To: T10 Technical Committee
From: Mark Overby, NVIDIA Corporation (moverby@nvidia.com)
Date: 20 February 2008
Subject: T10/08-018r1 SAT-2: NV Cache Translation

Revision History
Revision 0 - Initial draft of document
Revision 1 - Updated text to clarify certain actions as requested in December WG meeting. Revised text for clarity. Noted that this proposal will also need to be discussed in CAP for the mode page.

Related Documents
SAT-2 (T10/1699-D) Revision 1a

1 Overview

This proposal creates a translation from the ATA non-volatile cache feature set into a combination of the SBC non-volatile cache and SAT-specific mode pages to control the behavior of the cache.

Note - this draft is incomplete and is submitted for the purposes of discussion and early feedback. **In addition, this proposal will also need to be discussed by CAP for the assignment of a subpage code.**

2 Document Changes

2.1 Changes to SAT-2

2.1.1 Changes to parameters for SAT implementations

2.1.2 Non-Volatile Cache Log Page

The non-volatile cache mode page reports that a non-volatile cache is present and for how long the data remains volatile. If the attached ATA device reports support for the non-volatile cache feature set (see [Extended Inquiry VPD Page Translation]), the SATL shall implement the translation for this log page. As ATA non-volatile caches are required to remain non-volatile under all circumstances, the SATL shall report that the caches are indefinitely non-volatile as described in the following translation.

Table xx shows the translation of the fields for the returned parameter data. This translation shall be used for all defined non-volatile cache log parameters.

Table 1 — Non-volatile Cache Parameter Fields

Field	Changeable	Description or reference
PARAMETER LENGTH	n/a	Shall be set to 3h
REMAINING NON-VOLATILE TIME or MAXIMUM NON-VOLATILE TIME	n/a	Shall be set to FF_FFFFh

2.1.3 Changes to Vital product data parameters

2.1.4 Extended INQUIRY VPD Page

The SATL shall implement the Extended INQURY VPD Page if the ATA device reports support for the non-volatile cache feature set (i.e., IDENTIFY DATA word 214, bit 0 is set to 1) or the NCQ priority feature.

Table xx defines the translation of the Extended INQUIRY VPD Page.

Table 2 — Extended INQUIRY VPD Page

Field	Description or Reference
PERIPHERAL QUALIFIER	Shall be set to 000b
PERIPHERAL DEVICE TYPE	Shall be set to 00h
PAGE CODE	Shall be set to 86h
PAGE LENGTH	Shall be set to 3Ch
SPT	Unspecified (see 3.4.2)
GRD_CHK	Unspecified (see 3.4.2)
APP_CHK	Unspecified (see 3.4.2)
REF_CHK	Unspecified (see 3.4.2)
GROUP_SUP	Shall be set to zero
PRIOR_SUP	If the ATA device reports support for the NCQ priority feature (need reference), shall be set to one. Otherwise, shall be set to zero. See [NCQ Priority Translation].
HEADSUP	Unspecified (see 3.4.2)
ORDSUP	Unspecified (see 3.4.2)
SIMPSUP	Shall be set to one
COR_D_SUP	Shall be set to one if IDENTIFY DEVICE data word 119, bit 2 is set to one. Otherwise this field shall be set to zero.
NV_SUP	If the ATA device supports the non-volatile cache feature set (i.e. IDENTIFY DEVICE data word 214, bit 0 is set to one), this field shall be set to one. Otherwise, this field shall be set to zero.
V_SUP	If the ATA device write or read caches are enabled (i.e. IDENTIFY DATA word 85, bits 5 or 6 are set to one), this field shall be set to one. Otherwise, this field shall be set to zero.
LUICLR	Unspecified (see 3.4.2)

2.1.5 New SAT-specific SCSI extensions

2.1.6 ATA Non-Volatile Cache Control Mode Page

The ATA non-volatile cache control mode page provides ATA specific controls that allows an application client to control the range of logical block addresses that are resident in the non-volatile cache as well as query current characteristics of the non-volatile cache. This mode page shall be implemented by the SATL when the ATA device supports the non-volatile cache feature set, see [extended inquiry VPD page].

Table 3 — ATA Non-Volatile Cache Control Mode Page

Byte/Bit	7	6	5	4	3	2	1	0
0	PS	SPF (1B)	PAGE CODE (0AH)					
1	SUBPAGE CODE (F2H)							
2	(MSB)	PAGE LENGTH						
3								(LSB)
4		RANGE ALLOCATION LENGTH						(LSB)
5	(MSB)							
6	Reserved			ACTION			PWRMODE	ENABLE
7	Reserved							
8	COMMAND PARAMETER DATA							
n								

When processing a MODE SELECT command for this page, if the parameters savable (PS) bit is not 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 PARAMETER LIST.

The SPF bit (see SPC-4) shall be set to one to access this mode page.

