

To: T10 Technical Committee  
 From: Mark Overby, NVIDIA Corporation (moverby@nvidia.com)  
 Date: 14 January 2009  
 Subject: T10/08-018r4 SAT-3: NV Cache Translation

### **Revision History**

Revision 0 - Initial draft of document  
 Revision 1 - Revised proposal to use commands instead of doing through mode pages  
 Revision 2 - Beginning shift to command based mechanisms using MAINTANENCE IN (12)  
 Revision 3 - Added MAINTANENCE OUT (10) version as well as moved to MAINTANENCE IN (16)  
     Revised SPC-4 changes  
 Revision 4 - Moved to SERVICE ACTION only.  
     Incorporated requested review changes from Jan 2009 SAT WG meeting  
     Added missing SYNCHRONIZE CACHE (16) translation

### **Related Documents**

SAT-2 (T10/1699-D) Revision 6  
 SPC-4 (T10/1731-D) Revision 14

## **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.

## **2 Document Changes**

### **2.1 Changes to SPC-4 (r14)**

Replace Table D.6 with the following:

**Table D.6 — SERVICE ACTION IN (16) and SERVICE ACTION OUT (16) service actions**

<b>Service Action</b>	<b>Description</b>
SERVICE ACTION IN (16) [Operation Code 9Eh]	
00h - 0Eh	Reserved
0Fh	Restricted (See SAT-3)
10h - 1Fh	Restricted
SERVICE ACTION OUT (16) [Operation Code 9Fh]	
00h - 0Fh	Reserved
0Fh	Restricted (See SAT-3)
10h - 1Fh	Restricted

## 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 non-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 1 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

## 2.2.3 Changes to Vital product data parameters

### 2.2.4 Extended INQUIRY VPD Page

The SATL shall implement the Extended INQUIRY VPD Page (see table BB) if:

- a) The NV\_SUP bit is set to one; or
- b) The PRIOR\_SUP bit is set to one.

**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 (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

#### 2.2.6.1 Overview

The SAT NON-VOLATILE CACHE CONTROL IN 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 3 — SAT NON-VOLATILE CACHE CONTROL IN command**

Byte	7	6	5	4	3	2	1	0
Bit								
0	OPERATION CODE (9Eh)							
1	Reserved			SERVICE ACTION (0Fh)				
2	Reserved							
3								
4	(MSB)	STARTING LBA						(LSB)
9								
10	Reserved							
11								
12	(MSB)	RETURNED DATA LENGTH						(LSB)
13								
14	NON-VOLATILE CACHE SERVICE							
15	CONTROL							

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

**Table 4 — NON-VOLATILE CACHE SERVICE field**

Code	Description	Reference
00h	Query non-volatile cache misses service	XX
01h	Query non-volatile cache pinned set service	YY
All other code values reserved		

The RETURNED DATA LENGTH field specifies the number of 512-byte data units (i.e., a value of 1 means 512 bytes, two means 1 024 bytes, etc.) that shall be transferred in the data-in buffer for the requested service action. A value

of zero in this field indicates that 65 536 data units (i.e., 33 554 432 or 0200\_0000h bytes) are requested for transfer.

The STARTING LBA field contents vary based on the NON-VOLATILE CACHE SERVICE field (see table ZZ).

### 2.2.6.2 Non-volatile cache in services

#### 2.2.6.2.1 Query non-volatile cache misses service

The query non-volatile cache miss service 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.

If the NON-VOLATILE CACHE SERVICE field is set to 01h, then the returned data length field shall be set to 0001h. If the RETURNED DATA LENGTH field is set to any other value, then 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 shall be ignored for this non-volatile cache service action and should be set to 0000h by the application client.

The SATL shall issue an ATA QUERY NV CACHE MISSES command to the ATA device for this service and place the data from the ATA device into the data-in buffer.

#### 2.2.6.2.2 Query non-volatile cache pinned set service

The query non-volatile cache pinned set service 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 addresses pinned in the non-volatile cache shall be returned in ascending order (see ATA8-ACS).

For this service, 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 RETURNED DATA LENGTH field.

