

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com) and Jeff Wolford, HP (jeff.wolford@hp.com)
Date: 26 April 2007
Subject: 07-200r0 SAT-2 WRITE LONG to WRITE UNCORRECTABLE EXT

Revision history

Revision 0 (26 April 2007) First revision

Related documents

sat2r01 - SCSI/ATA Translation (SAT-2) revision 0
sbc3r09 - SCSI Block Commands - 3 (SBC-3) revision 9
ata8-acs-r3f - AT Attachment 8 - ATA/ATAPI Command Set (ATA8-ACS) revision 3f

Overview

The SCSI WRITE LONG (10)/(16) command with the WR_UNCOR bit set to one should be translated into the ATA WRITE UNCORRECTABLE EXT command.

The ATA WRITE UNCORRECTABLE EXT command is in the General feature set (not the 48-bit Address feature set) and supports these fields:

- a) Opcode 45h
- a) Feature bits 7:0 Uncorrectable options
 - A) 55h Pseudo-uncorrectable error with logging
 - B) 5Ah Pseudo-uncorrectable error without logging
 - C) A5h Flagged error with logging
 - D) AAh Flagged error without logging

Convention:

- A) 5xh means reads to any sectors within the physical block of the specified sector cause errors. Drive shall perform normal error recovery to the fullest extent.
- B) Axh means only the logical sector is marked as bad. Drive apparently does not perform error recovery (since it is not mentioned as in the 5xh case).
- C) x5h means the sectors are listed as failed in error logs and SMART reports
- D) xAh means the sectors are not listed as failed in error logs and SMART reports
- b) Count - number of sectors to be marked. 0000h means 65,536 sectors.
- c) LBA - address of first sector to be marked

The SCSI WRITE LONG (10)/(16) command with the WR_UNCOR bit set to zero could be mapped to the ATA SMART WRITE LOG and WRITE LOG (DMA) EXT commands to log page E1h with the SCT WRITE LONG service action (action code 0001h, function code 0002h), but that is not included in this proposal. The promise of WRITE UNCORRECTABLE is to obsolete the need for access to the long data, which may not follow a simple "user data + ECC bytes" format as disk drives use more sophisticated encoding techniques.

Suggested changes

2.2 Approved references

ISO/IEC 24739-1, AT Attachment with Packet Interface - 7 (ATA/ATAPI-7) [ANSI INCITS 397-2005]
ISO/IEC 14776-413, SCSI Architecture Model - 3 (SAM-3) [ANSI INCITS 402-2005]
ISO/IEC 14776-412, SCSI Architecture Model - 2 (SAM-2) [ANSI INCITS 366-2003]
ISO/IEC 14776-453, SCSI Primary Commands - 3 (SPC-3) [ANSI INCITS 408-2005]
ISO/IEC 14776-322, SCSI Block Commands - 2 (SBC-2) [ANSI INCITS 405-2005]
ISO/IEC 14776-151, Serial Attached SCSI - 1.1 (SAS-1.1) [ANSI INCITS 417-2006]
ANSI INCITS TR38-2005, SMART Command Transport (SCT)

2.3 References under development

At the time of publication, the following referenced standards were still under development. For information on the current status of the document, or regarding availability, contact the relevant standards body or other organization as indicated.

ISO/IEC 14776-861, AT Attachment-8 Architecture Model (ATA8-AAM) [T13/1700D]

ISO/IEC 14776-871, AT Attachment-8 ATA/ATAPI Command Set (ATA8-ACS) [T13/1699D]

ISO/IEC 14776-881, AT Attachment-8 Parallel Transport (ATA8-APT) [T13/1698D]

ISO/IEC 14776-454, SCSI Primary Commands - 4 (SPC-4) [T10/1731-D]

[ISO/IEC 14776-323, SCSI Block Commands - 3 \(SBC-3\) \[T10/1799-D\]](#)

Editor's Note 1: Start to migrate all SBC-2 references to SBC-3. SBC-2 reference should only remain for features that are obsolete in SBC-3 and only defined in SBC-2.

4 General

...

