

March 15, 1999

T10/99-144 revision 1



To: John Lohmeyer
NCITS Technical Committee T10
From: Bob Snively
Subject: Suggested changes to REPORT LUNS command for SPC-2

1 Introduction

Today's RAID devices and Fibre Channel bridge devices have a very large logical unit address space and potentially large numbers of logical units actually implemented. The REPORT LUNS command has become a necessary tool for probing for the presence of logical units in that large logical unit address space. Use of the command by operating systems and applications is very similar to the use of the INQUIRY command during the probe process, and is usually required before the INQUIRY command can be executed. For this reason, the REPORT LUNS command's definition should be modified to behave like an INQUIRY command with respect to reservations, Unit Attention conditions, and other similar functions that may inadvertently interfere with the required configuration probing functions.

SPC-2, revision 8, has now defined the ASC/ASCQ of REPORTED LUNS DATA HAS CHANGED, 3Fh 0Eh, for all devices. This is the expected Unit Attention to warn that a LUN has either come into existence, changed type, or been destroyed. Since EPI revision 16 is vague about how the Unit Attention shall be offered, a proposal for that mechanism is also included in this document.

The following changes in SPC-2, revision 8, and SAM-2, revision 9, are required to support the proposed behavior.

Revision 1 of this proposal corrects some wording in unrelated parts of SPC-2 concerning the REPORT LUNS command. In addition, some of the modified text is improved and clarified and a glossary entry in SAM-2 is provided to clarify the explanation of the REPORT LUNS command.

2 Changes required in SPC-2, revision 8.

The text of section 7.19, REPORT LUNS (pdf page 100 and 101) needs to be modified as shown below to include the requirement that the REPORT LUNS command be able to execute in the presence of check conditions and unit attention conditions. The text is similar to section 7.4 (pdf pages 62 and 63), describing the INQUIRY command.

7.19 REPORT LUNS command

The REPORT LUNS command (see table 60) requests that the peripheral device logical unit ~~inventory numbers of known logical units in the target~~ be sent to the application client. ~~The REPORT LUNS command shall return information about only those logical units to which commands may be sent. When the HISUP bit is one (see 7.4.1), the~~

~~device server shall support the REPORT LUNS command.~~ The logical unit inventory is a list that shall include the logical unit numbers of all logical units having a peripheral qualifier of 000. Logical unit numbers for logical units with a peripheral qualifier of 100, 101, 110, or 111 may optionally be included in the logical unit inventory. A SCSI device that is capable of supporting a LUN address other than zero shall ~~should~~ support the REPORT LUNS command on logical unit zero. Support of the REPORT LUNS command on logical units other than logical unit zero is optional. Support of the REPORT LUNS command on devices having only a single logical unit with the logical unit number of zero is optional.

The allocation length shall be at least 16 bytes. If the allocation length is less than 16 bytes, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense data shall be set to INVALID FIELD IN CDB.

~~The~~ If the allocation length is not sufficient to contain the logical unit number values for all configured logical units, the device server shall report as many logical unit number values as ~~will~~ fit in the specified allocation length. This shall not be considered an error.

The device server shall report ~~the logical unit numbers of those devices in the logical unit inventory configured logical units~~ using the format shown in table 61.

The LUN LIST LENGTH field shall contain the length in bytes of the LUN list that is available to be transferred. The LUN list length is the number of logical unit numbers ~~reported in the logical unit inventory~~ multiplied by eight. If the allocation length in the command descriptor block is too small to transfer information about ~~the entire logical unit inventory, all configured logical units,~~ the LUN list length value shall not be adjusted to reflect the truncation.

The REPORT LUNS command shall return CHECK CONDITION status only when the device server cannot return the requested report of the logical unit inventory.

If a REPORT LUNS command is received from an initiator with a pending unit attention condition (i.e., before the device server reports CHECK CONDITION status), the device server shall perform the REPORT LUNS command. If the unit attention condition was established because of a change in the logical unit inventory, that unit attention condition shall be cleared for that initiator by the REPORT LUNS command. Unit attention conditions established for other reasons shall not be cleared by the REPORT LUNS command (see SAM-2).

The REPORT LUNS data should be returned even though the device server is not ready for other commands. To minimize delays after a hard reset or power-up condition, the default report of the logical unit inventory should be available without incurring any media access delays. The default report of the logical unit inventory shall contain at least LUN 0.

If the logical unit inventory changes for any reason, including completion of initialization, removal of a logical unit, or creation of a logical unit, the device server shall generate a unit attention command for all initiators (see SAM-2). The device server shall set the additional sense code to REPORTED LUNS DATA HAS CHANGED.

