

From: Bob Sheffield – LSI Corporation
 To: T10 Committee
 Subject: T10/08-044r1 SBC-3 Set Pseudo Format Command
 Date: 10 June 2008

Related Documents: SBC-3r15

Revision History:

| Revision | Date | Description |
|----------|---------|---|
| r0 | 1/3/08 | Initial version |
| r1 | 5/5/08 | Clarify justification (4K initiator,...) and add per-command controls |
| r2 | 6/10/08 | Propose SET PSEUDO FORMAT command |

Overview

There has been recent momentum in the storage industry to support larger logical blocks. This offers the opportunity for improved disk reliability and/or other benefits for systems that use larger logical blocks. Interoperability with existing applications that largely depend on logical blocks having a size of 512 bytes prompts the need for storage virtualization solutions that emulate 512-byte logical blocks for the virtual logical units they expose to applications, but use disk drives formatted with larger logical blocks (e.g., 4096 bytes) as the raw storage for the virtual logical units.

With the move to larger disk block lengths, applications that access data using matching larger block lengths (rather than depending upon emulation) may access data more efficiently by avoiding latencies caused by read-modify-write operations in the emulation layer. So, it stands to reason that both traditional applications that access media using 512-byte blocks and new applications that access media using larger blocks (e.g., 4K-byte) will coexist, perhaps in the same platforms (e.g., virtualized servers).

Likewise, there will be a gradual transition from systems that use 512-byte formatted disks to larger block length media, and the larger block length media may include both media formatted to a larger logical block length as well as media formatted to a larger physical block length but that emulates 512-byte logical blocks.

This leads to a complex intermix of application expected block lengths and disk supported block lengths that, if not appropriately addressed in SCSI standards, leads to a variety of non-interoperability problems, particularly when protection information (PI) is used. The table below lists the relevant combinations, and whether interoperability holds for PI-enabled implementations.

| Case | Application client block length | Disk block length | | Comment |
|--------------|--|-------------------|----------|---|
| | | Logical | Physical | |
| 1 | T ^a | T | T | Traditional usage |
| 2 | L ^b | T | T | Not interoperable – mismatched PI field locations and intervals between app and disk. |
| 3 | T | L | T | |
| 4 | L | L | T | Interoperable, but not defined in SCSI |
| 5 | T | T | L | Emulation of traditional block length |
| 6 | L | T | L | Not interoperable – mismatched PI field locations and intervals between app and disk. |
| 7 | T | L | L | |
| 8 | L | L | L | Complete agreement on larger block lengths |
| ^a | T: Traditional block length (i.e., 512-bytes) | | | |
| ^b | L: Large block length (e.g., 4Kbyte – but could be something else like 2K or 8K) | | | |

In SPC-3, as PI is currently defined, when protection information is present, interoperability is achieved as long as disks and applications agree on the logical block length. When protection information is not used, a block virtualization layer may apply a reasonably straightforward mapping to emulate application block lengths that do not match the logical block lengths provided by the disks. This poses an insurmountable challenge, however, when protection information is involved. The reason is that the application expects PI fields to appear at intervals the disk is incapable of providing without violating SBC-3 (i.e., the requirement that PI fields must appear at logical block boundaries).

The existing definitions for the logical blocks per physical block exponent and lowest aligned logical block address fields in READ CAPACITY (16) parameter data provide interoperability for cases 1 and 5 in the table above, but do not address interoperability for the remaining 6 cases.

In virtualized server and external storage environments, block storage devices are likely to encounter the need to support an intermix of applications that require both traditional 512-byte access and larger block length access, simultaneously.

A solution to this problem should accommodate future increases in block length, and should provide the framework to establish interoperability with any application, regardless of the block length expected by the application.

The pseudo-format SCSI protocol extensions in this proposal define the framework to allow a disk to provide interoperability with different applications needing different block lengths, at the same time, regardless of the physical and logical block lengths provided by the disk.

