

Table TBD - Reference: relationship between SCC and SAF-TE states

SAF-TE State	SCC State
No Error Flag	Volume Set: Available, Verify in Progress Redundancy Group: Available, Verify in Progress Peripheral Device: Available Spare State: Available Component Device: Available
Device Faulty	Volume Set: Broken Redundancy Group: Exposed, Invalidated Protected Space, Partially Exposed Peripheral Device: Broken Spare State: Broken Component Device: Broken
Device Rebuilding	Volume Set: Rebuild, Recalculate, Protected Rebuild Redundancy Group: Protected Rebuild, Rebuild, Recalculate Peripheral Device: Rebuild
Failed Array	Volume Set: Broken, Data Lost Redundancy Group: Invalidated Protected Space
Critical Array	Volume Set: Exposed, Fractionally exposed, Protection disabled, rebuild, recalculate Redundancy Group: Protection Disabled, Rebuild, Recalculate
Parity Check	Volume Set: Verify in Progress Redundancy Group: Verify in progress
Predicted Fault	Reported as status
No Drive	Peripheral Device: Not available
Hot Spare	Peripheral Device: Available Spare State: Available, Present, Spare in use
Parity of array failed	Volume Set: Data lost Redundancy Group: Invalidated protected space (?)
Unconfigured Drive	Peripheral Device: Present Spare: Present Component Device: Present
Rebuild Abort	Reported as status
Self Identify	?



3) Proposed change in management of diagnostic code page lengths

At present, the SPC standard specifies that a page length that truncates a parameter will be posted as a CHECK CONDITION. The changes offered in this document modify that to allow truncation by an allocation length shorter than the required page length, as long as all the parameters provided are valid and fully defined and no required parameters are truncated.

4) Device type codes

The device type codes for SCC are a bit strange for enclosures. I have elected to include all the specified device types, but to give them codes in the order of most probable implementation in the enclosure services pages. I have then chosen to enforce that order of presentation. I have added the SCC component types not already covered using an address range of 08h to 15h but using the same low order bits

5) SCSI Slot parameters

I personally think that the collection of flags is overkill. I would like to see these reduced to three or four indicator flags at the most. Perhaps the following might be helpful:

- a) Device is okay
- b) Set device fault indicator
- c) Set device identification indicator
- d) Set device membership indicator (indicates device is member of an array)
- e) Set membership in recovering or failed array (indicates device is member of an array that

6) Temperature

The present SAF-TE specification indicates the temperature in degrees F + 10 degrees. The zero value is -10 degrees F and the maximum is 245 degrees F. I propose that the standard use degrees C + 20. The range is -20 degrees C to 235 degrees C, although most sensors will indicate overtemperature somewhere around 50 degrees C.

7) Power supply indicators

It may be useful to extend the power supply definitions so that overvoltage conditions, undervoltage conditions, and overcurrent conditions can be optionally separated. While two of those functions probably could be treated as power supply failures, the overcurrent condition may be caused by a failure of the enclosure, a failure of a device within the enclosure, or an invalid configuration of the enclosure.

8) Mapping of device indicators to SCC

Rod Dekonin of Symbios has provided a preliminary analysis of how the SCC indicators might map to this enclosure proposal. This will be related to item 5 and to the alternate slot/bay proposals.

Rod notes that there is some SCC information that is not mappable to an enclosure activity. As an example, the readying of a drive is indicated through the normal status process.



- The drive shall respond to any command that accesses the drive media with CHECK CONDITION status and with sense information indicating an ASC/ASCQ of Enclosure Failure.
- Outstanding tasks shall remain suspended until after the EFW indication is ended.
- Normal drive operation shall resume after the EFW signal is ended. Suspended operations are resumed.

If the post error according to informational exceptions page mode is supported by the SCSI device and is requested by the MODE SENSE/SELECT procedure, the SCSI device shall execute the following actions upon detecting an EFW indication.

- The SCSI device server shall not abnormally terminate any activity. All commands shall continue normal operation.
- The EFW check condition shall be posted under the rules of the notification procedure defined for background errors in the informational exceptions mode page. These rules typically require the posting of a CHECK CONDITION for a command the first time the EFW is detected, then periodically reposting the CHECK CONDITION as required by the setting of the MODE SELECT page as long as the EFW condition is present.
- Normal device operation continues after the EFW condition is terminated and no additional posting activity takes place.

2.6 Bibliography and references

A bibliography with an entry for the Small Form Factor Committee specification SFF-8045 (SCA-2 disk drive for Fibre Channel Loop Attachment) shall be created. The references need to include: Fibre Channel Physical and Signaling Interface (FC-PH), ANS X3.230-1995.

3 Issues and questions

- 1) Should INQUIRY indicate support of ESI

While it is possible to put an indicator bit to show that ESI is supported by a particular device in byte 5 or 6 of the INQUIRY command, it is not clear that this is necessary. If an application client actually cares, it can easily find out by performing a RECEIVE DIAGNOSTIC RESULTS command requesting the supported diagnostics page. ESI devices of course identify themselves as implementing the ESI commands.

- 2) Should variable length entries for elements be allowed.

At present, all element status and control descriptions are 4 bytes long. A pair of bits could be included to select 4-byte 8-byte, or 16-byte descriptions, but this seems like overkill, since all standard device types are already defined. A length descriptor could be included in the configuration table, but this would allow the creation of nonstandard element descriptions. I believe that 4 bytes is a good number and should be a fixed number.



The EFW (Enclosure Failure Warning) Mode describes the operations that the SCSI device will perform if the enclosure presents an enclosure failure warning to it. The enclosure may transmit such a warning when a new ESI status is available, when a partial or incipient failure is detected, when a complete failure is detected, or any combination of these. The application client uses the EFW Mode bits to adjust the behavior of the SCSI device to properly accommodate whichever type of warnings the enclosure is known to provide. The EFW indication from the enclosure to a device within the enclosure is optional.

