

To: T10 Technical Committee
 From: Bob Sheffield(robert.l.sheffield@intel.com)
 Date: 1 June 2006
 Subject: 06-179r1: SAT - Fix Task Management Functions

Revision history

Revision 0 (30 March 2006) First revision

Revision 1 (1 June 2006) Add descriptions for handling collateral aborts w/ queued & non-queued commands

Related documents

SAT-r08 - SCSI / ATA Translation revision 08

06-121r0 SAT-r08_LB_Comment_Resolution

Overview

The the comments in SAT subclause 6.3 Task Management Functions were extensive and to some degree conflicting. This proposal outlines a compromise that attempts to address the combined comments.

Suggested changes

Add the following definitions

3.1.1 ATA abort retry: A policy implemented by a SATL whereby the SATL retries ATA commands aborted by ATA collateral abort once.

3.1.2 ATA volatile settings: ATA device settings affecting the way an ATA device responds to ATA commands that are configurable using ATA commands (e.g., ATA SET FEATURES or ATA SET MAX EXT), and that are set by the SATL to correspond to SCSI mode parameters, log parameters, or INQUIRY data.

*Editor's Note 1: Perhaps what's really relevant here are ATA *volatile settings*? What the SATL needs to do is configure any ATA volatile settings that are needed to emulate the current settings of corresponding SCSI mode parameters, log parameters, or INQUIRY data.*

3.1.3 ATA collateral abort: An ATA command that is aborted as a side effect of a different command being aborted when an ATA device is processing queued commands (i.e., NCQ or TCQ).

3.1.4 ATA hardware reset: AThe routines performed by the ATA device server and the ATA device port in an ATA device after a hardware reset event occurs (see ATA8-AAM). The hardware reset routines performed by the ATA device include the actions performed by the ATA device for a software reset (see ATA-AAM), and the actions defined in ATA8-ACS and the applicable ATA transport standards.

3.1.5 ATA software reset: A reset that is triggered by a task management function request (see ATA8-AAM).

3.1.6 service response: The device service response or SCSI transport protocol specific service response returned to an application client by the SATL on completion of a SCSI transport protocol service request (see SAM-3).

Modify subclause 6.2.2 as follows:

6.2.2 Mapping of SCSI queued commands to ATA queued commands

A SATL that translates SCSI tagged tasks to an ATA device using **SATA** NCQ or **ATA** TCQ, whether or not the SATL also queues commands internally, shall either:

- a) **report indicate** support for the basic task management model in **SCSI** standard INQUIRY data (i.e., the BQUE bit is set to one and CMDQUE bit is set to zero), and follow the rules for the basic task management model (see SAM-3); or

- ~~b) report support for the full task management model in SCSI standard INQUIRY data (i.e., the BQUE bit is set to zero and CMDQUE bit is set to one), and report 01b in the Queue error management (QERR) field of the SCSI Control mode page.~~
- b) indicate support for the full task management model in standard INQUIRY data (i.e., the BQUE bit is set to zero and CMDQUE bit is set to one), and set the QERR (Queue error management) field of the Control mode page (see 10.1.4) as follows:
 - A) a value of 01b if the SATL does not reissue ATA queued commands aborted by the ATA device due to an error condition on any one of the ATA queued commands; or
 - B) a value other than 01b if the SATL reissues all other ATA queued commands (i.e. except the one in error) aborted by the ATA device due to an error condition on any one of the ATA queued commands

~~A SATL that supports SATA NCQ or ATA TCQ may report support for the full task management model with a QERR field set to a value other than 01b only if the SATL reissues all queued commands aborted by the ATA device due to an error condition reported by the ATA device on any one of the queued commands.~~

~~Error conditions with outstanding commands to an attached ATA device generally affect all outstanding commands being processed by the ATA device. See ATA/ATAPI-7 or SATA-2 for a description of how to determine the status of each command.~~

Editor's Note 2: The previous paragraph (strikeout) is being moved (modified) to subclause 6.2.5 below.

For each SCSI tagged task the SATL translates to an ATA device using SATA NCQ or ATA TCQ, the SATL shall allocate an ~~inactive~~ available tag value (e.g. for NCQ, the value corresponding to ~~an available bit in the reserved field of~~ the position of a bit set to zero in the SActive register). The SATL shall maintain a mapping between allocated NCQ or TCQ tags and the corresponding SCSI task tags.

The SATL shall detect the maximum queue depth supported by the ATA device (i.e., check word 75 in IDENTIFY DEVICE data), and may either:

- a) ~~report~~ return a status of TASK SET FULL in response to a SCSI command issued to the corresponding emulated SCSI ~~device~~ logical unit when the ATA device represented has the maximum number of queued commands outstanding; or
- b) queue the command internally and return TASK SET FULL status when the SATL exhausts ~~it's~~ internal queueing resources.

Add subclauses 6.2.5 and 6.2.6 as follows (new text):

6.2.5 Collateral abort with queued commands

Error conditions with outstanding commands to an ATA device terminate all outstanding ATA commands being processed by the ATA device. An ATA host determines the status and error for each outstanding ATA queued command affected by the error condition and which ATA command(s) caused the error(s) (see ATA8-ACS or SATA 2.5). The SATL shall process aborted ATA commands as shown in table 1.

