To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 17 March 2007
Subject: 07-066r1 SAM-4 SAS-2 QUERY TASK SET task management function

Revision history

Revision 0 (14 February 2007) First revision
Revision 1 (17 March 2007) Incorporated comments from March 2007 FCP and CAP WG. Split the FCP-4 changes into 07-143.

Related documents

sam4r08 - SCSI Architecture Model - 4 (SAM-4) revision 8
sas2r08 - Serial Attached SCSI - 2 (SAS-2) revision 8
fcp4r00 - Fibre Channel Protocol - 4 (FCP-4) revision 0
07-067 - SAM-4 SAS-2 QUERY UNIT ATTENTION task management function (Rob Elliott, HP)
07-072 - FCP-4 QUERY TASK task management function (Rob Elliott, HP)
07-143 - FCP-4 QUERY TASK SET task management function (Rob Elliott, HP)
07-144 - FCP-4 QUERY UNIT ATTENTION task management function (Rob Elliott, HP)

Overview

Just as ABORT TASK has a companion task management function called ABORT TASK SET that aborts all tasks in the task set from the specified I_T nexus, QUERY TASK should be joined by a task management function called QUERY TASK SET that reports FUNCTION SUCCEEDED if there is any task in the task set from the specified I_T nexus.

QUERY TASK SET can be used by an initiator to determine if all its tasks were aborted (e.g., due to a persistent reservations PREEMPT AND ABORT, an I_T nexus loss, CLEAR TASK SET, LOGICAL UNIT RESET, or a hard reset) without having to query each task one-by-one.

The proposed semantics parallel ABORT TASK SET rather than CLEAR TASK SET. QUERY TASK SET only queries tasks in the task set from the specified I_T nexus, not tasks from any I_T nexus in the task set.

Changes are proposed for SAM-4 and SAS-2. FCP-4 changes are proposed in 07-143. iSCSI could adopt this change when it is upgraded to SAM-4 from SAM-2. The ADT editor is including this in the ADT-3 work list.

Suggested changes to SAM-4

4.5.8 SCSI task router class

The SCSI task router class routes information (e.g., commands and task management functions) between a logical unit and a service delivery subsystem using the route task operation.

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

a) Commands addressed to a valid logical unit are routed to the task manager in the specified logical unit;
b) Commands addressed to an incorrect logical unit are handled as described in 5.8.4;
c) Task management functions with I_T_L nexus scope (e.g., ABORT TASK SET, CLEAR TASK SET, QUERY 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 are routed to the task manager in the specified logical unit;
d) 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
e) Task management functions with I_T_L nexus scope or I_T_L_Q nexus scope addressed to an incorrect logical unit are handled as described in 7.10.

In some transport protocols, the task router may check for overlapped task tags on commands (see 5.8.3).

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_L	7.6
LOGICAL UNIT RESET	I_T_L	7.7
QUERY TASK	I_T_L_Q	7.8
QUERY TASK SET	I_T_L	7.xx

...

One of the following SCSI transport protocol specific service 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 or QUERY TASK SET).

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 3.1.120) 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) A task management request of ABORT TASK, ABORT TASK SET, CLEAR ACA, I_T NEXUS RESET, ~~or~~ QUERY TASK, or QUERY TASK SET 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.

7.8 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.~~

The task manager in the specified logical unit shall return a service response as follows:

a) if the specified task is present in the task set, FUNCTION SUCCEEDED; and
b) if the specified task is not present in the task set, FUNCTION COMPLETE.

7.xx QUERY TASK SET

Request:

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

Description:

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

The task manager in the specified logical unit shall return a service response as follows:

a) if there is any task present in the task set from the specified I_T nexus, FUNCTION SUCCEEDED; and
b) if there is no task present in the task set from the specified I_T nexus, FUNCTION COMPLETE.

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_L_Q nexus (e.g., an ABORT TASK or QUERY TASK task management function (see SAM-4)) to any PL_PM state machine until this state has received one of the following messages for each Tx Frame message with the same I_T_L_Q 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).