Table qp - EFW Mode values

EFW value			Description
bit 2	bit 1	bit 0	
0	0	0	Default mode, SCSI device vendor unique.
0	0	1	Terminate at block boundary and reset
0	1	0	Terminate at command boundary and temporarily suspend.
0	1	1	Post error according to informational exceptions page
All others			Vendor-specific pages

The default EFW Mode allows a vendor unique recovery or error presentation action to be taken by the SCSI device.

If the terminate at block boundary and reset mode is supported by the SCSI device and is requested by the MODE SENSE/SELECT procedure, the SCSI device shall execute the following actions upon detecting an EFW indication.

- The SCSI device server shall terminate any active media modification (write) tasks at an appropriate block boundary ending point.
- All other tasks on the SCSI device shall end and the device shall return all parameters to their power on initial reset state.
- The device server shall respond to any new commands with CHECK CONDITION status and indicate the presence of an EFW indication as long as the EFW condition is still present.
- The SCSI transport mechanism remains active and initialized.
- After the EFW state is ended, the first command received shall indicate CHECK CONDITION status and UNIT ATTENTION sense information. Subsequent tasks will be executed normally.

If the terminate at command boundary and temporarily suspend mode is supported by the SCSI device and is requested by the MODE SENSE/SELECT procedure, the SCSI device shall execute the following actions upon detecting an EFW indication.

- The SCSI device server shall terminate any active media modification (write) tasks at an appropriate block boundary ending point and provide a CHECK CONDITION indication with sense information indicating an ASC/ASCQ of Enclosure Failure.
- All further tasks on the drive shall be temporarily suspended.
- The device server shall respond to any command that does not access the drive media without regard to the EFW state.



Enclosure Transfer Refused: This ASC/ASCQ indicates that the device or the enclosure services function indicated either an error or an invalid format in their communication.

Table TBD - Summary of new ASC/ASCQs required

ASC/ASCQ	SEND DIAGNOSTIC
TBD1/0	Enclosure failure
TBD2/0	Enclosure degraded warning
TBD3/0	Unspecified Enclosure Services Failure
TBD3/1	Unsupported Enclosure Function
TBD3/2	Enclosure Services Unavailable
TBD3/3	Enclosure Transfer Failure
TBD3/4	Enclosure Transfer Refused

2.5 Additional MODE SENSE/MODE SELECT pages for enclosure services.

The mode select pages defined for all devices should include the necessary definitions for support of devices that support enclosure services. This will be added as a new section to section 8.3, probably placed alphabetically between sections 8.3.2 and 8.3.3. In addition, Table 83 will need to have an additional entry to include the new optional page. The entry is obvious, and so will not be expanded here. The text of the new section will be as follows:

8.3.n Enclosure services management page.

The enclosure services management page (see table qq) provides controls over those SCSI features involving communication with the enclosure for a SCSI device.

Table qq” Enclosure services management page

Bits Bytes	7	6	5	4	3	2	1	0
0	PS	Rsrvd	Page Code (TBDh)					
1	Page length							
2	Reserved							
2						EFW Mode		
2 - 7	Reserved							



12 Parameters for enclosure services devices

12.1 Diagnostic Parameters

This subclause defines the descriptors and pages for diagnostic parameters used with enclosure services devices.

The diagnostic page codes for enclosure services devices are defined in table mm.

Table mm - Diagnostic page codes

Page Code	SEND DIAGNOSTIC	Mandatory/ Optional	RECEIVE DIAGNOSTIC RESULTS	Mandatory/ Optional
00h	N/A		Supported diagnostics	Mandatory
01h	N/A		Enclosure Services Config	Mandatory
02h	Enclosure Services Control	Mandatory	Enclosure Services Status	Mandatory
03h	N/A		Enclosure Services Help Text	Optional
04h	Enclosure Services String Out	Optional	Enclosure Services String In	Optional

2.4 Additional ASC/ASCQ for enclosure services function integrated into devices.

Seven new ASC/ASCQ values should be defined for devices that provide the enclosure services pages, whether they are enclosure service devices or other device models that provide an interface to the enclosure.

Enclosure Failure: This ASC/ASCQ is provided to indicate when a critical or an unrecoverable enclosure failure has been detected by the enclosure. This is provided using the Sense Key of **HARDWARE ERROR**. Further information may be available using the **REQUEST DIAGNOSTIC RESULTS** command and requesting the enclosure services in page.

Enclosure Degraded Warning: This ASC/ASCQ is provided to indicate that an informational condition or a non-critical failure has been detected by the enclosure. This is provided using the Sense Key of **RECOVERED ERROR** and may be managed by the Informational Exceptions Control mode page. Further information may be available using the **REQUEST DIAGNOSTIC RESULTS** command and requesting the enclosure services in page.

Unspecified Enclosure Services Failure: This ASC/ASCQ is provided to indicate that the enclosure services device has failed in an unknown manner.

Unsupported Enclosure Function: This ASC/ASCQ indicates that the device has been asked to invoke an enclosure services function that does not exist.

Enclosure Services Unavailable: This ASC/ASCQ indicates that the device has been asked to invoke an enclosure services function that is temporarily busy or unavailable.

Enclosure Transfer Failure: This ASC/ASCQ indicates that the device communication with the enclosure services function has failed.



and the enclosure services inbound pages to examine various status and warning information available from the enclosure. The pages and page formats are defined in section 8.1.

