

SCSI Accessed Fault-Tolerant Enclosures

Interface Specification

Copyright © nStor Corporation and Intel Corporation

Revision Intermediate Review R041497
April 14, 1997

**Intermediate
SAFTE Committee
Review Revision**

THIS DOCUMENT IS PROVIDED BY NSTOR CORPORATION AND INTEL CORPORATION “AS IS” AND NO MAINTENANCE OR SUPPORT SERVICES ARE PROVIDED BY NSTOR OR INTEL. NSTOR AND INTEL EXPRESSLY DISCLAIM ALL WARRANTIES, EITHER EXPRESSED OR IMPLIED, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR AGAINST INFRINGEMENT OF ANY COPYRIGHT, PATENT, TRADE SECRET OR OTHER INTELLECTUAL PROPERTY RIGHT OF ANY THIRD PARTY. INTEL AND NSTOR DISCLAIM ANY LIABILITY FOR ANY DAMAGES WHATSOEVER, WHETHER DIRECT, INDIRECT, INCIDENTAL, SPECIAL, OR CONSEQUENTIAL (INCLUDING BUT NOT LIMITED TO LOSS OF PROFITS, REVENUE, DATA OR USE) INCURRED BY YOU OR ANY THIRD PARTY IN CONNECTION WITH OR ARISING FROM THE USE OF THIS DOCUMENT, EVEN IF NSTOR AND INTEL HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

“SCSI ACCESSED FAULT-TOLERANT ENCLOSURES” AND “SAF-TE” ARE NOT TRADEMARKS NOR ARE THEY INTENDED TO BE TRADEMARKS, BUT RATHER ARE INDUSTRY TERMS OF ART THAT DESCRIBE TECHNOLOGY TO MONITOR FAULT-TOLERANT ENCLOSURES.

<p>Intermediate SAFTE Committee Review Revision</p>
--

Table Of Contents

1.0 INTRODUCTION	24
2.0 SCSI SPECIFICATION	2
2.1 SCSI Messages	3
2.2 SCSI Commands	3
INQUIRY	4
READ BUFFER	7
REQUEST SENSE	8
SEND DIAGNOSTIC	10
TEST UNIT READY	11
WRITE BUFFER - Write SEP Device Command	12
3.0 SAF-TE INTERFACE	13
3.1 READ BUFFER Commands	16
Read Enclosure Configuration (00h)	16
Read Enclosure Status (01h)	18
Read Usage Statistics (02h)	21
Read Device Insertions (03h)	22
Read Device Slot Status (04h)	23
Read Global Flags (05h)	25
3.2 WRITE BUFFER Commands	26
Write Device Slot Status (10h)	26
Set SCSI ID (11h)	29
Perform Slot Operation (12h)	30
Set Fan Speed (13h)	32
Activate Power Supply (14h)	33
Send Global Flags (15h)	34
4.0 POWER ON DIAGNOSTICS	36
5.0 THE SAF-TE ORGANIZATION	36

1.0 Introduction

The objective of the SAF-TE Interface is to provide a standard, non-proprietary way for third party disk and RAID controllers to be automatically integrated with peripheral packaging that supports status signals (LEDs, audible alarm, LCD, etc.), hot swapping of hard drives, and monitoring of enclosure components. From the system vendor's point of view, this allows quick integration of the "best" third party controllers, knowing that they will fully integrate with disk and peripheral packaging. It also allows a selected controller to work with a variety of expansion packaging, being fully able to sense status and drive enclosure indicators. The cost of a separate cable and interface for enclosure services is also eliminated.

SCSI is the underlying transport mechanism chosen for communicating enclosure information. This means that all standard SCSI host adapters will work. No special considerations, such as reserved signals on the SCSI bus, or additional cables are required. The SAF-TE interface can be implemented as an inexpensive SCSI target, using a simple SCSI part and an 8 bit microcontroller. In this document, the target devices that implement the SAF-TE interface are collectively referred to as the SAF-TE Processor (SEP) device.

All communication is initiated by the host. The SAF-TE Processor device acts only in a target role. Asynchronous Event Notification is not used. The SAF-TE Processor device should periodically be polled by the host to detect changes in status.

Drive failure indications are controlled by the host's driver software through this command set, because it is the host that knows if a drive has failed. Status indicators for other components, such as fans and power supplies, may be controlled automatically by the SEP device.

2.0 SCSI Specification

The SAF-TE protocol is implemented using the SCSI microprocessor device type. The SAF-TE Processor (SEP) device conforms to the ANSI SCSI-2 specification for processor devices. It must support the following six SCSI commands -WRITE BUFFER, READ BUFFER, INQUIRY, TEST UNIT READY, SEND DIAGNOSTIC, and REQUEST SENSE. Any unsupported CDB opcode will result in a check condition, with a sense key of 05h, Illegal Request. (ASC 20h, ASCQ 00h - Invalid CDB Operation Code).

The maximum recovery time (time until 00h status is returned from TEST UNIT READY) of the SEP device from a SCSI bus reset is 15 seconds. While the SEP device must be polled, this is not expected to be a performance impact due to the short duration and low frequency of polling required. If the host polls for enclosure status every ten seconds or so, then the overall system performance will be minimally impacted. This specification does not place any restrictions on polling frequency. It is expected that most implementations will poll the SEP once every 2 to 10 seconds.

The SEP device will always support LUN 0, which is responsible for managing device slots on the SCSI channel on which LUN 0 accessed. If the SEP is responsible for managing the device slots on only one channel, then only LUN 0 will be supported, and commands (other than INQUIRY or REQUEST SENSE) to LUNs other than 0 will result in a CHECK CONDITION, with a sense key of 05h, Illegal Request (ASC 20h, ASCQ 00h - Logical Unit Not Supported).

It is possible, in an enclosure that supports more than one SCSI channel, to use a single SCSI ID on one channel to implement multiple SEP devices. In this case, additional LUNs will be supported, one LUN per channel, to implement a logically separate SEP for each channel.