This state shall not send a Tx Frame message containing a TASK frame for a task that only affects an I_T_L nexus (e.g., an ABORT TASK SET, CLEAR TASK SET, QUERY TASK SET, CLEAR ACA, or LOGICAL UNIT RESET task management function (see SAM-4)) to any PL_PM state machine until this state has received one of the following messages for each Tx Frame message with the same I_T_L 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).

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-4)) to any PL_PM state machine 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 frame - Task Management Function information unit

Table 152 defines the TASK MANAGEMENT FUNCTION field.

Table 152 — TASK MANAGEMENT FUNCTION **field**

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-4). [a]
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-4). [a]
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-4). [a]
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-4). [a]
10h	I_T NEXUS RESET	no	no	The task manager shall perform the I_T NEXUS RESET task management function (see SAM-4). [a]
20h	Reserved			
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-4). [a]
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-4). [a]
81h	QUERY TASK SET	yes	no	The task manager shall perform the QUERY TASK SET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-4). [a]
82h	*QUERY UNIT ATTENTION (proposed in 07-xxxr0)*			
All others	Reserved			

[a] 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.

If the TASK MANAGEMENT FUNCTION field contains a reserved or unsupported value, the task manager shall return a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and its RESPONSE CODE field set to TASK MANAGEMENT FUNCTION NOT SUPPORTED.

If the TASK MANAGEMENT FUNCTION field is set to ABORT TASK or QUERY TASK, the TAG OF TASK TO BE MANAGED field specifies the TAG value from the COMMAND frame that contained the task to be aborted or checked. For all other task management functions, the TAG OF TASK TO BE MANAGED field shall be ignored.

10.2.1.1 SCSI transport protocol services overview

...

An application client requests the processing of a SCSI command by invoking SCSI transport protocol services, the collective operation of which is conceptually modeled in the following remote procedure call (see SAM-4):

Service response = Execute Command (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Data-In Buffer Size], [Data-Out Buffer], [Data-Out Buffer Size], [Task Priority]), OUT ([Data-In Buffer], [Sense Data], [Sense Data Length], Status))

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-4):

a) Service Response = ABORT TASK (IN (Nexus));
b) Service Response = ABORT TASK SET (IN (Nexus));
c) 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))~~.~~; and
h) Service Response = QUERY TASK SET (IN (Nexus)).

SSP defines the transport protocol services required by SAM-4 in support of the these remote procedure calls.

Table 180 describes the mapping of the remote procedure calls to transport protocol services and the SSP implementation of each transport protocol service.

Table 180 — SCSI architecture mapping

Remote procedure call	Type of transport protocol service	Transport protocol service interaction	Transport protocol service	I/T [a]	SSP implementation
Execute Command	Request/ Confirmation	Request	Send SCSI Command	I	COMMAND frame
		Indication	SCSI Command Received	T	Receipt of the COMMAND frame
		Response	Send Command Complete	T	RESPONSE frame
		Confirmation	Command Complete Received	I	Receipt of the RESPONSE frame or problem transmitting the COMMAND frame
	Data-In Transfer [b]	Request	Send Data-In	T	Read DATA frames
		Confirmation	Data-In Delivered	T	Receipt of ACKs for the read DATA frames
	Data-Out Transfer [b]	Request	Receive Data-Out	T	XFER_RDY frame
		Confirmation	Data-Out Received	T	Receipt of write DATA frames
	Terminate Data Transfer [b]	Request	Terminate Data Transfer	T	
		Confirmation	Data Transfer Terminated	T	
ABORT TASK, ABORT TASK SET, CLEAR ACA, CLEAR TASK SET, I_T NEXUS RESET, LOGICAL UNIT RESET, ~~and~~ QUERY TASK, and QUERY TASK SET	Request/ Confirmation	Request	Send Task Management Request	I	TASK frame
		Indication	Task Management Request Received	T	Receipt of the TASK frame
		Response	Task Management Function Executed	T	RESPONSE frame
		Confirmation	Received Task Management Function Executed	I	Receipt of the RESPONSE frame or problem transmitting the TASK frame

[a] I/T indicates whether the SSP initiator port (I) or the SSP target port (T) implements the transport protocol service.
[b] Data transfer transport protocol services for SCSI initiator ports are not specified by SAM-4.

