

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 12 October 2007
Subject: 07-099r0 SES-2 Clarify no null termination of strings

Revision history

Revision 0 (12 October 2007) First revision

Related documents

ses2r18 - SCSI Enclosure Services - 2 (SES-2) revision 18

Overview

SES-2 ASCII strings do not contain NUL (00h) termination characters, since they are always accompanied by precise length fields. SES-2 does not state this directly, though.

A new definition is proposed for "ASCII string" that only allows graphic characters (with an i.e., defining that as 20h through 7Eh) and thus prohibits inclusion of 00h. A similar definition is proposed for "text string" for use with strings that honor the Language element, with the same "containing only graphic characters" restriction (but no i.e., since the definition of graphic characters can vary based on the encoding).

Suggested changes to SES-2

2.2 Approved references

[ANSI INCITS 4-1986 \(R2002\), Information Systems - Coded Character Sets - 7-Bit American National Standard Code for Information Interchange \(7-Bit ASCII\)](#)

ISO 639-1:2002, Codes for the representation of names of languages - Part 1: Alpha-2 code

ISO/IEC 8859-1:1998, Information processing—8-bit single-byte coded graphic character sets - Part 1: Latin alphabet No. 1
ISO/IEC 10646-1:2000, Universal Multiple-Octet Coded Character Set (UCS) - Part 1: Architecture and Basic Multilingual Plane (BMP)

3.1 Definitions

3.1.xx ASCII string: [String of US-ASCII characters each encoded in 8 bits per ISO/IEC 8859-1 containing only graphic characters \(i.e., code values 20h through 7Eh\). ASCII strings do not contain the NUL character \(i.e., 00h\).](#)

3.1.xx text string: [String of characters using the character encoding and language specified by the Language element \(see 7.3.18\) containing only graphic characters. Text strings do not contain the NULL character \(i.e., 00h or 0000h\).](#)

6.1.2.4 Type descriptor text list

The type descriptor text is an optional text string from zero to 255 bytes for each type descriptor header. The text string, if it has a length greater than zero, may contain any descriptive information about the element type that may be useful to an application client that is displaying the configuration of the enclosure. The type descriptor texts shall be placed in the same order as the type descriptor headers, except that type descriptor texts of zero length shall be omitted.

Examples of information that may be included in the type descriptor text include the manufacturer's part number for a replacement element, a brief description of the element and its properties, or instructions about configuration limitations and redundancy requirements of the elements of that type.

The type descriptor text uses the character encoding and language specified by the Language element (see 7.3.18).

6.1.5 Help Text diagnostic page

The Help Text diagnostic page contains a ~~string of characters~~ [text string \(see 3.1.xx\)](#) from the primary subenclosure that describes the present state of the enclosure and provides text indicating what corrective actions, if any, are desirable to bring the enclosure to its fully operational state. The Help Text diagnostic page

is intended to allow the writing of enclosure independent application clients that return enclosure specific text describing the state of the enclosure and explaining enclosure dependent corrective actions that may be required. The page is optional. The language and character set of the help text are defined by the Language element (see 7.3.18). This page does not support subenclosures; the Subenclosure Help Text diagnostic page (see 6.1.14) does.

The Help Text diagnostic page is read by the RECEIVE DIAGNOSTIC RESULTS command with a PCV bit set to one and a PAGE CODE field set to 03h. If the parameter list for a SEND DIAGNOSTIC command contains a PAGE CODE field set to 03h, the command shall be treated as having an invalid field error (see 4.5).

Table 1 defines the Help Text diagnostic page.

Table 1 — Help Text diagnostic page

Byte\Bit	7	6	5	4	3	2	1	0	
0	PAGE CODE (03h)								
1	Obsolete								
2	(MSB)	PAGE LENGTH (n - 3)							
3								(LSB)	
4	PRIMARY SUBENCLOSURE HELP TEXT								
n									

The PAGE CODE field is set to 03h.

The PAGE LENGTH field indicates the length in bytes of the diagnostic parameters that follow.

The PRIMARY SUBENCLOSURE HELP TEXT field contains ~~text~~ [a text string \(see 3.1.xx\)](#) describing what corrective actions should be performed on the primary subenclosure to change it from its present state to a fully operational state. The text shall use the language and character set specified by the Language element (see 7.3.18).

6.1.10 Element Descriptor diagnostic page

The Element Descriptor diagnostic page returns a list of vendor-specific, variable-length ASCII strings, one for each element in the Enclosure Status diagnostic page (see 6.1.4).

...

Table 22 defines the overall descriptor and the element descriptor.

