT RESET task management function.](#)

...

5.7 Aborting tasks

5.7.1 Mechanisms that cause tasks to be aborted

...

The following actions affect only the task(s) received on the I_T nexus on which the action is transmitted:

- a) Completion of an ABORT TASK task management function directed to the specified task;
- b) Completion of an ABORT TASK SET task management function under the conditions specified in 7.3; ~~or~~
- c) [Completion of an I_T NEXUS RESET task management function; or](#)
- d) Completion of a command with a CHECK CONDITION status, without establishing an ACA condition (see 5.9.1.3) or establishing an ACA condition (see 5.9.2.2), while the Control mode page (see SPC-3) contains fields that are set as follows:
 - A) The QERR field set to 01b and the TST field set to 001b; or
 - B) The QERR field set to 11b.

...

5.9.4 Incorrect logical unit selection

The SCSI target device's response to a [command addressed to an](#) incorrect logical unit number is described in this subclause.

...

6 SCSI events and event notification model

6.1 SCSI events overview

[In figure 35, add "Processing an I_T NEXUS RESET task management function" yellow box input into the "I_T Nexus Loss" box.](#)

6.3.4 I_T nexus loss

I_T nexus loss is a SCSI device condition resulting from:

- a) A hard reset condition (see 6.3.2); ~~or~~
- b) An I_T nexus loss event (e.g., logout) indicated by a Nexus Loss event notification (see 6.4); or
- c) An I_T nexus loss event indicating that an I_T NEXUS RESET task management request (see 7.n) has been processed.

An I_T nexus loss event is an indication from the SCSI transport protocol to the SCSI application layer that an I_T nexus no longer exists. SCSI transport protocols may define I_T nexus loss events.

Each SCSI transport protocol standard that defines I_T nexus loss events should specify when those events result in the delivery of a Nexus Loss event notification to the SCSI applications layer.

...

7 Task management functions

7.1 Introduction

An application client requests the processing of a task management function by invoking the SCSI transport protocol services described in 7.8, the collective operation of which is modeled in the following procedure call:

Service Response = Function name (IN (nexus))

The task management function names are summarized in table 30.

Table 1 — Task Management Functions

Task Management Function	Nexus	Reference
ABORT TASK	I_T_L_Q	7.2
ABORT TASK SET	I_T_L	7.3
CLEAR TASK SET	I_T_L	7.4
CLEAR ACA	I_T_L	7.5
LOGICAL UNIT RESET	I_T_L	7.6
<u>I_T NEXUS RESET</u>	<u>I_T</u>	<u>7.n</u>
QUERY TASK	I_T_L_Q	7.7

...

Service Response:

One of the following SCSI transport protocol specific responses shall be returned:

...

INCORRECT LOGICAL UNIT NUMBER: An optional task ~~manager~~router response indicating that the function requested processing for an incorrect logical unit number.

...

The task manager response to task management requests is subject to the presence of access restrictions, as managed by ACCESS CONTROL OUT and ACCESS CONTROL IN commands (see SPC-3), as follows:

- a) A task management request of ABORT TASK, ABORT TASK SET, CLEAR ACA, I_T NEXUS RESET, or QUERY TASK shall not be affected by the presence of access restrictions;
- b) A task management request of CLEAR TASK SET or LOGICAL UNIT RESET received from a SCSI initiator port that is denied access to the logical unit, either because it has no access rights or because it is in the pending-enrolled state, shall not cause any changes to the logical unit; and
- c) The task management function service response shall not be affected by the presence of access restrictions.

...

7.n I_T NEXUS RESETRequest:

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

Description:

SCSI transport protocols may or may not support I_T NEXUS RESET and may or may not require logical units accessible through SCSI target ports using such transport protocols to support I_T NEXUS RESET.

Each logical unit accessible through the SCSI target port shall perform the I_T nexus loss functions specified in 6.3.4 for the I_T nexus on which the function request was received, then the SCSI target device shall return a FUNCTION COMPLETE response. After returning a FUNCTION COMPLETE response, the logical unit(s) and the SCSI target port shall perform any additional functions specified by the SCSI transport protocol.

7.8 Task management SCSI transport protocol services

...

Response from task manager:

Task Management Function Executed (IN (Nexus, Service Response))

...

Service Response: An encoded value representing one of the following:

...

