



where information lives®

To: INCITS Technical Committee T10

From: David L. Black, EMC

Email: black_david@emc.com

Date: Sep 4, 2008

Subject: SPC-4, SBC-3: Reporting Thin Provisioning Support (08-357r0)

1) Revision history

Revision 0 (Sep 4, 2008) First revision (r0)

2) Related documents

spc4r16 – SCSI Primary Commands – 4

sbc3r16 – SCSI Block Commands – 3

ssc3r04a – SCSI Sequential Commands – 3

08-149r1 – Thin Provisioning Commands

08-347r0 – TRIM: Behavior of subsequent READs

08-356r0 – TRIM bit for WRITE SAME

T13/e07154r6 – ATA8-ACS2 accepted TRIM proposal

3) Overview

08-149r1 proposes that a bit in the parameter data for READ CAPACITY (16) be used to report support for thin provisioning. This proposal suggests two alternative mechanisms.

There is a case to be made that support for thin provisioning should be reported in a fashion that has minimal impact on existing application client code. Thin Provisioning has been implemented in system software (e.g., filesystem) on hosts; there is utility to reporting use of thin provisioning on the storage so that the host can disable its thin provisioning, even if the host supports no other SCSI-related thin provisioning functionality. This suggests that device thin provisioning support should be reported by a mechanism already widely used by host software.

08-149r1 proposes that a bit in the parameter data for READ CAPACITY (16) be used to report support for thin provisioning. READ CAPACITY (10) appears to be far more widely used rather than READ CAPACITY (16), and EMC has seen code that uses an ALLOCATION LENGTH of 12 bytes with READ CAPACITY (16) and hence would not obtain the bit proposed in 08-149r1. Two widely used commands that could be used to report thin provisioning are INQUIRY and READ CAPACITY (10). This proposal includes an alternative for use of each command to reporting thin provisioning support; only one of these alternatives should be adopted.

Existing text is shown in **BLACK**, new text is shown in **RED**, and comments (not to be included) are shown in **BLUE**.

Proposal:

SPC-4: 6.4.2 Standard INQUIRY data (Alternative 1)

If thin provisioning is expected to be supported by command sets in addition to SBC-4, an INQUIRY is an appropriate mechanism to report thin provisioning support. The first 36 bytes of INQUIRY parameter data are always retrieved, and hence the bit should be located there. There are a limited number of possible locations.

The standard INQUIRY data (see table 130) shall contain at least 36 bytes.

Table 130 – Standard INQUIRY data format

Bit	7	6	5	4	3	2	1	0
Byte								
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	RMB	Reserved						
2	VERSION							
3	Obsolete	Obsolete	NormACA	HiSup	RESPONSE DATA FORMAT			
4	ADDITIONAL LENGTH (n-4)							
5	SCCS	ACC	TPGS		3PC	Reserved		PROTECT
6	Obsolete	EncServ	VS	MultiP	Obsolete	Obsolete	Obsolete	ADDR16 ^a
7	Obsolete	Obsolete	WBUS16 ^a	SYNC ^a	Obsolete	Obsolete	CmdQue	VS
8	(MSB) T10 VENDOR IDENTIFICATION (LSB)							
15								
16	(MSB) PRODUCT IDENTIFICATION (LSB)							
31								
32	(MSB) PRODUCT REVISION LEVEL (LSB)							
35								
Table Continues								
^a The meanings of these fields are specific to SPI-5 (see 6.4.3). For SCSI transport protocols other than the SCSI Parallel Interface, these fields are reserved.								

Possible unused (Reserved) locations for a TPE (Thin Provisioning Enabled) bit:

- Byte 1, bit 6: RMB (bit 7) indicates whether the medium is removable. Thin provisioning also affects all the logical blocks. Bits 0-6 were the Device-type modifier in SCSI-2, but have been Reserved since at least the original SPC standard.
- Byte 5, bit 1 or 2: Available - the two remaining reserved bits for features.

This is a scarce resource, but one of the bits in Byte 5 would be the best choice.

Possible reusable (Obsolete) locations for a TPE (Thin Provisioning Enabled) bit. These locations risk failing initiator checks that the obsolete feature is not supported:

- Byte 3: Not appropriate, the fields in this byte affect command processing.
- Byte 6, bits 7 & 3: Not available, these bits are defined in SPC-3 (BQue and MChngr).
- Byte 6, bits 2 & 1: May be available. These bits were SPI-specific (ACKREQQ and Addr32), in SPC, but were already Obsolete in SPC-2.
- Byte 7, bit 7: Probably not available. This was the RelAdr (Relative Addressing) bit for linked commands in SPC-2, and was Obsolete in SPC-3. Linked commands are Obsolete in the SPC-4 draft, but the LINKED bit (Byte 7, bit 3) is defined in SPC-3, and hence SPC-3-based code could reasonably be checking that this related bit is zero.
- Byte 7, bit 6: May be available. This bit was SPI-specific (WBus32) in SPC, and was already Obsolete in SPC-2.
- Byte 7, bit 3: Not available, defined in SPC-3 (LINKED).
- Byte 7, bit 2: May be available. This bit was SPI-specific (TranDis) in SPC, and was already Obsolete in SPC-2.

The best candidates for reuse are SPI-specific bits that have been Obsolete since SPC-2 (Byte 6, bits 2 & 1; Byte 7, bits 6 & 2).

T10 may wish to define a bit in the INQUIRY CDB that releases Obsolete parameter data bits for other uses (if the bit is set in the CDB, some or all of the Obsolete parameter data bits become Reserved and available for other uses without needing to be concerned about previous use).

08-149r1 defines the TPE bit as:

The TPE bit set to one indicates that this logical unit implements Thin Provisioning (see 4.4.1.2).

SBC-3: 5.12 READ CAPACITY (10) command (Alternative 2)

There is no room available in the current format for the READ CAPACITY (10) parameter data. There are 4 bytes used for the logical block length, which is beyond conceivable needs. One byte could be used for other purposes, such as reporting thin provisioning support. If this is done, the same approach should be applied to READ CAPACITY (16) instead of the bit defined in 08-149r1.

Table 44 — READ CAPACITY (10) command

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

Insert the following paragraph after Note 13 and before the paragraph defining the CONTROL byte:

The INFO bit specifies that the parameter data shall include device information in byte 4 (see 5.12.2).

5.12.2 READ CAPACITY (10) parameter data

If the INFO bit is set to zero, the READ CAPACITY (10) parameter data is defined in table 45. If the INFO bit is set to one, the READ CAPACITY (10) parameter data is defined in table 45a. Any time the READ CAPACITY (10) parameter data changes, the device server should establish a unit attention condition as described in 4.7.

Table 45 — READ CAPACITY (10) parameter data [INFO bit set to zero]

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

All of Table 45a is new - the red fields highlight the changes from Table 45.

Table 45a — READ CAPACITY (10) parameter data [INFO bit set to one]

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

Add the following text from 08-149r1 that defines the TPE bit to the end of 5.12.2:

The TPE bit set to one indicates that this logical unit implements Thin Provisioning (see 4.4.1.2).