



Knut Grimsrud
 Intel Corp. M/S JF2-53
 2111 NE 25th Ave
 Hillsboro, OR 97124

(503) 264-8419
 knut.s.grimsrud@intel.com

March 22, 2002

T13 Committee
 c/o Pete McLean
 Maxtor Corp.
 2190 Miller Drive
 Longmont, CO 80501

RE: Serial ATA II request for ATA opcode assignments and capabilities discovery

Dear Pete,

The Serial ATA II working group is striving to extend the capabilities of Serial ATA in order to better address the needs of server and networked storage market segments. We believe Serial ATA is an attractive solution for some portion of this market segment and that there are opportunities to make Serial ATA even more attractive for this use by utilizing some of the extended capabilities accommodated in the Serial ATA 1.0 spec as well as definition of new extensions in Serial ATA II.

The Serial ATA II working group would like to make a request of the T13 committee for some specific opcode assignments, signatures in identify device, and set features capabilities that are outlined in the remainder of this memo.

OpCode Assignment Requests

The Serial ATA II working group has developed a new command queuing scheme that takes advantage of the "first-party DMA" capability accommodated in the Serial ATA 1.0 specification. The command queuing scheme has low overhead and accommodates less than one interrupt per I/O for interrupt aggregation purposes. It requires no intermediate device driver software handshakes during the phases of command execution and accommodates out of order command execution. It does not require a revamped host controller interface to implement and primarily utilizes the existing defined facilities.

To support the new queuing model, two new opcodes are requested – one for a first-party DMA queued read command and one for a first-party DMA queued write command. Both commands would use the 48-bit addressing mode and no request is made for 28-bit addressing equivalents of these two commands.

An essential element of storage subsystems for the server and networked storage market segment is a complete enclosure services/management solution. Subsystems for these market segments today include controllers within the subsystem for monitoring and controlling environmental conditions (fan tachometer, overtemp, etc) as well as to provide operator notification of status (blink the right lights on the enclosure). Such enclosure services controllers are already readily available.

Because subsystem controllers are generally used in server and networked storage segments where SCSI is one of the most predominant interfaces, the existing enclosure service chips generally respond to SCSI commands passed in the form of CDBs.

Rather than attempt to define ATA equivalents for the myriad of control and configuration functions associated with the special requirements of storage subsystems, a means for covering such auxiliary service functions through a single simple and uniform mechanism is more attractive.

The Serial ATA II working group requests an opcode assignment for accommodating a CDB being passed directly in the taskfile. With the 48-bit addressing capability, there is sufficient space in the taskfile registers to directly convey a 10-byte CDB without the need for the multi-phased command delivery protocol used for the equivalent task in ATAPI. The CDB delivery mechanism is not intended for general data movement commands for target storage devices and is limited to mere delivery of a 10-byte or smaller CDB.

Capabilities Discovery Requests

With the addition of new capabilities, a means by which host software can determine the optional support of these facilities is needed. A request is made for allocation of space in the Identify Device data structure for conveying the device's capability and optional support for Serial ATA-specific features and capabilities. Specific capabilities that are expected to be conveyed would include:

- Identifying mark to indicate the device is a Serial ATA device
- Identification of which Serial ATA generation the device supports
- Identification of whether the device supports first-party DMA command queuing
- Identification of whether the device supports CDB pass-through

In addition, a means for identifying an enclosure services device as a new class of ATA device is requested through appropriate signature in the Identify Device data structure.

The Serial ATA II group would also like to request that the T13 committee's approach to indicating the device's native sector size not preclude the ability for specialty devices to have sector sizes that are not a multiple of 512 bytes. Specialty drives addressing the specific needs of server and networked storage solutions may have reason to use sector sizes of 520 bytes or 528 bytes (or multiples thereof) rather than the traditional 512 bytes. Because such drives are specialty products addressing specific point-needs, no provision for setting the sector size is requested – delivery of such specialty products would be negotiated directly between customers and suppliers.

Set Features Requests

A pair of Set Features values are requested in order to provide a means for the host to enable/disable Serial ATA specific capabilities on devices. One set features value would be used for enabling features (as identified in the Sector Count subcode field) and the second would be used for disabling features. The standard convention of having these two values be 80h apart is requested.

ReadLogExt Request

Serial ATA has error sources that have no equivalents in parallel ATA (for example a disparity error in the serial data stream). To better accommodate error handling, diagnostics, and fault isolation, a means for conveying Serial ATA-specific error information is required through the allocation of one ReadLogExt log page. This allocated log page is also anticipated to be used for error handling in the first-party DMA command queuing model in order to accommodate a means

to return race-free error status information that includes sufficient information to adequately identify the affected queued command and the error condition.

If you'd like additional clarification, background, or supporting data for these requests, please let me know and I would be happy to provide additional details.

Best Regards,

Knut Grimsrud
Intel Corp.