INCORRECT LOGICAL UNIT NUMBER: An optional task **manager**[router](#) response indicating that the function requested processing for an incorrect logical unit number.

...

Confirmation received by application client:

Received Task Management Function Executed (IN (Nexus, Service Response))

...

Service Response: An encoded value representing one of the following:

...

INCORRECT LOGICAL UNIT NUMBER: An optional task **manager**[router](#) response indicating that the function requested processing for an incorrect logical unit number.

...

Suggested changes to SPC-4**6.24 REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS command**

...

The format of the parameter data returned by the REPORT TASK MANAGEMENT FUNCTIONS command is shown in table 159.

Table 2 — REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS parameter data

Byte/Bit	0	1	2	3	4	5	6	7
0	ATS	ATSS	CACAS	CTSS	LURS	QTS	TRS	WAKES
1	Reserved							ITNRS
2	Reserved							
3	Reserved							

[An I_T NEXUS RESET supported \(ITNLS\) bit set to one indicates the I_T NEXUS RESET task management function \(see SAM-4\) is supported by the logical unit. An ITNLS bit set to zero indicates the I_T NEXUS RESET task management function is not supported.](#)

Suggested changes to SAS-2**8.2.2.3.6 PL_OC2:Overall_Control state frame transmission**

...

This state shall not send a Tx Frame message containing a TASK frame for a task that only affects an I_T nexus ([e.g., an I_T NEXUS RESET task management function \(see SAM-3\)](#)) until this state has received one of the following messages for each Tx Frame message with the same I_T nexus:

- a) Transmission Status (ACK Received);
- b) Transmission Status (NAK Received);
- c) Transmission Status (ACK/NAK Timeout); or
- d) Transmission Status (Connection Lost Without ACK/NAK).

...

9.2.2.2 TASK information unit

...

Table 3 defines the TASK MANAGEMENT FUNCTION field.

Table 3 — Task management functions

Code	Task management function	Uses LOGICAL UNIT NUMBER field	Uses TAG OF TASK TO BE MANAGED field	Description
01h	ABORT TASK	yes	yes	The task manager shall perform the ABORT TASK task management function with L set to the value of the LOGICAL UNIT NUMBER field and Q set to the value of the TAG OF TASK TO BE MANAGED field (see SAM-3).
02h	ABORT TASK SET	yes	no	The task manager shall perform the ABORT TASK SET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
04h	CLEAR TASK SET	yes	no	The task manager shall perform the CLEAR TASK SET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
08h	LOGICAL UNIT RESET	yes	no	The task manager shall perform the LOGICAL UNIT RESET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
10h	I_T NEXUS RESET	no	no	The task manager shall perform the I_T NEXUS RESET task management function (see SAM-3).
20h	Reserved ^a			
40h	CLEAR ACA	yes	no	The task manager shall perform the CLEAR ACA task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
80h	QUERY TASK	yes	yes	The task manager shall perform the QUERY TASK task management function with L set to the value of the LOGICAL UNIT NUMBER field and Q set to the value of the TAG OF TASK TO BE MANAGED field (see SAM-3).
All others	Reserved			
^a The TARGET RESET task management function defined in SAM-3 is not supported.				

[The task manager shall perform the specified task management function with the I and T arguments set to the initiator port and target port involved in the connection used to deliver the TASK frame.](#)

...

10.2 SCSI application layer

10.2.1 SCSI transport protocol services

10.2.1.1 SCSI transport protocol services overview

An application client requests the processing of a SCSI task management function by invoking SCSI transport protocol services, the collective operation of which is conceptually modeled in the following remote procedure calls (see SAM-3):

- Service Response = ABORT TASK (IN (Nexus));
- Service Response = ABORT TASK SET (IN (Nexus));
- Service Response = CLEAR ACA (IN (Nexus));

- d) Service Response = CLEAR TASK SET (IN (Nexus));
- e) [Service Response = I_T NEXUS RESET \(IN \(Nexus\)\)](#);
- f) Service Response = LOGICAL UNIT RESET (IN (Nexus)); and
- g) Service Response = QUERY TASK (IN (Nexus)).

[Add I_T NEXUS RESET to:](#)

[10.2.1.1 table 124 SCSI architecture mapping](#)

[10.2.1.10 table 133 Send Task Management Request transport protocol service arguments](#)

[10.2.1.11 table 134 Task Management Request Received transport protocol service arguments](#)