

Date: May 06, 2004

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: Changing Logical Block Sizes

1 Overview

It is common knowledge that the method used in SCSI to change the number of bytes in a logical block on a block device is to first issue a Mode Select command with the desired block size placed into the BLOCK LENGTH field of the mode parameter block descriptor, then issue a FORMAT command. The block device will then format the logical blocks on the media to the size that was specified in the BLOCK LENGTH field.

Unfortunately this common knowledge is not specified in either the SPC-3 nor the SBC-2 standards. Actually it has never been specified in any standard starting with SCSI-2.

This proposal defines this behavior.

2 SPC-3 changes

2.0.0.1 General block descriptor format

When the LONGLBA bit is set to zero (see 7.4.3), the mode parameter block descriptor format for all device types except direct-access is shown in table 1.

Table 1 — General mode parameter block descriptor

| Bit Byte | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
|-------------|--------------|---|---|---|------------------|---|---|-------|--|
| 0 | DENSITY CODE | | | | | | | | |
| 1 | (MSB) | | | | | | | | |
| 2 | | | | | NUMBER OF BLOCKS | | | | |
| 3 | | | | | | | | (LSB) | |
| 4 | Reserved | | | | | | | | |
| 5 | (MSB) | | | | | | | | |
| 6 | | | | | BLOCK LENGTH | | | | |
| 7 | | | | | | | | (LSB) | |

Block descriptors specify some of the medium characteristics for all or part of a logical unit. Support for block descriptors is optional. Each block descriptor contains a DENSITY CODE field, a NUMBER OF BLOCKS field, and a BLOCK LENGTH field. Block descriptor values are always current (i.e., saving is not supported). A unit attention condition (see 6.7 and SAM-2) shall be generated when any block descriptor values are changed. [Command set standards may place additional requirements on the general mode parameter block descriptor. Requirements in the command set standards that conflict with requirements defined in this subclause shall take precedence over the requirements defined in this subclause.](#)

The DENSITY CODE field is unique for each device type. Refer to the mode parameters subclause of the specific device type command standard (see 3.1.17) for definition of this field. Some device types reserve all or part of this field.

The NUMBER OF BLOCKS field specifies the number of logical blocks on the medium to which the DENSITY CODE and BLOCK LENGTH FIELDS apply. A value of zero indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.

If the number of logical blocks on the medium exceeds the maximum value that may be specified in the NUMBER OF BLOCKS field a value of FF FF FFh indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.

NOTES

- 1 There may be implicit association between parameters defined in the mode pages and block descriptors. In this case, the target may change parameters not explicitly sent with the MODE SELECT command. A subsequent MODE SENSE command may be used to detect these changes.
- 2 The number of remaining logical blocks may be unknown for some device types.

The BLOCK LENGTH field specifies the length in bytes of each logical block described by the block descriptor. For sequential-access devices, a block length of zero indicates that the logical block size written to the medium is specified by the transfer length field in the CDB (see SSC).

2.0.0.2 Direct-access device block descriptor format for ~~LONGLBA=0~~

~~When the LONGLBA bit is set to zero (see 7.4.3), the mode parameter block descriptor format for the direct-access device type is shown in table 2.~~

Table 2 — Direct-access device mode parameter block descriptor

| Bit Byte | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|------------------|---|---|---|---|---|---|-------|
| 0 | (MSB) | | | | | | | |
| 1 | | | | | | | | |
| 2 | NUMBER OF BLOCKS | | | | | | | |
| 3 | | | | | | | | (LSB) |
| 4 | DENSITY CODE | | | | | | | |
| 5 | (MSB) | | | | | | | |
| 6 | BLOCK LENGTH | | | | | | | |
| 7 | | | | | | | | (LSB) |

~~This block descriptor format shall apply only to direct-access devices. When the LONGLBA bit is set to zero (see 7.4.3), all other device types shall use the block descriptor format described in 2.0.0.1.~~

~~Block descriptors specify some of the medium characteristics for a logical unit. Support for block descriptors is optional. Each block descriptor contains a DENSITY CODE field, a NUMBER OF BLOCKS field, and a BLOCK LENGTH field. A unit attention condition (see 6.7 and SAM-2) shall be generated when any block descriptor values are changed.~~

~~The NUMBER OF BLOCKS field specifies the number of logical blocks on the medium to which the DENSITY CODE and BLOCK LENGTH fields apply. A value of zero indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.~~

