

To: T10 Technical Committee
From: Bob Sheffield(robert.l.sheffield@intel.com)
Date: 1 May 2006
Subject: 06-216r0: SAT - Block Mapping Issues

Revision history

Revision 0 (1 May 2006) First revision

Related documents

SAT-r08 - SCSI / ATA Translation revision 08
06-121r0 SAT-r08_LB_Comment_Resolution.pdf

Overview

Letter ballot comments received for SAT (documented in 06-121r0) reflect diverging opinions about the basic mapping of ATA logical sectors to SCSI logical blocks. Two related aspects of this are:

- a) the mapping of logical block addresses between ATA and SCSI command sets; and
- b) the mapping of the number of bytes within each block between ATA and SCSI command sets.

Item b is the more crucial element, with the result there driving the resolution of item a.

One school of thought is that because most ATA disks are formatted fixed 512-byte logical sector sizes, a SATL should implement a one-for-one mapping and support only a fixed 512-byte logical block size without protection information.

The other school of thought says that many factors drive the need to provide more flexibility in the mapping of ATA sectors to SCSI logical blocks:

- a) recent modifications to ATA provide for larger physical sector sizes (e.g., 4096);
- b) various methods exist for providing data protection capabilities in a SATL, and many of those methods involve something other than a one-for-one mapping.

This proposal is intended to define a common direction regarding resolution of the SAT LB comments affecting the mapping between ATA logical sectors and SCSI logical blocks, as well as identifying which changes are required in which subclauses in the SAT standard to resolve those LB comments. The LB comment resolution for each LB comment listed in 06-121 related to block mapping will simply refer to this proposal as the resolution.

The authors's preference is driven by the notion that the purpose of a standard is to provide interoperability where there is potential for cooperative implementations that operate under a common set of assumptions. But it is not the purpose of a standard to constrain implementations that may derive value by providing a level of flexibility that goes beyond a specific set of definitions. In that spirit, the author's recommendation is, as a first pass, to leave the specific mapping of ATA logical sectors to SCSI logical blocks undefined (i.e., allow implementations that do something other than a one-to-one mapping of 512-byte ATA logical sectors to 512-byte SCSI logical blocks), but to provide a reporting mechanism so that an application client can readily determine whether a SATL implements a simple one-for-one 512-byte mapping, or something more elaborate.

Suggested changes

Modify subclause 8.5.5 General mode parameter block descriptor fields as follows:

8.5.5 General mode Mode parameter block descriptor fields

Table 19 describes the translation of fields in the general short LBA mode parameter block descriptor and the short LBA mode parameter block descriptor fields.

Table 19 — ~~General mode~~ Mode parameter block descriptor fields

Field	Description or reference
NUMBER OF BLOCKS	Unspecified (see 3.4.3)
DENSITY CODE	For direct access devices, this field should be set to 00h.
BLOCK LENGTH	Describes the block length for the section of the LUN described by this mode parameter block descriptor. Since there is only one mode parameter block descriptor, this describes the block length of the entire LUN. For direct access devices, the block length is set to 200h or 512 bytes per block. Unspecified (see 3.4.3)

SATL shall support only the mode parameter block descriptor formats for direct access devices.

Remove subclause 8.5.6 mode sense block descriptor (8-byte format).

Modify subclause 9.1 Translating LBA and transfer length and ATA command use constraints as follows:

9.1 Translating LBA and transfer length and ATA command use constraints

A SATL emulates SCSI logical blocks. The SCSI BLOCK LENGTH IN BYTES field in the READ CAPACITY data (see 9.8.2 and 9.9.2) may not be equal to the Logical Sector Size of the ATA device (see ATA/ATAPI-7).

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 or may not be equal to the Logical Sector Size of the ATA device (see ATA/ATAPI-7) (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, but so that the result of a logical block address translated in one direction and then translated in the reverse direction yields the original logical block address.

NOTE 6 - The Logical Sector Size indicated by an ATA device is the number of words in a logical sector. The number of bytes in an ATA device logical sector is twice the value indicated in the Logical Sector Size.

NOTE 6 - The Logical Sector Size indicated by an ATA device is represented in words; therefore, the number of bytes in an ATA device logical sector is twice the value indicated in the Logical Sector Size.

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 attached ATA device and the ATA host within the SATL.