To: INCITS T10 Committee

From: Paul Entzel, Quantum

Date: 26 August 2005

Document: T10/05-320r0

Subject: SSC-3: Standardize processing of ERASE command with LONG=0

1 Revision History

Revision 0:
Posted to the T10 web site on 26 August 2005.

2 General

Changes in tape drive technology have made the need to control gap sizes on the tape obsolete. This fact was noted and addressed by proposal T10/05-066r0 which was approved for inclusion in SSC-3 by the SSC-3 working group on March 9, 2005 (see T10-05-111). This change leaves the concept of an ERASE command with the LONG bit set to 0 (henceforth referred to as a “short erase”) undefined. This proposal attempts to address this issue by defining actions to be taken by a device server that receives just such a command.

This problem is not new, just exasperated by the removal of the concept of gaps size from the standard. In fact, the concept of an erase gap as defined in the definition and model sections of SSC-3 have not been incorporated in most tape drive designs for years. The problem of how to deal with a short erase command on a device that does not place gaps between logical objects is far from new. The action taken on a short erase command varies by products, and this proposal attempts to standardize the actions taken while maintaining the flexibility required by the different technologies and implementations already existing. To do this, a mode parameter is introduced that allows the device to report how it will process a short erase. By using a mode parameter, a device server can support more than one option by making the field changeable via a MODE SELECT command.

The mode parameter proposed has been placed in a new mode page called the Device Configuration Extension mode page, a subpage of the Device Configuration mode page. This is a new mode page, but not a new concept, since this subpage has been proposed already, several times. See T10/05-140 and T10/05-154 for examples of other proposals that add this mode page. Unfortunately, none of these proposals have yet been approved for inclusion in SSC-3. My intention is not to conflict with these proposals, but simply to allocate some space from the same new page to add a new field.

One of the options for how a short erase is processed is to do it as defined by SSC-2. However, that is the only option that deals with the concept of erase gaps. Since the short erase command is the only place that requires a definition of a gap, the definitions will be altered to refer to SSC-2 or removed when possible. This should simplify the model section of the standard as well as catch it up with modern tape drive designs.

An alternative method of dealing with different implementations for short erase that was considered involved adding more bits to the ERASE CDB or changing the LONG bit to a multiple bit field with a value of 1 defined as what LONG means now. This approach has merit, but was not used due to the problems it would create for an application to discover what modes are supported by the drive. Using a mode parameter allow the drive to report what it does by default, and if the drive supports more than one option, the field can be reported as changeable so the application client can support its favorite option.

3 Changes to SSC-3

3.1 Clause 3, definitions

~~**3.1.26 gap:** A non-data object recorded on the medium. Gaps may be recorded between logical objects. The format and method of recording a gap may vary.~~

3.2 Clause 4.2.5.1, logical objects within a partition

~~Inter-block gaps, the gaps between logical objects, are introduced on the medium at the time a logical object is written without explicit action by the application client. Minimum and maximum lengths for inter-block gaps are defined by the recording format.~~

~~Erase gaps may be recorded on the medium through use of the ERASE command or device-initiated error recovery actions. Although explicitly recorded on the medium, there is normally no distinction between two contiguous erase gaps. An erase gap may be a length of erased medium or a recorded pattern not distinguishable as a logical object. Minimum and maximum lengths for erase gaps are defined by the recording format.~~

After writing data from BOP *x*, the medium is considered to be a contiguous grouping of logical objects ~~and gaps~~. ~~Certain American National Standards define gap lengths that, if exceeded, are to be considered as having reached blank medium (i.e., end-of-data within a partition).~~ Depending on the format, ~~this~~ blank medium may be treated as an end-of-data indication, an error recovery area, or an unrecoverable medium error causing an interchange error. Unrecorded volumes, new or erased, may exhibit blank medium characteristics if an attempt is made to read or space the volume before data has been written.

3.3 Clause 5.2, Erase (16) command

The ERASE(16) command (see table 13) causes part or all of the medium to be erased beginning at the logical object identifier and partition specified in the command descriptor block. ~~Erased means the medium shall be erased or a pattern shall be written on the medium that appears as a gap.~~ Prior to performing the erase operation, the device server shall perform a synchronize operation (see 4.2.8).

…

An immediate (IMMED) bit of zero specifies the device server shall not return status until the erase operation has completed. Interpretation of an IMMED bit of one depends on the value of the LONG bit, see below. However, for all values of the LONG bit, if CHECK CONDITION status is returned for an ERASE(16) command with an IMMED bit of one, the erase operation shall not be performed.