The execution of a REPORT LUNS command to any valid and installed logical unit shall clear the REPORTED LUNS DATA HAS CHANGED unit attention for all logical

units of that target to the requesting initiator. A valid and installed logical unit is one having a peripheral qualifier of 000 if the INQUIRY command were executed.

3 Changes required in SAM-2, revision 9

An additional glossary entry should be added to define “logical unit inventory”.
The following text is proposed:

3.1.x logical unit inventory: The list of the logical unit numbers that shall be reported by a REPORT LUNS command. The list shall include the logical unit numbers of all logical units with a peripheral qualifier of 000 and may include in a vendor specific manner the logical unit numbers for those logical units with a peripheral qualifier of 100, 101, 110, or 111, as defined by SPC-2.

Section 5.6.5 on pdf pages 72 and 73 needs to be modified in the following manner.
The modifications clarify when a unit attention may be caused and how it interacts with the REPORT LUNS command.

5.6.5 Unit Attention condition

Each logical unit shall generate a unit attention condition whenever the logical unit has been reset as described in 5.6.6 or by a power-on reset. In addition, a logical unit shall generate a unit attention condition for each initiator whenever one of the following events occurs:

- a) A removable medium may have been changed;
- b) The mode parameters in effect for this initiator have been changed by another initiator;
- c) The version or level of microcode has been changed;
- d) Tasks for this initiator were cleared by another initiator;
- e) INQUIRY data has been changed;
- f) The mode parameters in effect for the initiator have been restored from non-volatile memory;
- g) A change in the condition of a synchronized spindle; or
- h) Any other event requiring the attention of the initiator.

Logical units may queue unit attention conditions. After the first unit attention condition is cleared, another unit attention condition may exist (e.g., a power on condition followed by a microcode change condition).

A unit attention condition shall persist on the logical unit for each initiator until that initiator clears the condition as described in the following paragraphs.

If an INQUIRY or a REPORT LUNS command is received from an initiator to a logical unit with a pending unit attention condition (before the logical unit generates the auto contingent allegiance or contingent allegiance condition), the logical unit shall perform the INQUIRY command. If the unit attention condition was established because of a change in the logical unit inventory, the unit atten-

tion condition shall be cleared for that initiator for all logical units by a REPORT LUNS command. In all other cases, the INQUIRY or REPORT LUNS command and shall not clear the unit attention condition.

If a request for sense data is received from an initiator with a pending unit attention condition (before the logical unit establishes the auto contingent allegiance or contingent allegiance condition), then the logical unit shall either:

- a) Report any pending sense data and preserve the unit attention condition on the logical unit; or,
- b) Report the unit attention condition.

If the second option is chosen (reporting the unit attention condition), the logical unit may discard any pending sense data and may clear the unit attention condition for that initiator.

If the logical unit has already generated the auto contingent allegiance or contingent allegiance condition for the unit attention condition, the logical unit shall perform the second action listed above. If NACA for the REQUEST SENSE command is zero and the command is untagged the contingent allegiance condition shall be cleared.

If an initiator issues a command other than INQUIRY, REPORT LUNS, or REQUEST SENSE while a unit attention condition exists for that initiator (prior to generating the auto contingent allegiance or contingent allegiance condition for the unit attention condition), the logical unit shall not perform the command and shall report ACA ACTIVE (NACA=1, see 5.1.2) or BUSY (NACA=0) status.

If a logical unit successfully sends an asynchronous event report informing the initiator of the unit attention condition, then the logical unit shall clear the unit attention condition for that initiator on the logical unit (see 5.6.4.1).

4 Additional problems identified during study of this proposal

4.1 Incomplete definition of task completion, SAM-2 revision 9

Section 5.4 of SAM-2 (pdf page 65) indicates that one of the conditions marking the end of a task is:

- b) A unit attention condition with one of the following additional sense codes:
 - a) COMMANDS CLEARED BY ANOTHER INITIATOR (if in reference to the task set containing the task);
 - b) POWER ON;
 - c) RESET; or
 - d) TARGET RESET.

This is really incomplete. Other text in this section indicates that service responses of task complete (including GOOD status or CHECK CONDITION status) for the referenced task indicate task completion. However, it is not made explicit that the



case described above is for any unit attention for any task for any initiator. I would propose that item b) be changed to read:

- | b) A Notification of a unit attention condition with one of the following additional sense codes:
 - a) COMMANDS CLEARED BY ANOTHER INITIATOR (if in reference to the task set containing the task);
 - b) POWER ON;
 - c) RESET; or
 - d) TARGET RESET.