This proposal defines a method by which a disk may support, in addition to its native logical block length, one or more pseudo-formats for blocks on media. Applications may specify, on a per command basis, the pseudo-format that matches the application block format required. The SET PSEUDO FORMAT command specifies one or more pseudo-formats to be used to process read and write commands that may (or may not) be supported by the drive. The SET PSEUDO FORMAT command specifies the number of disk logical blocks that correspond to a pseudo-block, and whether protection information is appended at logical block boundaries or pseudo logical block boundaries.

This document also proposes using two currently reserved bits in the CDBs for read and write type commands, referred to as the pseudo-format ID or PFID field, in the same byte as the GROUP NUMBER field, to specify a pseudo-format to be used for processing that command.

This proposal allows the application of protection information to coincide with the pseudo logical block boundaries as specified in the CDB.

- For type-1 data protection, the logical block reference tag contains the lower four bytes of the LBAs specified to be transferred in the CDB.
- For type-2 data protection, the CDB specifies the value of the first logical block reference tag transferred and the value increments by one for each subsequent pseudo logical block transferred.

This addresses both the need to provide raw block capacity for virtual devices emulating legacy block sizes, and also avoids diluting the effectiveness of the CRC field as logical block sizes increase.

| | | | | | | | | | | | | | | | | | | | |
|--------------------------------|---|-----|---|-----|--------------------|-----|---|-----|--------------------------------|--------------------|---|-----|---|-----|--------------------|--|---|--|---|
| 512 | X | 512 | 8 | 512 | X | 512 | 8 | 512 | X | 512 | 8 | 512 | X | 512 | 8 | | | | |
| 1024-byte pseudoLB | | | | 8 | 1024-byte pseudoLB | | | | 8 | 1024-byte pseudoLB | | | | 8 | 1024-byte pseudoLB | | | | 8 |
| 512 | 8 | 512 | 8 | 512 | 8 | 512 | 8 | 512 | 8 | 512 | 8 | 512 | 8 | 512 | 8 | | | | |
| 2048-byte pseudo logical block | | | | | | | | 8 | 2048-byte pseudo logical block | | | | | | | | 8 | | |
| 512 | X | 512 | X | 512 | X | 512 | X | 512 | X | 512 | X | 512 | X | 512 | 8 | | | | |
| 4096-byte pseudo logical block | | | | | | | | | | | | | | | 8 | | | | |

Figure 1 – Examples mapping logical blocks to pseudo logical blocks

Suggested Changes to SBC-3

3.1.40 protection information: Fields appended to each logical block or pseudo logical block that contain a cyclic redundancy check (CRC), an application tag, and a reference tag. See 4.17.

3.1.xx pseudo logical block: A unit of user data (see 3.1.50) composed of two or more adjacent logical blocks (see 3.1.26) that may be protected by one or more instances of protection information (see 3.1.40).

3.1.xx pseudo logical block address (PLBA): The value used to reference a pseudo logical block (see 3.1.xx).

Add the following to the list of symbols and abbreviations:

PLBA pseudo logical block address (see 3.1.xx)

Modify subclause 4.14.1 Error reporting overview to specify reporting the PLBA instead of the LBA in the INFORMATION field when applicable.

When a command attempts to access or reference an invalid LBA, the first invalid LBA shall be returned in the INFORMATION field of the sense data (see SPC-4).

When a command attempts to access or reference an invalid PLBA, the first invalid PLBA shall be returned in the INFORMATION field of the sense data.

When a recovered read error is reported for a command with the PFID field set to zero, the INFORMATION field of the sense data shall contain the LBA of the last logical block accessed by the command on which a recovered read error occurred.

When a recovered read error is reported for a command with the PFID field set to a non-zero value, the INFORMATION field of the sense data shall contain the PLBA of the last pseudo logical block accessed by the command on which a recovered read error occurred.

When an unrecovered error is reported for a command with the PFID field set to zero, the INFORMATION field of the sense data shall contain the LBA of the logical block on which the unrecovered error occurred.