The PAGE CODE field shall be set to 0Ah.

The SUBPAGE CODE field shall be set to F2h.

The PAGE LENGTH field shall be set to 8h plus the contents of the RANGE ALLOCATION LENGTH field.

The RANGE ALLOCATION LENGTH field shall be set to the length, in 512-byte units, of additional data transmitted with the mode page beginning at offset 8h. The field is in little-endian format as specified in ATA8-ACS.

When processing a MODE SENSE command for this page, ENABLE shall be set to one if the ATA device IDENTIFY DEVICE data word 214, bit 4 is set to one.

When processing a MODE SELECT command for this page, if ENABLE is set to zero, the SATL shall issue an NV CACHE DISABLE command to the device and all other fields in the mode page shall be ignored. If

ENABLE is set to one, and IDENTIFY DEVICE data word 214, bit 4 is set to zero, the SATL shall issue an NV CACHE ENABLE command to the device. action shall be set to 000b when changing the state of enable.

When processing a MODE SENSE command for this page, PWRMODE shall be set to one if the ATA device IDENTIFY DEVICE data word 214, bit 1 is set to one.

When processing a MODE SELECT command for this page, if PWRMODE is set to zero and IDENTIFY DEVICE data word 214, bit 1 is set to one, the SATL shall issue a RETURN FROM NV CACHE POWER MODE command to the device. If PWRMODE is set to one and IDENTIFY DEVICE data word 214, bit 1 is set to zero the SATL shall issue a SET NV CACHE POWER MODE command to the device with the minimum high-power time parameter set to the value as described in table xx. The application client passes a command parameter data block as described in table xx in command parameter data containing the minimum high-power time value for the NV cache power mode (see ATA8-ACS).

When processing a MODE SENSE command for this page, ACTION shall be set to 000b.

When processing a MODE SELECT command, ACTION shall be set to a value from table xx. This coding represents an action that the SATL and ATA device shall take for managing the contents of the non-volatile cache.

Table 4 — ACTION field coding

ACTION	Description
000b	Perform actions specified by ENABLE or PWRMODE
001b	Add LBA ranges to non-volatile cache, immediate
010b	Remove LBA ranges from non-volatile cache
011b	Query cache misses
100b	Query cache LBA list
101b	Unpin all LBAs from the non-volatile cache
110b	Add LBA ranges to non-volatile cache
All other codes reserved	

2.1.6.1 Add LBA ranges to non-volatile cache, immediate action

When this action is performed, the application client provides one or more non-volatile cache LBA elements (see table 5) in COMMAND PARAMETER DATA. Each element represents a LBA and a number of sequential LBAs that are to be added to the non-volatile cache pinned set. The application client shall not provide non-volatile cache LBA elements that represent LBA ranges that overlap.

The SATL shall issue an ATA ADD LBA(S) TO NON VOLATILE CACHE PINNED SET command with bit zero of the ATA LBA field set to one. The data sent with this command is the list of LBA ranges provided in the COMMAND PARAMETER DATA. If the number of ranges provided is not a multiple of 512 bytes, the SATL shall pad the end of the data with the required number of bytes, set to zero, that causes the length to become a multiple of 512 bytes.

Table 5 — Non-Volatile Cache LBA Element

Byte/Bit	7	6	5	4	3	2	1	0
0		LBA						(LSb)
6	(MSb)							
7		RANGE LENGTH						(LSb)
8	(MSb)							

2.1.6.2 Add LBA ranges to non-volatile cache action

When this action is performed, the SATL shall perfom the same actions as specified in 2.1.6.1, except the SATL shall set bit 0 of the ATA LBA field to zero when issuing the ATA ADD LBA(S) TO NON VOLATILE CACHE command.

2.1.6.3 Remove LBA ranges from non-volatile cache action

2.1.6.4 Query cache misses action

2.1.6.5 Query cache LBA list action

2.1.6.6 Unpin all LBAs action

2.1.6.7 Power mode command parameter data

Table 6 — NV Cache Power Mode command parameter data

Byte\Bit	7	6	5	4	3	2	1	0
0	(LSB)	MINIMUM HIGH-POWER TIME						
1								(MSB)

MAXIMUM HIGH-POWER TIME specifies the time, in seconds, that the device shall remain fully powered after accessing the media to satisfy a request.

2.1.7 Changes to SYNCHRONIZE CACHE (10) and SYNCHRONIZE CACHE (16)

Change Tables 43 and , field SYNC_NV as follows:

Table 7 — SYNCHRONIZE CACHE (10) CDB field translations

SYNC_NV	If sync_nv is set to one, and the SATL has reported a non-volatile cache in the extended INQUIRY VPD page, the SATL shall issue a FLUSH NV CACHE command before issuing an ATA flush command as required elsewhere in the translation for SYNCHRONIZE CACHE.