

To: T10 Technical Committee
From: Bob Sheffield(robert.l.sheffield@intel.com)
Date: 18 August 2006
Subject: 06-378r0: SAT - Miscellaneous changes

Revision history

Revision 0 (18 August 2006) First revision

Related documents

SAT-r08c - SCSI / ATA Translation revision 08c
06-121r1 SAT-r08_LB_Comment_Resolution.pdf

Overview

I received (by e-mail) a number of comments on SAT-r08c. Most of them were editorial, but some of them (with suggested resolutions herein) need concurrence of the WG for integration.

Suggested changes

3.5.7 Use of field names defined in ATA standards and specifications

This standard discusses fields and values defined in other standards and specifications, in particular the ATA8-APT, ATA8-ACS, ATA8-AAM, and SCT standards developed by T13, and the SATA 2.5 specification. Such fields and values discussed in this standard are shown using the same notation conventions used in the standards where those fields and values are defined.

When this standard uses terms defined in T13 ATA standards or the SATA-2.5 specification, the following conventions apply:

- a) The names of abbreviations, commands, and acronyms used as signal names are in all uppercase (e.g., IDENTIFY DEVICE). Fields containing only one bit are usually referred to as the "name" bit instead of the "name" field.
- b) Names of device registers, fields in data structures, and other defined terms begin with an upper-case letter (e.g., LBA Mid register) and may be represented in mixed-case (e.g. PhyRdy).
- c) The expression "word n" or "bit n" shall be interpreted as indicating the content of word n or bit n.
- d) Bit names are shown in all uppercase letters ~~except where a lowercase n precedes a bit name. If there is no preceding n, then when BIT is set to one the meaning of the bit is true, and when BIT is cleared to zero the meaning of the bit is false. If there is a preceding n, then when nBIT is cleared to zero the meaning of the bit is true and when nBIT is set to one the meaning of the bit is false.~~
- e) Bit (n:m) denotes a set of bits, for example, bits (7:0).

5.6 Overview

Clause 5 defines SCSI / ATA translation elements that impact the representation of the storage domains defined in SAM-3 and ATA8-AAM. Figure 4 shows a SATL providing a communication path between a SCSI application client and an ATA device.