The SATL shall place the returned data from the ATA device into the data-in buffer.

### 2.2.7 SAT NON-VOLATILE CACHE CONTROL OUT command

#### 2.2.7.1 Overview

The SAT NON-VOLATILE CACHE CONTROL OUT command (see Table JJ) transfers data from the application client to the ATA device to control the ATA non-volatile cache.

**Table 5 — SAT NON-VOLATILE CACHE CONTROL OUT command**

Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (9Fh)							
1	Reserved	NVC_CTL	SERVICE ACTION (0Fh)					
2	Reserved							
3								
4	(MSB)	NVC SERVICE DATA						(LSB)
9								
10	Reserved							
11								
12	(MSB)	PARAMETER LIST LENGTH						(LSB)
13								
14	NON-VOLATILE CACHE SERVICE							
15	CONTROL							

The NON-VOLATILE CACHE SERVICE field (see table 6) specifies which non-volatile cache service is being performed.

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

Code	Description	Reference
00h	Enable non-volatile cache service	2.2.7.2.1
01h	Disable non-volatile cache service	2.2.7.2.2
02h	Enable non-volatile cache power mode service	2.2.7.2.3
03h	Disable non-volatile cache power mode service	2.2.7.2.4
04h	Add LBAs to pinned set service	2.2.7.2.5
05h	Remove LBAs pinned set service	2.2.7.2.6
06h	Reserve non-volatile cache space service	2.2.7.2.7
All other code values reserved		

The behavior of the NVC\_CTL bit and NVC SERVICE DATA field depends on the non-volatile cache service field (see table 6).

The PARAMETER LIST LENGTH field specifies the number of 512-byte data units (i.e., a value of 1 means 512 bytes, two means 1 024 bytes, etc.) that shall be transferred from the data-out buffer for the requested service action. A

value of zero in this field indicates that 65 536 data units (i.e., 33 554 432 or 0200\_0000h bytes) are requested for transfer.

### **2.2.7.2 Non-volatile cache out services**

#### **2.2.7.2.1 Enable non-volatile cache service**

This service requests that the SATL enable the non-volatile cache on an ATA device. No data is transferred from the data-out buffer for this command (i.e., the PARAMETER LIST LENGTH field and the data-out buffer shall be ignored).

The NVC\_CTL bit and NVC SERVICE DATA field are ignored for this service.

For this service, the SATL shall issue an ATA NV CACHE ENABLE command to the ATA device.

#### **2.2.7.2.2 Disable non-volatile cache service**

This service requests that the SATL disable the non-volatile cache on an ATA device. No data is transferred from the data-out buffer for this command (i.e., the PARAMETER LIST LENGTH field and the data-out buffer shall be ignored).

The NVC\_CTL bit and NVC SERVICE DATA field are ignored for this service.

For this service, the SATL shall issue an ATA NV CACHE DISABLE command to the ATA device.

#### **2.2.7.2.3 Enable non-volatile cache power mode service**

This service requests that the SATL enable the non-volatile cache power savings mode (see ATA8-ACS) on an ATA device. No data is transferred from the data-out buffer for this command (i.e., the PARAMETER LIST LENGTH field and the data-out buffer shall be ignored).

The NVC\_CTL bit is ignored for this service.

The NVC SERVICE DATA field contains the minimum value, in seconds, that the device shall stay in the ATA active or idle power states after entering the active state to access its media when the non-volatile cache power savings mode is active.

For this service, the SATL shall issue an ATA SET NV CACHE POWER MODE command to the ATA device with the Count field set to the value in the NVC SERVICE DATA field.

#### **2.2.7.2.4 Disable non-volatile cache power mode service**

This service requests that the SATL disable the non-volatile cache power savings mode (see ATA8-ACS) on an ATA device. No data is transferred from the data-out buffer for this command (i.e., the PARAMETER LIST LENGTH field and the data-out buffer shall be ignored).

The NVC\_CTL bit and NVC SERVICE DATA field are ignored for this service.

For this service, the SATL shall issue an ATA RETURN FROM NV CACHE POWER MODE command to the ATA device.