When an unrecovered error is reported for a command with the PFID field set to a non-zero value, the INFORMATION field of the sense data shall contain the PLBA of the pseudo logical block on which the unrecovered error occurred.

Add subclause 4.6 to describe virtual logical blocks

4.6 Pseudo logical blocks

An application client may reference logical blocks using a pseudo logical block address (PLBA) by specifying a non-zero value in the PFID field. A PLBA references a pseudo logical block (see 3.1.xx) that is composed of two or more adjacent logical blocks, as specified by the SET PSEUDO FORMAT command for the specified pseudo format ID. The LBA of the first logical block in the pseudo logical block referenced by a PLBA value of zero shall be the LBA of the lowest aligned logical block as indicated in the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field returned in the READ CAPACITY (16) parameter data. The LBA of the first logical block in the next pseudo logical block shall be the LBA of the first logical block in the previous pseudo logical block plus n, where n is the value of the NUMBER OF LOGICAL BLOCKS IN A PSEUDO LOGICAL BLOCK field as specified in the logical blocks per pseudo block field of the SET PEDUDO FORMAT CDB.

The pattern continues until the number of remaining LBAs, as determined by the returned LOGICAL BLOCK ADDRESS field in READ CAPACITY (16) parameter data, is less than n.

A non-zero value in the PFID field specifies that the application client is specifying a pseudo logical block address in the LOGICAL BLOCK ADDRESS FIELD of the CDB. If a non-zero value is specified in the PFID field, the field in the CDB specifying the number of logical blocks to process (e.g., the TRANSFER LENGTH field or the NUMBER OF LOGICAL BLOCKS field) shall specify the number of pseudo logical blocks processed by the command.

If the device server does not support pseudo block addressing and the PFID field contains a value other than zero, the command shall be terminated with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the PFID field contains a value other than zero that is not supported by the device server, or if the device server has not successfully processed a SET PSEUDO FORMAT command to establish the specified pseudo format, the command shall be terminated with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

4.17.3 Protection information format

Table 7 defines the **placement format** of protection information in a logical block or pseudo logical block.

Table a – user data and protection information format

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-------------------------------|---|---|---|---|---|---|---|
| 0 | USER DATA | | | | | | | |
| n - 1 | USER DATA | | | | | | | |
| n | LOGICAL BLOCK GUARD | | | | | | | |
| n + 1 | LOGICAL BLOCK GUARD | | | | | | | |
| n + 2 | LOGICAL BLOCK APPLICATION TAG | | | | | | | |
| n + 3 | LOGICAL BLOCK APPLICATION TAG | | | | | | | |
| n + 4 | LOGICAL BLOCK REFERENCE TAG | | | | | | | |
| n + 7 | LOGICAL BLOCK REFERENCE TAG | | | | | | | |

The USER DATA field shall contain user data. The contents of the USER DATA field shall be used to generate and check the CRC contained in the LOGICAL BLOCK GUARD field.

The LOGICAL BLOCK GUARD field contains the CRC (see 4.17.4) of the contents of the USER DATA field.

The LOGICAL BLOCK APPLICATION TAG field is set by the application client. If the device server detects a:

- a) LOGICAL BLOCK APPLICATION TAG field set to FFFFh and type 1 protection (see 4.17.2.3) or type 2 protection (see 4.17.2.4) is enabled; or
- b) LOGICAL BLOCK APPLICATION TAG field set to FFFFh, LOGICAL BLOCK REFERENCE TAG field set to FFFF FFFFh, and type 3 protection (see 4.17.2.5) is enabled,

then the device server disables checking of all protection information for the logical block when reading from the medium. Otherwise, the contents of the logical block application tag are not defined by this standard.

The LOGICAL BLOCK APPLICATION TAG field may be modified by a device server if the ATO bit is set to zero in the Control mode page (see SPC-4).

The contents of the LOGICAL BLOCK APPLICATION TAG field shall not be used to generate or check the CRC contained in the LOGICAL BLOCK GUARD field.

