To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com) and Gerry Houlder, Seagate (gerry.houlder@seagate.com)
Date: 9 March 2007
Subject: 07-113r0 SBC-3 Maximum transfer sizes for XPWRITE XDWRITE XDREAD PRE-FETCH

Revision history

Revision 0 (9 March 2007) First revision

Related documents

sbc3r08a - SCSI Block Commands - 3 (SBC-3) revision 8a
94-111r9 - XOR Commands on SCSI Disk Drives (Gerry Houlder, Seagate)(incorporated into sbc-r02)
07-110r0 - SBC-3 Update Block Limits VPD page for ORWRITE (Roger Cummings, Symantec)

Overview

1. The XOR Control mode page MAXIMUM XOR WRITE SIZE field should not apply to the XPWRITE command.

This field allows the device server to indicate that it has a restriction on the amount of data that it can store in a buffer between running an XDWRITE command and an XDREAD command. There is likely some bound; a drive cannot be expected to hold FFFFFFFFh logical blocks of data between commands.

The field should not apply to the XPWRITE command (or its new descendent ORWRITE), however. With those commands, the device server can request data in whatever chunk size it can support (one XFER_RDY/R2T at a time), perform the XOR (or OR) operation, and write the result to the medium before requesting another chunk of data. It doesn't have to buffer the entire transfer size for the command.

2. The Block Limits VPD page MAXIMUM TRANSFER SIZE field should not apply to the XDWRITE and XDREAD commands. Those commands are likely to have different limits than other commands (which might not have any limits). The XOR Control mode page MAXIMUM XOR WRITE SIZE field covers the transfer size for the XDWRITE (and XDREAD) commands, but is not required to be implemented. If the mode page is not implemented, then a new VPD page field is proposed to report the maximum transfer size for these commands.

3. The Block Limits VPD page MAXIMUM TRANSFER SIZE field should not apply to PRE-FETCH command. The PRE-FETCH command reads data from the medium and stores it into a buffer that is likely limited in size. A new VPD page field is proposed to report this limit. The field proposed is shared by XDWRITE, XDREAD, and PRE-FETCH.

Suggested changes to SBC-3

6.3.7 XOR Control mode page

The XOR Control mode page (see table 128) provides the application client with the means to obtain or modify certain XOR operating parameters of the logical unit.

Table 128 — XOR Control mode page

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	Reserved	PAGE CODE (10h)					
1	PAGE LENGTH (16h)							
2	Reserved						XORDIS	Reserved
3	Reserved							
4	(MSB)	MAXIMUM XOR WRITE SIZE						
7								(LSB)
8	Reserved							
11								
12	Obsolete							
19								
20	Reserved							
21								
22	Obsolete							
23								

The parameters savable (PS) bit is only used with the MODE SENSE command. This bit is reserved with the MODE SELECT command. A PS bit set to one indicates that the device server is capable of saving the mode page in a non-volatile vendor-specific location.

An XOR disable (XORDIS) bit set to zero specifies that the device server shall enable processing of XOR commands (i.e., the XDREAD commands (see 5.38 and 5.39), the XDWRITE commands (see 5.40 and 5.41), the XDWRITEREAD commands (see 5.42 and 5.43), and the XPWRITE commands (see 5.44 and 5.45)). An XORDIS bit set to one shall disable processing of XOR commands. If the XORDIS bit is set to one and an XOR command is received, the device server shall terminate the XOR command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE.

The MAXIMUM XOR WRITE SIZE field specifies the maximum transfer length in blocks that the device server accepts for a single XDWRITE ~~or XPWRITE~~ command or XDREAD command. Requests for transfer lengths exceeding this limit result in CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB. The MAXIMUM PREFETCH XDREAD XDWRITE TRANSFER LENGTH field in the Block Limits VPD page (see 6.4.2) indicates the maximum value of the MAXIMUM XOR WRITE SIZE field supported by the device server.

6.4.2 Block Limits VPD page

The Block Limits VPD page (see table 129) provides the application client with the means to obtain certain operating parameters of the logical unit.

Table 129 — Block Limits VPD page

Byte\Bit	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	PAGE CODE (B0h)							
2	Reserved							
3	PAGE LENGTH (~~0Ch~~10h)							
4	Reserved							
5								
6	(MSB)	OPTIMAL TRANSFER LENGTH GRANULARITY						
7								(LSB)
8	(MSB)	MAXIMUM TRANSFER LENGTH						
11								(LSB)
12	(MSB)	OPTIMAL TRANSFER LENGTH						
15								(LSB)
16	(MSB)	MAXIMUM PREFETCH XDREAD XDWRITE TRANSFER LENGTH						
19								(LSB)

The PERIPHERAL QUALIFIER field and the PERIPHERAL DEVICE TYPE field are defined in SPC-3.

The PAGE CODE field shall be set to B0h.

The PAGE LENGTH field is defined in SPC-3.

The OPTIMAL TRANSFER LENGTH GRANULARITY field indicates the optimal transfer length granularity in blocks for a single PRE-FETCH command, READ command, VERIFY command, WRITE command, WRITE AND VERIFY command, XDREAD command, XDWRITE command, XDWRITEREAD command, or XPWRITE command. Transfers with transfer lengths not equal to a multiple of this value may incur significant delays in processing.

The MAXIMUM TRANSFER LENGTH field indicates the maximum transfer length in blocks that the device server accepts for a single ~~PRE-FETCH command,~~ READ command, VERIFY command, WRITE command, WRITE AND VERIFY command, ~~XDREAD command, XDWRITE command,~~ XDWRITEREAD command, or XPWRITE command.

Requests for transfer lengths exceeding this limit result in CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB. A MAXIMUM TRANSFER LENGTH field set to zero indicates that there is no reported limit on the transfer length.

The OPTIMAL TRANSFER LENGTH field indicates the optimal transfer length in blocks for a single PRE-FETCH command, READ command, VERIFY command, WRITE command, WRITE AND VERIFY command, XDREAD command, XDWRITE command, XDWRITEREAD command, or XPWRITE command. Transfers with transfer lengths exceeding this value may incur significant delays in processing.

The MAXIMUM PREFETCH XDREAD XDWRITE TRANSFER LENGTH field indicates:

a) the maximum transfer length in blocks that the device server accepts for a single PRE-FETCH command;

b) if the XOR Control mode page (see 6.3.7) is implemented, the maximum value supported by the MAXIMUM XOR WRITE SIZE field in the XOR Control mode page; and
c) if the XOR Control mode page is not implemented, the maximum transfer length in blocks that the device server accepts for a single XDWRITE command or XDREAD command.

The device server should set the MAXIMUM PREFETCH XDREAD XDWRITE TRANSFER LENGTH field to less than or equal to the MAXIMUM TRANSFER LENGTH field.