

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 11 April 2006
 Subject: 06-187r0 SAS-2 Self-configuring expander status

Revision history

Revision 0 (11 April 2006) First revision

Related documents

sas2r03 - Serial Attached SCSI - 2 (SAS-2) revision 3

Overview

Self-configuring expander devices may encounter errors while performing the discover and configuration process. Possible errors include:

- a) Physical layer problems. The expander device tries to talk to a SAS address through one of its expander ports. There could be errors on specific phys within the expander port.
 - A) The phy(s) being used lose dword synchronization (either the receiver lost it outright, or there were too many invalid dwords)
- b) Connection request problems. The OPEN address frame results in:
 - A) Open Timeout timer expiration
 - B) An abandon-class OPEN_REJECT (e.g., BAD DESTINATION, PROTOCOL NOT SUPPORTED, ZONE VIOLATION, STP RESOURCES BUSY, WRONG DESTINATION)
 - C) Too many retry-class OPEN_REJECTs (e.g. RETRY, PATHWAY BLOCKED)
 - D) I_T nexus loss (OPEN_REJECT (NO DESTINATION) for longer than the STP SMP I_T Nexus Loss time timer)
 - E) BREAK received
- c) Other SMP link layer issues. The expander device can retry a few times, but eventually needs to give up if the errors keep occurring.
 - A) BREAK occurs
 - B) SMP response frame has a CRC error
- d) Port layer issues.
 - A) No SMP response frame within the 2 ms SMP connection time (enforced by the port layer)
- e) Transport layer issues
- f) Application layer issues. The expander device must parse REPORT GENERAL and DISCOVER responses.
 - A) SMP response frame is too short (not SAS-1.1 compliant)
 - B) SMP response frame contains fields with illegal values
 - C) SMP response frame isn't coherent (e.g. the upstream attached SAS address reported in DISCOVER for an expander several levels away doesn't match the SAS address that the expander device at the higher level)
 - D) The CONFIGURING bit in REPORT GENERAL is set to one. This is informational, not an error. The discovery process should not care - it should just continue through the expander and handle OPEN_REJECT (NO DESTINATION)s per the I_T nexus loss rules. However, if multiple passes keep finding the CONFIGURING bit stuck at one, something is wrong.
- g) Expander has no more room in its routing table or is otherwise out of resources

A new SMP REPORT SELF-CONFIGURATION STATUS function is proposed to report these errors to interested management software. It returns an expander-wide list, reporting the phy that was involved in the problem in each status descriptor.

Additionally, a SELF-CONFIGURATION STATUS field is added to the DISCOVER response to report the latest status related to that specific phy. As a management application client is running the discover process, it sends DISCOVER to query about each phy in an expander; this will let it know if it should send REPORT SELF-CONFIGURATION STATUS to the expander to investigate problems in the self-configuration process.

Editor's Note 1: Could create another proposal to define an SMP configuration function to instruct the expander device how to respond to certain errors. For example, how many retries shall it

attempt (or how long shall it retry) if CRC errors occur? Which errors are considered retryable?
This is all vendor-specific right now.

Editor's Note 2: The configuration proposal should include a way to tell a self-configuring expander to stop using a specified phy to perform self-configuration, and also to tell it to stop self-configuring altogether.

Editor's Note 3: 06-078 by Steve Johnson (LSI) is defining a new mechanism to read the route table from a self-configuring expander. The SAS-1.1 expander route index rules implied by REPORT ROUTE INFORMATION need not be followed by the new function; just return all the SAS addresses that are routed through the selected expander phy.

Editor's Note 4: How should self-configuration status entries be cleared? Don't want errors to be reported fleetingly, disappearing the next time the expander runs the discover process. Don't want multi-initiator issues stemming from a shared CLEAR STATUS function, though. Should the function just report the Last N errors (FIFO like behavior) at all times, with a fixed maximum number than can be reported?

Suggested changes

10.4.3.1 SMP function request frame format

...

The FUNCTION field specifies which SMP function is being requested and is defined in table 1. If the value in the FUNCTION field is not supported by the SMP target port, it shall return a function result of UNKNOWN SMP FUNCTION as described in table 197.