~~If the number of logical blocks on the medium exceeds the maximum value that may be specified in the NUMBER OF BLOCKS field a value of FFFF FFFFh indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.~~

~~If the SCSI device doesn't support changing its capacity by changing the NUMBER OF BLOCKS field using the MODE SELECT command, the value in the NUMBER OF BLOCKS field is ignored. If the device supports~~

changing its capacity by changing the NUMBER OF BLOCKS field, then the NUMBER OF BLOCKS field is interpreted as follows:

- a) If the number of blocks is set to zero, the device shall retain its current capacity if the block size has not changed. If the number of blocks is set to zero and the block size has changed, the device shall be set to its maximum capacity when the new block size takes effect;
- b) If the number of blocks is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and L_T nexus losses;
- c) If the number of blocks field is set to a value greater than the maximum capacity of the device and less than FFFF FFFFh, then the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings; or
- d) If the number of blocks is set to FFFF FFFFh, the device shall be set to its maximum capacity. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and L_T nexus losses.

NOTE 3 - There may be implicit association between parameters defined in the mode pages and block descriptor. For direct access devices, the block length affects the optimum values (i.e., the values that achieves the best performance) for the sectors per track, bytes per physical sector, track skew factor, and cylinder skew factor fields in the format parameters mode page. In this case, the target may change parameters not explicitly sent with the MODE SELECT command. A subsequent MODE SENSE command may be used to detect these changes.

The DENSITY CODE field is unique for each device type. Refer to the mode parameters subclause of the specific device type command standard (see 3.1.17) for the definition of this field. Some device types reserve all or part of this field.

The BLOCK LENGTH field specifies the length in bytes of each logical block described by the block descriptor.

2.0.0.3 Long LBA block descriptor format

When the LONGLBA bit is set to one (see 7.4.3), the mode parameter block descriptor format for all device types is shown in table 3.

Table 3 — Long LBA mode parameter block descriptor

| Bit Byte | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-------------|------------------|---|---|---|---|---|-------------|---|
| 0 | (MSB) _____ | | | | | | | |
| 7 | NUMBER OF BLOCKS | | | | | | _____ (LSB) | |
| 8 | DENSITY CODE | | | | | | | |
| 9 | Reserved | | | | | | | |
| 10 | Reserved | | | | | | | |
| 11 | Reserved | | | | | | | |
| 12 | (MSB) _____ | | | | | | | |
| 15 | BLOCK LENGTH | | | | | | _____ (LSB) | |

Block descriptors specify some of the medium characteristics for all or part of a logical unit. Support for block descriptors is optional. Each block descriptor contains a DENSITY CODE field, a NUMBER OF BLOCKS field, and a BLOCK LENGTH field. Block descriptor values are always current (i.e., saving is not supported). A unit attention condition (see 6.7 and SAM-2) shall be generated when any block descriptor values are changed. Command set standards may place additional requirements on the long LBA mode parameter block descriptor. Require-

~~ments in the command set standards that conflict with requirements defined in this subclause shall take precedence over the requirements defined in this subclause.~~

~~The NUMBER OF BLOCKS field specifies the number of logical blocks on the medium to which the DENSITY CODE and BLOCK LENGTH fields apply. A value of zero indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.~~

~~If the SCSI device doesn't support changing its capacity by changing the NUMBER OF BLOCKS field using the MODE SELECT command, the value in the NUMBER OF BLOCKS field is ignored. If the device supports changing its capacity by changing the NUMBER OF BLOCKS field, then the NUMBER OF BLOCKS field is interpreted as follows:~~

- ~~a) If the number of blocks is set to zero, the device shall retain its current capacity if the block size has not changed. If the number of blocks is set to zero and the block size has changed, the device shall be set to its maximum capacity when the new block size takes effect;~~
- ~~b) If the number of blocks is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses and shall take effect on successful completion of the MODE SELECT command;~~
- ~~c) If the number of blocks field is set to a value greater than the maximum capacity of the device and less than FFFF FFFF FFFF FFFFh, then the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings; or~~
- ~~d) If the number of blocks is set to FFFF FFFF FFFF FFFFh, the device shall be set to its maximum capacity. If the block size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses.~~

~~NOTE 4 - There may be implicit association between parameters defined in the mode pages and block descriptor. For direct access devices, the block length affects the optimum values (i.e., the values that achieves the best performance) for the sectors per track, bytes per physical sector, track skew factor, and cylinder skew factor fields in the format parameters mode page. In this case, the target may change parameters not explicitly sent with the MODE SELECT command. A subsequent MODE SENSE command may be used to detect these changes.~~