Any field marked *Reserved* should not be used and should not be interpreted by the target.

This specification is mainly concerned only with those elements of the SCSI protocol required to define the SAF-TE Interface. It is not the intent of this specification to be a complete description of the SCSI specification for processor devices.

For more information concerning the SCSI commands, messages, and phases mentioned in this specification please refer to the appropriate ANSI SCSI specification.

2.1 SCSI Messages

An SEP device will support all MESSAGES which are mandatory for a SCSI-2 processor device type. Additional MESSAGES may be supported as required by a target SEP implementation. Any unsupported message sent to the SEP device will cause it to send a MESSAGE REJECT message back to the initiator.

BUS DEVICE RESET (0Ch) - Out to Target - Receiving this message should not cause the target SEP to modify the current status of any external enclosure signals, such as LEDs, speaker, etc. Otherwise the SEP should respond normally as a SCSI-2 target by entering a UNIT ATTENTION condition and preparing SENSE information indicating that the device has been reset.

2.2 SCSI Commands

Any command other than INQUIRY or REQUEST SENSE, sent to an unsupported LUN on the SEP device, will result in CHECK CONDITION status with sense information indicating ILLEGAL REQUEST, Sense Key 05h (ASC 25h, ASCQ 00h - LOGICAL UNIT NOT SUPPORTED)

The INQUIRY and REQUEST SENSE commands will return valid data if sent to an unsupported LUN. This data will indicate that there is no device connected to the particular LUN, as described below in the command descriptions.

INQUIRY

Bit	7	6	5	4	3	2	1	0
0	Operation Code (12h)							
1	Logical Unit Number				Reserved			0
2	00h							
3	Reserved							
4	Allocation Length							
5	00h							

The INQUIRY command requests that the SEP device return parameter information, the format of which is illustrated below.

Logical Unit Number - Specifies which SCSI logical unit (0 through 7) of the selected SCSI ID is the target of this command. If the LUN specified in this field is not supported, then the response data will indicate such, as described below in the description of the *Peripheral Qualifier* and *Peripheral Device Type* in the response data format.

Allocation Length - Specifies the number of bytes that the initiator has allocated for the response data. The SEP device transfers the number of bytes specified up to the maximum available. An allocation length of 00h indicates that no data will be transferred. This condition is not considered an error, GOOD status should be returned without entering a DATA IN phase.

The response data format for the INQUIRY command is defined as follows.

Bit	7	6	5	4	3	2	1	0
0	Peripheral Qualifier				Peripheral Device Type			
1	00h							
2	0	0	0	0	0	ANSI -Approved Version (02h)		
3	0	0	0	0	Response Data Format (02h)			
4	Additional Length							
5	Reserved							
6	Reserved							
7	00h							

8			
-----		Vendor Identification	
15			
16			
-----		Product Identification	
31			
32			
-----		Firmware Revision Level	
35			
36			
-----		Enclosure Unique Identifier	
42			
43			
-----		Channel Identifier	
44			
-----		SAF-TE Interface Identification String	
49			
50			
-----		SAF-TE Specification Revision Level	
53			
54			
-----		Reserved By SAF-TE	
55			
-----		Reserved By SAF-TE	
56			
-----		Reserved By SCSI	
95			
96			
-----		Vendor Unique	
n			

Peripheral Qualifier - Indicates whether the selected LUN is a valid SCSI device. This field will be 000b if a supported LUN is selected, and 011b if an unsupported LUN is selected.

000b indicates that the device specified in the Peripheral Device Type is currently connected to this LUN. It does not imply that the specified logical device is ready for access by the initiator (as in TEST UNIT READY).

011b indicates that the target is not capable of supporting a device on this logical unit.

Peripheral Device Type - Indicates the type of SCSI device at the selected LUN. If a valid LUN is selected, this field will be 03h (SCSI Processor Device), if an invalid LUN is selected this field will be 1Fh (no device type).

ANSI-Approved Version - This field is 02h to indicate compliance with the ANSI SCSI-2 specification.

Response Data Format - This field is 02h to indicate that the format of the INQUIRY response data is as defined in the ANSI SCSI-2 specification.

Additional Length - This field indicates the number of bytes of additional INQUIRY command parameters available for transfer, beginning with byte 05h. This value is not adjusted if the *Allocation Length* in the CDB is too large or too small to accommodate the entire response.

Vendor Identification - This 8 byte ASCII string identifies the product vendor.

Product Identification - This 16 byte ASCII string specifies the product ID.

Firmware Revision Level - This 4 byte ASCII string identifies the current firmware revision of the SEP device.

Enclosure Identifier - This field holds a 58 bit identification number. This field should be unique for each individual unit produced by a particular manufacturer. This makes it possible to positively distinguish between two or more RAID devices connected to the same host system.

The combination of this field, along with the *Vendor Identification* and *Product Identification* fields, will uniquely identify any peripheral unit from any manufacturer.

Channel Identifier - This field is used to distinguish between separate channels supported by a single enclosure. The value in this field will be unique for each channel. It does not necessarily bear any relation to the channel number of the host adapter used to access the enclosure channel.

SAF-TE Interface Identification String - This 6 byte field holds the constant ASCII string "SAF-TE". This serves to identify that the enclosure is compliant with the SAF-TE Interface Specification.

SAF-TE Specification Revision Level - This 4 byte field holds an ASCII string of the format x.xx which identifies the revision of the SAF-TE Interface Specification to which this device claims compliance. ASCII string data is stored with the most significant (leftmost) character stored at the lowest byte offset of the field.

READ BUFFER

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (3Ch)							
1	Logical Unit Number			Reserved		Mode (01h)		
2	Buffer ID							
3	00h							
6								
7	Transfer Length (MSB)							
8	Transfer Length (LSB)							
9	00h							

The READ BUFFER command is used to receive data from the SEP device in a DATA IN phase. The format of the returned data is discussed in Section 3 of this document.