The LOGICAL BLOCK REFERENCE TAG field of the first logical block [or pseudo logical block](#) in the data-in buffer and/or data-out buffer shall contain the value specified in table 8.

Table 8 – Content of the LOGICAL BLOCK REFERENCE TAG field of the first logical block [or pseudo logical block](#) in the data-in buffer and/or data-out buffer

| Protection Type | PSID | Content of the LOGICAL BLOCK REFERENCE TAG field of the first logical block or pseudo logical block in the data-in buffer and/or data-out buffer |
|-------------------------------------|------|--|
| Type-1 protection (see 4.17.2.3) | = 0 | The least significant four bytes of the LBA contained in the LOGICAL BLOCK ADDRESS field of the command. |
| | ≠ 0 | See 4.17.4 |
| Type-2 protection (see 4.17.2.4) | = 0 | The value in the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field of the command. |
| | ≠ 0 | See 4.17.4 |
| Type-3 protection (see 4.17.2.5) | n/a | Not defined in this standard. |

The LOGICAL BLOCK REFERENCE TAG field subsequent logical blocks in the data-in buffer and/or data-out buffer shall be set as specified in table 9.

Table 9 – Setting of the LOGICAL BLOCK REFERENCE TAG field of the subsequent logical block in the data-in buffer and/or data-out buffer

| Protection Type | PSID | The content of the LOGICAL BLOCK REFERENCE TAG field of the subsequent logical block or pseudo logical block in the data-in buffer and/or data-out buffer |
|--|------|---|
| Type-1 protection (see 4.17.2.3) and Type-2 protection (see 4.17.2.4) | = 0 | The logical block reference tag of the previous logical block plus one. |
| | ≠ 0 | See 4.17.4 |
| Type-3 protection (see 4.17.2.5) | n/a | Not defined in this standard. |

The contents of the LOGICAL BLOCK REFERENCE TAG field shall not be used to generate or check the CRC contained in the LOGICAL BLOCK GUARD field.

Add subclause 4.17.4 to describe application of protection information with pseudo logical blocks

4.17.4 Protection information with pseudo logical blocks

4.17.4.1 Overview

[For commands specifying a non-zero value in the PFID field, the eight bytes of protection information are appended either to each pseudo logical block transferred, or to each logical block within each pseudo logical block, depending on the setting of the APIPB bit for the pseudo format specified in the PFID field. Media shall be formatted to accommodate \$2^n\$ x 8 bytes of protection information, where n is the value of the NUMBER OF LOGICAL BLOCKS IN A PSEUDO LOGICAL BLOCK EXPONENT field specified in the SET PSEUDO FORMAT CDB \(i.e., space is provided on media for one instance of protection information for each logical block contained within the pseudo logical block\).](#)

4.17.4.2 Protection information alignment with pseudo logical block

Alignment of protection information with pseudo block boundaries is specified using a SET PSEUDO FORMAT command with the APIPB bit set to one for the specified pseudo format. The last protection information field available on media contained within a pseudo logical block shall contain protection information that pertains to the pseudo logical block.

The LOGICAL BLOCK GUARD FIELD shall contain the CRC calculated over all user data contained within the pseudo logical block.

The LOGICAL BLOCK REFERENCE TAG field of the first logical block or pseudo logical block in the data-in buffer and/or data-out buffer shall contain the value specified in table 8.

Table 10 – Content of the LOGICAL BLOCK REFERENCE TAG field of the first pseudo logical block in the data-in buffer and/or data-out buffer

| <u>Protection Type</u> | <u>Content of the LOGICAL BLOCK REFERENCE TAG field of the first pseudo logical block in the data-in buffer and/or data-out buffer</u> |
|---|---|
| <u>Type-1 protection (see 4.17.2.3)</u> | <u>The least significant four bytes of the PLBA contained in the LOGICAL BLOCK ADDRESS field of the command.</u> |
| <u>Type-2 protection (see 4.17.2.4)</u> | <u>The value in the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field of the command.</u> |
| <u>Type-3 protection (see 4.17.2.5)</u> | <u>Not defined in this standard.</u> |

The LOGICAL BLOCK REFERENCE TAG field subsequent logical blocks in the data-in buffer and/or data-out buffer shall be set as specified in table 9.

