

T10/07-402 Revision 1

Date: October 17, 2007

To: T10 Committee

From Brad Besmer, LSI

Subject: SAT-2 Add Translation for SATA NCQ Priority

Overview

There is a new capability in SATA NCQ command to specify a command as either a “normal” or “high” priority command. This is done using the PRIO bit defined in the READ FPDMA QUEUED and WRITE FPDMA QUEUED commands (refer to SATA specification revision 2.6).

Background

SAM-4:

4.11 Tasks and task tags

A task is represented by an I_T_L_Q nexus (see 4.12) and is composed of:

- a) A definition of the work to be performed by the logical unit in the form of a command or a group of linked commands;
- b) A task attribute (see 8.6) that allows the application client to specify processing relationships between various tasks in the task set; and
- c) Optionally, a task priority (see 8.7).

8.7 Task priority

Task priority specifies the relative scheduling importance of a task having a SIMPLE task attribute in relation to other tasks having SIMPLE task attributes already in the task set. If the task has a task attribute other than SIMPLE, the task priority is not used. Task priority is a value in the range of 0h through Fh. A task with either no task priority or a task priority set to 0h has a vendor-specific level of scheduling importance. A task with a task priority set to 1h has the highest scheduling importance, with increasing task priority values indicating decreasing scheduling importance. A task with a task priority set to Fh has the lowest scheduling importance.

If the Task Priority argument is set to zero or is not contained within the Send SCSI Received SCSI transport protocol service indication (see 5.4.2) and a priority has been assigned to the I_T_L nexus, the device server shall use that priority as the task priority. A priority may be assigned to an I_T_L nexus by a SET PRIORITY command (see SPC-3) or by the INITIAL PRIORITY field in the Control Extension mode page (see SPC-3). If no priority has been assigned to the I_T_L nexus using the SET PRIORITY command and the logical unit does not support the INITIAL PRIORITY field in the Control Extension mode page the device server shall set the task priority to 0h (i.e., vendor specific) or the task shall have no task priority.

A task manager may use task priority to determine an ordering to process tasks with the SIMPLE task attribute within the task set. A difference in task priority between tasks may not override other scheduling considerations (e.g., different times to access different logical block addresses) or vendor specific scheduling considerations. However, processing of a collection of tasks with different task priorities should cause the subset of tasks with the higher task priorities to return status sooner in aggregate than the same subset would if the same collection of tasks were submitted under the same conditions but with all task priorities being equal.

SAT-2 Changes

New section following section 6.2 Queued Commands:

6.3 Task Priority

A SATL that supports SATA native command queuing (NCQ) feature set may also support the SAM-4 Task Priority feature. SAM-4 Task Priority supports 16 priorities (0-15), whereas SATA NCQ only supports 2 priorities via the PRIO bits in the READ FPDMA QUEUED and WRITE FPDMA QUEUED commands. The SATL shall translate SAM-4 Task Priorities to SATA NCQ priority as shown in Table 4.

Table 4 —

SAM-4 Task Priority	SATA NCQ PRIO
0	0
1-3	1
4-15	0