This standard refers to behaviors for SCSI devices defined in ~~SBC-2~~[SBC-3](#) and SPC-3. Unless otherwise specified, any behaviors that are optional in ~~SBC-2~~[SBC-3](#) or SPC-3 are optional for devices implementing SCSI / ATA Translation. Any optional behaviors referred to in this standard and implemented by the SATL shall be implemented as described in this standard.

...

7 Summary of SCSI / ATA command mappings

7.1 Translated and emulated commands

In the event of a discrepancy between the contents of this clause and the description of individual commands, description of individual commands shall apply.

Clause 7, clause 8, and clause 9 describe the SCSI to ATA command mapping for ATA devices. Translation for ATAPI devices is described in ~~Annex~~ Annex A.

The SATL shall not issue more than one ATA command to the ATA device representing the logical unit with the exception of ATA queued commands (see 3.1.19). The SATL shall queue received SCSI commands as necessary to enforce this.

Table 8 lists the SCSI / ATA command mappings defined in this standard. A SATL may implement commands defined in SPC-3 and ~~SBC-2~~[SBC-3](#), but not listed in table 8. Translation of commands not listed in table 8 is vendor-specific.

Table 8 — Summary of SCSI / ATA Command Mapping

SCSI command	ATA command(s)	Reference
...		
WRITE LONG (10)	WRITE UNCORRECTABLE EXT	9.xx
WRITE LONG (16)		9.xy
...		

9 SCSI Block Commands (SBC) mapping

9.1 Translating LBA and transfer length and ATA command use constraints

9.1.1 Overview

A SATL may implement a direct block mapping of ATA logical sectors to SCSI logical blocks (see 9.1.2), or the SATL may implement indirect block mapping translation (see 9.1.3).

9.1.2 Direct block mapping model

If the SATL implements direct block mapping (see 3.1.33), the logical block size indicated by the BLOCK LENGTH IN BYTES field in the READ CAPACITY data (see 9.8.2 and 9.9.2) shall equal the ATA logical sector Size (see 3.1.15). The ATA LBA of an ATA logical sector shall equal the logical block address of the corresponding SCSI logical block.

9.1.3 Indirect block mapping model

If the SATL implements indirect block mapping (see 3.1.38), the constraints of the direct block mapping model do not apply. The logical block size indicated by the BLOCK LENGTH IN BYTES field in the READ CAPACITY data (see 9.8.2 and 9.9.2) may not equal the ATA logical sector size (see 3.1.15) (e.g., SCSI logical block size of 520 bytes with an ATA Logical Sector Size of 512 bytes). The SATL translates between the SCSI LOGICAL BLOCK ADDRESS field and the ATA LBA in a vendor-specific manner. The result of a logical block address translated in one direction and then translated in the reverse direction shall yield the original logical block address.

9.1.4 Selection of ATA block commands

The ATA commands the SATL may use to implement the functions specified by SCSI block commands depend upon:

- a) the value of the LOGICAL BLOCK ADDRESS and TRANSFER LENGTH fields specified in the SCSI CDB; and
- b) the capabilities of the ATA device and the ATA host within the SATL.

Table 26 relates selection conditions to allowable ATA commands used to implement SCSI block storage data transfer commands. ATA commands listed in the Allowed ATA commands column shall not be used in the emulation of a SCSI block command if the prerequisite conditions listed in Selection Prerequisites columns are not met (i.e., the word 'yes' in a Selection Prerequisites column means the prerequisite shall be met before

the SATL may use an ATA command listed in that row, and the word 'no' indicates the prerequisite need not be met for the SATL to use the ATA command listed).

Table 9 — ATA commands used for SCSI block command translations

Selection Prerequisites					Allowed ATA commands
Highest ATA logical sector accessed	ATA feature sets required to be supported and enabled				
Required that the logical sector address is < 2 ²⁸	48-bit Address	DMA	Overlap	SATA 2.5 NCQ	
no	no	no	no	no	FLUSH CACHE WRITE UNCORRECTABLE EXT
...					
no	yes	no	no	no	READ MULTIPLE EXT READ SECTOR(S) EXT READ VERIFY SECTOR(S) EXT WRITE MULTIPLE EXT WRITE MULTIPLE FUA EXT WRITE SECTOR(S) EXT
...					