#### 2.2.7.2.5 Add LBAs to pinned set service

This service requests that the SATL add a list of logical block addresses and ranges, provided in the data-out buffer, to be permanently added to the non-volatile cache and that portion of the non-volatile cache be treated as media. This is referred to as the non-volatile cache pinned set (see ATA8-ACS). The format of the data in data-out buffer is described in ATA8-ACS.

The application client should ensure that the data-out buffer is a multiple of 512-bytes and padded with zeros to achieve this alignment. Failure to comply with the zero-value padding requirements may result in the attached ATA device to fail the operation.

If the size of the data-out buffer is not a multiple of 512, the SATL shall terminate the service with a CHECK CONDITION status with a sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.

The NVC SERVICE DATA field is ignored for this service.

The NVC\_CTL bit specifies, when set to one, that the ATA device not complete the ADD LBAS TO NV CACHE PINNED SET command until all requested LBAs have been made resident in the non-volatile cache.

For this service, the SATL shall issue an ATA ADD LBA(S) TO NV CACHE PINNED SET command to the ATA device with the ATA Count field set to the value in the PARAMETER LIST LENGTH field and bit zero of the ATA LBA field set to the value of the NVC\_CTL bit.

#### 2.2.7.2.6 Remove LBAs from pinned set service

This service requests that the SATL remove a list of logical block addresses and ranges, provided in the data-out buffer, to be removed from the non-volatile cache and that portion of the non-volatile cache no longer be treated as media. The format of the data in data-out buffer is described in ATA8-ACS.

The application client should ensure that the data-out buffer is a multiple of 512-bytes and padded with zeros to achieve this alignment. Failure to comply with the zero-value padding requirements may result in the attached ATA device to fail the operation.

If the size of the data-out buffer is not a multiple of 512, the SATL shall terminate the service with a CHECK CONDITION status with a sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.

The NVC SERVICE DATA field is ignored for this service.

The NVC\_CTL bit specifies, when set to one, that the ATA device remove all logical block addresses from the non-volatile cache pinned set. When this bit is set, the parameter list length field is ignored and no data-out buffer is used.

For this service, the SATL shall issue an ATA REMOVE LBA(S) FROM NV CACHE PINNED SET command to the ATA device with the ATA Count field set to the value in the PARAMETER LIST LENGTH field and bit zero of the ATA LBA field set to the value of the NVC\_CTL bit.

#### 2.2.7.2.7 Reserve non-volatile cache space service

This service requests that the SATL reserve a multiple of the logical block-size region in the non-volatile cache to add logical block addresses to the pinned set (see ATA8-ACS). No data is transferred from the data-out buffer for this command (i.e., the PARAMETER LIST LENGTH field and the data-out buffer shall be ignored).



The NVC\_CTL bit shall be ignored for this service.

The NVC SERVICE DATA field contains the number of logical blocks for which space is to be reserved in the non-volatile cache. The ATA device flushes logical blocks from its unpinned set, if necessary, in order to satisfy this request (see ATA8-ACS).

For this service, the SATL shall issue an ATA FLUSH NV CACHE command to the ATA device with the ATA LBA field set to the value in the NVC SERVICE DATA field.

**2.2.8 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	<p>If the SYNC_NV bit is set to one, and the SATL has reported a non-volatile cache in the extended INQUIRY VPD page, then the SATL shall:</p> <ol style="list-style-type: none"> <li>1) Issue an ATA FLUSH NV CACHE command to the ATA device with the ATA LBA field set to FFFF_FFFFh;</li> <li>2) Issue an ATA FLUSH NV CACHE command to the ATA device with the ATA LBA field set to the value in the ATA LBA field returned from the previous FLUSH NV CACHE command;</li> <li>3) Issue an ATA FLUSH NV CACHE command to the ATA device with the ATA LBA field set to zero; and</li> <li>4) Continue further processing as required by the SYNCHRONIZE CACHE (10) translation.</li> </ol>
---------	--

**Table 8 — SYNCHRONIZE CACHE (16) CDB field translations**

SYNC_NV	As defined in SYNCHRONIZE CACHE (10) (see x.xx).
---------	--