

To: T10 Technical Committee
From: Luben Tuikov, Vitesse Semiconductor (ltuikov@vitesse.com)
Date: 9 August 2006
Subject: 06-368r0 SAM-4 QUERY TASK TMF Extended Response

Revision history

Revision 0 (9 August 2006) First revision

Related documents

sam4r07 - SCSI Architecture Model - 4, revision 07
05-284r2 - Self-Describing Command Timeouts, revision 2

Overview

Recently there had been a need for an Application Client to properly determine the timeout of commands sent to a LU. This is partly addressed in document 05-284r2, Self Describing Command timeouts, revision 2. Such a mechanism would allow an Application Client to properly set a command timeout for commands sent to an LU, and not enter an error recovery procedure prematurely causing "error recovery thrashing".

A LU can only guarantee an upper bound of the execution time, and thus a command timeout, only when given factors which depend on the LU itself, such as mechanical characteristics or device media access algorithms including internal error recovery. A LU cannot guarantee how long a task would spend in the Dormant or Blocked state, since this may depend on the behavior of other Initiators (Application Clients). Thus, generally, a LU can guarantee an upper bound of the execution time when the task is a "current task".

This proposal extends the Function Result of the QUERY TASK Task Management Function, such that the Application client can dynamically make a decision whether to abort the task or wait for its completion for the timeout value obtained as specified in 05-284r2, Self-Describing Command Timeouts, revision 2. In particular, an Application Client using the QUERY TASK Task Management Function extended service response may decide what kind of command timeout value to use and when to start it. For this, it is expected that Application Clients would issue a QUERY TASK Task Management Function at least once after submitting the task to the LU for execution.

Suggested changes

(next page)

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 7.10, 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 34.

Table 34 — 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 ACA	I_T_L	7.4
CLEAR TASK SET	I_T_L	7.5
I_T NEXUS RESET	I_T	7.6
LOGICAL UNIT RESET	I_T_L	7.7
QUERY TASK	I_T_L_Q	7.8

Argument descriptions:

- Nexus:** An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.7) identifying the task or tasks affected by the task management function.
- I_T Nexus:** A SCSI initiator port and SCSI target port nexus (see 4.7).
- I_T_L Nexus:** A SCSI initiator port, SCSI target port, and logical unit nexus (see 4.7).
- I_T_L_Q Nexus:** A SCSI initiator port, SCSI target port, logical unit, and task tag nexus (see 4.7).

Service Response:

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

- FUNCTION COMPLETE:** A task manager response indicating that the requested function is complete. Unless another response is required, the task manager shall return this response upon completion of a task management request supported by the logical unit or SCSI target device to which the request was directed.
- FUNCTION SUCCEEDED:** An optional task manager response indicating that the requested function is supported and completed successfully. This task manager response shall only be used by functions that require notification of success (e.g., QUERY TASK).
- TASK ENABLED:** [An optional QUERY TASK Task Management Function task manager response indicating that the task referenced by the I T L Q nexus is in TASK ENABLED state. Implies FUNCTION SUCCEEDED task manager response.](#)
- TASK BLOCKED:** [An optional QUERY TASK Task Management Function task manager response indicating that the task referenced by the I T L Q nexus is in TASK BLOCKED state. Implies FUNCTION SUCCEEDED task manager response.](#)
- TASK DORMANT:** [An optional QUERY TASK Task Management Function task manager response indicating that the task referenced by the I T L Q nexus is in TASK DORMANT state. Implies FUNCTION SUCCEEDED task manager response.](#)

TASK ENDED:	An optional QUERY TASK Task Management Function task manager response indicating that the task referenced by the I_T_L_Q nexus is in TASK ENDED state. Implies FUNCTION SUCCEEDED task manager response.
FUNCTION REJECTED:	An task manager response indicating that the requested function is not supported by the logical unit or SCSI target device to which the function was directed.
INCORRECT LOGICAL UNIT NUMBER:	An optional task router response indicating that the function requested processing for an incorrect logical unit number.
SERVICE DELIVERY OR TARGET FAILURE:	The request was terminated due to a service delivery failure (see 3.1.122) or SCSI target device malfunction. The task manager may or may not have successfully performed the specified function.

Each SCSI transport protocol standard shall define the events comprising each of these service responses.

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 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;
- 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
- The task management function service response shall not be affected by the presence of access restrictions.

7.2 ABORT TASK

Request:

Service Response = ABORT TASK (IN (I_T_L_Q Nexus))

Description:

This function shall be supported by all logical units.

The task manager shall abort the specified task, if any, as described in 5.6.2. Previously established conditions, including MODE SELECT parameters, reservations, and ACA shall not be changed by the ABORT TASK function.

A response of FUNCTION COMPLETE shall indicate that the task was aborted or was not in the task set. In either case, the SCSI target device shall guarantee that no further requests or responses are sent from the task.