Table 22 — Overall descriptor format and element descriptor format

Byte\Bit	7	6	5	4	3	2	1	0	
0	Reserved								
1									
2	(MSB)	DESCRIPTOR LENGTH (m - 3)							
3								(LSB)	
4	DESCRIPTOR								
m									

The DESCRIPTOR LENGTH field indicates the length in bytes of the DESCRIPTOR field. A DESCRIPTOR LENGTH of zero indicates that no DESCRIPTOR field is contained in the overall descriptor or element descriptor.

The DESCRIPTOR field indicates [an ASCII string \(see 3.1.xx\) reporting](#) vendor-specific information about the element. The information ~~shall use ASCII characters and~~ shall not be modified by the Language element (see 7.3.18).

6.1.14 Subenclosure Help Text diagnostic page

The Subenclosure Help Text diagnostic page contains a ~~string of characters~~ [text string \(see 3.1.xx\)](#) from an enclosure that describes the present state of the enclosure and provides text indicating what corrective actions, if any, are desirable to bring the enclosure to its fully operational state. The Subenclosure Help Text diagnostic page is intended to allow the writing of enclosure independent application clients that return enclosure specific text describing the state of the enclosure and explaining enclosure dependent corrective actions that may be required. The page is optional. The language and character set of the help text are defined by the Language element (see 7.3.18).

The Subenclosure Help Text diagnostic page is read by the RECEIVE DIAGNOSTIC RESULTS command with a PCV bit set to one and a PAGE CODE field set to 0Bh. If the parameter list for a SEND DIAGNOSTIC command contains a PAGE CODE field set to 0Bh, the command shall be treated as having an invalid field error (see 4.5).

Table 1 defines the Subenclosure Help Text diagnostic page.

Table 23 — Subenclosure Help Text diagnostic page

Byte\Bit	7	6	5	4	3	2	1	0
0	PAGE CODE (0Bh)							
1	NUMBER OF SUBENCLOSURES							
2	(MSB)	PAGE LENGTH (n - 3)						(LSB)
3								
4	(MSB)	GENERATION CODE						(LSB)
7								
Subenclosure help text list								
8	Subenclosure help text (primary subenclosure)(see table 24)							
...								
n	Subenclosure help text (last subenclosure)(see table 24)							

The PAGE CODE field is set to 0Bh.

The NUMBER OF SUBENCLOSURES field specifies the number of separate subenclosure help texts that are included, not including the primary subenclosure. The NUMBER OF SUBENCLOSURES value shall be the same as the number of subenclosures value in the Configuration diagnostic page.

The PAGE LENGTH field indicates the length in bytes of the diagnostic parameters that follow.

The GENERATION CODE field contains the same value as the GENERATION CODE field in the Configuration diagnostic page (see 6.1.2).

Table 24 defines the format of each subenclosure help text. The first subenclosure help text shall be for the primary subenclosure; subenclosure help text for the remaining subenclosures may follow in any order.

Table 24 — Subenclosure help text format

Byte\Bit	7	6	5	4	3	2	1	0	
0	Reserved								
1	SUBENCLOSURE IDENTIFIER								
2	(MSB)	SUBENCLOSURE HELP TEXT LENGTH (m - 3)							
3								(LSB)	
4	SUBENCLOSURE HELP TEXT								
m									

The SUBENCLOSURE IDENTIFIER field indicates the subenclosure identifier to which the help text applies.

The SUBENCLOSURE HELP TEXT LENGTH field indicates the number of bytes in the SUBENCLOSURE HELP TEXT field. If a subenclosure has no help text, the SUBENCLOSURE HELP TEXT LENGTH field shall contain 0000h.

The SUBENCLOSURE HELP TEXT field contains ~~text~~ [a text string \(see 3.1.xx\)](#) describing what corrective actions should be performed on the subenclosure to change it from its present state to a fully operational state. The text shall use the language and character set specified by the Language element (see 7.3.18).

6.1.18 Language element

The Language element manages the language used for visual displays.

The format of the control field for a Language element is defined in table 25.

Table 25 — Language element for control-type diagnostic pages

Byte\Bit	7	6	5	4	3	2	1	0	
0	COMMON CONTROL								
1	RQST IDENT	Reserved							
2	(MSB)	LANGUAGE CODE							
3								(LSB)	

The COMMON CONTROL field is specified in 7.2.2.

The RQST IDENT bit is defined in the Power Supply element (see 7.3.4).

The LANGUAGE CODE field specifies the language and character encoding to be used in all fields that are defined as being modified by the Language element. The enclosure should provide external indications in the requested language.