Some enclosures do not have a node that can be addressed as a logical unit. However, some SCSI devices have a special interface to the enclosure that can be addressed through their normal SCSI connection. Those devices shall use the same SEND DIAGNOSTIC and RECEIVE DIAGNOSTIC RESULTS commands that would be used by an enclosure services device. Support for the enclosure services pages is determined by using the RECEIVE DIAGNOSTIC RESULTS supported diagnostics page.

11.2 Commands for enclosure services devices

The commands for enclosure services devices shall be as shown in table nn. All remaining operation codes are reserved for future standardization.

Table nn - Commands for enclosure services devices

Command name	Operation Code	Type	Subclause
INQUIRY	12h	M	7.5
MODE SELECT (6)	15h	O	7.8
MODE SELECT (10)	55h	O	7.9
MODE SENSE (6)	1Ah	O	7.10
MODE SENSE (10)	5Ah	O	7.11
PERSISTENT RESERVE IN	5Eh	O	TBD
PERSISTENT RESERVE OUT	5Fh	O	TBD
RECEIVE DIAGNOSTIC RESULTS	1Ch	M	7.16
RELEASE(6)	17h	O	7.17
RELEASE(10)	57h	O	7.18
REQUEST SENSE	03h	M	7.20
RESERVE(6)	16h	O	7.21
RESERVE(10)	56h	O	7.22
SEND DIAGNOSTIC	1Dh	M	7.23
TEST UNIT READY	00h	M	7.24
WRITE BUFFER	3Bh	0	7.25
Key: M = Command implementation is mandatory O = Command implementation is optional.			



Code 0Fh, SCSI port/transceiver element typ

Table bag - Element status bytes for SCSI port/transceiver

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									Disabled

Table bah - Element control bytes for SCSI port/transceiver

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									Disable

Byte 3, Bit 0

Disabled (Status)

1 = Indicates transceiver is disabled

Disable (Control)

1 = Instructs transceiver to disable or remain disabled.

0 = Instructs transceiver to enable or remain enabled.

2.3 Definition of enclosure services device model

A new section is added to SPC or to the appropriate document to describe the enclosure services device model. The model uses a new device type code defined in the INQUIRY command.. The model is added as section 11 of SPC revision 8. The description of parameters for devices obeying the model is added as section 12 of SPC revision 8.

11 Model and commands for enclosure services devices

11.1 Model for enclosure services devices

Enclosure services devices define a mechanism to communicate across the SCSI bus with an intelligent node in an enclosure that contains other SCSI devices attached to the same bus. An application client may address the enclosure services logical unit and use the SEND DIAGNOSTIC command and the enclosure services outbound pages to set various indicators and states in the enclosure, allowing the enclosure to provide the most appropriate environment for the SCSI devices contained within it. Similarly, the application client may request information from the enclosure services logical unit using the RECEIVE DIAGNOSTIC RESULTS command



Code 0Dh, Key pad entry device element type

Table bac - Element status bytes for key pad entry

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table bad - Element control bytes for key pad entry

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined

Code 0Eh, UPS battery element type

Table bae - Element status bytes for UPS battery

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table baf - Element control bytes for UPS battery

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined



Code 0Bh, Uninterruptible power supply element type

Table by - Element status bytes for uninterruptible power supply

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table bz - Element control bytes for uninterruptible power supply

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined

Code 0Ch, Display element type

Table baa - Element status bytes for display element

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table bab - Element control bytes for display element

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined



Code 08h, SCC controller electronics element type

Table bu - Element status bytes for SCC controller electronics

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table bv - Element control bytes for SCC controller electronics

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined

Code 09h, Nonvolatile cache element type

Table bw - Element status bytes for nonvolatile cache

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table bx - Element control bytes for nonvolatile cache

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined



Table br - Element control bytes for speaker

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									Turn on

Byte 3, Bit 0

On (Status)

1 = Indicates speaker is on

Unlock (Control)

1 = Instructs speaker to turn on or remain on.

0 = Instructs speaker to turn off or remain off.

Code 07h, ESI controller electronics element type

Table bs - Element status bytes for ESI controller electronics

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Table bt - Element control bytes for ESI controller electronics

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									

Bits to be defined



Code 05h, Door lock element type

Table bo - Element status bytes for door lock

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									Unlocked

Table bp - Element control bytes for door lock

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									Unlock

Byte 3, Bit 0

Unlocked (Status)

1 = Indicates door lock is unlocked

Unlock (Control)

1 = Instructs door latch to unlock or remain unlocked.

0 = Instructs door latch to lock or remain locked.

Code 06h, Speaker element type

Table bq - Element status bytes for speaker

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2									
3									On



Byte 3, Bit 6

Fail (Status)

1 = Indicates that fan failure indicator is on or has been set on.

Set Fail Indicator (Control)

1 = Instructs the enclosure to turn the fan failure indicator on.

0 = Instructs the enclosure to turn the fan failure indicator off unless turned on by internal circuitry.

Code 04h, Temperature sensor element type

Table bm - Element status bytes for temperature sensors

Bytes	Bits	7	6	5	4	3	2	1	0
1									
2		Temperature (degrees Celsius + 20)							
3						OT, Failure	OT, warning		

Table bn - Element control bytes for temperature sensors

Bytes	Bits	7	6	5	4	3	2	1	0
1-3		Reserved							

Byte 2, bits 0-7

Temperature (Status)

value = temperature at sensor in degrees Celsius + 20. Range is from -20 to +245 degrees Celsius

Byte 3, bit 2

Over Temperature (OT) Warning (Status)

1 = temperature at sensor is higher than normal operating temperature but will not damage the enclosure.

Byte 3, bit 3

Over Temperature (OT) Failure (Status)

1 = temperature at sensor is higher than a safe operating temperature and enclosure damage may occur.