A LONG bit of one specifies all remaining medium ~~in the current partition~~ ~~shall be erased~~ beginning at the specified logical object identifier and partition shall be erased or over-written with a format specific pattern that renders any over-written data on the medium unreadable. If the format on the medium specifies a recorded indication of EOD (see 3.1.6), the erase operation shall establish an EOD indication at the specified location as part of the erase operation. If the IMMED bit is one, the device server shall return status as soon as all buffered logical objects have been written to the medium and the command descriptor block of the ERASE(16) command has been validated. The logical position following an ERASE(16) command with a LONG bit of one is not specified by this standard.

NOTE 7 Some logical units may reject an ERASE(16) command ~~with the LONG bit set to one~~ if the logical object identifier is not zero.

A LONG bit of zero specifies shall perform the action specified by the SHORT ERASE MODE field in the Device Configuration Extension mode page (see 8.x) at the logical object identifier and partition specified in the command ~~a vendor-specific erase gap shall be written to the medium~~. The logical position following a ERASE(16) command with a LONG bit of zero shall be at the specified logical object identifier and partition. If the IMMED bit is one, the device server shall return status as soon as the command descriptor block has been validated. ~~Erase gaps may be used in application client controlled error recovery or update in place applications.~~

~~Editors Note 1 - DAP: Since the gap size field has been obsoleted the ability for application client controlled error recovery or update in place applications is suspect. Suggest removing the sentence.~~

3.4 Clause 6.2 ERASE (6) command

The ERASE(6) command (see table 20) causes part or all of the medium to be erased beginning at the current position. ~~Erased means the medium shall be erased or a pattern shall be written on the medium that appears as a gap~~. Prior to performing the erase operation, the device server shall perform a synchronize operation (see 4.2.8).

…

An immediate (IMMED) bit of zero specifies the device server shall not return status until the erase operation has completed. Interpretation of an IMMED bit of one depends on the value of the LONG bit, see below. However, for all values of the LONG bit, if CHECK CONDITION status is returned for an ERASE(6) command with an IMMED bit of one, the erase operation shall not be performed.

A LONG bit of one specifies all remaining medium in the current partition ~~shall be erased~~ beginning at the current logical position shall be erased or over-written with a format specific pattern that renders any over-written data on the medium unreadable. If the format on the medium specifies a recorded indication of EOD (see 3.1.6), the erase operation shall establish an EOD indication at the specified location as part of the erase operation. If the IMMED bit is one, the device server shall return status as soon as all buffered logical objects have been written to the medium and the command descriptor block of the ERASE(6) command has been validated. The logical position following an ERASE(6) command with a LONG bit of one is not specified by this standard.

NOTE 16 Some logical units may reject an ERASE(6) command ~~with the LONG bit set to one~~ if the logical unit is not at beginning-of-partition.

A LONG bit of zero specifies ~~a vendor-specific erase gap shall be written to the medium~~ shall perform the action specified by the SHORT ERASE MODE field in the Device Configuration Extension mode page (see 8.3.y). If the IMMED bit is one, the device server shall return status as soon as the command descriptor block has been validated. ~~Erase gaps may be used in application client controlled error recovery or update in place applications.~~

3.5 Clause 8.3.y Device Configuration Extension mode page

The Device Configuration Extension mode page (see table X), a subpage of the Device Configuration mode page (see 8.3.3), provides control over SCSI features specific to Sequential Access devices. If a device server supports the Device Configuration Extension mode page, then the device server shall provide access to the mode page using the shared mode page policy (see SPC-3).

Table X Device Configuration Extension mode page

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	SPF (1b)	PAGE CODE (10h)					
1	SUBPAGE CODE (01h)							
2	(MSB)	PAGE LENGTH (1Ch)						
3								(LSB)
4	Reserved							
5	Reserved				SHORT ERASE MODE			
6	Reserved							
31								

The SHORT ERASE MODE field controls the action to be taken by the device server when an ERASE (16) or ERASE (6) command with the LONG bit set to 0 is processed. The values for the SHORT ERASE MODE field are defined in table X+1. A device server shall support at least of the listed options, but support for more than one or any specific one is optional.

Table X+1 SHORT ERASE MODE field description

Value	Description
0h	The erase operation shall be performed as defined in SPC-2.
1h	The erase operation shall have no effect on the medium.
2h	The device server shall record an EOD indication at the specified location on the medium.
3h – Fh	Reserved