Table 1 — SATL processing of ATA commands aborted by ATA collateral abort

Association between the aborted ATA command and the ATA command that caused the error		Value of the QERR field set in the Control mode page	Method applied by the SATL for processing the aborted ATA command
I_T_L_Q nexus	I_T nexus		
same		00b	The SATL shall terminate the affected I_T_L_Q nexus with CHECK CONDITION status with the sense key and sense code set according to the reported ATA error as described in link-> clause 11.
		01b	
different	same	01b	The SATL shall terminate the affected I_T_L_Q nexus, but the SATL shall neither return status for the I_T_L_Q nexus affected by the aborted ATA command, nor retry the aborted ATA command.
		00b	The SATL shall reissue the ATA command and continue processing the corresponding I_T_L_Q nexus.
different		01b	The SATL shall terminate the affected I_T_L_Q nexus and establish a unit attention condition (see SAM-3) for the affected initiator port with the additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

6.2.6 Collateral abort with non-queued commands

The SATL should not send an ATA command to an ATA device while the ATA device is processing a non-queued command unless the intent is to abort a non-queued command being processed by the ATA device. Doing so causes the ATA device to abort both the issued ATA command and the non-queued command already being processed.

If an ATA non-queued command terminates with an error, the SATL shall terminate the affected I_T_L_Q nexus with CHECK CONDITION status with the sense key and sense code set according to the reported ATA error as described in [link->](#)clause 11.

Replace subclause 6.3 Task Management Functions with the following text:

6.3 Task management functions

6.3.1 Task management functions overview

6.3 describes the translation of SCSI task management functions to ATA equivalents.

6.3.2 Aborting ATA Commands

Some task management functions processed by the SATL may result in ATA commands aborted by ATA collateral abort (see 3.1.2) affecting an I_T_L_Q nexus other than the I_T_L_Q nexus(es) specified in the task management function request. The subclause defining the translation for each task management function defines how the SATL processes the I_T_L_Q nexuses affected by the task management function.

Processing some task management functions requires the SATL to abort one or more ATA commands being processed by an ATA device.

The SATL shall abort an ATA queued command being processed by an ATA device by sending an ATA CHECK POWER MODE command to the ATA device.

6.3.3 ABORT TASK

The service request for the ABORT TASK task management function is (see SAM-3):

Service Response = ABORT TASK (IN (I_T_L_Q nexus)).

If no ATA commands associated with the I_T_L_Q nexus specified in the ABORT TASK task management function are outstanding to the ATA device, then the SATL shall abort the task for the specified I_T_L_Q nexus from the SATL internal context and respond to the ABORT TASK task management function with a service response of FUNCTION COMPLETE (see SAM-3).

If the ATA device is processing one or more ATA commands that are related to the specified I_T_L_Q nexus, then the SATL shall either

- a) allow the ATA command(s) to complete as follows:
 - 1) wait until the ATA device returns command complete for the ATA command(s);
 - 2) if the completed ATA command completes processing of the specified I_T_L_Q nexus, then return completion status for the I_T_L_Q nexus; and
 - 3) return a service response of FUNCTION COMPLETE for the ABORT TASK task management function regardless of whether or not completion status was returned for the I_T_L_Q nexus;
 or
- b) abort the ATA command(s) (see 6.3.2) for the specified I_T_L_Q nexus and respond to the ABORT TASK task management function with a service response of FUNCTION COMPLETE.

If aborting the ATA commands related to the specified I_T_L_Q nexus results in one or more other ATA commands being aborted by ATA collateral abort (see 3.1.2), then the SATL shall:

- a) if the SATL supports ATA abort retry (see 3.1.1), then re-issue all ATA commands aborted by ATA collateral abort (see 3.1.2) and continue processing the affected I_T_L_Q nexuses; or

Editor's Note 3: Deleted the list item describing what to do if the TAS bit is set to one because the description of the Control mode page requires the TAS bit to be set to zero.

- b) if the SATL does not support ATA abort retry, then for each I_T nexus affected by an ATA command aborted by ATA collateral abort:
 - 1) complete processing of all but one of the SCSI tasks without returning a function result; and
 - 2) complete processing of the last SCSI task by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

6.3.4 ABORT TASK SET

The service request for the ABORT TASK SET task management function (see SAM-3) is:

Service Response = ABORT TASK SET (IN (I_T_L nexus)).

If the ATA device is not processing ATA commands for SCSI tasks associated with the specified I_T_L nexus, then the SATL shall abort all tasks for the specified I_T_L nexus from the SATL internal context and respond to the ABORT TASK SET task management function with a service response of FUNCTION COMPLETE.

If the ATA device is processing any ATA commands related to the specified I_T_L nexus, then the SATL shall either:

- a) allow the ATA command(s) to complete as follows:
 - 1) wait until the ATA device returns command complete for the ATA command(s);
 - 2) if the completed ATA command completes processing a SCSI task in the task set, return completion status for the SCSI task; and
 - 3) after all ATA commands return completion status, return a service response of FUNCTION COMPLETE for the ABORT TASK SET task management function;
- or
- b) abort outstanding ATA command(s) for the specified I_T_L nexus, and respond to the ABORT TASK SET task management function with a service response of FUNCTION COMPLETE.

