

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com) and Mark Williams, HP (mark.d.williams@hp.com)
 Date: 20 October 2007
 Subject: 07-457r0 SES-2 Bypass reason code

Revision history

Revision 0 (20 October 2007) First revision

Related documents

T10/ses2r18 - SCSI Enclosure Services - 2 (SES-2) revision 18
 T11/FC-FS-3-r0.20 - Fibre Channel Framing and Signaling - 3 (FC-FS-3) revision 0.20
 SNIA/Storage Management Initiative Specification (SMI-S) revision 1.1 - see the CIM_FCPortStatistics error counters for a Fibre Channel switch

Overview

This proposal adds a way for the enclosure services process to report the reason that it has bypassed a Fibre Channel port (if the bypass is due to the enclosure, rather than the application client or the device itself).

Reasons are based on the Fibre Channel Link Error Status Block (LESB) fields.

Suggested changes to SES-2

6.1.13.2 Additional Element Status descriptor protocol-specific information for Fibre Channel

The Additional Element Status descriptor is used to describe a Device element or an Array Device element that may contain a Fibre Channel device, or a SCSI Initiator Port, SCSI Target Port, or Enclosure Services Controller Electronics element that is a Fibre Channel device.

Table 28 defines the Additional Element Status descriptor protocol-specific information for Fibre Channel devices (see FCP-4) with the EIP bit set to one.

Table 28 — Additional Element Status descriptor protocol-specific information for Fibre Channel with the EIP bit set to one

Byte\Bit	7	6	5	4	3	2	1	0	
0	NUMBER OF PORTS								
1	Reserved								
2	BAY NUMBER								
3	BAY NUMBER								
4	(MSB)	NODE NAME						(LSB)	
11	NODE NAME								
Port descriptor list									
12	Port descriptor (first)(see table 30)								
27	Port descriptor (first)(see table 30)								
...									
y - 15	Port descriptor (last)(see table 30)								
y	Port descriptor (last)(see table 30)								

Table 29 defines the Additional Element Status descriptor protocol-specific information for Fibre Channel devices (see FCP-4) with the EIP bit set to zero.

Fibre Channel device. This format does not include the two extra bytes that are in table 28

Table 29 — Additional Element Status descriptor protocol-specific information for Fibre Channel with the EIP bit set to zero

Byte\Bit	7	6	5	4	3	2	1	0
0	NUMBER OF PORTS							
1	Reserved							
2	(MSB)	NODE NAME						(LSB)
9								
Port descriptor list								
10	Port descriptor (first)(see table 30)							
25	...							
y - 15	Port descriptor (last)(see table 30)							
y								

The NUMBER OF PORTS field indicates how many Fibre Channel ports are in the port descriptor list. There is one port descriptor for each port.

The BAY NUMBER field indicates the number of the bay (i.e., the device slot for a Device or Array Device element) represented by the element.

The NODE NAME field contains the node Name_Identifier of the corresponding Fibre Channel node.

Table 30 defines the port descriptor.

Table 30 — Port descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	PORT LOOP POSITION							
1	Reserved BYPASS REASON							
2	Reserved							
3								
4	PORT REQUESTED HARD ADDRESS							
5	(MSB)	N_PORT IDENTIFIER						(LSB)
7								
8	(MSB)	N_PORT_NAME						(LSB)
15								

The PORT LOOP POSITION field indicates the position of the corresponding Fibre Channel port on a Fibre Channel Arbitrated Loop.

The **BYPASS REASON** field indicates the reason the corresponding Fibre Channel port is being bypassed, if it is being bypassed, and is defined in table 31.

Table 31 — BYPASS REASON field [all new]

Code	Description
00h	Either: a) the port is not being bypassed; or b) the port is being bypassed and no reason is available (e.g., it is being bypassed by request of an application client or the device).
01h - 0Fh	Reserved
10h	Link failure rate is too high
11h	Loss-of-synchronization rate is too high
12h	Loss-of-signal rate is too high
13h	Primitive sequence protocol error rate is too high
14h	Invalid transmission word rate is too high
15h	CRC error rate is too high
16h - 1Fh	Reserved
20h	Link failure count is too high
21h	Loss-of-synchronization count is too high
22h	Loss-of-signal count is too high
23h	Primitive sequence protocol error count is too high
24h	Invalid transmission word count is too high
25h	CRC error count is too high
26h - 2Fh	Reserved
30h - BFh	Reserved
C0h - FFh	Vendor specific

Editor's Note 1: The FC-FS Link Error Status Block (LESB), returned in response to an RLS (Read Link Error Status Block) ELS, records:

word 0: link failure (loss of signal, loss of sync > timeout)

word 1: loss-of-synchronization

word 2: loss-of-signal

word 3: primitive sequence protocol error

word 4: invalid transmission word

word 5: invalid CRC

The reasons proposed correspond to the LESB counters.

Editor's Note 2: SMI-S page 741 defines these properties related to FC port statistics:

CRCErrors

LinkFailures

PrimitiveSeqProtocolErrCount

LossOfSignalCounter

InvalidTransmissionWords

LIPCount
 NOCount
 ErrorFrames
 DumpedFrames
 LossOfSyncCounter

The reasons proposed correspond to the LESB counters rather than the SMI-S properties, which are poorly defined (just the names in most cases).