Table 11 – Setting of the LOGICAL BLOCK REFERENCE TAG field of the subsequent pseudo logical block in the data-in buffer and/or data-out buffer

| <u>Protection Type</u> | <u>The content of the LOGICAL BLOCK REFERENCE TAG field of the subsequent pseudo logical block in the data-in buffer and/or data-out buffer</u> |
|--|--|
| <u>Type-1 protection (see 4.17.2.3) and Type-2 protection (see 4.17.2.4)</u> | <u>The logical block reference tag of the previous pseudo logical block plus one.</u> |
| <u>Type-3 protection (see 4.17.2.5)</u> | <u>Not defined in this standard.</u> |

When a command specifying a nonzero value in the PFID field causes data to be written to media, the unused protection information fields for each logical block shall be written as follows:

- a) the value stored in the LOGICAL BLOCK GUARD field shall be the one's complement of the CRC calculated over the data in the corresponding logical block,
- b) the value stored in the LOGICAL BLOCK REFERENCE TAG field shall be the one's complement of the least significant four bytes of the LBA of the logical block corresponding to that protection information field, and
- c) the value stored in the LOGICAL BLOCK APPLICATION TAG field shall be zero.

When a command reads data from media with the PFID field set to a value that is different from the value used to write the data to media (i.e., the read uses a different pseudo block format than the previous write to the same area), if protection information checking is enabled, the device server detects an error in the protection information and reports CHECK condition status with the sense key set to ABORTED COMMAND and the additional sense code set to LOGICAL BLOCK GUARD CHECK FAILED.

4.17.4.3 Protection information alignment with logical block

If the APIPB bit is set to zero for a specified non-zero pseudo format ID, the read or write CDB shall specify data transferred as a range of pseudo logical blocks (i.e., the LOGICAL BLOCK ADDRESS field specifies the starting PLBA and the TRANSFER LENGTH field specifies a count of pseudo logical blocks). However, protection information shall be applied to the data in terms of logical blocks. Protection information is present for each logical block contained within each pseudo block. The contents of the protection information field for each logical block is the same as it would be for a command specifying the transfer of the same range of data with the PFID field set to zero.

Modify the READ CAPACITY (16) command as follows

5.13 READ CAPACITY (16) command

5.13.1 READ CAPACITY (16) command overview

The READ CAPACITY (16) command (see table 46) requests that the device server transfer parameter data describing the capacity and medium format of the direct-access block device to the data-in buffer. This command is mandatory if the logical unit supports protection information (see 4.17) and is optional otherwise. This command is implemented as a service action of the SERVICE ACTION IN operation code (see A.2). This command may be processed as if it has a HEAD OF QUEUE task attribute (see 4.12).

Table 46 – READ CAPACITY (16) command

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-----------------------|---|----------------------|----------------------|---|---|---|-----|
| 0 | OPERATION CODE (9Eh) | | | | | | | |
| 1 | Reserved | | | SERVICE ACTION (10h) | | | | |
| 2 | LOGICAL BLOCK ADDRESS | | | | | | | |
| 9 | ALLOCATION LENGTH | | | | | | | |
| 10 | | | | | | | | |
| 13 | | | | | | | | |
| 14 | Reserved | | PFID | Reserved | | | | PMI |
| 15 | CONTROL | | | | | | | |

The OPERATION CODE field and SERVICE ACTION field are defined in SPC-4 and shall be set to the values defined in table 46.

See the READ CAPACITY (10) command (see 5.12) for definitions of the LOGICAL BLOCK ADDRESS field and the PMI bit.