Mode - This 3 bit field should hold (001b) to indicate that the data buffer is in the SAF-TE command format.

Buffer ID - This byte determines the content and format of the data buffer to be transferred to the initiator during the data phase. (See Section 3.2)

Transfer Length - The transfer length, in bytes, of the data to be transferred in the data phase of this command. A transfer length of 0 is not considered an error; the SEP device will return GOOD status without entering the DATA IN PHASE. If an transfer length is sent which is greater than the available data length for a particular *Buffer ID* on the SEP device, then only the available number of bytes will be transferred and no error condition will be indicated.

REQUEST SENSE

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (03h)							
1	Logical Unit Number				Reserved			
2	00h							
3	00h							
4	Allocation Length							
5	00h							

The REQUEST SENSE command requests that the SEP device transfer sense data to the initiator. Note that the sense data pertains to the transport of data over the SCSI bus. Sense data is not used to report enclosure specific failures and status, such as a cooling failure.

Allocation Length - This field specifies the number of bytes that the initiator has allocated for the response data. The SEP device transfers the number of bytes specified, up to the maximum available for transfer by the target. An allocation length of zero indicates that no data will be transferred. This is not considered an error. If more bytes are requested than are available, then only the available bytes will be transferred and no error condition will be indicated.

The sense data format is as follows.

Bit	7	6	5	4	3	2	1	0
Byte								
0	F0h							
1	00h							
2	Reserved	Reserved	0	Reserved	Sense Key			
3	Reserved							
6	Reserved							
7	Additional Sense Length							
8	00h							
11	Reserved							
12	Additional Sense Code (ASC)							
13	Additional Sense Code Qualifier (ASCQ)							

Additional Sense Length - Indicates the number of additional bytes that are available after Byte 7. A minimum of 13 bytes must be available for transfer.

SCSI-2 targets must support the Additional Sense Code field (ASC). If no additional sense information is available, the ASC (and ASCQ if returned) must be set to 00h, indicating No Additional Sense Information.

Some possible sense information fields returnable by an SEP are defined as follows. Other SCSI-2 standard Sense Keys and Additional Sense codes may be supported as required by an SEP implementation. All values listed are in hexadecimal.

<i>Sense Key</i>	<i>ASC</i>	<i>ASCO</i>	<i>Error Condition</i>
00			No Sense, No error condition
02			Not Ready
	04	00	Cause Not Reportable
	04	01	In The Process Of Becoming Ready
04			Hardware Error
	40	81	Failed ROM Checksum Test
	40	82	Failed RAM Checksum Test
	40	83	Failed RAM Read/Write Test
	40	84	Failed Processor Register Test
	44	00	Internal Target Failure (Undetermined Fatal Error)
05			Illegal Request
	20	00	Invalid Command Operation Code
	24	00	Invalid Field In CDB
	25	00	Logical Unit Not Supported
	26	02	Invalid SEP Command In Write Buffer Data
06			Unit Attention
	29	00	Power-On, Reset, or Bus Device Reset Occurred
	3F	01	Microcode Has Been Changed
09	80	xx	Reserved for SAF-TE specific error conditions
	81	xx	Reserved for SAF-TE specific error conditions
0B			Aborted Command
	48	00	Initiator Detected Error Message Received
	4A	00	Command Phase Error
	4B	00	Data Phase Error

SEND DIAGNOSTIC

Bit	7	6	5	4	3	2	1	0
0	Operation Code (1Dh)							
1	Logical Unit Number			Reserved				
2	00h							
3	00h							
4	00h							
5	00h							

The SEND DIAGNOSTIC command requests that the SEP device perform diagnostic self-tests on itself and its memory resources (ROM and RAM). These are the same tests that are conducted upon power up and when a BUS RESET DEVICE message is received. GOOD status is returned if no errors are detected. CHECK CONDITION status is returned if an error is detected. If an error condition is indicated then sending REQUEST SENSE immediately following the SEND DIAGNOSTIC will return sense information that provides more details about the error.

If there are no enclosure components (such as non-volatile RAM) that require diagnostic tests, then this CDB may be treated as a “no operation”, and must simply return GOOD status.

Note that this command does not return any information about the enclosure or any of its components (fans, power supplies, etc.) other than the SEP device. The status of other enclosure components are returned through the use of the READ BUFFER command, described elsewhere in this section.

TEST UNIT READY

Bit	7	6	5	4	3	2	1	0
0	Operation Code (00h)							
1	Logical Unit Number				Reserved			
2	00h							
3	00h							
4	00h							
5	00h							

The TEST UNIT READY command provides the means to check if the SEP device is ready to accept other commands. If it is ready then GOOD status will be returned. If the SEP device is not ready to accept new commands, CHECK CONDITION status will be returned with a sense key of NOT READY.

WRITE BUFFER - Write SEP Device Command

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (3Bh)							
1	Logical Unit Number			Reserved		Mode (01h)		
2	00h							
3	00h							
4	00h							
5	00h							
6	00h							
7	Transfer Length (MSB)							
8	Transfer Length (LSB)							
9	00h							

The WRITE BUFFER command, with a mode value of 01h, is used to send a WRITE type command data buffer to the SEP device. These data buffers define operations to be performed by the SEP device. The functionality and format of each WRITE BUFFER commands is described in Section 3.2 of this document.

Mode - This 3 bit field should hold (001b) to indicate that the data buffer is in the SAF-TE command format.

Transfer Length - The transfer length, in bytes, of the data buffer to be sent in the data phase of this command. An allocation length of 0 is not considered an error; the SEP device will return GOOD status without entering a DATA OUT PHASE.

3.0 SAF-TE Interface

The SAF-TE Interface is a set of processor device commands (rather than SCSI commands) that the host system may use to request specific actions of the processor.