All SCSI transport protocol standards shall support the ABORT TASK task management function.

7.3 ABORT TASK SET

Request:

Service Response = ABORT TASK SET (IN (I_T_L Nexus))

Description:

This function shall be supported by all logical units.

The task manager shall abort all tasks in the task set that were received on the specified I_T nexus as described in 5.6. Tasks received on other I_T nexuses or in other task sets shall not be aborted. This task management function performed is equivalent to a series of ABORT TASK requests.

Other previously established conditions, including MODE SELECT parameters, reservations, and ACA shall not be changed by the ABORT TASK SET function.

All SCSI transport protocol standards shall support the ABORT TASK SET task management function.

7.4 CLEAR ACA

Request:

Service Response = CLEAR ACA (IN (I_T_L Nexus))

Description:

This function shall be supported by a logical unit if it supports ACA (see 5.2).

For the CLEAR ACA task management function, the task set shall be the one defined by the TST field in the Control mode page (see SPC-3).

An application client requests a CLEAR ACA using the faulted I_T nexus (see 3.1.39) to clear an ACA condition from the task set serviced by the logical unit. The state of all tasks in the task set shall be modified as described in 8.8. For a task with the ACA attribute (see 8.6.5) receipt of a CLEAR ACA function shall have the same effect as receipt of an ABORT TASK function (see 7.2) specifying that task. If successful, this function shall be terminated with a service response of FUNCTION COMPLETE.

If the task manager clears the ACA condition, any task within that task set may be completed subject to the requirements for task set management specified in clause 8.

The service response for a CLEAR ACA request received from an I_T nexus other than the faulted I_T nexus shall be FUNCTION REJECTED.

All SCSI transport protocol standards shall support the CLEAR ACA task management function.

7.5 CLEAR TASK SET

Request:

Service Response = CLEAR TASK SET (IN (I_T_L Nexus))

Description:

This function shall be supported by logical units.

For the CLEAR TASK SET task management function, the task set shall be the one defined by the TST field in the Control mode page (see SPC-3).

All tasks in the task set shall be aborted as described in 5.6.

All pending status and sense data for the task set shall be cleared. Other previously established conditions, including MODE SELECT parameters, reservations, and ACA shall not be changed by the CLEAR TASK SET function.

All SCSI transport protocol standards shall support the CLEAR TASK SET task management function.

7.6 I_T NEXUS RESET

Request:

Service Response = I_T NEXUS RESET (IN (I_T_L 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.

7.7 LOGICAL UNIT RESET

Request:

Service Response = LOGICAL UNIT RESET (IN (I_T_L Nexus))

Description:

This function shall be supported by all logical units.

Before returning a FUNCTION COMPLETE response, the logical unit shall perform the logical unit reset functions specified in 6.3.3.

NOTE 1 - Previous versions of this standard only required LOGICAL UNIT RESET support in logical units that supported hierarchical logical units.

All SCSI transport protocol standards shall support the LOGICAL UNIT RESET task management function.

7.8 QUERY TASK

Request:

Service Response = QUERY TASK (IN (I_T_L_Q Nexus))

Description:

SCSI transport protocols may or may not support QUERY TASK and may or may not require logical units accessible through SCSI target ports using such transport protocols to support QUERY TASK.

The task manager shall return a response of FUNCTION SUCCEEDED if the specified task is present in the task set, or FUNCTION COMPLETE if the specified task is not present in the task set.

Optionally, the task manager may return a response of TASK ENABLED, TASK BLOCKED, TASK DORMANT or TASK ENDED as shown in table 35.

Table 35 — QUERY TASK Task Management Function extended service response

TASK ENABLED	The task referenced by the I T L Q nexus is in the Enabled state. Implies FUNCTION SUCCEEDED service response.
TASK BLOCKED	The task referenced by the I T L Q nexus is in the Blocked state. Implies FUNCTION SUCCEEDED service response.
TASK DORMANT	The task referenced by the I T L Q nexus is in the Dormant state. Implies FUNCTION SUCCEEDED service response.
TASK ENDED	The task referenced by the I T L Q nexus is in the Ended state. Implies FUNCTION SUCCEEDED service response.

7.9 Task management function lifetime

The task manager shall create a task management function upon receiving a Task Management Request Received indication (see 7.10). The task management function shall exist until:

- the task manager sends a SCSI transport protocol service response for the task management function;
- an I_T nexus loss (see 6.3.4);
- a logical unit reset (see 6.3.3);
- a hard reset (see 6.3.2); or
- a power on condition (see 6.3.1).

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

- A service response of FUNCTION COMPLETE, FUNCTION SUCCEEDED, FUNCTION REJECTED, or SERVICE DELIVERY OR TARGET FAILURE is received for that task management function;
- Notification of a unit attention condition with 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_T NEXUS LOSS OCCURRED); or
- Notification of a unit attention condition with an additional sense code of COMMANDS CLEARED BY POWER LOSS NOTIFICATION.