Code 03h, Fan element type

Table bk - Element status bytes for fan

Bytes	7	6	5	4	3	2	1	0
1								
2								
3		Fail	Rqsted On				Actual Speed Code	

Table bl - Element control bytes for fan

Bytes	7	6	5	4	3	2	1	0
1								
2								
3		Set Fail Indicator	Rqst On				Requested Speed Code	

Byte 3, Bit 1 and 0

Actual Speed Code (Status)

Code indicates actual speed of fan

Requested Speed Code (Control)

Code indicates requested speed of fan

Table bla - Speed Code

Speed Code	Description
00b	Fan stopped (Actual Speed Code only)
01b	Fan at lowest speed
10b	Fan at intermediate speed
11b	Fan at highest speed

Byte 3, Bit 5

Requested on (Status)

1 = Indicates that the fan has been requested to be in the operating state.

Request on (Control)

1 = Instructs the fan to turn on or remain on.

0 = Instructs the fan to turn off or remain off.



Table bj - Element control bytes for power supply

Bytes	7	6	5	4	3	2	1	0
1								
2								
3		Set Fail Indicator	Rqst On					

Byte 3, Bit 0

DC Fail (Status)

1 = Indicates power supply cannot provide specified DC power.

Byte 3, Bit 1

AC Fail (Status)

1 = Indicates power supply is not receiving specified AC power.

Byte 3, Bit 2

Temp warning (Status)

1 = Indicates temperature is higher than normal operating temperature, but has not yet reached a temperature that requires the power supply to shut down.

Byte 3, Bit 3

Overtemp fail (Status)

1 = Indicates temperature is higher than a safe operating temperature and the power supply is or soon will be shutdown.

Byte 3, Bit 5

Requested on (Status)

1 = Indicates that the power supply has been requested to be in the operating state.

Request on (Control)

1 = Instructs the power supply to turn on or remain on.

0 = Instructs the power supply to turn off or remain off.

Byte 3, Bit 6

Fail (Status)

1 = Indicates that failure indicator is on or has been set on.

Set Fail Indicator (Control)

1 = Instructs the enclosure to turn the power supply failure indicator on.

0 = Instructs the enclosure to turn the power supply failure indicator off unless turned on by internal circuitry.



Table bh - Element control bytes for SCSI device bay/slot

Bytes	7	6	5	4	3	2	1	0
2	Array	Do not remove			Ready to Insert	Rdy to remove	Identify Fault	Sense Fault
3				Drive Off	Enable Port A	Enable Port B	Port A Enabled	Port B Enabled

[Editor's comment:

A mapping of the SCC states to the states defined for SAF-TE has been made and is shown in section 3 of this proposal, entitled issues and questions.]

Code 02h, Power Supply element type

Table bi - Element status bytes for power supply

Bytes	7	6	5	4	3	2	1	0
1								
2								
3		Fail	Rqsted On		Overtmp Fail	Temp warning	AC Fail	DC Fail



Table bh - Element control bytes for SCSI device bay/slot

Bytes	7	6	5	4	3	2	1	0
1					Prepare to Insert	Prep. to remove	Prep. to operate	
2	Uncon-fig Drive	Pred Fault	Parity Check	In Crit. Array	In Failed Array	Rebuild	Identify Fault	Sense Fault
3	OK	Hot Spare	Rebuild Abort	Drive Off	Enable Port A	Enable Port B	Port A Enabled	Port B Enabled

[Editor's comment:

A simplified alternative proposal is shown below. It focuses on revealing to the observer in front of an array only that information which will actually be useful to him in deciding whether or not a drive can be removed. The final bit descriptions will be provided in the same manner as for the other element types when they are resolved.]

Alternate Implementation: Code 01h, SCSI device bay/slot element type

Table bg - Element status bytes for SCSI device bay/slot

Bytes	7	6	5	4	3	2	1	0
1								
2	Array	Do not remove			Ready to Insert	Rdy to remove	Identify Fault	Sense Fault
3				Drive Off	Enable Port A	Enable Port B	Port A Enabled	Port B Enabled



The select bit is set to one if the control function defined by this element control entry is to be performed. The select bit is set to zero if this element control entry is to be ignored by the enclosure. The select bit allows individual element control entries from the entire list of control entries to be selected for execution.

Unless otherwise specified, all status and control bits are optional. The enclosure is not required to present any optional status bit. The enclosure is not required to act on any optional control bit and may ignore any control bit if required to maintain a proper operating environment in the enclosure.

Code 00h, Unspecified element type

The element status bits are reserved for the status entry for the unspecified element type. The element control bits are reserved for the control entry for the unspecified element type.

Code 01h, SCSI device bay/slot element type

Table bg - Element status bytes for SCSI device bay/slot

Bits Bytes	7	6	5	4	3	2	1	0
1				Inserted	Ready to Insert	Rdy to remove	Rdy to Operate	
2	Uncon-fig Drive	Pred Fault	Parity Check	In Crit. Array	In Failed Array	Rebuild	Identify Fault	Sense Fault
3	OK	Hot Spare	Rebuild Abort	Drive Off	Enable Port A	Enable Port B	Port A Enabled	Port B Enabled



8.1.8 Global and element status and control formats

The element status format for all elements is provided in table bd.

Table bd - General format of global and element status entries

Bytes	Bits	7	6	5	4	3	2	1	0
0	Resrvd								Status Code
1	Element status 1								
2	Element status 2								
3	Element status 3								

The status codes and the definition of the status codes are shown in table be.

