To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 15 December 2003
Subject: 03-387r1 SBC-2 Data protection usage detection

Revision history
Revision 0 (1 November 2003) First revision
Revision 1 (15 December 2003) Incorporated comments from November 2003 CAP WG - moved the indication to READ CAPACITY (16).

Related documents
sbc2r11 - SCSI Block Commands - 2 revision 11
03-365 SPC-3 SBC-2 End-to-end data protection (George Penokie)

Overview
03-365 defines data protection for block devices. 03-365r0 requested a two-bit PROTECT field to standard INQUIRY data indicating if the media has been formatted with protection information or not.

This introduced a few problems, especially for tapes:

a) Standard INQUIRY data doesn't normally change more than once after a reset. With the PROTECT field, a unit attention can now be created whenever a FORMAT UNIT completes.
b) Standard INQUIRY data is normally not dependent on the media. With the PROTECT field, a unit attention can now be created whenever media is changed. The media serial number feature (01-027) was moved away from an INQUIRY VPD page because of this (May 2001).
c) Standard INQUIRY data is rather full and should avoid hosting device type-specific fields.
d) Tapes with data protection support are probably not going to be "formatted to include protection information." They will likely store a bit indicating whether protection information is in use on-the-fly on a file-by-file or partition-by-partition basis. The proposed values won't apply very well.
e) Tapes will probably require a mode page to enable/disable inclusion of protection information to instead of modifying the FORMAT UNIT command (which suffices for disks).

At the November 2003 CAP meeting, 03-365 was revised to just add a 1-bit protect field to INQUIRY indicating the device server supports data protection. This revision of this proposal expands the READ CAPACITY (16) data to indicate whether the medium was formatted with protection information (revision 0 proposed using the Format Status log page).

Suggested changes to SBC-2

4.5 Protection information model

4.5.1 Protection information overview

This data protection model provides for protection of the data while it is being transferred between a sender and a receiver. Protection information is generated at the application layer and may be checked by any object along the I_T_L nexus. Once received, protection information is retained (e.g., write to media, store in non-volatile memory, recalculate on read back) by the device server until overwritten (e.g., power loss, hard reset, logical unit reset, and I_T nexus loss have no effect on the retention of protection information).

Support for protection information shall be indicated in the PROTECT field of standard INQUIRY data (see SPC-3).

...

4.5.9 FORMAT UNIT command

4.5.9.1 FORMAT UNIT command overview

The FORMAT UNIT command (see table 3) formats the medium into application client addressable logical blocks per the 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.

...

A format protection information (FMTPINFO) bit of zero specifies that the device server shall format the medium to the block length specified in the mode parameter block descriptor of the mode parameter header (see SPC-3). A FMTPINFO bit of one specifies that the device server shall format the medium to the block length specified in the mode parameter block descriptor of the mode parameter header plus eight (e.g., if the block length equals 512 the formatted block length is 520). A successful format that changes whether protection information (see 4.5) is included shall cause the ~~PROTECT~~PROT_EN ~~-~~field in the READ CAPACITY (16) data to be changed.

...

5.2.11 READ CAPACITY (10) command

The READ CAPACITY (10) command (see table 1) provides a means for the application client to request information regarding the capacity of the block device. This command may be processed as if it has a HEAD OF QUEUE task attribute (see 4.2.1.6). If the logical unit supports protection information (see 4.5), the READ CAPACITY (16) command should be used instead of READ CAPACITY (10).

Table 1 — READ CAPACITY (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (25h)							
1	Reserved							Obsolete
2	(MSB)							
3								
4				LOGICAL BLOCK ADDRESS				
5								(LSB)
6								
7				Reserved				
8	Reserved							PMI
9	CONTROL							

See 4.2.1.9 for reservation requirements for this command. See the LOCK UNLOCK CACHE (10) command (see 5.2.3) for a definition of the LOGICAL BLOCK ADDRESS field.

The LOGICAL BLOCK ADDRESS field shall be zero if the PMI bit is zero. If the PMI bit is zero and the LOGICAL BLOCK ADDRESS field is not set to zero, the device server shall return a CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST with the additional sense code set to ILLEGAL FIELD IN CDB.

A partial medium indicator (PMI) bit of zero specifies that the device server return information on the last logical block on the block device.

A PMI bit of one specifies that the device server return information on the last logical block after that specified in the LOGICAL BLOCK ADDRESS field before a substantial delay in data transfer may be encountered.

NOTE 1 - This function is intended to assist storage management software in determining whether there is sufficient space on the current track, cylinder, etc., to contain a frequently accessed data structure, such as a file directory or file index, without incurring an access delay.

The short read capacity data (see table 2) shall be sent during the data-in buffer transfer of the command.

Table 2 — Short read capacity data

Byte\Bit	7	6	5	4	3	2	1	0
0	(MSB)			RETURNED LOGICAL BLOCK ADDRESS				
3								(LSB)
4	(MSB)			BLOCK LENGTH IN BYTES				
7								(LSB)

If the number of logical blocks exceeds the maximum value that may be specified in the RETURNED LOGICAL BLOCK ADDRESS field, the device server shall transfer FFFFFFFFh in the RETURNED LOGICAL BLOCK ADDRESS field. The initiator should then issue a READ CAPACITY (16) command.

If the PMI bit is zero, the RETURNED LOGICAL BLOCK ADDRESS field contains the LBA of the last logical block on the block device or FFFFFFFFh, whichever is lower.

If the PMI bit is one, the RETURNED LOGICAL BLOCK ADDRESS field contains the last LBA after that specified in the LOGICAL BLOCK ADDRESS field of the CDB before a substantial delay in data transfer may be encountered or FFFFFFFFh, whichever is lower. The RETURNED LOGICAL BLOCK ADDRESS shall be greater than or equal to that specified by the LOGICAL BLOCK ADDRESS field in the CDB.

The BLOCK LENGTH IN BYTES field contains the number of bytes of user data in the logical block indicated by the RETURNED LOGICAL BLOCK ADDRESS field. This value does not include protection information or additional information (e.g., ECC bytes) recorded on the medium.

0.0.1 READ CAPACITY (16) command

The READ CAPACITY (16) command (see table 3) provides a means for the application client to request information regarding the capacity of the block device. This command is implemented as a service action of the SERVICE ACTION IN opcode. This command may be processed as if it has a HEAD OF QUEUE task attribute (see 4.2.1.6)

Table 3 — READ CAPACITY (16) command

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

See 4.2.1.9 for reservation requirements for this command. See the LOCK UNLOCK CACHE (10) command (see 5.2.3) for a definition of the LOGICAL BLOCK ADDRESS field. See the READ CAPACITY (10) command (see) for a description of the other fields in this command.

The long read capacity data (see table 4) shall be sent during the data-in buffer transfer of the command. The maximum value that shall be returned in the RETURNED LOGICAL BLOCK ADDRESS field is FFFFFFFF FFFFFFFEh.

Table 4 — Long read capacity data

Byte\Bit	7	6	5	4	3	2	1	0
0	(MSB)			RETURNED LOGICAL BLOCK ADDRESS				
7								(LSB)
8	(MSB)			BLOCK LENGTH IN BYTES				
11								(LSB)
12	Reserved							PROT_EN
13				Reserved				
31								

The RETURNED LOGICAL BLOCK ADDRESS field and BLOCK LENGTH IN BYTES field of the long read capacity data are the same as in the short read capacity data described in the READ CAPACITY (10) command (see).

A PROT_EN bit set to one indicates that the medium was formatted with protection information (see 4.5) enabled. A PROT_EN bit set to zero indicates that the medium was not formatted with protection information enabled.