

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 27 February 2004
Subject: 04-076r0 SBC-2 PRE-FETCH and errors

Revision history

Revision 0 (27 February 2004) First revision

Related documents

sbc2r12 - SCSI Block Commands - 2 revision 12

Overview

SBC-2 defines a few rules for the PRE-FETCH command, which is the only command left in SCSI returning the "CONDITION MET" status:

- a) IMMED=1, cache has sufficient room => CONDITION MET status
- b) IMMED=1, cache does not have sufficient room => GOOD status
- c) IMMED=0, complete success => CONDITION MET

However, it does not mention what to do if IMMED=0 and there is not complete success transferring the requested blocks into the cache. Errors could range from not enough room in the cache (relatively harmless) to an unrecoverable read error (a major error).

To parallel the IMMED=1 behavior, it should return GOOD if the cache isn't big enough and CHECK CONDITION for other reasons (that would generate deferred errors under IMMED=1).

Alternative: obsolete PRE-FETCH altogether.

Suggested changes

5.6 PRE-FETCH (10) command

The PRE-FETCH (10) command (see table 27) requests that the device server transfer the specified logical blocks from the medium to the cache memory. No data shall be transferred to the application client.

...

An immediate (IMMED) bit set to zero specifies that status shall be returned after the operation is complete. An IMMED bit set to one specifies that status shall be returned as soon as the CDB has been validated.

See the LOCK UNLOCK CACHE (10) command (see 5.4) for a definition of the LOGICAL BLOCK ADDRESS field.

The PREFETCH LENGTH field specifies the number of contiguous logical blocks of data that shall be transferred to the block device's cache memory. A PREFETCH LENGTH of zero specifies that the contiguous logical blocks up to and including the last logical block shall be transferred to the cache memory. Any other value specifies the number of logical blocks that shall be transferred. The device server may elect to not transfer logical blocks that already are contained in the cache memory.

~~If the IMMED bit is set to zero and the specified logical blocks were successfully transferred to the cache memory, the device server shall return CONDITION MET status. If the LINK bit is set to one in the CONTROL byte (see SPC-3), the device server shall return INTERMEDIATE CONDITION MET status.~~

~~If the IMMED bit is set to one, and the unlocked cache memory has sufficient capacity to accept all of the specified logical blocks, the device server shall return CONDITION MET status. If the LINK bit is set to one in the CONTROL byte (see SPC-3), and the unlocked cache memory has sufficient capacity to accept all of the specified logical blocks, the device server shall return INTERMEDIATE CONDITION MET status.~~

~~If the IMMED bit is set to one, and the unlocked cache memory does not have sufficient capacity to accept all of the specified logical blocks, the device server shall return GOOD status. The device server shall transfer to cache memory as many logical blocks that fit. If the LINK bit is set to one, the device server shall return INTERMEDIATE status.~~

If the IMMED bit is set to zero and the specified logical blocks were successfully transferred to the cache memory, the device server shall return:

- a) CONDITION MET status if the LINK bit is set to zero in the CONTROL byte (see SPC-3); or
- b) INTERMEDIATE-CONDITION MET status if the LINK bit is set to one.

If the IMMED bit is set to zero and the unlocked cache memory does not have sufficient capacity to accept all of the specified logical blocks, the device server shall transfer to the cache memory as many of the specified logical blocks that fit. If these logical blocks are transferred successfully it shall return:

- a) GOOD status if the LINK bit is set to zero in the CONTROL byte (see SPC-3); or
- b) INTERMEDIATE status if the LINK bit is set to one.

If the IMMED bit is set to one and the unlocked cache memory has sufficient capacity to accept all of the specified logical blocks, the device server shall return:

- a) CONDITION MET status if the LINK bit is set to zero in the CONTROL byte (see SPC-3); or
- b) INTERMEDIATE-CONDITION MET status if the LINK bit is set to one.

If the IMMED bit is set to one and the unlocked cache memory does not have sufficient capacity to accept all of the specified logical blocks, the device server shall return:

- a) GOOD status if the LINK bit is set to zero in the CONTROL byte (see SPC-3); or
- b) INTERMEDIATE status if the LINK bit is set to one.

If the IMMED bit is set to zero and one or more of the specified logical blocks were not successfully transferred to the cache memory for reasons other than lack of cache memory capacity, the device server shall return CHECK CONDITION status with the appropriate sense key and additional sense code. If the IMMED bit is set to one and one or more of the specified logical blocks were not successfully transferred to the cache memory for reasons other than lack of cache memory capacity, the device server shall report a deferred error (see SPC-3).