

To: T10 Technical Committee
From: Steve Johnson, LSI Logic (sjohnson@lsil.com)
Date: 10 August, 2005
Subject: 05-108r1 SAT-Task Management

Revision history

Revision 0 (9 July 2005) First revision

[Revision 1 \(18 July 2005\) Cleaned up text](#)

[Revision 2 \(10 August 2005\) Incorporated changes discussed in July 26 conference call](#)

Related documents

satr04 - SCSI/ATA Translations 1.0 revision 4

Overview

Add support for ATA Queued feature set, Informative ~~bus-resents~~[BUS RESET](#), QUERY TASK, and address editors notes for task management section.

Suggested changes

See text changes below.

1 Task Management Model

1.1 Overview

SAT implementations may support the full task management model or the basic task management model as defined by SAM-3 (~~see~~), depending on whether the SATL ~~makes use of~~ supports SATA native command queuing (NCQ), ~~and ATA queued feature set (TCQ), and/or~~ the tagged task management capabilities of the SATL itself. ~~Elements-Features~~ of the task management model (e.g. SIMPLE task attribute and ORDERED task attribute) may be provided through the use of SATA NCQ, ATA TCQ, and/or queueing within the SATL itself, ~~or a combination thereof~~.

1.2 Queued ~~commands~~ Commands

1.2.1 Overview

SCSI provides for various modes of command queuing via task tags and ORDERED, SIMPLE, HEAD OF QUEUE, and ACA task attributes. SCSI devices do not report the maximum number of commands that may be queued, but report a status of BUSY or TASK SET FULL when a command is received and the ~~queue has no space~~ device does not have enough resources to process another command. ~~SATA NCQ supports up to 32 commands at a time. The SATL has the option of mapping SCSI queued commands directly to SATA native queued commands and/or queueing commands internally.~~

1.2.2 Mapping of SCSI ~~queued commands~~ tagged tasks to SATA ~~native queued commands~~ NCQ and ATA TCQ

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

- a) report support for the basic task management model in SCSI standard INQUIRY data (BQue bit is 1 and CmdQue bit is 0), and follow the rules for the basic ~~SCSI command queuing~~ task management model (see SAM-3); or
- b) report support for the full task management model in SCSI standard INQUIRY data (BQue bit is 0 and CmdQue bit is 1), and report a value of 01b in the Queue error management (QErr) field of SCSI ~~standard INQUIRY data~~. Control Mode page

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

~~Editor's Note 1: Is there a problem if data has already been transferred for an outstanding command that was aborted by the SATA device due to an error in another command?~~

Error conditions with outstanding commands to an attached ATA device generally effect 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.

For each ~~queued~~ SCSI ~~command~~ tagged task the SATL ~~forwards~~ translates to a ~~SATA~~ ATA device using ~~NCQ~~ SATA NCQ or ATA TCQ, the SATL shall allocate an inactive ~~NCQ~~ tag value (e.g. for ~~NCQ~~, corresponding to an available bit in the reserved field of the ~~Set Device Bits FIS~~ Active 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 attached SATA devices (word 75 in IDENTIFY DEVICE), and ~~shall~~ may return TASK SET FULL STATUS in response to a SCSI command issued to the corresponding emulated SCSI device when the SATA device represented has the maximum number of queued commands outstanding; or the SATL may queue the command internally and return TASK SET FULL status when the SATL exhausts it's internal queueing resources.

~~The mapping of NCQ tags to SCSI task tags may involve mapping more than one NCQ tag to a single SCSI task tag, for example, in cases where processing of a SCSI command involves processing of multiple ATA commands by the attached non-packet device.~~

~~Editor's Note 2: May want to add text to address how to handle sending ATA commands that cannot be queued when the SATA device has queued commands outstanding (i.e. have to wait for queued commands to complete before issuing the non-NCQ command). Does this impose an implicit requirement that the SATL queue commands (at least one anyway) internally?~~

1.2.3 Commands the SATL queues internally

~~The SATL may only use SATA NCQ for READ and WRITE commands. Other new commands require are received by the SATL that cannot be queued by the ATA device the SATL to shall queue the commands internally, or report return a TASK SET FULL or BUSY condition to avoid receiving more commands status until the non-READ or WRITE command is complete queued commands have been completed. Furthermore, if the The SATL has outstanding READ or WRITE commands in process under NCQ, the SATL shall defer processing of the newly received non-READ or WRITE type command queued commands until the queued commands complete processing. The SATL shall perform task management in accordance with the reported task management model supported in accordance with the requirements in SAM-3.~~

If the SATL supports queued commands and the translation requires a mix of queued and non-queued commands the SATL shall defer processing of subsequent commands, complete processing of all outstanding queued commands, process the non-queued command, then continue processing the previously deferred commands.