Table be - Element status codes

Type Code	Name	Condition	Mandatory Optional
00h	Unsupported	element status detection is not implemented for this element	Optional
01h	OK	element is installed and operating normally	Mandatory
02h	Critical	critical element failure is detected	Optional
03h	Noncritical	noncritical element failure is detected	Optional
04h	Unrecoverable	unrecoverable element failure is detected	Optional
05h	Not installed	element is detected to be not installed in enclosure	Mandatory
06h	Unknown	element sensor has failed or element status is not available	Optional

The element control format for all elements is provided in table bf.

Table bf - General format of global and element control entries

Bytes	Bits	7	6	5	4	3	2	1	0
0	Select	Reserved							
1	Element control 1								
2	Element control 2								
3	Element control 3								



Table ee - Enclosure services string in page

Bytes	Bits	7	6	5	4	3	2	1	0
0		Page Code (04h)							
1		Reserved							
2		(MSB) Page Length (= n-3) (LSB)							
3									
4-n		Binary string provided by enclosure to application client							

8.1.7 Enclosure services string out page

The enclosure services string out page transmits an enclosure dependent binary string from the application client to the enclosure. The string may contain bits describing indicator states, text or graphic display information, or control information outside the context of the enclosure elements defined in the enclosure services configuration page. The format is not specified by this standard. The format of the string may be established by an application client using manufacturer name and model from the INQUIRY command and any applicable descriptive text in the RECEIVE DIAGNOSTIC RESULTS enclosure services configuration page. Table ff describes the enclosure services string out page transmitted using the SEND DIAGNOSTIC command.

Table ff - Enclosure services string out page

Bytes	Bits	7	6	5	4	3	2	1	0
0		Page Code (04h)							
1		Reserved							
2		(MSB) Page Length (= n-3) (LSB)							
3									
4-n		Binary string provided by application client to enclosure							



8.1.5 Enclosure services help text page

The enclosure services help text page transmits an ASCII string of data from the 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 help text page has the standard enclosure services diagnostic page header. The help text page is intended to allow the writing of enclosure independent SCSI drivers that can 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 enclosure services help text page shall be read by the RECEIVE DIAGNOSTIC RESULTS command. The transmission of a page with page code 03 using a SEND DIAGNOSTIC command is an error, causing the presentation of CHECK CONDITION and ASC/ASCQ of INVALID FIELD IN PARAMETER LIST (26/00).

Some very simple devices are not capable of reporting any enclosure service page except the short status page, defined in section 8.1.3. Such devices shall always provide the short status page, regardless of which enclosure services page was requested.

Table dd describes the enclosure services help text page.

Table dd - Enclosure services help text page

Bytes	Bits	7	6	5	4	3	2	1	0
0		Page Code (03h)							
1		Reserved							
2	(MSB)	Page Length (= n-3)							
3	(LSB)								
4 - n		Enclosure services ASCII help text string							

8.1.6 Enclosure services string in page

The enclosure services string in page transmits an enclosure dependent binary string from the enclosure to the application client. The string may contain bits describing keyboard states, switch states, or the content of other information provided through or by the enclosure to the application client. The format is not specified by this standard. The format of the string may be determined by an application client using manufacturer name and model from the INQUIRY command and any applicable descriptive text in the RECEIVE DIAGNOSTIC RESULTS enclosure services configuration page. Table ee describes the enclosure services string in page received using the RECEIVE DIAGNOSTIC RESULTS command.



services diagnostic page header. The page shall be implemented if the device supports enclosure services. The control page is transmitted by the SEND DIAGNOSTIC command. The request of a page using RECEIVE DIAGNOSTIC RESULTS, page code 2 is defined as the request for an enclosure services status page. Table cc describes the enclosure services control page.

Table cc - Enclosure services control page

Bytes	Bits	7	6	5	4	3	2	1	0
0		Page Code (02h)							
1		Reserved							
2	(MSB) Page Length (= n-3) (LSB)								
3									
4-7		Global control for first element type							
8-15		Element control for first element of first element type							
...									
4 bytes		Element control for last element of first element type							
4 bytes		Global control for second element type							
4 bytes		Element control for first element of second element type							
...									
n-3 to n		Element control for last element of last element type							

The global control entry for each element type has the same format as the corresponding element control entry. There is exactly one global control entry for each type descriptor in the configuration page. The global control entry provides collective control for all elements described in the element control entries. Control values may be applied to either the global control entry or the element control entry.

Following the global control entry, there shall be n element control entries for each type descriptor in the configuration page, where n shall be identical to the number of elements field in the type descriptor entry. Each element control entry optionally contains control information for the specified element. Each standard element type has a standard fixed format for its control entry. In most cases, the status entries and control entries are very similar and contain complementary functions.

The global status and element status formats and the corresponding element control formats for each element type are described in section 8.1.8



The critical condition (Crit) bit shall be set to 1 to indicate that an element is not operating correctly or is operating outside its specification and that the failure of the element has affected or will soon affect the continued normal operation of the enclosure. Other elements of the enclosure may continue to operate normally. The critical condition bit shall continue to be presented until all critical conditions have been corrected.

The unrecoverable condition (Unrecov) bit shall be set to 1 to indicate that an element is not operating correctly and that the failure of the element has disabled some functions of the enclosure. The enclosure may be incapable of recovering or redundantly bypassing the failure and will normally require repairs to correct the condition. The unrecoverable condition bit shall continue to be presented until all repairs necessary to correct unrecoverable conditions are completed.

If the optional short status bit is set to 1, the page format is fixed. The enclosure failure warning bit (EFW) is set to one to indicate that the enclosure is about to fail and to provide warning information to those SCSI devices able to make use of that to set themselves in a state preparatory for the enclosure failure. The EFW signal is optional. The action taken when a SCSI device detects the enclosure failure warning is set by the MODE SELECT command and described in subclause 8.3.n. The remaining 6 bits of byte 1 contain vendor unique enclosure condition indicator bits.