The ALLOCATION LENGTH field specifies the maximum number of bytes that the application client has allocated for returned parameter data. An allocation length of zero indicates that no data shall be transferred. This condition shall not be considered as an error. The device server shall terminate transfers to the data-in buffer when the number of bytes specified by the ALLOCATION LENGTH field have been transferred or when all available data has been transferred, whichever is less. The contents of the parameter data shall not be altered to reflect the truncation, if any, that results from an insufficient allocation length.

The contents of the CONTROL byte are defined in SAM-4.

[A non-zero value in the PFID field requests information returned in the READ CAPACITY \(16\) parameter data that pertains to the specified pseudo logical block format.](#)

5.13.2 READ CAPACITY (16) parameter data

The READ CAPACITY (16) parameter data is defined in table 47. Any time the READ CAPACITY (16) parameter data changes, the device server should establish a unit attention condition as described in 4.7.

Table 46 – READ CAPACITY (16) command

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|----------|--------------------------------|--------------------------------------|--|---|---|---------|-------|
| 0 | (MSB) | RETURNED LOGICAL BLOCK ADDRESS | | | | | | (LSB) |
| 7 | | | | | | | | |
| 8 | (MSB) | LOGICAL BLOCK LENGTH IN BYTES | | | | | | (LSB) |
| 11 | | | | | | | | |
| 12 | Reserved | | | P_TYPE | | | PROT_EN | |
| 13 | Reserved | | | LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT | | | | |
| 14 | Reserved | | LOWEST ALIGNED LOGICAL BLOCK ADDRESS | | | | | |
| 15 | | | | | | | | |
| 16 | Reserved | PFID | APIPB | LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT | | | | |
| 17-31 | Reserved | | | | | | | |
| 31 | | | | | | | | |

The RETURNED LOGICAL BLOCK ADDRESS field and LOGICAL BLOCK LENGTH IN BYTES field of the READ CAPACITY (16) parameter data are the same as the in the READ CAPACITY (10) parameter data (see 5.12). The maximum value that shall be returned in the RETURNED LOGICAL BLOCK ADDRESS field is FFFF_FFFF_FFFF_FFFEh.

The protection type (P_TYPE) field and the protection enable (PROT_EN) bit (see table 48) indicate the logical unit's current type of protection.

Table 48 – P_TYPE field and PROT_EN bit

| PROT_EN | P_TYPE | Description |
|---------|-------------|--|
| 0 | xxxb | The logical unit is formatted to type 0 protection (see 4.17.2.2). |
| 1 | 000b | The logical unit is formatted to type 1 protection (see 4.17.2.3). |
| 1 | 001b | The logical unit is formatted to type 2 protection (see 4.17.2.4). |
| 1 | 010b | The logical unit is formatted to type 3 protection (see 4.17.2.5). |
| 1 | 011b – 111b | Reserved |

The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is defined in table 49.

Table 49 – LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field and

| Code | Description |
|--|--|
| 0 | One or more physical blocks per logical block ^a |
| n > 0 | 2 ⁿ logical blocks per physical block |
| ^a The number of physical blocks per logical block is not reported | |

The LOWEST ALIGNED LOGICAL BLOCK ADDRESS field indicates the LBA of the first logical block that is located at the beginning of a physical block (see 4.5).

NOTE 14 - The highest LBA that the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field supports is 3FFFh (i.e., 16 383).

The PFID field shall have the same pseudo format ID value as the PFID field specified in the READ CAPACITY (16) CDB.

If the PFID field is non-zero, the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field indicates the number of logical blocks contained in each pseudo logical block for the specified pseudo format ID as specified in a previous SET PSEUDO FORMAT command.

If the PFID field is non-zero, the APIPB field indicates whether protection information is included for each logical block or for each pseudo logical block for the specified pseudo format ID as specified in a previous SET PSEUDO FORMAT command.

If the PFID field is zero, the APIPB field and the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field are reserved.

Add subclause 5.xx (following 5.18 REASSIGN BLOCKS) describing the SET PSEUDO FORMAT command

5.xx SET PSEUDO FORMAT command

The SET PSEUDO FORMAT command (see table x) specifies use of an alternate method to address logical blocks associated with a pseudo format identifier that may be specified in a read or write command in the PFID field. This command is implemented as a service action of the SERVICE ACTION OUT operation code (see A.2). Support for this command is optional.