There are two types of SEP commands. Those commands that request some action to be performed in the enclosure, such as indicating a drive failure, are sent to the SEP device with a WRITE BUFFER operation. Those commands that request information from the SEP device, such as the status of the cooling fans, are sent to the SEP device using a READ BUFFER command.

SAF-TE commands use a system of one byte opcodes similar in concept to the opcodes used in SCSI CDBs. Opcodes in the range of 00h to 7Fh are reserved for standardized SAF-TE commands common to all vendor's implementations and should not be used for any vendor unique commands. Opcodes in the range of 80h to FFh are open for vendor specific use.

Each of the SAF-TE commands outlined in this section is described as being either mandatory or optional for a target implementation.

Command Support for an Initiator - This specification will not explicitly list mandatory commands for an initiator application. The commands required to perform various functions will vary between applications. A hardware RAID host adapter, for example, would probably not need to use the Send Global Command (15h) to explicitly turn on or off an audible enclosure alarm, whereas a software utility for enclosure maintenance would probably want to use the command to allow a user to silence an alarm from the host console.

A few essential commands will likely be used by any host application to provide the basics of communication to the enclosure.

Read Enclosure Configuration (00h) returns a list of the enclosure hardware resources.

Read Enclosure Status (01h) returns the operating status of hardware resources and provides a map between SCSI IDs and device slot numbers.

Read Device Slot Status (04h) to read logical and physical status of device slots, and determine if a drive is physically inserted in a slot.

The Send Global command is intended for use only by high level array management software such as those components of a RAID system that interact with the user. It is assumed that low level system components, such as RAID adapter firmware, automatically drive any appropriate enclosure signals by sending array information in the Write Device Slot Status command. Likewise, the Set SCSI ID command will probably be useful only at the higher levels of array management.

Mandatory Target SEP Commands - A mandatory command is one that must be implemented by a target SEP device to adhere to a minimal implementation of this specification. The mandatory commands provide the basics of communication between a host and an enclosure, and standardize the inventory, status, and state indication of standard enclosure elements such as fans, hard drives, power supplies, etc.

Optional Target SEP Commands - An optional command need not necessarily be implemented on an enclosure adhering to this specification. These are commands that may not be needed or desired on a particular vendor's implementation. However, even though not all vendor's will support these commands, they are not vendor unique opcodes. This is because it is still desirable to have them standardized with a defined opcode and parameter format among those implementations that do support them. If the enclosure supports such functionality as settable device slot SCSI IDs, then the appropriate optional SEP command(s) should be supported. A vendor may elect to augment their implementation with additional vendor-specific commands.

Unsupported Commands - If an SEP device implementation receives an unsupported opcode in byte 2 of a READ BUFFER CDB, then a CHECK CONDITION occurs with sense data indicating Illegal Request, Sense Key 05h. (ASC 24h, ASCQ 00h - Invalid Field in CDB).

If an invalid opcode is received as the first byte of a WRITE BUFFER data buffer, then a CHECK CONDITION occurs with sense data indicating Illegal Request, Sense Key 05h. (ASC 26h, ASCQ 02h - Invalid SEP Command).

If a target SEP returns CHECK CONDITION status for any WRITE BUFFER command, it should not perform any further action for that command, even if some or all of the parameter fields contain valid values.

Default Status Values - For each device slot in a RAID enclosure, the SAF-TE protocol uses four bytes to represent the current status. These flags contained in these bytes are fully described in Section 3.0 in the description of the two commands that use them - Read Device Slot Status (04h) and Write Device Slot Status (10h). Bytes 0 and 1 contain flags describing the logical status of a device as an element of a RAID array, Byte 2 is Reserved, and Byte 3 contains flags describing the physical status of the device and slot.

In general, the contents of the slot status bytes 0 and 1 are set by the initiator through the Write Device Slot Status command. This allows the initiator to inform the SEP of the logical status of each of the device slots. The contents of the first two bytes should not normally be modified by the SEP. There are, however, two events which are exceptions to this rule.

1) Upon being powered on, the SEP must set the device slot status bytes for all slots in the enclosure to the following initial state.

Device Slot Status Byte 0 - 00h if no device is inserted
80h if a device is inserted

Device Slot Status Byte 1 - 00h

Device Slot Status Byte 2 - 00h

Device Slot Status Byte 3 - If a device is present in this slot, then the device should be prepared for operation by default, and the Prepare For Operation flag set.

- 2) When a drive is physically inserted into a slot while power is supplied to the enclosure, the SEP should set the Unconfigured flag in status byte 0 for that slot. Removing a drive should not cause the SEP to modify device slot status bytes 0, 1 or 2, but flags in byte 3 will be changed by the SEP when a drive is removed.

Slot SCSI IDs - An enclosure may allow the SCSI IDs of device slots to be set through the optional Set SCSI IDs command, or by other means such as an integrated enclosure control panel. In any case, the enclosure should not assert the new IDs to the slots until the next time it is turned off and then back on again.

A SAF-TE initiator can assume that the slot SCSI IDs will not change without a power cycle occurring.

3.1 READ BUFFER Commands

This section defines the format of the data returned by READ BUFFER commands, and defines the Buffer ID field that is Byte 2 of the READ BUFFER CDB. For each READ BUFFER command detailed below, the hex value is given that should be placed in the Buffer ID field of the CDB.

For all READ BUFFER commands, it should not be considered an error condition for the initiator to request fewer than the maximum number of bytes available for transfer. An allocation length of zero bytes should result in GOOD status being returned without entering the DATA PHASE.

Read Enclosure Configuration (00h)

Mandatory

The Read Enclosure Configuration command is used by the host to inquire about the system components of the RAID enclosure. If the unit supports multiple SCSI channels, only those devices connected to the current channel are included in the returned data.

To send this command, set the Buffer ID field of the READ BUFFER CDB to 00h. The following table illustrates the data returned. The data length is 64 bytes if no vendor specific fields are used.