If the LANGUAGE CODE field contains:

- a) 0000h: the enclosure services process shall use the default language of English with the US-ASCII character set encoding as defined by ISO/IEC 8859-1 (i.e., encoded as 8-bit characters with each MSB set to zero);
- b) two characters containing the ISO 639-1 two-letter code for a language that is supported by the enclosure services process (e.g., “en” for English, “fr” for French, “de” for German, or “jp” for Japanese) expressed as US-ASCII characters as defined by ISO/IEC 8859-1 (i.e., encoded as 8-bit characters each with its MSB set to zero): the enclosure services process shall use UCS-2 as defined by ISO 10646-1 (i.e., encode using 16-bit characters); or

- c) a value other than 0000h or the two-letter code of a language that is supported by the enclosure services process: the enclosure services process shall use the default language of English with the US-ASCII character set encoded as defined by ISO/IEC 8859-1 (i.e., encoded as 8-bit characters with each MSB set to zero) and shall report an invalid field error (see 4.5).

The format of the status field for a Language element is defined in table 26.

Table 26 — Language element for status-type diagnostic pages

Byte\Bit	7	6	5	4	3	2	1	0
0	COMMON STATUS							
1	IDENT	Reserved						
2	(MSB)	LANGUAGE CODE						
3								(LSB)

The COMMON STATUS field is specified in 7.2.3.

The IDENT bit is defined in the Power Supply element (see 7.3.4). Since the Language element may not represent a physical element, the visual indication may be an indication of the language being used.

The LANGUAGE CODE field indicates the language and character encoding that the enclosure services process uses for those fields that have the capability of being modified by the Language element. A LANGUAGE CODE field set to 0000h indicates the enclosure services process is using the default language of English and the US-ASCII character set encoded as defined by ISO/IEC 8859-1 (i.e., encoded as 8-bit characters with each MSB set to zero). A LANGUAGE CODE field set to an ISO 639-1 two-letter code (e.g., “en” for English, “fr” for French, “de” for German, or “jp” for Japanese) expressed as US-ASCII characters as defined by ISO/IEC 8859-1 (i.e., encoded as 8-bit characters each with its MSB set to zero) indicates the enclosure services process is using the indicated language and is using UCS-2 as defined by ISO/IEC 10646-1 (i.e., encoding using 16-bit characters).

6.1.20 Subenclosure Nickname Control diagnostic page

The Subenclosure Nickname Control diagnostic page transmits ~~an ASCII~~ a text string (see 3.1.xx) to the enclosure services process to serve as the nickname for the specified subenclosure. The nickname is saved to non-volatile storage (e.g., a flash ROM) so it may be retrieved after future hard resets.

Table 27 describes the Subenclosure Nickname Control diagnostic page.

Table 27 — Subenclosure Nickname Control diagnostic page

Byte\Bit	7	6	5	4	3	2	1	0
0	PAGE CODE (0Fh)							
1	SUBENCLOSURE IDENTIFIER							
2	(MSB)	PAGE LENGTH (24h)						
3								(LSB)
4	(MSB)	GENERATION CODE						
7								(LSB)
8	SUBENCLOSURE NICKNAME							
39								

The PAGE CODE field is set to 0Fh.

The SUBENCLOSURE IDENTIFIER field specifies the vendor-specific identifier for the subenclosure to which the application client is sending the subenclosure nickname. If the SUBENCLOSURE IDENTIFIER value does not

match a SUBENCLOSURE IDENTIFIER value found in the Configuration diagnostic page (see 6.1.2), the enclosure services process shall set the SUBENCLOSURE NICKNAME STATUS field to 80h in the Subenclosure Nickname Status diagnostic page.

The PAGE LENGTH field specifies the length in bytes of the remainder of the diagnostic page. If the PAGE LENGTH field value does not match the length of the page, the enclosure services process shall not change the subenclosure nickname and set the SUBENCLOSURE NICKNAME MICROCODE STATUS field to 80h in the Subenclosure Nickname Status diagnostic page.

The GENERATION CODE field specifies the value expected to be found in the GENERATION CODE field of the Configuration diagnostic page (see 6.1.2). To prevent the misinterpretation of the microcode data, the enclosure services process shall verify that the value of the GENERATION CODE field matches the generation code value known by the enclosure services process. If there is a mismatch, the enclosure services process shall not change the subenclosure nickname and set the SUBENCLOSURE NICKNAME MICROCODE STATUS field to 80h in the Subenclosure Nickname Status diagnostic page.

The SUBENCLOSURE NICKNAME field specifies the subenclosure nickname. If a Language element (see 7.3.18) is present, the SUBENCLOSURE NICKNAME field shall contain [a text string \(see 3.1.xx\) containing](#) characters using the language and character set indicated by the Language element and the enclosure services process shall store the language code value indicated by the Language element along with the subenclosure nickname. If a Language element is not available, the SUBENCLOSURE NICKNAME field shall contain ~~US-ASCII characters~~[an ASCII string](#) and the enclosure services process shall store the language code value of 0000h along with the subenclosure nickname.