~~The DENSITY CODE field is unique for each device type. Refer to the mode parameters subclause of the specific device type command standard (see 3.1.17) for the definition of this field. Some device types reserve all or part of this field.~~

~~The BLOCK LENGTH field specifies the length in bytes of each logical block described by the block descriptor.~~

3 SBC-2 changes

3.0.1 FORMAT UNIT command

3.0.1.1 FORMAT UNIT command overview

The FORMAT UNIT command (see table 13) formats the medium into application client addressable logical blocks ~~per~~ as specified in the number of blocks and the block length values received in the last MODE SELECT commands mode header block descriptor (see SPC-3). ~~application-client-defined options.~~ In addition, the medium may be certified and control structures may be created for the management of the medium and defects. The degree that the medium is altered by this command is vendor-specific.

If a device server receives a FORMAT UNIT command before receiving a MODE SELECT command with a block descriptor the device server shall use the number of blocks and block length at which the logical unit is currently formatted (i.e., no change is made to the number of blocks and the block length of the logical unit during the format operation).

3.1 Mode parameters

3.1.1 Mode parameters overview

This subclause defines the descriptors and pages for mode parameters used with direct-access device types.

The mode parameter list, including the mode parameter header ~~and mode block descriptor are~~ [is](#) described in SPC-3.

The MEDIUM TYPE field in the mode parameter header (see SPC-3) shall be set to 00h.

The DEVICE-SPECIFIC PARAMETER field (see table 4) is contained in the mode parameter header (see SPC-3).

Table 4 — DEVICE-SPECIFIC PARAMETER field

| Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-----|----|----------|---|--------|----------|---|---|---|
| | WP | Reserved | | DPOFUA | Reserved | | | |

When used with the MODE SELECT command the write protect (WP) bit is not defined.

When used with the MODE SENSE command a WP bit set to zero indicates that the medium is write enabled. A WP bit set to one indicates that the medium is write protected.

When used with the MODE SELECT command, the DPOFUA bit is reserved.

When used with the MODE SENSE command, a DPOFUA bit set to zero indicates that the device server does not support the DPO and FUA bits. When used with the MODE SENSE command, a DPOFUA bit set to one indicates that the device server supports the DPO and FUA bits (see 4.9).

3.1.2 Block descriptor format for `LONGLBA=0` ([this entire section is new to SBC-3](#))

When the LONGLBA bit is set to zero (see SPC-3), the mode parameter block descriptor format is as shown in table 5.

Table 5 — Mode parameter block descriptor

| Bit Byte | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-------------|--|---|---|---|---|---|---|-------|
| 0 | (MSB) | | | | | | | |
| 1 | NUMBER OF BLOCKS | | | | | | | |
| 2 | | | | | | | | |
| 3 | | | | | | | | (LSB) |
| 4 | DENSITY CODE Reserved | | | | | | | |
| 5 | (MSB) | | | | | | | |
| 6 | BLOCK LENGTH | | | | | | | |
| 7 | | | | | | | | |
| | | | | | | | | (LSB) |

Block descriptors specify some of the medium characteristics for a logical unit. Support for block descriptors is optional. Each block descriptor contains a ~~DENSITY CODE field~~, a NUMBER OF BLOCKS field and a BLOCK LENGTH field. A unit attention condition (see [SPC-3](#) and [SAM-3](#)) shall be generated when any block descriptor values are changed.

[On a MODE SELECT command](#) the NUMBER OF BLOCKS field specifies the number of logical blocks on the medium to which the ~~DENSITY CODE and~~ BLOCK LENGTH fields apply. A value of zero indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.

A device server shall respond to a MODE SENSE command (see SPC-3) by reporting the number of blocks specified in the NUMBER OF BLOCKS field sent in the last MODE SELECT command that contained a mode page descriptor. If no MODE SELECT command with a block descriptor has been received then the current number of blocks is returned. An application client shall ignore the value reported in the NUMBER OF BLOCKS field. To determine the number of blocks at which the logical unit is currently formatted the application client shall use the READ CAPACITY command (see x.x.x).

If the number of logical blocks on the medium exceeds the maximum value that may be specified in the NUMBER OF BLOCKS field a value of FFFF FFFFh indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.