Table x – SET PSEUDO FORMAT command

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-----------|---|-------------|--------------|---|---|---|---|---|
| <u>0</u> | <u>OPERATION CODE (9Fh)</u> | | | | | | | |
| <u>1</u> | <u>Reserved</u> | | | <u>SERVICE ACTION (10h)</u> | | | | |
| <u>2</u> | <u>LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT</u> | | | | | | | |
| <u>3</u> | <u>Reserved</u> | <u>PFID</u> | <u>APIPB</u> | <u>LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT</u> | | | | |
| <u>4</u> | <u>Reserved</u> | | | | | | | |
| <u>14</u> | <u>Reserved</u> | | | | | | | |
| <u>15</u> | <u>CONTROL</u> | | | | | | | |

The OPERATION CODE field and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values defined in table x. If the device server does not support the SET PSEUDO FORMAT command, the device server shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE.

When processing a read or write command that specifies the corresponding pseudo format identifier, the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field specifies that, the device server shall:

- a) calculate the LBA of the first logical block to transfer as the product of the LOGICAL BLOCK ADDRESS field of the read or write command times 2^n , where n is the value in the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field plus the value of the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data; and
- b) calculate the number of logical blocks to transfer as the product of the TRANSFER LENGTH field of the read or write command times 2^n , where n is the value in the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field for the specified pseudo format.

If the device server does not support the value specified in the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field, the device server shall terminate the command with CHECK CONDITION status

with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The PSEUDO FORMAT IDENTIFIER (PFID) field specifies the pseudo format identifier value the device server shall associate with the LOGICAL BLOCKS PER PSEUDO BLOCK EXPONENT field and APIPB BIT values specified in the CDB, and the device server shall apply the pseudo format for the specified pseudo format identifier when that value appears in the PFID field of a read or write command. A value of zero in the PFID field is invalid. If the device server receives a SET PSEUDO FORMAT command with the PFID field set to zero, the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The ALIGN PROTECTION INFORMATION ON PSEUDO BLOCK (APIPB) bit specifies whether to align protection information, if present, with logical block boundaries or pseudo block boundaries.

An APIPB bit set to one specifies that protection information fields align with pseudo logical block boundaries. If the protection information is type-1, the LOGICAL BLOCK REFERENCE TAG field for the first pseudo logical block transferred contains the lower four bytes of the PLBA specified in the LOGICAL BLOCK ADDRESS field of the read or write command. If the protection information is type-2, the LOGICAL BLOCK REFERENCE TAG field in the protection information of the first pseudo logical block contains the value specified in the EXPECTED LOGICAL BLOCK REFERENCE TAG field of the read or write CDB. The LOGICAL BLOCK REFERENCE TAG field increments by one with each subsequent pseudo logical block transferred. The LOGICAL BLOCK GUARD field is calculated over the combined contents of the logical blocks contained within a pseudo logical block.

An APIPB bit set to zero specifies that protection information fields align with logical block boundaries. If the protection information is type-1, the LOGICAL BLOCK REFERENCE TAG field contains the lower four bytes of the LBA specified in the LOGICAL BLOCK ADDRESS field of the read or write command. If the protection information is type-2, the LOGICAL BLOCK REFERENCE TAG field in the protection information of the first pseudo logical block contains the value specified in the EXPECTED LOGICAL BLOCK REFERENCE TAG field of the read or write CDB. The LOGICAL BLOCK REFERENCE TAG field increments by one with each subsequent logical block transferred.

If the device server does not support the value specified in the APIPB bit, the device server shall terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to INVALID FIELD IN CDB.