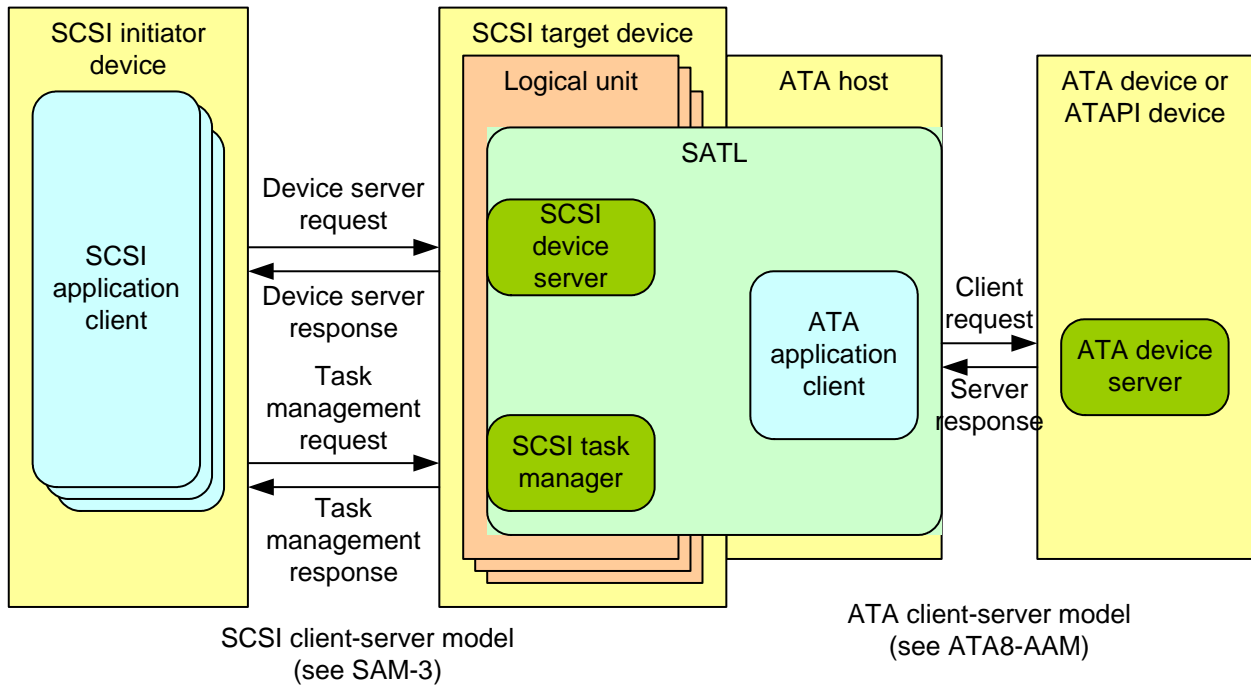


Figure 4 — Example of a SATL between a SCSI application client and an ATA or ATAPI device

Editor's Note 1: Replaced figure-4. This is more consistent with figures in SAM-3 and ATA8-AAM. Credit goes to Rob Elliott for creating this diagram.

The SATL provides the communication path between a SCSI application client and an ATA device or ATAPI device by:

- emulating a SCSI logical unit;
- integrating an ATA host; and
- providing the translation elements to link them together.

This standard defines SCSI / ATA translation using SCSI and ATA command sets. This standard does not define the mapping of transport capabilities as defined at the SCSI transport protocol layer and the ATA protocol interconnect layer.

An implementation utilizing a SATL may include a SCSI transport. A SATL may appear in different configurations:

EXAMPLE 1 - Figure 5 shows a SATL contained within a SCSI to ATA protocol bridge, where the ATA device is being accessed by an ATA host port, and the SATL is being accessed with a SCSI target port using a SCSI transport protocol (e.g, FCP-3 over Fibre Channel).

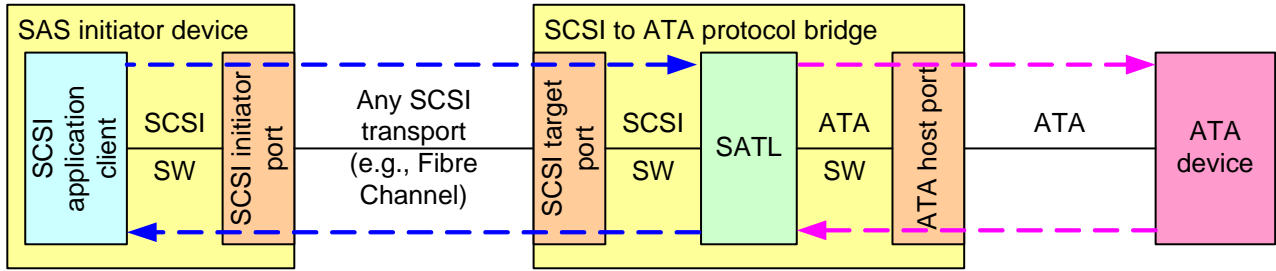


Figure 5 — SATL contained within a SCSI to ATA protocol bridge

EXAMPLE 2 - Figure 6 shows an ATA Host Bus Adapter (HBA) directly connected to an ATA device. The SATL provides SCSI transport protocol layer services to a SCSI application client in accordance with SAM-3.

