



To: INCITS Technical Committee T10
From: Fred Knight, Network Appliance
Email: knight@netapp.com
Date: Nov 17, 2008
Subject: SBC-3 Host Alignment Detection

1) Revision history

Revision 0 (Nov 17, 2008)

2) Related documents

sbc3r15 – SCSI Block Commands – 3

3) Overview

The Physical Blocks section of SBC makes it clear that the standard does not define a method to set the physical to logical block associations (number of logical blocks per physical block, or number of physical blocks per logical blocks). Any such method is outside the scope of SBC.

However, the application client can use information in the READ CAPACITY Parameter data structure to determine that information. In addition, the BLOCK LIMITS VPD page provides additional granularity and transfer size hints. This proposal adds a note that application clients can take advantage of this information to obtain optimal performance.

Existing text is shown in **BLACK**, new text is shown in **RED**, comments in **BLUE**.

Proposal:

4.5 Physical blocks

A physical block is a set of data bytes on the medium accessed by the device server as a unit. A physical block may contain:

- a) a portion of a logical block (i.e., there are multiple physical blocks in the logical block)(e.g., a physical block length of 512 bytes with a logical block length of 2 048 bytes);
- b) a single complete logical block; or

- c) more than one logical block (i.e., there are multiple logical blocks in the physical block)(e.g., a physical block length of 4 096 bytes with a logical block length of 512 bytes).

Each physical block includes additional information not normally accessible to the application client (e.g., an ECC) that the device server uses to manage storage and retrieval.

If the device server supports the COR_DIS bit and/or the WR_UNCOR bit in a WRITE LONG command (see 5.35 and 5.36), then the device server shall have the capability of marking individual logical blocks as containing pseudo uncorrectable errors with correction enabled (see 3.1.45) or with correction disabled (see 3.1.46).

Logical blocks may or may not be aligned to physical block boundaries. A mechanism for establishing the alignment is not defined by this standard.

Note x: The LOWEST ALIGNED LOGICAL BLOCK ADDRESS field and LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT in the READ CAPACITY (16) parameter data (see 5.13.2) may be used by application clients to determine the logical block alignment. Application clients may use this information along with fields in the Blocks Limits VPD Page (see 6.4.2) (e.g., OPTIMAL TRANSFER LENGTH, OPTIMAL TRANSFER LENGTH GRANULARITY) to obtain optimal performance.

Figure 2 shows examples of logical blocks and physical blocks, where LBA 0 is aligned to a physical block boundary.

LOGICAL BLOCKS PER PHYSICAL BLOCK field set to 0h

(indicating one or more physical blocks per logical block):

4 physical blocks per logical block:

LBA 0				LBA 1				...
PB	PB	PB	PB	PB	PB	PB	PB	...

3 physical blocks per logical block:

LBA 0			LBA 1			LBA 2			...
PB	PB	PB	PB	PB	PB	PB	PB	PB	...

2 physical blocks per logical block:

LBA 0		LBA 1		LBA 2		LBA 3		LBA 4		...
PB	PB	PB	PB	PB	PB	PB	PB	PB	PB	...

1 physical block per logical block:

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7	LBA 8	LBA 9	LBA 10	...
PB	PB	PB	PB	PB	PB	PB	PB	PB	PB	PB	...

LOGICAL BLOCKS PER PHYSICAL BLOCK field set to a non-zero value

(indicating more than one logical block per physical block):

LOGICAL BLOCKS PER PHYSICAL BLOCK field set to 1h (indicating 2^1 logical blocks per physical block):

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7	LBA 8	LBA 9	LBA 10	LBA 11	...
PB		PB		PB		PB		PB		PB		...

LOGICAL BLOCKS PER PHYSICAL BLOCK field set to 2h (indicating 2^2 logical blocks per physical block):

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7	LBA 8	LBA 9	LBA 10	LBA 11	...
PB				PB				PB				...

LOGICAL BLOCKS PER PHYSICAL BLOCK field set to 3h (indicating 2^3 logical blocks per physical block):

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7	...
PB								...

Key:

LBA n = logical block with LBA n


PB = physical block

Figure 2 — Logical blocks and physical blocks examples


Figure 3 shows examples of logical blocks and physical blocks, where various LBAs are aligned to the physical block boundaries.

LOGICAL BLOCKS PER PHYSICAL BLOCK field set to 1h (indicating 2^1 logical blocks per physical block):

LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 0:




LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 1:




LOGICAL BLOCKS PER PHYSICAL BLOCK field set to 2h (indicating 2^2 logical blocks per physical block):


LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 0:




LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 1:



LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 2:



LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 3:



Key:

LBA n = logical block with LBA n

PB = physical block

Figure 3 — Logical block to physical block alignment examples

When there are more than one logical block per physical block, not all of the logical blocks are aligned to the physical block boundaries. When using medium access commands, application clients should:

- specify an LBA that is aligned to a physical block boundary; and
- access an integral number of physical blocks, provided that the access does not go beyond the last LBA on the medium.