Modify the descriptions of the applicable block commands (those that contain a GROUP NUMBER field) as follows

Table x – 10-byte CDB format

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-------------------|---|--------------|---|---|---|---|----------|
| 0 | OPERATION CODE | | | | | | | |
| 1 | Command dependent | | | | | | | Obsolete |
| 2 | (MSB) | LOGICAL BLOCK ADDRESS | | | | | | (LSB) |
| 5 | | | | | | | | |
| 6 | Reserved | Reserved PFID | GROUP NUMBER | | | | | |
| 7 | (MSB) | TRANSFER LENGTH (or other length field) | | | | | | (LSB) |
| 8 | | | | | | | | |
| 9 | CONTROL | | | | | | | |

A non-zero value in the PSEUDO FORMAT ID (PFID) field specifies that the LOGICAL BLOCK ADDRESS field contains the PLBA of the first pseudo logical block to be transferred, and that the TRANSFER LENGTH field specifies the number of pseudo logical blocks to transfer, where each pseudo logical block contains 2ⁿ logical blocks, and n is the value specified by the SET PSEUDO FORMAT command defining the pseudo format for the specified pseudo format ID.

Table x – 12-byte CDB format

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-------------------|---|--------------|---|---|---|---|----------|
| 0 | OPERATION CODE | | | | | | | |
| 1 | Command dependent | | | | | | | Obsolete |
| 2 | (MSB) | LOGICAL BLOCK ADDRESS | | | | | | (LSB) |
| 5 | | | | | | | | |
| 6 | (MSB) | TRANSFER LENGTH (or other length field) | | | | | | (LSB) |
| 9 | | | | | | | | |
| 10 | Reserved | Reserved PFID | GROUP NUMBER | | | | | |
| 11 | CONTROL | | | | | | | |

Table x – 16-byte CDB format

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-------------------|---|--------------|---|---|---|---|----------|
| 0 | OPERATION CODE | | | | | | | |
| 1 | Command dependent | | | | | | | Obsolete |
| 2 | (MSB) | LOGICAL BLOCK ADDRESS | | | | | | (LSB) |
| 9 | | | | | | | | |
| 10 | (MSB) | TRANSFER LENGTH (or other length field) | | | | | | (LSB) |
| 13 | | | | | | | | |
| 14 | Reserved | Reserved PFID | GROUP NUMBER | | | | | |
| 15 | CONTROL | | | | | | | |

Editor's note: Some of the commands show the high-order bit of the vBLE field as "Restricted for MMC-4" because some MMC-4 commands define it as the "streaming" bit. This bit is available to direct-access storage devices because MMC devices (e.g., optical media) have no need to use the streaming bit (which implies tolerance of bit-errors for video applications – not appropriate for direct-access storage devices).

Table x – 32-byte CDB format (example)

| Byte/Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-----------------------------|--|--------------|-----|-----|----------|--------|----------|
| 0 | OPERATION CODE | | | | | | | |
| 1 | CONTROL | | | | | | | |
| 2 | Reserved | | | | | | | |
| 5 | Reserved | | | | | | | |
| 6 | Reserved | Reserved PFID | GROUP NUMBER | | | | | |
| 7 | ADDITIONAL CDB LENGTH (18h) | | | | | | | |
| 8 | Reserved | | | | | | | |
| 9 | SERVICE ACTION | | | | | | | |
| 10 | XXPROTECT | | | DPO | FUA | Reserved | FUA_NV | Reserved |
| 11 | Reserved | | | | | | | |
| 12 | (MSB) | LOGICAL BLOCK ADDRESS | | | | | | (LSB) |
| 19 | Reserved | | | | | | | |
| 20 | (MSB) | EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG | | | | | | (LSB) |
| 23 | Reserved | | | | | | | |
| 24 | (MSB) | EXPECTED LOGICAL BLOCK APPLICATION TAG | | | | | | (LSB) |
| 25 | Reserved | | | | | | | |
| 26 | (MSB) | LOGICAL BLOCK APPLICATION TAG MASK | | | | | | (LSB) |
| 27 | Reserved | | | | | | | |
| 28 | (MSB) | TRANSFER LENGTH (or other length field) | | | | | | (LSB) |
| 31 | Reserved | | | | | | | |