To: T10 Technical Committee
From: Mark Overby, NVIDIA Corporation (moverby@nvidia.com)
Date: 6 May 2008
Subject: T10/08-018r2 SAT-2: NV Cache Translation

Revision History
Revision 0 - Initial draft of document
Revision 1 - Revised proposal to use commands instead of doing through mode pages

Related Documents
SAT-2 (T10/1699-D) Revision 4

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 command codes for the operations to manage the NV cache.

2 Document Changes

2.1 Changes to SPC-4

Add the following to table D.4 (MAINTANENCE (IN))

Table 1 —

11h	SAT NON-VOLATILE CACHE CONTROL IN

Add the following to table D.4 (MAINTANENCE (OUT))

Table 2 —

11h	SAT NON-VOLATILE CACHE CONTROL OUT

2.2 Changes to SAT-2

2.2.1 Changes to parameters for SAT implementations

2.2.2 Non-Volatile Cache Log Page

The non-volatile cache log 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 3 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 3 — 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

2.2.3 Changes to Vital product data parameters

2.2.4 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 4 — 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)

Table 4 — Extended INQUIRY VPD Page

Field	Description or Reference
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 (i.e., IDENTIFY DEVICE data word 76, bit 12 is set to one), then this field shall be set to one. Otherwise, this field 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 or IDENTIFY DEVICE data word 214, bit 4 is set to one), then 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)

2.2.5 New SAT-specific SCSI extensions

2.2.6 SAT NON-VOLATILE CACHE CONTROL IN command

The SAT NON-VOLATILE CACHE CONTROL command (see Table ZZ) transfers data from the ATA device to the application client to obtain status or configuration information about the ATA non-volatile cache in the device.

Table 5 — SAT NON-VOLATILE CACHE CONTROL IN command

Byte / Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (A3h)							
1	Reserved			SERVICE ACTION (11h)				
2	NON-VOLATILE CACHE SERVICE ACTION							
3		RANGE ALLOCATION LENGTH						
4								
5		STARTING LBA						
10								

The NON-VOLATILE CACHE SERVICE ACTION field (see table ZZ) specifies which non-volatile cache operation is being used.

The RANGE ALLOCATION LENGTH field specifies the number of 512-byte data units that shall be transferred from the SATL to the application client for the requested service action. A value of zero in this field indicates that 65 536 512-byte data blocks are requested for transfer.

Table 6 — NON-VOLATILE CACHE SERVICE ACTION **field**

Code	Description	Reference
01h	Query Non-Volatile Cache Misses	XX
02h	Query Non-Volatile Cache Pinned Set	YY
All other code values reserved		

2.2.6.1 Non-volatile cache service actions

2.2.6.1.1 Query Non-Volatile Cache Misses

This non-volatile cache service action queries the non-volatile cache in the ATA drive for a list of the last 64 logical block addresses that were accessed and were not present in the non-volatile cache. When this service action is requested the RANGE ALLOCATION LENGTH field shall be set to 1 or the command shall be terminated with CHECK CONDITION status with a sense key of ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.

The STARTING LBA field is ignored for this non-volatile cache service action.

The SATL shall issue an ATA QUERY NV CACHE MISSES command to the ATA device and return the data from the ATA device.

2.2.6.1.2 Query Non-Volatile Cache Pinned Set

This non-volatile cache service action queries the non-volatile cache in the ATA drive for a list of the current logical block addresses that are resident in the non-volatile cache.

The STARTING LBA FIELD indicates the first LBA from which a scan for resident LBAs shall begin. The list of logical block addressed pinned in the non-volatile cache is always returned in ascending order (see ATA8-ACS).

For this service action, the SATL shall issue an ATA QUERY NV CACHE PINNED SET command to the ATA device with the ATA LBA field set to the value in the STARTING LBA field and with the ATA Count field set to the value in the RANGE ALLOCATION LENGTH field.

The SATL shall return the returned data from the ATA device to the host upon completion of this command.

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

Change Tables 43 and the 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: 1) Issue a FLUSH NV CACHE command to the ATA device with the LBA field set to FFFF_FFFFh; 2) Issue a FLUSH NV CACHE command to the ATA device with the LBA field set to the value in the ATA LBA field returned from the previous FLUSH NV CACHE command; 3) Issue a FLUSH_NV_CACHE command to the ATA device with the LBA field set to zero; and 4) Continue further processing as required by the SYNCHRONIZE CACHE (10) translation.