10.2.1.12 Send Task Management Request transport protocol service

An application client uses the Send Task Management Request transport protocol service request to request that an SSP initiator port transmit a TASK frame.

Send Task Management Request (IN (Nexus, Function Identifier, [Association]))

Table 191 shows how the arguments to the Send Task Management Request transport protocol service are used.

Table 191 — Send Task Management Request transport protocol service arguments

Argument	SAS SSP implementation
Nexus	I_T_L nexus or I_T_L_Q nexus (depending on the Function Identifier), where: a) I specifies the initiator port to send the TASK frame; b) T specifies the target port to which the TASK frame is sent; c) L specifies the LOGICAL UNIT NUMBER field in the TASK frame header; and d) Q (for an I_T_L_Q nexus) specifies the TAG OF TASK TO BE MANAGED field in the TASK frame header.
Function Identifier	Specifies the TASK MANAGEMENT FUNCTION field in the TASK frame. Only these task management functions are supported: a) ABORT TASK (Nexus argument specifies an I_T_L_Q Nexus); b) ABORT TASK SET (Nexus argument specifies an I_T_L Nexus); c) CLEAR ACA (Nexus argument specifies an I_T_L Nexus); d) CLEAR TASK SET (Nexus argument specifies an I_T_L Nexus); e) I_T NEXUS RESET (Nexus argument specifies an I_T Nexus); f) LOGICAL UNIT RESET (Nexus argument specifies an I_T_L Nexus); ~~and~~ g) QUERY TASK (Nexus argument specifies an I_T_L_Q Nexus)~~.~~; and h) QUERY TASK SET (Nexus argument specifies an I_T_L Nexus).
[Association]	Specifies the TAG field in the TASK frame header.

10.2.1.13 Task Management Request Received transport protocol service

An SSP target port uses the Task Management Request Received transport protocol service indication to notify a task manager that it has received a TASK frame.

Task Management Request Received (IN (Nexus, Function Identifier, [Association]))

Table 192 shows how the arguments to the Task Management Request Received transport protocol service are determined.

Table 192 — Task Management Request Received transport protocol service arguments

Argument	SAS SSP implementation
Nexus	I_T_L nexus or I_T_L_Q nexus (depending on the Function Identifier), where: a) I indicates the initiator port that sent the TASK frame; b) T indicates the target port that received the TASK frame; c) L indicated by the LOGICAL UNIT NUMBER field in the TASK frame header; and d) Q (for an I_T_L_Q nexus) indicated by the TAG OF TASK TO BE MANAGED field in the TASK frame header.
Function Identifier	Indicates the TASK MANAGEMENT FUNCTION field in the TASK frame. Only these task management functions are supported: a) ABORT TASK (Nexus argument ~~specifies~~indicates an I_T_L_Q Nexus); b) ABORT TASK SET (Nexus argument ~~specifies~~indicates an I_T_L Nexus); c) CLEAR ACA (Nexus argument ~~specifies~~indicates an I_T_L Nexus); d) CLEAR TASK SET (Nexus argument ~~specifies~~indicates an I_T_L Nexus); e) I_T NEXUS RESET (Nexus argument ~~specifies~~indicates an I_T Nexus); f) LOGICAL UNIT RESET (Nexus argument ~~specifies~~indicates an I_T_L Nexus); ~~and~~ g) QUERY TASK (Nexus argument ~~specifies~~indicates an I_T_L_Q Nexus)~~.~~; and h) QUERY TASK SET (Nexus argument indicates an I_T_L Nexus).
[Association]	Indicates the TAG field in the TASK frame header.