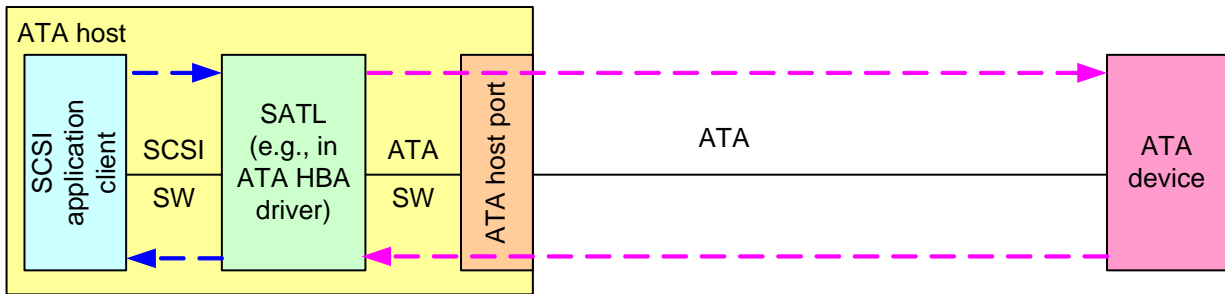


Figure 6 — SATL contained within an ATA host

EXAMPLE 3 - Figure 7 shows an ATA device accessed by a SAS STP initiator port (see SAS-1.1) through a SAS interconnect. The SAS initiator device includes a SATL to provide the SCSI transport protocol layer services to the application client in accordance with SAM-3.

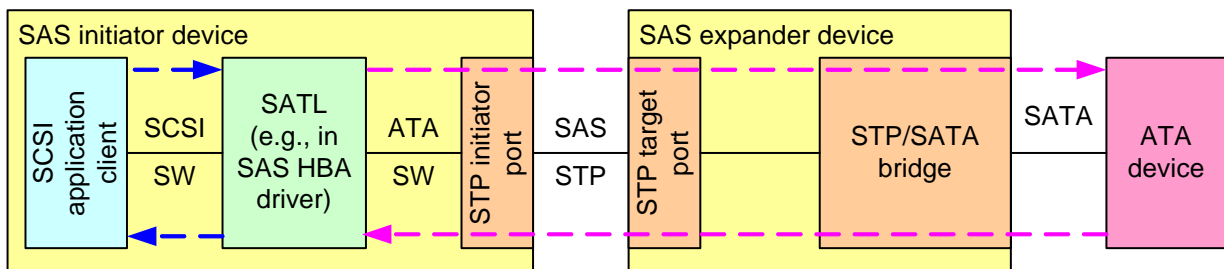


Figure 7 — SATL contained in a SAS initiator device

6.2.1 Comparison of SCSI and ATA queuing

Some differences between SCSI and ATA queuing methods are shown in table 5.

Table 5 — Comparison of SCSI and ATA queuing methods

| Feature | SCSI | NCQ | TCQ |
|--|---|--|--|
| Ordering | Specified by task attributes (e.g. SIMPLE, ORDERED) associated with each command. | Always at the discretion of the device. | Always at the discretion of the device. |
| Queue depth | Indeterminate | Fixed at 1 to 32 commands as reported in the ATA IDENTIFY DEVICE data word 75. | Fixed at 1 to 32 commands as reported in the ATA IDENTIFY DEVICE data word 75. |
| Queue full reporting | TASK SET FULL status | n/a | n/a |
| Queue full management | Device manages and indicates via TASK SET FULL status. | ATA host managed. | ATA host managed. |
| Queued commands | Queuing is applicable to all commands. | Limited to READ FPDMA QUEUED and WRITE FPDMA QUEUED commands. | Limited to READ DMA QUEUED, READ DMA QUEUED EXT, WRITE DMA QUEUED, WRITE DMA QUEUED EXT and WRITE DMA QUEUED FUA EXT commands, or a NOP command with a non-zero subcommand code. |
| Handling of non-queued commands received while one or more queued commands are being processed | n/a; all commands are queued commands. | Receipt of any command other than a READ FPDMA QUEUED command or a WRITE FPDMA QUEUED command is an error. | Receipt of any command other than a NOP command with a non-zero subcommand code, a SERVICE command, a READ DMA QUEUED command, a READ DMA QUEUED EXT command, a WRITE DMA QUEUED command, a WRITE DMA QUEUED EXT command, or a WRITE DMA QUEUED FUA EXT command is an error. |
| Error handling | Controlled with mode parameters. | Any error aborts all queued commands. | Any error aborts all queued commands. |