The global status entry for each element type has the same format as the corresponding element status entry. There is exactly one global status entry for each type descriptor in the configuration page. The global status optionally indicates a summary of the status for all of the elements of that type collectively. The global status also may be used to indicate the status of those elements whose individual status is not available, but that do have a measurable global status. An example might be an enclosure with 3 temperature sensors that cannot record the individual temperatures from those sensors. The global status entry might report the average of the 3 sensors.

One or more element status entries are provided immediately after the global status entry for that element type. There are n element status entries for each type descriptor in the configuration page, where n shall be identical to the number of elements field in the type descriptor entry. Each element status entry optionally indicates the status for the particular element. Each standard element type has a fixed format for its status.

The global status and element status formats and the corresponding element control formats for each element type are described in section 8.1.8

8.1.4 Enclosure services control page

The enclosure services control page transmits a global control entry for each type of element included in the enclosure services configuration page. Below each global control entry, there is an element control entry for each of the elements of that type that have been allowed for by the element count value of the configuration page. The control page has the standard enclosure



Table bc - Enclosure services short status page

Bytes	Bits	7	6	5	4	3	2	1	0
0	Page Code (02h)								
1	Short (= 1)	EFW	VU	VU	VU	VU	VU	VU	VU
2	(MSB) Page Length (= 0) (LSB)								
3									

Note: The short status page is provided when a device that is capable of supporting full enclosure services determines that the enclosure attached to it is only capable of providing a single 7-bit data field. Such a device is defined by the Small Form Factor specification SFF-8045, describing a single connector SCSI FCP FC-AL disk drive with a simple enclosure sensing mechanism.

The short bit is 1 if the enclosure services short status page is being returned. The short status page may be returned in response to a RECEIVE DIAGNOSTIC RESULTS page with any enclosure services page code. The short bit is 0 if the normal enclosure services status page is being returned.

The summary bits are returned in normal status pages and are mandatory. The bits can be recovered with a very short allocation length and can be examined by an enclosure polling procedure to determine if any special or unusual conditions requiring more exhaustive sensing procedures are required. The bits are set independently and can be presented in any combination.

The information bit (Info bit) shall be set to 1 if one or more conditions have been detected in the enclosure since the last time a enclosure services status page has been presented to the application client that should be provided for the information of the application client, but that do not reflect any error or abnormal condition. Such conditions include temperature and battery status that do not reflect a failure or reduction of availability. The information bit shall be provided once as an indication to the application client that interesting information is available and not provided again until some new condition requires it to be set.

The noncritical condition (Non-Crit) shall be set to 1 to indicate that an element is not operating correctly or is operating outside its specification, but that the failure of the element will not affect continued normal operation of the enclosure and all SCSI devices in the enclosure. The degree of operational redundancy may be reduced by a noncritical condition. The noncritical condition bit shall continue to be presented until all noncritical conditions have been corrected.



8.1.3 Enclosure services status page

The enclosure services status page returns a generic information entry for each type of element included in the enclosure services configuration page. Below each generic information entry, there is a type specific status entry for each of the elements of that type that have been allowed for by the element count value of the configuration page. The status page has the standard enclosure services diagnostic page header. This page shall be implemented if the device supports enclosure services. The status page is read by the RECEIVE DIAGNOSTIC RESULTS command. The transmission of a page with page code 2 is defined as the transmission of an enclosure services control page. Table bb describes the enclosure services status page.

Table bb - Enclosure services status page

Bytes	Bits	7	6	5	4	3	2	1	0
0	Page Code (02h)								
1	Short					Info	Non-Crit	Crit	Unrecov
2	(MSB) Page Length (= n-3) (LSB)								
3									
4-7	Global status for first element type								
8-15	Element status for first element of first element type								
•••									
4 bytes	Element status for last element of first element type								
4 bytes	Global status for second element type								
4 bytes	Element status for first element of second element type								
•••									
n-3 to n	Element status for last element of last element type								

Some very simple devices are not capable of reporting any enclosure service page except the short status page, defined in table bc. Such devices shall always provide the short status page, regardless of which enclosure services page was requested. It shall not be an error to respond with a short status page when another enclosure services page has been requested by a RECEIVE DIAGNOSTIC RESULTS command.



Table ac - Element type codes

Type Code	Type of element
00h	unspecified
01h	SCSI device bay/slot
02h	Power supply
03h	Fan
04h	Temperature Sensors
05h	Door Lock
06h	Speaker
07h	ESI controller electronics
08h	SCC controller electronics
09h	Nonvolatile cache
0Ah	Reserved
0Bh	Uninterruptible Power Supply
0Ch	Display
0Dh	Key pad entry device
0Eh	UPS battery
0Fh	SCSI port/transceiver
10-7Fh	Reserved
80h-FFh	Vendor-specific codes

The number of elements field indicates the number of elements of that type that can be installed in the enclosure.

The type descriptor text length field specifies the length of the type descriptor text for the particular element. Vendor unique element types shall have a descriptor length field that is nonzero and shall have a type descriptor text adequate to identify the element to a driver program.



The type descriptor header is a 4-byte header that indicates the type of element being described, the number of such elements, and the length of an ASCII text descriptor of the element type. The elements of an enclosure will be listed in the same order in the configuration page, the type descriptor text of the configuration page, the status page, and the control page. The format of the type descriptor header is shown in table ab.

The type descriptor text is an optional ASCII text string from 0 to 256 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 might be useful to a driver that is displaying the configuration of the enclosure. Examples 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.