If the SCSI device doesn't support changing its capacity by changing the NUMBER OF BLOCKS field using the MODE SELECT command (see SPC-3), the value in the NUMBER OF BLOCKS field is ignored. If the device supports changing its capacity by changing the NUMBER OF BLOCKS field, then the NUMBER OF BLOCKS field is interpreted as follows:

- a) If the NUMBER OF BLOCKS field is set to zero, the device shall retain its current capacity if the block size has not changed. If the NUMBER OF BLOCKS field is set to zero and the contents of the BLOCK SIZE field (i.e., new block size) is different than the current block size has changed, the device shall be set to its maximum capacity when the new block size takes effect (i.e., after a successful FORMAT UNIT command);
- b) If the NUMBER OF BLOCKS field is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the contents of the BLOCK SIZE field is the same as the current block size block-size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses. If the contents of the BLOCK LENGTH field is the same as the current block length this capacity setting shall take effect on successful completion of the MODE SELECT command. If the contents of the BLOCK LENGTH field (i.e., new block size) is different than the current block length this capacity setting shall take effect when the new block size takes effect (i.e., after a successful FORMAT UNIT command);
- c) If the NUMBER OF BLOCKS field is set to a value greater than the maximum capacity of the device and less than FFFF FFFFh, then the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings; or
- d) If the NUMBER OF BLOCKS field is set to FFFF FFFFh, the device shall be set to its maximum capacity. If the contents of the BLOCK SIZE field is the same as the current block size block-size has not changed, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses If the contents of the BLOCK LENGTH field is the same as the current block length this capacity setting shall take effect on successful completion of the MODE SELECT command. If the contents of the BLOCK LENGTH field (i.e., new block size) is different than the current block length this capacity setting shall take effect when the new block size takes effect (i.e., after a successful FORMAT UNIT command).

~~NOTE 5 - There may be implicit association between parameters defined in the mode pages and block descriptor. For direct access devices, the block length affects the optimum values (i.e., the values that achieves the best performance) for the sectors per track, bytes per physical sector, track skew factor, and cylinder skew factor fields in the format parameters mode page. In this case, the target may change parameters not explicitly sent with the MODE SELECT command. A subsequent MODE SENSE command may be used to detect these changes.~~

~~The DENSITY CODE field is unique for each device type. Refer to the mode parameters subclause of the specific device type command standard (see 3.1.17) for the definition of this field. Some device types reserve all or part of this field.~~

The BLOCK LENGTH field specifies the length in bytes of each logical block described by the block descriptor. No change shall be made to any logical blocks on the medium until a format operation (see 5.3xx) is initiated by an application client

A device server shall respond to a MODE SENSE command (see SPC-3) by reporting the ~~actual~~ size of the logical blocks as specified in the BLOCK LENGTH field sent in the last MODE SELECT command that contained a mode page descriptor. If no MODE SELECT command with a block descriptor has been received then the current block size is returned. (e.g., if the block size is 512 bytes and a MODE SENSE command occurs with the BLOCK LENGTH field set to 520 bytes, any MODE SELECT commands would return 512 in the BLOCK LENGTH field until a FORMAT UNIT command has completed). ~~An application client shall ignore the value reported in the BLOCK LENGTH field.~~ To determine the block size at which the logical unit is currently formatted the application client shall use the READ CAPACITY command (see x.x.x).

3.1.2.1 Long LBA block descriptor format (this entire section is new to SBC-3)

When the LONGLBA bit is set to one (see SPC-3), the mode parameter block descriptor format is as shown in table 6.

Table 6 — Long LBA mode parameter block descriptor

| Bit Byte | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-------|----------------------------------|---|---|---|---|---|-------|
| 0 | (MSB) | NUMBER OF BLOCKS | | | | | | (LSB) |
| 7 | | | | | | | | |
| 8 | | DENSITY CODE Reserved | | | | | | |
| 9 | | Reserved | | | | | | |
| 10 | | Reserved | | | | | | |
| 11 | | Reserved | | | | | | |
| 12 | (MSB) | BLOCK LENGTH | | | | | | (LSB) |
| 15 | | | | | | | | |

Block descriptors specify some of the medium characteristics for all ~~or part~~ of a logical unit. Support for block descriptors is optional. Each block descriptor contains a ~~DENSITY CODE~~ field, a NUMBER OF BLOCKS field and a BLOCK LENGTH field. Block descriptor values are always current (~~i.e., saving is not supported~~). A unit attention condition (see ~~6.7~~ SPC-3 and SAM-2) shall be generated when any block descriptor values are changed.