6.2.2 Mapping of SCSI queued commands to ATA queued commands

A SATL that translates SCSI tagged commands to an ATA device using NCQ or TCQ, whether or not the SATL also queues commands internally, shall either:

- a) indicate support for the basic task management model in standard INQUIRY data (i.e., the BQUE bit is set to one and CMDQUE bit is set to zero), and follow the rules for the basic task management model (see SAM-3); or
- b) indicate support for the full task management model in standard INQUIRY data (i.e., the BQUE bit is set to zero and CMDQUE bit is set to one), and set the QERR (i.e., queue error management) field of the Control mode page (see 10.1.5) as follows:
 - A) a value of 01b if the SATL does not reissue ATA queued commands aborted by the ATA device due to an error condition on any one of the ATA queued commands; or
 - B) a value other than 01b if the SATL reissues all other ATA queued commands (i.e., except the one in error) aborted by the ATA device due to an error condition on any one of the ATA queued commands.

For each SCSI tagged command that the SATL translates to an ATA [queued commands \(see 3.1.20\) device-using-NCQ-or-TCQ](#), the SATL shall allocate an available tag value (e.g., for NCQ, the value corresponding to the position of a bit set to zero in the SActive register). The SATL shall maintain a mapping between allocated [NCQ-or-TCQ ATA queued command](#) tags and the corresponding SCSI task tags.

The SATL shall use the maximum queue depth supported by the ATA device (i.e., indicated by IDENTIFY DEVICE data word 75), and may either:

- a) return a status of TASK SET FULL in response to a SCSI command issued to the corresponding emulated SCSI logical unit when the ATA device represented has the maximum number of ATA queued commands outstanding; or
- b) queue the command internally and return TASK SET FULL status when the SATL exhausts internal queuing resources.

6.2.3 Standard INQUIRY data

Table 10 describes the standard INQUIRY data fields supported by the SATL.

Table 10 — Standard INQUIRY data fields

| Field | Description or reference |
|--|---|
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |
| VERSION DESCRIPTOR 1 to VERSION DESCRIPTOR 8 | The SATL shall include version descriptors (see SPC-3) for: <ul style="list-style-type: none"> a) the SCSI Architecture Model standard (e.g., SAM-3); b) this standard; c) the SCSI Primary Commands standard (e.g., SPC-3); d) the SCSI Block Commands standard (e.g., SBC-2); e) if the SATL receives SCSI commands through a SCSI target port (see figure 5 in 5.1), the version of the transport protocol to which the SCSI target port was designed; f) if the SATL sends ATA commands through a SAS STP initiator port (see figure 7 in 5.1), the version of SAS (e.g., SAS-1.1) to which the SAS STP initiator port was designed; and g) the version(s) of ATA/ATAPI standards (e.g., ATA/ATAPI-7 ATA8-ACS and ATA8-AAM) or to which the ATA device claims compliance in the ATA IDENTIFY DEVICE data word 80 (i.e., Major version number) and ATA IDENTIFY DEVICE data word 81 (i.e., Minor version number).^d |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

9.19 WRITE (10) command

The WRITE (10) command is used to request the SATL to transfer user data from the application client to the ATA device. Data may be written to the medium or to the cache of the ATA device.

Table 47 shows the translation of fields specified in the WRITE (10) CDB.