The values returned by this command are used to determine the lengths of the data buffers returned by other READ BUFFER commands.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Number of Fans (<i>f</i>)							
1	Number of Power Supplies (<i>p</i>)							
2	Number of Device Slots (<i>d</i>)							
3	Door Lock Installed							
4	Number of Temperature Sensors (<i>t</i>)							
5	Audible Alarm Installed							
6	Celsius	Reserved			Number of Thermostats (<i>s</i>)			
7	Reserved							
62	Reserved							
63	Number of Vendor Specific Bytes (<i>v</i>)							
64	Vendor Specific							
xx	Vendor Specific							

Number of Fans - The maximum number of SEP manageable cooling fans. This value is referred to as *f* in other command descriptions in this document.

Number of Power Supplies - The maximum number of SEP manageable power supplies. This value is referred to as *p* in other command descriptions in this document.

Number of Device Slots - The total number of device slots available on this channel, whether or not they currently have a device inserted. This value is referred to as *d* in other command descriptions in this document.

Door Lock Installed - This field will be either 0 to indicate that no host controllable enclosure door lock is installed, or 1 to indicate that one is installed.

Number of Temperature Sensors - This field indicates the number of temperature sensors installed in the enclosure, and is limited to a maximum value of 15. This value is referred to as *t* in other command descriptions in this specification, and is used to determine the number of integer *Temperature* fields returned by the Read Enclosure Status command.

Audible Alarm Installed - This field will be either 0 to indicate that the enclosure has no audible alarm or 1 to indicate that an alarm is available.

Celsius - This flag is set to indicate that the temperature fields returned by the Read Enclosure Status command are in Celsius. If this flag is cleared then the temperatures returned will be in Fahrenheit.

Number of Thermostats - This field indicates the number of temperature sensors which returns a binary signal, under/over preset threshold. These sensors are not included in the Number of Temperature Sensors field. This value is referred to as *s* in other command descriptions in this specification.

Bytes 7 through 62 are reserved and should be zero.

Number of Vendor Unique Bytes - This field indicates the number of Vendor Unique bytes available for transfer. If no vendor unique fields are implemented then this field will be zero.

Read Enclosure Status (01h)

Mandatory

The Read Enclosure Status command is used by the host to find the operational status of the components of the RAID enclosure.

To send this command, set the Buffer ID field of the READ BUFFER CDB to 01h. The following table illustrates the data returned during the data phase.

The number of status bytes available for transfer to the initiator will depend on the number of installed components. Use the *Read Enclosure Configuration* command to find the values of the variables that determine the length of the data returned. Vendor unique fields may extend the length of the data.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Fan 0 Status							
$f-1$	Fan $f-1$ Status							
	Power Supply 0 Status							
$f+p-1$	Power Supply $p-1$ Status							
	Device Slot 0 SCSI ID							
$f+p+d-1$	Device Slot $d-1$ SCSI ID							
$f+p+d$	Door Lock Status							
$f+p+d+1$	Speaker Status							
$f+p+d+2$	Temperature 0							
$f+p+d+t+1$	Temperature $t-1$							
$f+p+d+t+2$	Temperature Out Of Range Flags 1							
$f+p+d+t+3$	Temperature Out Of Range Flags 2							
$f+p+d+t+4$	Number of Vendor Specific Bytes (v)							
$f+p+d+t+5$	Vendor Specific							

Fan Status fields - These fields will contain one of the following values -

- 00h - Fan is operational.
- 01h - Fan is malfunctioning.
- 02h - Fan is not installed.
- 80h - Unknown Status, or Status Not Reportable

Power Supply Status fields - These fields will contain one of the following values -

- 00h - Power supply is operational and on
- 01h - Power supply is operational and off
- 10h - Power supply is malfunctioning and commanded on
- 11h - Power supply is malfunctioning and commanded off
- 20h - Power supply is not present
- 21h - Power supply is present (This status may be used instead of codes 00h - 11h if failure detection and On/Off status are not detectable)
- 80h - Unknown Status, or Status Not Reportable

Device Slot SCSI ID fields - These fields provide an associative map between physical slots and device SCSI IDs. The SCSI IDs are represented using their integer ID value, not the bit field representation of the ID used on the SCSI bus. SCSI ID 6, for example, is represented in this field as 00000110b - not 01000000b. In order to maintain compatibility with older target implementations, any SCSI ID field that is associated with an empty or otherwise unused slot may optionally contain FFh instead of the SCSI ID asserted at that slot.

Door Lock Status - This field will contain one of the following values -

- 00h - Door is currently locked
- 01h - Door is currently unlocked, or no SEP controllable door lock installed
- 80h - Unknown Status, or Status Not Reportable

Speaker Status - This field will contain one of the following values -

- 00h - Speaker is currently off, or no speaker installed
- 01h - Speaker is currently on

Temperature - Each of these fields contains an unsigned integer indicating internal enclosure temperature at this sensor in degrees Fahrenheit. This value, from 0 to 255, is based on a scale starting at -10 degrees. This gives an effective reportable temperature range from -10 degrees to 245 degrees.

There will be *t* *Temperature* fields, with *t* being determined with the Read Enclosure Configuration command. Note that if no temperature sensors are installed, no bytes will be dedicated to these fields, so the byte following speaker status will be the *Temperature Out Of Range Flags* field.

Temperature Out Of Range Flags - These fields contain 16 flags which indicate whether or not the SEP device has detected a temperature that is out of normal operating range. Fifteen flags are provided for individual sensors, and one flag is provided for an overall temperature alert. The heat thresholds at which the SEP device will set these flags may either be preset in firmware or user configurable.

Unlike the *Temperature* fields, these two fields will always be present and valid even if the number of temperature sensors *t*, returned by the Read Enclosure Configuration command is zero. One flag is used for each temperature sensor and each thermostat. The first *t* flags, starting from *t0*, are used for the temperature sensors, and the next *s* flags are used for the thermostats.