Table ab - Configuration page type descriptor header format

Bytes	Bits	7	6	5	4	3	2	1	0
0	Type of element								
1	Number of elements (number of that type possible in enclosure)								
2	Reserved								
3	Type descriptor text length								

The type of element field in each type descriptor indicates the type of element being described in that particular header. Standard element types and reserved element type values are shown in Table ac. More than one type descriptor header may exist for a particular type of element. As an example, there may be two power supplies that provide +12 volts, and five power supplies that provide +5 volts. In this case, separate headers may be used for each of the power supply types, since they may have separate text descriptions.



CONDITION and an ASC/ASCQ of INVALID FIELD IN PARAMETER LIST (26/00). Table aa describes the enclosure services configuration page.

Table aa - Enclosure services configuration page

Bytes	Bits	7	6	5	4	3	2	1	0
0		Page Code (01h)							
1		Reserved							
2	(MSB)	Page Length (n-3) (LSB)							
3									
4		Global descriptor length (m, allowable values 0-252, 4-byte granularity)							
5		Reserved							
6		Number of unit types supported (T, allowable values 0-255)							
7		Reserved							
8 - 15		Enclosure Logical Identifier (World Wide Name or other identifier) (optional)							
16 - (7+m)		Vendor Unique global descriptor information (optional)							
(8+m) through (7+m+4T)		Type descriptor header 4 bytes for each type of element in enclosure							
(8+m+4T) through n		Type descriptor text specified text string for each element in enclosure optional for each element							

The global descriptor length specifies the number of bytes that follow the 4-byte global descriptor header. The value shall be a multiple of 4, having allowed values between 0 and 252.

The number of unit types supported field defines the number of 4-byte type descriptor fields that will follow the global descriptor.

The enclosure logical identifier field is optional. The first 8 bytes of the global descriptor, if allowed by the global descriptor length field, shall always contain either the logical identifier of the enclosure or a value of 0. The enclosure logical identifier shall use the world wide name format defined by X3.230-1995.

The vendor unique global descriptor information is optional. It may contain vendor unique identification of the enclosure profile, vendor unique identification of the enclosure's software and hardware level, and vendor unique text strings describing the enclosure's characteristics.



results in the truncation of a parameter value, the device server shall terminate the command with the CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN PARAMETER LIST. The page length and/or allocation length may be set to a value that transmits less than the full page of data bytes without causing an error if the bytes not transmitted are optional or not required to perform the operations specified by the diagnostic page.

[Editor's note: See section 3 of this proposal, issues and questions, item 3 for consideration of the treatment of length described above.]

For a RECEIVE DIAGNOSTIC RESULTS command, it is not an error for the allocation length to specify a value shorter than the total length of the diagnostic page. The page length field shall always specify the correct length of a complete diagnostic page.

The last paragraph of subclause 8.1.1 is modified to indicate that, if enclosure services pages are supported by a SCSI device, then the pages are indicated as supported regardless of the temporary or permanent ability of the enclosure to provide the necessary information to the device.

The supported page list field shall contain a list of all diagnostic page codes implemented by the target in ascending order beginning with page code 00h. If the device is capable of accessing a diagnostic function or enclosure function that may temporarily or permanently be unavailable to the device, the page code associated with that information shall be included in the list. Check conditions are only posted by the target when a diagnostic command is executed and fails, not when the list of supported pages is generated.

Additional subclauses, 8.1.2 - 8.1.8, are added to define the contents of the Enclosure Services pages.

8.1.2 Enclosure services configuration page

The enclosure services configuration page returns a list of components in an enclosure. This page shall be implemented if the device supports enclosure services. The component list shall include all components for which component status or controls are defined and may list any other components in the enclosure. The configuration page provides global enclosure information and parameters. The configuration page optionally provides descriptive text which can be used by drivers to identify components in more detail. The configuration page may be read by the RECEIVE DIAGNOSTIC RESULTS command. The transmission of a page with page code 01 using a SEND DIAGNOSTIC command is an error, causing the presentation of CHECK

If the target supports the page code option and a page value of zero is specified by the RECEIVE DIAGNOSTIC RESULTS command, the page code field and parameters sent in the previous SEND DIAGNOSTIC command specifies the format of the returned data. If no previous SEND DIAGNOSTIC command was provided from the same application client, the Supported Diagnostic Pages (page code 0) is returned as defined in 8.1.1. If the page code is set to a value other than zero in the RECEIVE DIAGNOSTIC RESULTS command, the command shall return the specified page if it supports the page code option. If the specified page code is not supported or if the page code option is not implemented by the device server, the target shall return CHECK CONDITION status with a sense key of ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.

If the allocation length is greater than the length of the fields to be returned, only the data specified by the page formats is returned. Device servers shall not adjust the content of the returned data to reflect truncation if the allocation length is less than the length specified by the page formats.

[paragraphs including: “A reservation conflict shall occur....page format definitions.” remain unchanged.

Section 8.1 Table 66 is modified to define a generic page code for Enclosure Services as shown below. The mandatory or optional nature of each page is defined in the text for each page, consistent with the other subclauses of clause 8.

Table 66 - Diagnostic page codes

Page Code	SEND DIAGNOSTIC	RECEIVE DIAGNOSTIC RESULTS	Subclause
00h	N/A	Supported diagnostics	8.1.1
01h	N/A	Enclosure Services Configuration	8.1.2
02h	Enclosure Services Control	Enclosure Services Status	8.1.4, 3
03h	N/A	Enclosure Services Help Text	8.1.5
04h	Enclosure Services String Out	Enclosure Services String In	8.1.6, 7
02h-3Fh	Reserved (applies to all device type pages)	Reserved (applies to all device type pages)	
40h-7Fh	See specific device type for definition	See specific device type for definition	
80h-FFh	Vendor-specific pages	Vendor-specific pages	

The next to the last paragraph of subclause 8.1 is modified to indicate that truncation may not necessarily cause a CHECK CONDITION.