Table 47 — WRITE (10) CDB field translations

| Field | Description or reference |
|---|--|
| OPERATION CODE | 9.17.1 and 9.17.2 |
| WRPROTECT | Unspecified (see 3.4.2) |
| DPO | Unspecified (see 3.4.2) |
| FUA | 9.17.2 |
| FUA_NV | <p>The SATL may ignore the FUA_NV bit, or the SATL may implement the FUA_NV bit as defined in SBC-2.</p> <p>If the SATL does not support the Extended INQUIRY Data VPD page (see SPC-3), the FUA bit is set to zero, and the FUA_NV bit is set to one, then 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. Otherwise, the SATL may implement the FUA_NV bit as defined in SBC-2.</p> <hr/> <p>Editor's Note 4: Make the same changes for READ(10/12/16) and WRITE(12/16).</p> <hr/> |
| 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 write command (see 3.1.25) equal to the value specified in the LOGICAL BLOCK ADDRESS field. Otherwise, the mapping is unspecified (see 3.4.2). |
| GROUP NUMBER | Unspecified (see 3.4.2) |
| TRANSFER LENGTH ^a | The transfer length is used to set the ATA Sector Count (see 3.1.21), as defined by 9.1. If the SATL implements direct block mapping (see 3.1.33), then the SATL shall set the ATA Sector Count in the ATA write command (see 3.1.25) equal to the value specified in the TRANSFER LENGTH field. Otherwise, the mapping is unspecified (see 3.4.2). The SATL shall issue as many ATA write commands (see 3.1.25) as needed to satisfy the transfer length specified by the WRITE (10) command. |
| CONTROL | 6.4 |
| ^a A transfer length of zero specifies that a data transfer shall not take place. | |

Editor's Note 5: The argument for this is that since SBC-2 defines the behavior for the case FUA=0, FUA_NV = 1 in terms of the NV_SUP bit, that a device server cannot "support" the combination FUA=0, FUA_NV=1 unless the Extended INQUIRY Data VPD page is supported. Then, since the bit combination (FUA, FUA_NV) = (0,1) represents a non-zero reserved code value in the field comprising the pair of bits, the device server must return a CHECK CONDITION status for that combination if the Extended INQUIRY Data VPD page is not supported.

10.3 Vital product data parameters

10.3.1 Vital product data parameters overview

Table 72 provides a summary of the VPD page translations defined in this standard.

Table 72 — Summary of SCSI / ATA VPD page mapping

| SCSI VPD page | Reference |
|--|---|
| Supported VPD Pages VPD page (i.e., 00h) | 10.3.2 |
| Unit Serial Number VPD page (i.e., 80h) | 10.3.3 |
| Device Identification VPD page (i.e., 83h) | 10.3.4 |
| Mode Page Policy VPD page (i.e., 87h) | 10.3.5 |
| ATA Information VPD page (i.e., 89h) | 12.4.2 |
| All others | See SPC-3 and SBC-2 Unspecified (see 3.4.2) |

Modify the operation code field descriptions for all commands to include both the code value and a brief description.

Table 9 — INQUIRY CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 12h. The SATL shall issue an ATA IDENTIFY DEVICE command to the ATA device. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 11 — LOG SENSE CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 4Dh. The SATL shall implement support for this field by returning the log page data for the particular page requested. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 14 — MODE SELECT (6) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 15h. The SATL emulates the function of this command. Some operational parameters in individual pages are provided via ATA. See clause 10 for specific requirements. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 15 — MODE SENSE (6) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 1Ah. The SATL emulates the function of this command. Some operational parameters in individual pages are gathered by issuing ATA commands (see 10.1). |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 16 — READ BUFFER CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 3Ch. The SATL shall issue the ATA READ BUFFER command to the ATA device. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 19 — REQUEST SENSE CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 03h. The SATL emulates the function of this command. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 20 — SEND DIAGNOSTIC CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 1Dh. The SATL emulates the function of this command. This field value is specific to the SEND DIAGNOSTIC command. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 23 — TEST UNIT READY CDB field translations

| Field | Description or reference |
|----------------|---|
| OPERATION CODE | Set to 00h. See 8.10.2. |
| CONTROL | 6.4 |