Temperature Out Of Range Flags 1

Bit	7	6	5	4	3	2	1	0
	ETA	t14	t13	t12	t11	t10	t9	t8

Temperature Out Of Range Flags 2

Bit	7	6	5	4	3	2	1	0
	t7	t6	t5	t4	t3	t2	t1	t0

The *t0* through *t14* flags indicate whether an abnormal temperature has been detected on sensors 0 through 14. A value of 0 indicates normal temperature, or temperature sensor not supported. A value of 1 indicates that an abnormal temperature has been detected on the sensor. These 15 flags are optional, and should be 0 if not supported.

The *ETA* flag indicates whether there is an enclosure temperature alert on any sensor within the enclosure. This flag should be set if any of the *t0* through *t14* flags are set.

Number of Vendor Unique Bytes - This field indicates the number of Vendor Unique bytes available for transfer. If no vendor unique fields are implemented then this field will be zero.

Read Usage Statistics (02h)

Optional

The Read Usage Statistics command is used to receive information on total usage time and number of power-on cycles of the RAID device.

To send this command, set the Buffer ID field of the READ BUFFER CDB to 02h. The following table illustrates the data returned during the data phase of the Read command. The length of the data will be 16 bytes if no vendor unique fields are implemented.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Total Number Of Power On Minutes MSB							
1	Total Number Of Power On Minutes							
2	Total Number Of Power On Minutes							
3	Total Number Of Power On Minutes LSB							
4	Total Number of Power On Cycles MSB							
5	Total Number of Power On Cycles							
6	Total Number of Power On Cycles							
7	Total Number of Power On Cycles LSB							
8	Reserved							
14	Reserved							
15	Number of Vendor Specific Bytes (<i>v</i>)							
	Vendor Specific							

Total Number Of Power On Minutes - This 4 byte field holds the total number of minutes that the RAID device has been powered on. This count is cumulative over the life of the device.

Total Number of Power On Cycles - This 4 byte field holds a count of the number of times that the RAID device has been powered on. This count is cumulative over the life of the device.

Bytes 5 through 14 are reserved.

Number of Vendor Unique Bytes - This field indicates the number of Vendor Unique bytes available for transfer. If no vendor unique fields are implemented then this field will be zero.

Read Device Insertions (03h)

Optional

The Read Device Insertions command returns information indicating how many times a device has been inserted into each slot in the RAID system while the unit was powered on.

To send this command, set the *Control* field of the READ BUFFER CDB to 03h. The following table illustrates the data returned during the data phase.

The total length of the data buffer returned is determined by the *Number Of Device Slots* value returned by the Read Enclosure Configuration command. The data buffer will be $2d$ bytes long.

Bit	7	6	5	4	3	2	1	0
0	Number Of Insertions In Slot 0 MSB							
1	Number Of Insertions In Slot 0 LSB							
$2d-2$	Number Of Insertions In Slot $d-1$ LSB							
$2d-1$	Number Of Insertions In Slot $d-1$ MSB							

Number of Insertions fields - These 16 bit fields hold the number of times a drive has been inserted into a particular slot while the RAID unit was powered on. These values are cumulative over the life of the device.

Read Device Slot Status (04h)

Mandatory

The Read Device Slot Status command returns information on the current state of each drive/slot. There are 4 bytes returned for each device slot contained in the enclosure.

To send this command, set the *Control* field of the READ BUFFER CDB to 04h. The following table illustrates the data returned during the data phase.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Slot 0 Byte 0							
1	Slot 0 Byte 1							
2	Slot 0 Byte 2							
3	Slot 0 Byte 3							
	Slot d-1 Byte 0							
	Slot d-1 Byte 1							
	Slot d-1 Byte 2							
d*4-1	Slot d-1 Byte 3							
d*4	Number of Vendor Specific Bytes (v)							
d*4+1	vendor specific							

Each group of four bytes is a set of bit flags. The first three bytes returned are a copy of the three bytes of flags most recently sent by the initiator for the respective drive/slot using the Write Device Slot Status command (See Section 3.2). The bit definitions are identical to those used with that command. Refer to the description of Write Device Slot Status for a definition of these three bytes.

The fourth byte contains flags that indicate other conditions of interest on the drive slot. Flag bytes are defined with bit 0 being the least significant bit (00000001) and bit 7 the most significant bit (10000000). The bit fields are defined as follows (Bit 0 is the least significant bit).

Byte 3

Bit 0 Device Inserted Flag - Indicates whether or not there is a physical device inserted in this slot. This does not imply that the device is ready for access on the SCSI bus. 0 - no device inserted, 1 - device inserted.

Bit 1 Ready for Insertion/Removal Flag - This flag indicates whether or not the slot is ready for the physical insertion or removal of a device. 0 - not ready, 1 - ready.

Bit 2 Prepared For Operation Flag - This flag indicates whether or not the slot is activated so that the inserted drive may be accessed on the SCSI bus. 0 - not activated, 1 - activated.

Bits 3-7 Reserved

Read Global Flags (05h)

Optional

The Read Global Flags command is used to read from the SEP the most recent state of the global flags received by the SEP in the Send Global Flags command. Sending this command will not modify the state of any global flag.

The following table illustrates the format of the data returned from the SEP device during the data phase.

Bit	7	6	5	4	3	2	1	0
0	Global Flags 1							
1	Global Flags 2							
2	Global Flags 3							
3	Reserved							
15								

The format of each global flags field is described fully in the description of the Send Global Flags command.

3.2 WRITE BUFFER Commands

This section defines the format of the data buffers used in WRITE BUFFER commands. For each command, the first byte of the data will be the Operation Code.

If more bytes are transferred during the DATA OUT phase than actually required or implemented by the SEP device, then the extra bytes should simply be ignored and this should not be considered an error condition.