On a MODE SELECT command the NUMBER OF BLOCKS field specifies the number of logical blocks on the medium to which the ~~DENSITY CODE and~~ BLOCK LENGTH fields apply. A value of zero indicates that all of the remaining logical blocks of the logical unit shall have the medium characteristics specified.

A device server shall respond to a MODE SENSE command (see SPC-3) by reporting the number of blocks specified in the NUMBER OF BLOCKS field sent in the last MODE SELECT command that contained a mode page descriptor. If no MODE SELECT command with a block descriptor has been received then the current number of blocks is returned. ~~An application client shall ignore the value reported in the NUMBER OF BLOCKS field.~~ To determine the number of blocks at which the logical unit is currently formatted the application client shall use the READ CAPACITY command (see x.x.x).

If the SCSI device doesn't support changing its capacity by changing the NUMBER OF BLOCKS field using the MODE SELECT command (see SPC-3), the value in the NUMBER OF BLOCKS field is ignored. If the device supports changing its capacity by changing the NUMBER OF BLOCKS field, then the NUMBER OF BLOCKS field is interpreted as follows:

- a) If the NUMBER OF BLOCKS field is set to zero, the device shall retain its current capacity if the block size has not changed. If the NUMBER OF BLOCKS field is set to zero and the contents of the BLOCK SIZE field (i.e., new block size) is different than the current block size has changed, the device shall be set to its maximum capacity when the new block size takes effect (i.e., after a successful FORMAT UNIT command);

- b) If the NUMBER OF BLOCKS field is greater than zero and less than or equal to its maximum capacity, the device shall be set to that number of blocks. If the contents of the BLOCK SIZE field is the same as the current block size ~~block size has not changed~~, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses. If the contents of the BLOCK LENGTH field is the same as the current block length this capacity setting shall take effect on successful completion of the MODE SELECT command. If the contents of the BLOCK LENGTH field (i.e., new block size) is different than the current block length this capacity setting shall take effect when the new block size takes effect (i.e., after a successful FORMAT UNIT command);
- c) If the NUMBER OF BLOCKS field is set to a value greater than the maximum capacity of the device and less than FFFF FFFFh, then the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST. The device shall retain its previous block descriptor settings; or
- d) If the NUMBER OF BLOCKS field is set to FFFF FFFF FFFF FFFFh, the device shall be set to its maximum capacity. If the contents of the BLOCK SIZE field is the same as the current block size ~~block size has not changed~~, the device shall not become format corrupted. This capacity setting shall be retained through power cycles, hard resets, logical unit resets, and I_T nexus losses If the contents of the BLOCK LENGTH field is the same as the current block length this capacity setting shall take effect on successful completion of the MODE SELECT command. If the contents of the BLOCK LENGTH field (i.e., new block size) is different than the current block length this capacity setting shall take effect when the new block size takes effect (i.e., after a successful FORMAT UNIT command).

~~NOTE 6 - There may be implicit association between parameters defined in the mode pages and block descriptor. For direct access devices, the block length affects the optimum values (i.e., the values that achieves the best performance) for the sectors per track, bytes per physical sector, track skew factor, and cylinder skew factor fields in the format parameters mode page. In this case, the target may change parameters not explicitly sent with the MODE SELECT command. A subsequent MODE SENSE command may be used to detect these changes.~~

~~The DENSITY CODE field is unique for each device type. Refer to the mode parameters subclause of the specific device type command standard (see 3.1.17) for the definition of this field. Some device types reserve all or part of this field.~~

The BLOCK LENGTH field specifies the length in bytes of each logical block described by the block descriptor. No change shall be made to any logical blocks on the medium until a format operation (see 5.3xx) is initiated by an application client

A device server shall respond to a MODE SENSE command (see SPC-3) by reporting the actual size of the logical blocks as specified in the BLOCK LENGTH field sent in the last MODE SELECT command that contained a mode page descriptor. If no MODE SELECT command with a block descriptor has been received then the current block size is returned. (e.g., if the block size is 512 bytes and a MODE SENSE command occurs with the BLOCK LENGTH field set to 520 bytes, any MODE SELECT commands would return 512 in the BLOCK LENGTH field until a FORMAT UNIT command has completed). ~~An application client shall ignore the value reported in the BLOCK LENGTH field.~~ To determine the block size at which the logical unit is currently formatted the application client shall use the READ CAPACITY command (see x.x.x).