Table 24 — WRITE BUFFER CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 3Bh. The SATL shall: <ul style="list-style-type: none"> a) issue an ATA WRITE BUFFER command to the ATA device; b) issue an ATA DOWNLOAD MICROCODE command to the ATA device; or c) emulate the specified function (i.e., if supported); depending on the values in the BUFFER ID field and MODE field (see 8.11.2.1). |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 27 — FORMAT UNIT CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 04h. If no defect list header is provided or a defect list header is provided with the DCRT bit set to one, then the SATL shall return GOOD status without issuing any commands to the ATA device. If the SATL supports certification of media and a defect list header is provided with the DCRT bit set to zero, then the SATL shall certify the media as described in 9.2.5 |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 30 — READ (6) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 08h. See 9.3.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 31 — READ (10) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 28h. See 9.3.1 and 9.3.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 32 — READ (12) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to A8h. See 9.3.1 and 9.3.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 33 — READ (16) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 88h. See 9.3.1 and 9.3.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 34 — READ CAPACITY (10) CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 25h. The SATL shall use ATA IDENTIFY DEVICE data to compute the ATA device's maximum user addressable medium capacity of the ATA device. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 36 — READ CAPACITY(16) CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Shall be set Set to 9E/10h. The SATL shall use ATA IDENTIFY DEVICE data to compute the maximum user addressable medium capacity of the ATA device. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 38 — REASSIGN BLOCKS CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 07h. See 9.10.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 39 — START/STOP UNIT CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 1Bh. Commands issued to the ATA device depend upon the other values in the CDB as described in 9.11.3. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 41 — SYNCHRONIZE CACHE (10) CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 35h. The SATL shall issue an ATA flush command (see 3.1.11) in accordance with the constraints described in 9.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 42 — SYNCHRONIZE CACHE (16) CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 91h. The SATL shall issue an ATA flush command (see 3.1.11) in accordance with the constraints described in 9.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 43 — VERIFY (10) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 2Fh. The SATL shall issue an ATA verify command (see 3.1.23) in accordance with the constraints defined in 9.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 44 — VERIFY (12) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to AFh. The SATL shall issue an ATA verify command (see 3.1.23) in accordance with the constraints defined in 9.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 45 — VERIFY (16) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 8Fh. The SATL shall issue an ATA verify command (see 3.1.23) in accordance with the constraints defined in 9.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 46 — WRITE (6) CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 0Ah. See 9.17.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 47 — WRITE (10) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 2Ah. See 9.17.1 and 9.17.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 48 — WRITE (12) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to AAh. See 9.17.1 and 9.17.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 49 — WRITE (16) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 8Ah. See 9.17.1 and 9.17.2. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 50 — WRITE AND VERIFY (10) CDB field translations

| Field | Description or reference |
|--|---------------------------------------|
| OPERATION CODE | Set to 2Eh. See 9.22. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 51 — WRITE AND VERIFY (12) CDB field translations

| Field | Description or reference |
|--|---------------------------------------|
| OPERATION CODE | Set to AEh. See 9.22. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 52 — WRITE AND VERIFY (16) CDB field translations

| Field | Description or reference |
|--|---------------------------------------|
| OPERATION CODE | Set to 8Eh. See 9.22. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 53 — WRITE SAME (10) CDB field translations

| Field | Description or reference |
|--|---|
| OPERATION CODE | Set to 41h. If the ATA device supports the ATA SCD Write Same command (see ATA8-ACS), the SATL should issue the ATA SCT Write Same command to repeatedly write the data block transferred from the application client to the ATA device. If the ATA device does not implement ATA SCD Write Same command then the SATL shall issue ATA write commands as defined in 9.17.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |

Table 55 — WRITE SAME (16) CDB field translations

| Field | Description or reference |
|--|--|
| OPERATION CODE | Set to 93h . If the ATA device supports the ATA SCD Write Same command (see ATA8-ACS), the SATL should issue the ATA SCT Write Same command to repeatedly write the data block transferred from the application client to the ATA device. If the ATA device does not implement ATA SCD Write Same command then the SATL shall issue ATA write commands as defined in 9.17.1. |
| <i>Parts of table not affected are not shown here, but remain unchanged.</i> | |