Table 1 — SMP functions (FUNCTION field)

Code	SMP function	Description	Reference
00h	REPORT GENERAL	Return general information about the device	10.4.3.3
01h	REPORT MANUFACTURER INFORMATION	Return vendor and product identification	10.4.3.4
02h	READ GPIO REGISTER	See SFF-8485	
03h	REPORT SELF-CONFIGURATION STATUS	Return self-configuring expander device status	10.4.3.x
03h 04h - 0Fh	Reserved for general SMP input functions		
10h	DISCOVER	Return information about the specified phy	
11h	REPORT PHY ERROR LOG	Return error logging information about the specified phy	10.4.3.6
12h	REPORT PHY SATA	Return information about a phy currently attached to a SATA phy	10.4.3.7
13h	REPORT ROUTE INFORMATION	Return route table information	10.4.3.8
14h	REPORT PHY EVENT INFORMATION	Return phy event information for the specified phy	
15h - 1Fh	Reserved for phy-based SMP input functions		
20h - 3Fh	Reserved for SMP input functions		
...	...		

10.4.3.5 DISCOVER function

The DISCOVER function returns the physical link configuration information for the specified phy. This SMP function provides information from the IDENTIFY address frame received by the phy and additional phy-specific information. This SMP function shall be implemented by all SMP target ports.

Table 2 defines the request format.

Table 2 — DISCOVER request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (10h)								
2	Reserved								
3	REQUEST LENGTH (02h)								
4	Reserved								
8	Reserved								
9	PHY IDENTIFIER								
10	Reserved								
11	Reserved								
12	(MSB)	CRC							
15							(LSB)		

The SMP FRAME TYPE field shall be set to 40h.

The FUNCTION field shall be set to 10h.

The REQUEST LENGTH field shall be set to 02h. For compatibility with previous versions of this standard, a REQUEST LENGTH field set to 00h specifies that there are 2 dwords before the CRC field.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for the link configuration information being requested.

The CRC field is defined in 10.4.3.1.

Table 3 defines the response format.

Table 3 — DISCOVER response (part 1 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (10h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (0Eh)							
4	Reserved							
8	Reserved							
9	PHY IDENTIFIER							
10	Reserved							
11	Reserved							
12	Reserved	ATTACHED DEVICE TYPE			Reserved			
13	Reserved				NEGOTIATED PHYSICAL LINK RATE			

Table 3 — DISCOVER response (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
14	Reserved				ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	ATTACHED SATA HOST
15	ATTACHED SATA PORT SELECTOR	Reserved			ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET	ATTACHED SATA DEVICE
16	SAS ADDRESS							
23	SAS ADDRESS							
24	ATTACHED SAS ADDRESS							
31	ATTACHED SAS ADDRESS							
32	ATTACHED PHY IDENTIFIER							
33	Reserved							
39	Reserved							
40	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
41	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
42	PHY CHANGE COUNT							
43	VIRTUAL PHY	Reserved			PARTIAL PATHWAY TIMEOUT VALUE			
44	Reserved				ROUTING ATTRIBUTE			
45	Reserved	CONNECTOR TYPE						
46	CONNECTOR ELEMENT INDEX							
47	CONNECTOR PHYSICAL LINK							
48	Reserved							
49	Reserved							
50	Vendor specific							
51	Vendor specific							
52	ATTACHED DEVICE NAME							
59	ATTACHED DEVICE NAME							
60	SELF-CONFIGURATION STATUS							
61	SELF-CONFIGURATION LEVELS COMPLETED							
62	Reserved							
63	Reserved							
64	Reserved							
71	SELF-CONFIGURATION SAS ADDRESS							
6072	(MSB)				CRC			
6375					(LSB)			

The SMP FRAME TYPE field shall be set to 41h.

...

The SELF-CONFIGURATION STATUS field indicates the status of a self-configuring expander device pertaining to the specified phy and is defined in table 4.

Table 4 — SELF-CONFIGURATION STATUS field

<u>Code</u>	<u>Description</u>
00h	No status available
01h - FFh	As defined for the STATUS TYPE field in the self-configuration status descriptor in the REPORT SELF-CONFIGURATION STATUS response (see table xx in 10.4.3.x)

The SELF-CONFIGURATION LEVELS COMPLETED field indicates the number of levels beyond the expander port containing the specified phy for which the self-configuring expander device has completed the discovery and configuration process. A SELF-CONFIGURATION LEVELS COMPLETED field set to 00h indicates the expander device has performed no discovery and configuration process through the expander port containing the specified phy.

The SELF-CONFIGURATION SAS ADDRESS field indicates the SAS address of the SMP target port to which the self-configuring expander device established a connection or attempted to establish a connection using the specified phy and resulted in the status indicated by the SELF-CONFIGURATION STATUS field.

The CRC field is defined in 10.4.3.2.

10.4.3.x REPORT SELF-CONFIGURATION STATUS function [all new section]

The REPORT SELF-CONFIGURATION STATUS function returns self-configuration expander device status. This SMP function shall be implemented by the SMP target port in self-configuring expander devices and shall not be implemented by any other SMP target ports.

Table 5 defines the request format.