The PORT REQUESTED HARD ADDRESS field contains the Fibre Channel Arbitrated Loop requested hard address of the corresponding Fibre Channel port.

The N_PORT IDENTIFIER field contains the address identifier of the corresponding Fibre Channel port. Applications may compare the lower 8 bits of this field with the PORT REQUESTED HARD ADDRESS field to determine whether the port was assigned its requested address.

The N_PORT_NAME field contains the Name_Identifier of the corresponding Fibre Channel port.

7.3.2 Device element

The Device element manages a SCSI device (e.g., a disk drive) in the enclosure.

Additional information about a Device element may be reported in the Additional Element Status diagnostic page (see 6.1.13).

The format of the control field for a Device element in the Enclosure Control diagnostic page (see 6.1.3) is defined in table 32.

Table 32 — Device element for the Enclosure Control diagnostic page

Byte\Bit	7	6	5	4	3	2	1	0
0	COMMON CONTROL							
1	Reserved							
2	RQST ACTIVE	DO NOT REMOVE	Reserved	RQST MISSING	RQST INSERT	RQST REMOVE	RQST IDENT	Reserved
3	Reserved		RQST FAULT	DEVICE OFF	ENABLE BYP A	ENABLE BYP B	Reserved	

...

An ENABLE BYP A (enable bypass A) bit set to one specifies that port A for the device be bypassed. An ENABLE BYP A bit set to zero specifies that, if there is no other cause for the port to be bypassed, the port bypass shall be disabled and the device shall be included on the device interface.

An ENABLE BYP B (enable bypass B) bit set to one specifies that port B for the device be bypassed. An ENABLE BYP B bit set to zero specifies that, if there is no other cause for the port to be bypassed, the port bypass shall be disabled and the device shall be included on the device interface.

...

The format of the status field for a Device element in the Enclosure Status diagnostic page (see 6.1.4) is defined in table 33.

Table 33 — Device element for the Enclosure Status diagnostic page

Byte\Bit	7	6	5	4	3	2	1	0
0	COMMON STATUS							
1	SLOT ADDRESS							
2	APP CLIENT BYPASSED A	DO NOT REMOVE	ENCLOSURE BYPASSED A	ENCLOSURE BYPASSED B	READY TO INSERT	RMV	IDENT	REPORT
3	APP CLIENT BYPASSED B	FAULT SENSED	FAULT REQSTD	DEVICE OFF	BYPASSED A	BYPASSED B	DEVICE BYPASSED A	DEVICE BYPASSED B

...

A BYPASSED A bit set to one indicates that Port A has been bypassed by request of the application client, the device, or the enclosure. A BYPASSED A bit set to zero indicates that the port bypass is disabled and the device is included on the device interface.

A BYPASSED B bit set to one indicates that Port B has been bypassed by request of the application client, the device, or the enclosure. A BYPASSED B bit set to zero indicates that the port bypass is disabled and the device is included on the device interface.

~~A BYP A ENABLED (bypass A enabled) bit set to one indicates that port A of the device is bypassed under control of the device. The device may be removed, turned off, not operational, or controlling the bypass signals under control of the device server. A BYP A ENABLED bit set to zero indicates that Port A is not being bypassed under control of the device. The device may still be bypassed under control of the enclosure services process.~~

~~A BYP B ENABLED (bypass B enabled) bit set to one indicates that port B of the device is bypassed under control of the device. The device may be removed, turned off, not operational, or controlling the bypass signals under control of the device server. A BYP B ENABLED bit set to zero indicates that Port B is not being bypassed under control of the device. The device may still be bypassed under control of the enclosure services process.~~

An ENCLOSURE BYPASSED A bit set to one indicates that Port A has been bypassed by request of the enclosure services process. An ENCLOSURE BYPASSED A bit set to zero indicates that Port A is not being bypassed under control of the enclosure services process. The device may still be bypassed under control of the application client or the device.

An ENCLOSURE BYPASSED B bit set to one indicates that Port B has been bypassed by request of the enclosure services process. An ENCLOSURE BYPASSED B bit set to zero indicates that Port B is not being bypassed under control of the enclosure services process. The device may still be bypassed under control of the application client or the device.

An APP CLIENT BYPASSED A (application client bypassed [Port A](#)) bit set to one indicates that Port A has been bypassed by request of an application client. An APP CLIENT BYPASSED A bit indicates that Port A is not being bypassed under control of an application client. The device may still be bypassed under control of the enclosure services process or the device.

An APP CLIENT BYPASSED B (application client bypassed [Port B](#)) bit set to one indicates that Port B has been bypassed by request of an application client. An APP CLIENT BYPASSED B bit indicates that Port B is not being bypassed under control of an application client. The device may still be bypassed under control of the enclosure services process or the device.

A DEVICE BYPASSED A bit set to one indicates that Port A has been bypassed by request of the device. A DEVICE BYPASSED A bit indicates that Port A is not being bypassed by request of the device. When set to one,

the device may be removed, turned off, not operational, or controlling the bypass signals under control of the device server. When set to zero, the device may still be bypassed under control of the enclosure services process or the application client.

A DEVICE BYPASSED B bit set to one indicates that Port B has been bypassed by request of the device. A DEVICE BYPASSED B bit indicates Port B is not being bypassed by request of the device. When set to one, the device may be removed, turned off, not operational, or controlling the bypass signals under control of the device server. When set to zero, the device may still be bypassed under control of the enclosure services process or the application client.