The page length field specifies the length in bytes of the diagnostic parameters that follow this field. For a SEND DIAGNOSTIC command, if the application client sends a page length that



2.1 Additions to INQUIRY command

Those devices using the enclosure services device type model use a peripheral device type code of 0Dh. The following entries must be added or modified in Table 21, Peripheral device type. Only those rows that are added or modified are included in the following table.

Table 21 - Peripheral device type

Code	Doc.	Description
0Dh	here	Enclosure services device
0Eh - 1Eh		Reserved

2.2 Definition of SEND DIAGNOSTIC and RECEIVE DIAGNOSTIC RESULTS

Section 7.23 of SPC defines the SEND DIAGNOSTIC command. No change is required.

Section 7.16 of SPC defines the RECEIVE DIAGNOSTIC RESULTS command. The first paragraph and the parameter descriptions are modified to allow a particular page to be requested and to allow the command to be executed without a preceding SEND DIAGNOSTIC command as follows. Note that this text also corrects a discrepancy in the present definition of the command, where the definition of the returned page is not specified clearly for the case where no SEND DIAGNOSTIC precedes the RECEIVE DIAGNOSTIC RESULTS command.

7.16 RECEIVE DIAGNOSTIC RESULTS command.

The RECEIVE DIAGNOSTIC RESULTS command (see table 41) requests that analysis data be sent to the application client after completion of a SEND DIAGNOSTIC command (see 7.23). The command optionally specifies which diagnostic page is to be returned by setting the page code to a value other than zero.

Table 41: RECEIVE DIAGNOSTIC RESULTS command

Bits Bytes	7	6	5	4	3	2	1	0
0	Operation Code (1Ch)							
1	reserved							
2	page code							
3	(MSB) Allocation Length (LSB)							
4								
5	Control							

Enclosure Services proposal

1 Description of Enclosure Services Interface and proposal:

1.1 Overview

Using SCSI-3 technologies, disk drives and other SCSI devices may reside in subsystem enclosures some distance from the host initiators. To maintain a single communication interface with the host computer, it is important that any status information in the subsystem enclosure be available through the same SCSI interface that provides communications with the SCSI devices inside the enclosure. In addition, indicators, locks, and displays on the subsystem must be managed by the host through the SCSI interface. Some enclosures may make this information available through special ports to the standard SCSI devices. An example of such a port is the Enclosure Service Interface defined for FC-AL SCA-2 devices by the Small Form Factor document SFF-8067. More sophisticated enclosures may actually create a special logical unit with its own or a shared target id that provides only enclosure services. The SAF-TE proposal by Intel and Conner Peripherals defined a series of bits that are appropriate for most enclosure services functions. This document uses those definitions and clarifies and expands them for more general cases.

1.2 Enclosure Services capabilities

Enclosure services can be implemented as a component of a standard SCSI device or as a separate LUN. The enclosure services commands are optional for any SCSI device. The enclosure services model and the accompanying commands are defined in SPC or a follow-on document to SPC. The commands used to transmit enclosure service information to or from the device are SEND DIAGNOSTIC and RECEIVE DIAGNOSTIC RESULTS. New diagnostic pages are defined to contain the transmitted and received information. A page code field is provided in RECEIVE DIAGNOSTIC RESULTS to explicitly request the desired enclosure status.

A set of functionality is made available to SCSI devices from the enclosure. One of these functions is the Enclosure Failure Warning, which may be generated by an enclosure to indicate to a SCSI device that a failure condition, including loss of primary power or cooling, has occurred. The action that the drive shall take upon detecting EFW, together with certain other enclosure related states, is set by a new MODE SELECT/MODE SENSE page.

The enclosure services device model defines a very simple device that implements certain SCSI-3 primary commands and that transmits and receives enclosure services information.

2 Proposed Modifications to SPC

Table numbers and section numbers, where applicable, are referenced using SPC X3T10/995D revision 8.



February 5, 1996



John Lohmeyer
Chairperson, X3T10
Symbios Logic Inc.
1635 Aeroplaza Drive
Colorado Springs, Colorado 80916

Subject: Proposal for Enclosure Services Interface implementation

Dear Mr. Chairperson:

A broad class of SCSI device enclosures now provide significant internal intelligence to monitor the state of the power supplies, fans, and external switches as well as to set the state of reconfiguration circuits and indicators. In the Small Form Factor committee definitions for single connector attachment SCSI and Fibre Channel drives, these enclosure services can be accessed across the SCA connector using SCSI commands. In other implementations, the enclosure may actually provide a special target that performs these services. The SCSI committee has requested that I contribute the text that would define these capabilities in the SPC document. In addition, the SAF-TE specification has been proposed as a mechanism for performing similar functions. After consultation with the SAF-TE principals, Intel, Conner Peripherals, Adaptec, and Mylex, revision 2 of X3T10/95-324 combines, extends, and refines the definitions of both the SCSI ESI and the SAF-TE specifications in a single document.

This document will serve as the working paper for a meeting on ESI to be held on February 16, 1996, Sheraton San Jose, 1801 Barber Lane, Milpitas, California, 94035. The phone number of the hotel is 408-943-0600. The meeting will begin at 9 am and last till 5 P.M. A continental breakfast will be available at 8:30 am.

Since this document is a major revision, no change bars are included.

Sincerely,

Robert N. Snively
Sun Microsystems
Mail Stop MPK 12-204
2550 Garcia Ave.
Mountain View, CA 94043-1100
phone: 415-786-6694
e-mail: bob.snively@sun.com