NOTE 2 - The names of the unit attention conditions listed in the subclause (e.g., SCSI BUS RESET OCCURRED) are based on usage in previous versions of this standard. The use of these unit attention condition names is not to be interpreted as a description of how the unit attention conditions are represented by any given SCSI transport protocol.

7.10 Task management SCSI transport protocol services

The SCSI transport protocol services described in this subclause are used by a SCSI initiator device and SCSI target device to process a task management procedure call. The following arguments are passed:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.7).

Function Identifier: Argument encoding the task management function to be performed.

All SCSI transport protocol standards shall define the SCSI transport protocol specific requirements for implementing the **Send Task Management Request**, the **Task Management Request Received** indication, the **Task Management Function Executed** response, and the **Received Task Management Function Executed** confirmation SCSI transport protocol services described in this subclause.

A SCSI transport protocol standard may specify different implementation requirements for the **Send Task Management Request** SCSI transport protocol service for different values of the Function Identifier argument.

All SCSI initiator devices shall implement the **Send Task Management Request** and the **Received Task Management Function Executed** confirmation SCSI transport protocol services as defined in the applicable SCSI transport protocol standards.

All SCSI target devices shall implement the **Task Management Request Received** indication and the **Task Management Function Executed** response SCSI transport protocol services as defined in the applicable SCSI transport protocol standards.

Request sent by an application client:

Send Task Management Request (IN (Nexus, Function Identifier))

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.7).

Function Identifier: Argument encoding the task management function to be performed.

Indication received by the task manager:

Task Management Request Received (IN (Nexus, Function Identifier))

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.7).

Function Identifier: Argument encoding the task management function to be performed.

Response from task manager:

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

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.7).

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

FUNCTION COMPLETE: The requested function has been completed.

FUNCTION SUCCEEDED: The requested function is supported and completed successfully.

TASK ENABLED,
TASK BLOCKED,
TASK DORMANT,
TASK ENDED
[See table 35.](#)

FUNCTION REJECTED: The task manager does not implement the requested function.

INCORRECT LOGICAL UNIT NUMBER:	An optional task router response indicating that the function requested processing for an incorrect logical unit number.
SERVICE DELIVERY OR TARGET FAILURE:	The request was terminated due to a service delivery failure (see 3.1.122) or SCSI target device malfunction. The task manager may or may not have successfully performed the specified function.

Confirmation received by application client:

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

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.7).

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

FUNCTION COMPLETE:	The requested function has been completed.
FUNCTION SUCCEEDED:	The requested function is supported and completed successfully.
	TASK ENABLED, TASK BLOCKED, TASK DORMANT, TASK ENDED
	See table 35.
FUNCTION REJECTED:	The task manager does not implement the requested function.
INCORRECT LOGICAL UNIT NUMBER:	An optional task router response indicating that the function requested processing for an incorrect logical unit number.
SERVICE DELIVERY OR TARGET FAILURE:	The request was terminated due to a service delivery failure (see 3.1.122) or SCSI target device malfunction. The task manager may or may not have successfully performed the specified function.

Each SCSI transport protocol shall allow a **Received Task Management Function Executed** confirming completion of the requested task to be associated with the corresponding **Send Task Management Request**.

7.11 Task management function example

Figure 42 shows the sequence of events associated with a task management function.

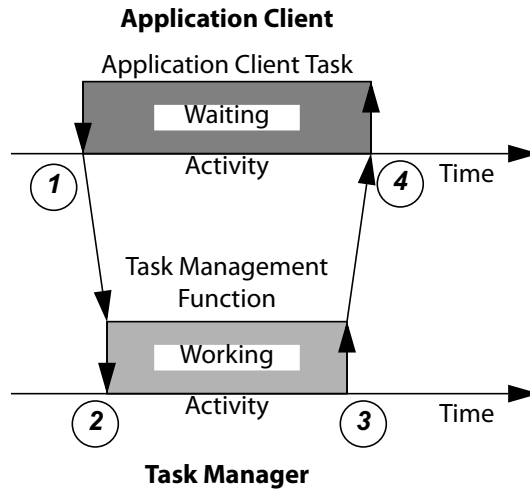


Figure 42 — Task management processing events

The numbers in figure 42 identify the events described as follows:

- 1) The application client task issues a task management request by invoking the **Send Task Management Request** SCSI transport protocol service.
- 2) The task manager is notified through a **Task Management Request Received** and begins processing the function.
- 3) The task manager performs the requested function and responds by invoking the **Task Management Function Executed** SCSI transport protocol service to notify the application client. The service response argument is set to a value of FUNCTION COMPLETE.
- 4) A **Received Task Management Function Executed** confirmation is received by the application client task.