If aborting ATA commands for the specified I_T_L nexus results in ATA commands aborted by ATA collateral abort that are related to processing SCSI commands in an I_T_L nexus other than the specified I_T_L nexus, then:

- a) if the SATL supports ATA abort retry, then the SATL shall re-issue all ATA commands aborted by ATA collateral abort and continue processing of the affected I_T_L_Q nexuses; or

Editor's Note 4: Deleted the list item describing what to do if the TAS bit is set to one because the description of the Control mode page requires the TAS bit to be set to zero.

- b) if the SATL does not support ATA abort retry, then for each I_T_L nexus other than the specified I_T_L nexus that had one or more SCSI tasks affected due to ATA commands aborted by ATA collateral abort, the SATL shall abort all commands for each affected I_T_L nexus and establish a UNIT ATTENTION condition with the additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

6.3.5 CLEAR ACA

The service request for the CLEAR ACA task management function (see SAM-3) is:

Service Response = CLEAR ACA (IN (I_T_L nexus)).

The SATL shall process the CLEAR ACA task management function as defined in SAM-3.

6.3.6 CLEAR TASK SET

The service request for the CLEAR TASK SET task management function (see SAM-3) is:

Service Response = CLEAR TASK SET (IN (I_T_L nexus)).

The SATL shall process the CLEAR TASK SET task management function in accordance with the full task management model (see SAM-3), and in accordance with a single task set that includes SCSI tasks for all I_T_L nexuses (i.e., the TST field in the Control mode page is set to 000h, see 10.1.4).

If the ATA device is processing any ATA commands, then the SATL shall either:

- a) allow the ATA command(s) to complete as follows:
 - 1) wait until the ATA device returns command complete for the ATA command(s); and
 - 2) if the completed ATA command completes processing a SCSI task in the task set, return completion status for the SCSI task;
- or
- b) abort all outstanding ATA command(s).

The SATL shall abort all tasks in the task set and respond to the CLEAR TASK SET task management function with a service response of FUNCTION COMPLETE.

If the SATL aborts tasks in the task set for an I_T_L nexus other than the specified I_T_L nexus, then for each other I_T_L nexus, the SATL shall establish a UNIT ATTENTION condition with the additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

[Editor's Note 5: Deleted the list item describing what to do if the TAS bit is set to one because the description of the Control mode page requires the TAS bit to be set to zero.](#)

6.3.7 LOGICAL UNIT RESET

The service request for the LOGICAL UNIT RESET task management function (see SAM-3) is:

Service Response = LOGICAL UNIT RESET (IN (I_T_L nexus)).

The SATL shall:

- 1) issue an ATA software reset (see 3.1.5) or an ATA hardware reset (see 3.1.4) to the ATA device associated with the logical unit for the specified I_T_L nexus;
- 2) abort all tasks in the task set from the SATL internal context;
- 3) restore ATA volatile settings (see 3.1.2) to values consistent with the emulation of current values of mode parameters, log parameters, and INQUIRY data; and
- 4) return a service response of FUNCTION COMPLETE for the LOGICAL UNIT RESET task management function.

NOTE 1 - If more than one PATA device is present on a PATA bus, issuing an ATA software reset causes both devices to be reset.

6.3.8 QUERY TASK

The service request for the QUERY TASK task management function (see SAM-3) is:

Service Response = QUERY TASK (IN (I_T_L_Q nexus)).

If the SATL supports the QUERY TASK task management function, the SATL shall return a service response of FUNCTION SUCCEEDED if the specified I_T_L_Q nexus is in the task set, or the SATL shall return a service response of FUNCTION COMPLETE if the specified I_T_L_Q nexus is not in the task set.

If the SATL does not support the QUERY TASK task management function the SATL shall return a service response of FUNCTION REJECTED.

6.3.9 Obsolete reset task management functions

The obsolete TARGET RESET (a.k.a., BUS DEVICE RESET) task management function is sometimes used by a SCSI application client to cause a hard reset (i.e., similar to a power-on condition) for each logical unit of a specified target device. The SATL may process the TARGET RESET task management function by issuing an ATA hard reset (see 3.1.9) to the ATA device(s) associated with the target device.

Replace the description of the QErr bit in table 65 - Control mode page fields in subclause 10.1.4 Control mode page as follows:

~~If the SATL resubmits queued commands that did not fail to the drive on behalf of any I_T Nexus, this field shall be set to 00b. Otherwise, the SATL shall set this field to 01b and comply with the unit attention requirements for a task completed with CHECK CONDITION status (see SPC-3)~~

If the SATL supports the full task management model and ATA abort retry (see 3.1.x) of ATA queued commands (see 3.1.x) aborted by ATA collateral abort (see 3.1.x), the SATL shall set this field to 00b. Otherwise, the SATL shall set this field to 01b and comply with the unit attention requirements for a task completed with CHECK CONDITION status (see SPC-3).