Write Device Slot Status (10h)

Mandatory

The Write Device Slot Status command is used to inform the SEP device of the state of each of its associated slots and the devices potentially inserted. This information is used to drive the enclosure status signals (LEDs, LCD, audible alarm, etc.) to some meaningful state, depending on the vendor's implementation.

The length of this data buffer depends on the number of device slots (d) on this channel. The number of device slots on a particular channel may be determined with the READ ENCLOSURE CONFIGURATION command (see Section 3.1). There are three bytes of data for each drive slot on the channel.

The following table illustrates the format of the data sent from the host to the SEP device during the data phase of the SCSI WRITE BUFFER operation.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (10h)							
1	Slot 0 Byte 0							
2	Slot 0 Byte 1							
3	Slot 0 Byte 2							
$3d$	Slot $d-1$ Byte 0							
$3d+1$	Slot $d-1$ Byte 1							
$3d+2$	Slot $d-1$ Byte 2							
$3d+3$	Vendor Specific							

Operation Code - is 10h for this command.

Flag bytes are defined with bit 0 being the least significant bit (00000001) and bit 7 the most significant bit (10000000). The 3 bytes for each drive/slot are a set of bit flags defined as follows.

Byte 0

Bit 0 No Error Flag. Set if there are no outstanding error conditions on this device/slot.

Bit 1 Device Faulty Flag. Set if this device has exhibited some hardware or data fault.

Bit 2 Device Rebuilding Flag. Set if the device is being rebuilt. Other drives in this array remain In Critical Array until the rebuild operation is complete. Non-redundant configurations such as RAID 0 are inherently non-fault-tolerant, and should not make use of this flag.

Bit 3 In Failed Array Flag. Set if this device is in an array with an error on any member device.

Bit 4 In Critical Array Flag. Set if this device is in an array that was previously fault-tolerant and has become non-fault-tolerant. Non-redundant configurations such as RAID 0 are inherently non-fault-tolerant, and should not make use of this flag.

Bit 5 Parity Check Flag. Set if the device is in an array which is undergoing a parity check operation.

Bit 6 Predicted Fault Flag. Set if a device has been tagged by a fault prediction algorithm as being likely to fail in the near future.

Bit 7 Unconfigured Flag. Set if a device is currently not configured as an element of an array or as a hot spare.

Byte 1

Bit 0 Hot Spare Flag. Set if a device in this slot is configured as a hot spare. This bit will be set in conjunction with the No Error flag.

Bit 1 Rebuild Stopped Flag. Set for a slot containing a device which was rebuilding, but the rebuild terminated abnormally or unsuccessfully. This flag and the Device Rebuilding flag cannot both be set for the same slot.

Bits 2 through 7 Reserved.

Byte 2

Bits 0 through 7 Reserved.

If no flags are set in any byte for a device slot this is a NO CHANGE FROM CURRENT STATE indication. This allows an initiator to change the state of one particular device slot without having to be aware of the current state of all device slots. Setting one or more flags requires that all flags be written to the correct binary value. Thus, changes should be preceded by a read of the corresponding device slot status and the Write Device Slot Status implemented as a “read-modify-write” operation.

Enclosures are free to alias these states as appropriate. For instance, one enclosure may have an amber LED driven constantly on to indicate that a drive is faulty, while the same LED might be blinking to indicate a different drive state, or off if there was no error. Another enclosure may use a green LED to indicate no error, and amber to indicate other drive states, with an LCD panel to print a textual description of drive states. Yet another vendor may elect to have a separate LED for each possible drive state. This command provides a common method of providing enclosures with the needed drive state information, regardless of vendor specific enclosure implementations.

Vendor specific fields may be implemented, and should be ignored by a target SEP device that is not aware of them.

Set SCSI ID (11h)

Optional

The Set SCSI ID command is used to set the SCSI bus ID of any target (hard drive, tape drive, or the SEP device) on the current channel. This command is optional as not all implementations will allow the host to change device SCSI IDs. The new IDs should not be asserted on the slots until the next power cycle on the enclosure.

An attempt to set the SCSI ID of the SEP device to an ID that is already used by another target on the bus, or an attempt to set another device's ID to the SCSI ID already used by the SEP device, will result in check condition status with sense information indicating INVALID SEP COMMAND IN WRITE BUFFER DATA. This will also be returned if a SCSI ID which is invalid for the device slot is selected.

The following table illustrates the format of the data sent from the host to the processor during the data phase of the Set SCSI ID operation. The data is 64 bytes long if no vendor specific fields are used.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (11h)							
1	Device Slot Number							
2	SCSI ID							
3	Reserved							
63	Reserved							
64	Vendor Specific							

Operation Code - 11h for the Set SCSI ID command.

Device Slot Number - The slot number of the device on which to set the SCSI ID. This field should be set to FFh to set the ID of the SEP device receiving this command.

SCSI ID - This field holds the SCSI ID that should be set.

Bytes 3 through 63 are reserved and should be zero.

Vendor specific fields may be implemented, but will be ignored by a target SEP device that is not aware of them.

Perform Slot Operation (12h)

Mandatory

The Perform Slot Operation command is used to perform various operations on device slots, such as to prepare a device slot for the insertion or removal of a hard drive. Although this advanced notice may not be required by all hardware implementations, this is a mandatory command to avoid having a CHECK CONDITION status returned. In those implementations that do not have to electrically or mechanically prepare the drive slot for drive insertion or removal, this command should be treated as a “no operation” with GOOD status returned.

The following table illustrates the format of the data sent from the host to the SEP device during the data phase of the Prepare Slot operation. The data is 64 bytes long if no vendor specific fields are used.

Bit	7	6	5	4	3	2	1	0
0	Operation Code (12h)							
1	Slot Number							
2	Operation Flags							
3	Reserved							
63								
64	Vendor Specific							

Operation Code - 12h for the Perform Slot Operation command.

Slot Number - This field indicates which drive slot to prepare to perform the given operation on.