1.2.4 Multi-initiator and multi-port command queuing

If the SATL is accessed through a SCSI target port the SATL may be accessible by more than one initiator port and through more than one SCSI target port. As specified in SAM-3, the task tags maintained in the SATL mapping of task tags to NCQ tags or TCQ tags shall be qualified by the ~~initiator ID and port ID~~ I_T nexus from which the command was received. When translating from an NCQ tag or TCQ tag to the corresponding SCSI task tag, the SATL shall determine the correct SCSI ~~port and initiator~~ I_T nexus using the qualification information associated with the SCSI task tag. The SATL may report TASK SET FULL even if the ~~SATA-ATA~~ device has available NCQ tags or TCQ tags in order to maintain tags available for other initiators.

1.3 Task Management Functions

This section describes the translation of SCSI task management functions to ATA or SATA equivalents.

NOTE 1 - Due to architectural differences, not all task management functions can be successfully translated to ATA commands or control operations.

1.3.1 ABORT TASK

The SATL may process the ABORT TASK service request ~~an~~ in any of the following ways:

- a) If no commands have yet been issued to the non-packet device for the processing of the specified SCSI task tag the SATL shall delete the specified task from the SATL internal context and respond to the ABORT TASK request with FUNCTION COMPLETE.
- b) If the only command(s) being processed in the non-packet device are related to the SCSI task specified by the ABORT TASK request, then the SATL may abort the ~~SATA-ATA~~ command(s) ~~(e.g. by issuing a NOP)~~, and respond with FUNCTION COMPLETE.
- c) If the non-packet device is processing one or more commands related to the SCSI task specified by the ABORT TASK request, and is also processing at least one other command related to a different SCSI task, and the SATL reports the full task management model and a QErr bit value other than 01b, then the SATL shall abort the command in a manner that does not prevent the processing of other

commands being processed by the attached non-packet device, and shall return FUNCTION COMPLETE status.

~~NOTE 2 – It may be that the first place the SATL can abort the command without preventing the processing of other commands is when the attached non-packet device returns completion status for the ATA command. In this case the SATL would not provide a SCSI response for the aborted command, even though processing completed in the ATA domain, but instead return FUNCTION COMPLETE for the ABORT TASK request.~~

- d) If the SATL reports support of the basic task management model, or reports support for the full task management model and a QErr field value of 01b, then the SATL ~~may abort outstanding ATA commands (e.g. by issuing a NOP), and shall respond to the ABORT TASK request with FUNCTION COMPLETE. The SATL shall return TASK ABORTED status for other SCSI tasks terminated because the corresponding ATA commands were aborted along with the one specified in the ABORT TASK request.~~

~~Editor's Note 3: Are there other conditions besides QErr that affect the behavior here?~~

- A) ~~If the SATL cannot abort the requested command, it shall Abort outstanding ATA commands and respond to the ABORT TASK request with SERVICE DELIVERY OR TARGET FAILURE.FUNCTION COMPLETE.~~

~~Editor's Note 4: Still need to specify non-packet device state after processing (e.g. restore-configuration options,...)~~

- B) Return TASK ABORTED status for other SCSI tasks terminated due to the corresponding ATA commands being aborted along with the one specified in the ABORT TASK request.

1.3.2 ABORT TASK SET

The SATL may handle the ABORT TASK SET service request differently depending on whether the SATL provides multiple initiators access to the emulated SCSI device or not. ~~If the SATL provides multiple initiator access to the emulated SCSI device, the SATL shall process the ABORT TASK SET service request as follows:~~

If the SATL does not provide multiple initiators access to the emulated SCSI device, the SATL shall process the service request as follows:

- a) If no commands have yet been issued to the non-packet device for the processing of the specified SCSI task tag the SATL shall delete all tasks in the task set from the SATL internal context and respond to the ABORT TASK SET request with FUNCTION complete.
- b) If the only command(s) being processed in the non-packet device are related to the SCSI tasks in the task set, then the SATL ~~may shall~~ abort the SATA command(s) ~~(e.g. by issuing a NOP),~~ and respond with FUNCTION complete.
- c) ~~If the~~ The SATL ~~reports support of the basic task management model, or reports support for the full task management model and a QErr field value of 01b, then the SATL may shall~~ abort outstanding ATA ~~commands~~ command(e.g. by issuing a NOP), s and ~~shall~~ respond to the ABORT TASK SET request with FUNCTION complete.
- d) ~~If the SATL cannot abort tasks in the task set, it shall respond to the service request with SERVICE DELIVERY OR TARGET FAILURE.~~

If the SATL provides multiple initiators access to the emulated SCSI device, the SATL shall process the service request as follows:

- a) ~~If the SATL supports the basic task management model, the SATL may abort outstanding ATA commands (e.g. by issuing a NOP), and respond with FUNCTION complete.~~
- b) If no commands have yet been issued to the non-packet device for the processing of the specified SCSI task tag the SATL shall delete the specified task from the SATL internal context and respond to the ABORT TASK request with FUNCTION COMPLETE.
- c) If the SATL supports the full task management model and reports a QErr value ~~of other than 01b~~, the SATL may either abort the outstanding ATA commands ~~(e.g. by issuing a NOP), or not~~ and ~~respond with FUNCTION complete~~ shall return SERVICE DELIVERY or TARGET FAILURE in response to the CLEAR TASK SET service request.
- d) ~~If the SATL supports the full task management model and reports a QErr value other than 01b, the SATL may either abort the outstanding ATA commands or not, and shall return SERVICE DELIVERY OR TARGET FAILURE in response to the CLEAR TASK SET service request.~~

NOTE 3 - Check with SAM-3 for behaviour when the SATL is processing commands for more than one I T L nexus.

1.3.3 CLEAR ACA

The SATL shall not support auto-contingent allegiance. The SATL shall indicate ACA is not supported by reporting a value of 0 in the NORMACA bit in standard INQUIRY ~~data. The SATL shall~~ data and respond to a CLEAR ACA service request with FUNCTION REJECTED.

1.3.4 CLEAR TASK SET

If the SATL reports support of the basic task management model or reports support of the full task management model with a QErr field value of 01b, the SATL shall reject the service request returning a status of FUNCTION REJECTED.

If the SATL reports support of the full task management model and a QErr value other than 01b, the SATL may ~~either~~ abort the outstanding ATA commands or not, and shall return SERVICE DELIVERY ~~OR~~ or TARGET FAILURE in response to the CLEAR TASK SET service request. The SATL may reset the device to clear any outstanding commands.

~~Editor's Note 5: Or perhaps the SATL should just reset the device?~~

1.3.5 LOGICAL UNIT RESET

LOGICAL UNIT RESET shall cause the SATL to issue a software reset (i.e., set the SRST bit to one in the Device Control register, then set the bit to zero) to the non-packet device representing the specified logical unit. Any persistent behaviors shall be reestablished by the SATL afterwards, including any behaviors related to saveable mode parameters.

~~Editor's Note 6: ATAPI may need special handling, but that issue is being deferred for now, and applies to all task management functions since ATAPI doesn't provide a standard way to send task management requests.~~

NOTE 4 - BUS RESET is commonly used by SCSI application layers to hard reset each device mapped to a target ID on a given SCSI bus. The SATL may translate the BUS RESET by issuing a protocol specific HARD RESET to each target device (e.g. SATA COM RESET or SAS PHY HARD RESET).

1.3.6 QUERY TASK

The SATL shall not support the QUERY TASK service request and shall return FUNCTION REJECTED to a QUERY TASK service request.

~~Editor's Note 7: Suggestion – Query task shouldn't be prohibited. At least one SCSI transport, SAS, requires a target device to support QUERY TASK, and so if the SATL is behind a SAS target port, there would be conflicting requirements.~~

~~Editor's Note 8: The original proposal included BUS RESET here. SAM-3 does not define "BUS RESET". This is a concept carried over from earlier SPI standards. The intent is to represent a HARD RESET condition, which is generated by transport specific means. SAS and SATA already define the behavior of HARD RESET. BUS RESET should be addressed in an informative annex. The editor is awaiting a formal proposal to add informative text for handling BUS RESET.~~

QUERY TASK shall cause the SATL to search for the specified task and if found, return a response of FUNCTION SUCCEEDED. If the specified task is not found the SATL shall return a response of FUNCTION COMPLETE.

1.4 SCSI Control Byte

1.4.1 CONTROL byte overview

~~Every SCSI CDB contains a CONTROL byte (see SAM-3).~~ Table 1 describes SATL handling of the CDB CONTROL byte. See SAM-3 for control byte details.

Table 1 — Control byte fields

Field	SATType ^a	Description or reference
Vendor specific(12h)	<u>U</u>	The SATL may use this field for vendor-specific purposes.
NACA	U	If set to one, the SATL shall return a <u>terminate the command with</u> CHECK CONDITION <u>status</u> with the sense key set to ILLEGAL REQUEST and <u>the</u> additional sense code set to INVALID FIELD IN CDB if a value of 1 is specified in this field <u>CDB</u> .
LINK	U	If set to one, the SATL shall return a CHECK CONDITION with the sense key set to ILLEGAL REQUEST and additional sense code set to INVALID FIELD IN CDB if a value of 1 is specified in this field <u>CDB</u> .

^a Key: I = implemented, E = emulated, U = unspecified (see 3.4)