

To: T10 Technical Committee
From: Steve Johnson, LSI Logic (sjohnson@lsil.com)
Date: 21 September, 2005
Subject: 05-108r3 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

Revision 3 (21 September 2005) Incorporated final changes as approved by working group September 13.

Related documents

satr05 - SCSI/ATA Translations 1.0 revision 5

Overview

Add support for ATA Queued feature set, Informative BUS RESET, QUERY TASK, and address editors notes for task management section.

Suggested changes

1 Task Management Model

1.1 Overview

SAT implementations may support the full task management model or the basic task management model as well as specific features of the task management model (e.g. SIMPLE and ORDERED task attributes) depending on the task management and queueing capabilities of the SATL and whether the SATL supports SATA native command queueing (NCQ) or the ATA queued feature set (TCQ).

1.2 Queued Commands

1.2.1 Overview

SCSI provides for various modes of command queueing 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 device does not have enough resources to process another command.

1.2.2 Mapping of SCSI tagged tasks to SATA NCQ and ATA TCQ

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 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 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 Control Mode page

A SATL that supports 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 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 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 SCSI tagged task the SATL translates to an ATA device using SATA NCQ or ATA TCQ, the SATL shall allocate an inactive tag value (e.g. for NCQ, corresponding to an available bit in the reserved field of 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 (word 75 in IDENTIFY DEVICE), and may either:

- a) report a status of TASK SET FULL in response to a SCSI command issued to the corresponding emulated SCSI device 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 its internal queueing resources.

1.2.3 Commands the SATL queues internally

When queued commands are outstanding to the ATA device and new commands are received by the SATL that cannot be queued by the ATA device the SATL shall queue the commands internally, or return a TASK SET FULL or BUSY status until at least one queued command has been completed. The SATL shall defer processing of the newly received non 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 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 I_T nexus using the qualification information associated with the SCSI task tag. The SATL may report TASK SET FULL even if the 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 in any of the following ways:

- a) If commands have not 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 a service response of 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, the SATL may abort the ATA command(s), and respond with a service response of FUNCTION COMPLETE.
- c) If the ATA device is processing commands for SCSI tasks in addition to task specified by the ABORT TASK request than the SATL shall:
 - A) abort outstanding ATA commands and respond to the ABORT TASK request with a service response of FUNCTION COMPLETE (see SAM3); and
 - B) for each initiator port associated with an I_T_Nexus that had a task aborted, the SATL shall complete at least one command for that I_T_Nexus with CHECK CONDITION status with the sense key set to UNIT ATTENTION and additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

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

- a) If commands have not 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 a service response of FUNCTION COMPLETE.
- b) If the only command(s) being processed in the non-packet device are associated to the SCSI tasks in the task set, then the SATL shall abort the SATA command(s) and respond with a service response of FUNCTION COMPLETE.
- c) The SATL shall abort outstanding ATA command(s) and respond to the ABORT TASK SET request with a service response of FUNCTION COMPLETE.

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

- a) If commands have not 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 SET request with a service response of FUNCTION COMPLETE.
- b) for each initiator port associated with an I_T_Nexus that had a task aborted, the SATL shall complete at least one command for that I_T_Nexus with CHECK CONDITION status with the sense key set to UNIT ATTENTION and additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

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 and respond to a CLEAR ACA service request with a service response of FUNCTION REJECTED.

1.3.4 CLEAR TASK SET

The SATL may handle the CLEAR TASK SET service request differently depending on whether the SATL provides multiple initiators access to the emulated SCSI device or not.

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 commands have not 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 CLEAR TASK SET request with a service response of FUNCTION COMPLETE.
- b) If the only command(s) being processed in the non-packet device are associated to the SCSI tasks in the task set, then the SATL shall abort the SATA command(s) and respond with a service response of FUNCTION COMPLETE.
- c) The SATL shall abort outstanding ATA command(s) and respond to the CLEAR TASK SET request with a service response of FUNCTION COMPLETE.

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

- a) If commands have not 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 CLEAR TASK SET request with a service response of FUNCTION COMPLETE.
- b) for each initiator port associated with an I_T_Nexus that had a task aborted, the SATL shall complete at least one command for that I_T_Nexus with CHECK CONDITION status with the sense key set to UNIT ATTENTION and additional sense code set to COMMANDS CLEARED BY ANOTHER INITIATOR.

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.

NOTE 2 - 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 COMRESET or SAS PHY HARD RESET).

NOTE 3 - If more than one device is present on a parallel ATA bus, issuing a soft reset causes both devices to be reset.

1.3.6 QUERY TASK

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

1.4 SCSI Control Byte

1.4.1 CONTROL byte overview

Table 4 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	U	The SATL may use this field for vendor-specific purposes.
NACA	U	If set to one, the SATL shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
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.
^a Key: I = implemented, E = emulated, U = unspecified (see 3.4)		

Editor's Note 1: Reminder to update Control byte fields table. This proposal made no modifications to this section