Operation Flags - This field is a set of 8 flags that indicate what operation is to be performed on this slot. There may be only one flag set at any time. If more than one flag is set then the SEP device will return CHECK CONDITION with sense data indicating INVALID SEP COMMAND IN WRITE BUFFER DATA.

The flags are defined as follows -

Bit 0 Prepare For Operation Flag. - This flag indicates that the SEP device should take any action required to make the device in this slot available on the SCSI bus.

Bit 1 Prepare For Insertion Or Removal Flag. - This bit indicates that the SEP device should take any action required to allow the safe physical insertion or removal of a device from this slot.

Bit 2 Identify Flag. - This flag indicates that the SEP device should indicate the physical location of this slot in some manner, to indicate to the operator which slot is being addressed. This is usually done by flashing an adjacent LED in a unique pattern.

Bits 3-7 Reserved

Bytes 3 through 63 are reserved and should be zero.

Vendor specific fields may be implemented, but will be ignored by a target SEP device that is not aware of them.

Set Fan Speed (13h)

Optional

The Set Fan Speed command is used to tell the SEP device at what speed each fan should be running. The unit of speed used is quarter parts of the fan's full speed, so that a fan may be off, quarter speed, half speed, three quarters speed, and full speed.

The following table illustrates the format of the data sent from the host to the SEP device during the data phase.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (13h)							
1	Fan Number							
2	Fan Speed							
3	Reserved							
15								

Operation Code - 13h for the Set Fan Speed command.

Fan Number - Indicates which fan on which to set the speed. This can range from 0 to $f-1$.

Speed for Fan - This field holds the value for the desired speed setting. It should hold one of the following values -

- 0 - turn fan off
- 1 - run fan at quarter capacity
- 2 - run fan at half capacity
- 3 - run fan at three quarters capacity
- 4 - run fan at full capacity

Activate Power Supply (14h)

Optional

The Activate Power Supply command is used to turn a power supply on or off. Note that most enclosures will have all power supplies activated by default after power on.

The following table illustrates the format of the data sent from the host to the SEP device during the data phase.

Bit	7	6	5	4	3	2	1	0
0	Operation Code (14h)							
1	Power Supply Number							
2	Power State							
3	Reserved							
15								

Operation Code - 14h for the Activate Power Supply command.

Power Supply Number - Indicates which power supply to turn on or off. This can range from 0 to $p-1$.

Power State - This field should hold one of the following values -

00h - turn power supply off

01h - turn power supply on

Send Global Flags (15h)

Mandatory

The Send Global Flags is used to send commands that apply to the enclosure globally rather than to a specific channel, or device slot. The actions caused by these commands flags may be implemented differently on different vendor's enclosures. Conditions such as Cooling Failure or Global Failure may optionally be used to drive alarms, LEDs, LCD messages, etc.

Support for this command is mandatory for a target SEP device, but not all flags may be applicable. For example, an enclosure that has no speaker will ignore the Audible Alarm Control flag. An enclosure that does not implement a method of remotely turning enclosure power on and off, will ignore the Enclosure Power flag.

This command is optional for an initiator because the responsibility of driving the enclosure signals belongs to the SEP device using information received from the Write Device Slot Status command. This command is provided for an Initiator that desires to override the default enclosure signals.

The following table illustrates the format of the data sent from the host to the SEP device during the data phase.

Bit	7	6	5	4	3	2	1	0
0	Operation Code (15h)							
1	Global Flags 1							
2	Global Flags 2							
3	Global Flags 3							
4	Reserved							
15								

Operation Code - 15h for Send Global Flags command.

Flag bytes are defined with bit 0 being the least significant bit (00000001) and bit 7 the most significant bit (10000000).

Global Flags 1

Bit 0 Audible Alarm Control. Set to sound an audible alarm within the enclosure. Reset to turn off the alarm.

Bit 1 Global Failure Indication. Set to request that the enclosure perform some actions (alarm, LEDs, LCD message) that indicate a global error condition.

- Bit 2 Global Warning Indication. Set to request that the enclosure perform some actions (alarm, LEDs, LCD message) that indicate a global warning condition.
- Bit 3 Enclosure Power. Set to turn off power to the enclosure. Clear to turn on power to the enclosure.
- Bit 4 Cooling Failure. Set to indicate that a fan/cooling failure has occurred.
- Bit 5 Power Failure. Set to indicate that system power has been lost. This allows the enclosure to perform any emergency shutdown procedure required before battery power is lost.
- Bit 6 Drive Failure. Set to drive any global enclosure indicator that a drive has failed.
- Bit 7 Drive Warning. Set to drive any global enclosure indicator that a potential error has or may have occurred on a drive.

Global Flags 2

- Bit 0 Array failure. Set to drive any global enclosure indicator that an array has failed.
- Bit 1 Array Warning. Set to drive any global enclosure indicator that an array is non fault-tolerant.
- Bit 2 Enclosure Lock. Set to lock the enclosure drive bay access door. Cleared to unlock the door.
- Bit 3 Identify Enclosure. Set to drive any global enclosure identify signal, such as a unique LED pattern. Clear to turn off the enclosure identify signal.
- Bits 4 through 7 Reserved.

Global Flags 3

- Bits 0 through 7 Reserved.

4.0 Power On Diagnostics

The SEP microprocessor device will perform diagnostic self-tests immediately upon being powered on. Such power on self tests (POSTs) may include memory checksums, NVRAM read/write tests, and SCSI bus access initialization. If no errors are detected, all enclosure status signals should be driven to some default state until changed by a command received from a host initiator. If a POST error is detected, then this condition is indicated by some signal (perhaps an audible alarm and amber LED state) for a predetermined amount of time before the default power-on signal state is assumed.

The existence of POST errors should cause CHECK CONDITION status to be returned from a SEND DIAGNOSTIC command.

5.0 The SAF-TE Organization

This document, and additional information about SAF-TE is available on the internet at <http://www.safte.org>. Email can be sent to safte@safte.org.