...

The SATL may use the ATA commands listed in table 26 in the translation of SCSI read commands (see 3.1.77), SCSI write commands (see 3.1.81), SCSI write and verify commands (see 3.1.82), SCSI verify commands (see 3.1.80), and SCSI synchronize cache commands (see 3.1.78) if the prerequisites defined for the command as shown in table 26 are satisfied. The translations for specific SCSI block commands in clause 9 further constrain the use of the available ATA commands in implementing the translation.

The SATL emulation of the READ (6) command and the WRITE (6) command in which the TRANSFER LENGTH field is set to zero, shall translate the transfer length to 256, and issue ATA commands that operate on the ATA logical sectors corresponding to the specified 256 SCSI logical blocks.

In all other cases, the SATL shall transfer or operate on the ATA logical sectors corresponding to the number of logical blocks specified by the SCSI command.

9.xx WRITE LONG (10) command [\[all new\]](#)

The WRITE LONG (10) command (see SBC-3) requests that the SATL mark a logical block or physical block as containing an error.

Table 10 shows the translation of fields specified in the WRITE LONG (10) CDB.

Table 10 — WRITE LONG (10) CDB field translations

Field	Description or reference
OPERATION CODE	Set to 3Fh. If the ATA device does not support the WRITE UNCORRECTABLE EXT command (see ATA8-ACS), 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 OPERATION CODE.
COR_DIS	If the COR_DIS bit is not set to zero, 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.
WR_UNCOR	See table 3
PBLOCK	See table 3
LOGICAL BLOCK ADDRESS	The logical block address shall be used to set the ATA LBA (see 3.1.14) as defined by 9.1. If the SATL implements direct block mapping (see 3.1.33), then the SATL shall set the ATA LBA in the ATA read command (see 3.1.20) equal to the value specified in the LOGICAL BLOCK ADDRESS field. Otherwise, the mapping is unspecified (see 3.4.2).
BYTE TRANSFER LENGTH	If the BYTE TRANSFER LENGTH field is not set to zero, 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.
CONTROL	6.4

Editor's Note 2: The COR_DIS bit set to one could almost be mapped to AAh and 5Ah (errors without logging). In SCSI, the bit means the device server performs no error recovery, automatic reallocation, and does not report the errors as information exceptions. However, it requires that media access commands get CHECK CONDITION/MEDIUM ERROR/READ ERROR - LBA MARKED BAD BY APPLICATION CLIENT. The ATA device does not return any such indication, so SATL will always translate UNC=1 into CHECK CONDITION/MEDIUM ERROR/UNRECOVERABLE READ ERROR.

The interaction of the WR_UNCOR bit and PBLOCK bit are defined in table 11.

Table 11 — WR_UNCOR bit and PBLOCK bit

WR_UNCOR	PBLOCK	Description
0	0 or 1	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.
1	0	The SATL shall issue an ATA WRITE UNCORRECTABLE EXT command with: a) the Feature bits set to A5h (i.e., flagged error with logging); and b) the Count field set to 0001h.
1	1	The SATL shall issue an ATA WRITE UNCORRECTABLE EXT command with: a) the Feature bits set to 55h (i.e., pseudo-uncorrectable error with logging); and b) the Count field set to 0001h.

9.xy WRITE LONG (16) command [\[all new\]](#)

The WRITE LONG (16) command (see SBC-3) requests that the SATL mark a logical block or physical block as containing an error.

Table 12 shows the translation of fields specified in the WRITE LONG (16) CDB.

Table 12 — WRITE LONG (16) CDB field translations

Field	Descriptor or reference
OPERATION CODE/SERVICE ACTION	Set to 9Fh/11h.
COR_DIS	As defined in WRITE LONG (10) (see 9.xx)
WR_UNCOR	As defined in WRITE LONG (10) (see 9.xx)
PBLOCK	As defined in WRITE LONG (10) (see 9.xx)
LOGICAL BLOCK ADDRESS	As defined in WRITE LONG (10) (see 9.xx)
BYTE TRANSFER LENGTH	As defined in WRITE LONG (10) (see 9.xx)
CONTROL	6.4