Table 5 — REPORT SELF-CONFIGURATION STATUS request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (03h)								
2	Reserved								
3	REQUEST LENGTH (01h)								
4	Reserved								
6	Reserved								
7	PAGE NUMBER								
8	(MSB)	CRC							
11							(LSB)		

The SMP FRAME TYPE field shall be set to 40h.

The FUNCTION field shall be set to 03h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 1).

The PAGE NUMBER field specifies the page of status information that is being requested. This function is capable of returning more data than fits in a single SMP response frame, so the response data is divided into pages starting with page number 00h. If the PAGE NUMBER field specifies a page which contains no status

information, the SMP target port shall return a response with the NUMBER OF SELF-CONFIGURATION STATUS DESCRIPTORS field set to zero.

Editor's Note 5: "Page" is an overloaded term in SCSI; another name might be advisable

The CRC field is defined in 10.4.3.1.

Table 6 defines the response format.

Table 6 — REPORT SELF-CONFIGURATION STATUS response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (03h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
4	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5								
6	Reserved							
7	PAGE NUMBER							
8								
13	Reserved							
14	TOTAL NUMBER OF SELF-CONFIGURATION STATUS DESCRIPTORS							
15	NUMBER OF SELF-CONFIGURATION STATUS DESCRIPTORS							
16								
n - 4	Self-configuration status descriptor(s)							
n - 3	(MSB)	CRC						(LSB)
n								

The SMP FRAME TYPE field shall be set to 41h.

The FUNCTION field shall be set to 03h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field contains the number of dwords that follow, not including the CRC field.

The EXPANDER CHANGE COUNT field is defined in the SMP REPORT GENERAL function response (see 10.4.1.3). If the SMP initiator port detects a change in the value of this field while retrieving multiple pages, it should start again because the status information returned is incomplete and inconsistent.

The PAGE NUMBER field indicates the page of status information being returned, and is set to the same value as the PAGE NUMBER field in the SMP request. frame

The TOTAL NUMBER OF SELF-CONFIGURATION STATUS DESCRIPTORS field indicates how many self-configuration status descriptors are available on all pages.

Editor's Note 6: better to return the total number of descriptors or the total number of pages? Each descriptor is the same size. The 1024 byte frame size holds the 16 byte header + (16 bytes x 63 descriptors)

The NUMBER OF SELF-CONFIGURATION STATUS DESCRIPTORS field indicates how many self-configuration status descriptors follow.

Each self-configuration status descriptor follows the format defined in table 7.

Table 7 — Self-configuration status descriptor

Byte\Bit	7	6	5	4	3	2	1	0	
0	STATUS TYPE								
1	Reserved						FINAL	VITAL	
2	Reserved								
2	Reserved								
3	PHY IDENTIFIER								
4	Reserved								
7	Reserved								
8	(MSB)	SAS ADDRESS							
15							(LSB)		

The STATUS TYPE field indicates the type of status being reported and is defined in table 8.

Table 8 — STATUS TYPE field (part 1 of 2)

Code	Description
00h	Reserved
01h	The expander device currently has a connection or is currently attempting to establish a connection with the SMP target port with the indicated SAS address.
02h	Expander route table is full. The expander device was not able to add the indicated SAS address to the expander route table.
03h	Expander device is out of resources (e.g., it discovered too many SAS addresses while performing the discover/configuration process through a subtractive port). This does not affect the expander route table.
04h - 1Fh	Reserved for status not related to specific layers
Status reported by the phy layer	
20h	All phys in the expander port containing the indicated phy lost dword synchronization
21h - 3Fh	Reserved for status reported by the phy layer

Table 8 — STATUS TYPE field (part 2 of 2)

Code	Description
Status reported by the link layer	
40h	Connection request failed: Open Timeout timer expired
41h	Connection request failed: Received an abandon-class OPEN_REJECT (e.g., BAD DESTINATION, PROTOCOL NOT SUPPORTED, ZONE VIOLATION, STP RESOURCES BUSY, WRONG DESTINATION)
42h	Connection request failed: Received too many retry-class OPEN_REJECTs (e.g. RETRY, PATHWAY BLOCKED)
43h	Connection request failed: I_T nexus loss occurred (e.g., OPEN_REJECT (NO DESTINATION) for longer than the time specified by the STP SMP I_T NEXUS LOSS TIME field in the CONFIGURE GENERAL function
44h	Connection request failed: Received BREAK
45h	Connection established: SMP response frame had a CRC error
46h - 5Fh	Reserved for status reported by the link layer
Status reported by the port layer	
60h	During an SMP connection, there was no SMP response frame within the maximum SMP connection time [i.e. 2 ms]
61h - 7Fh	Reserved for status reported by the port layer
Status reported by the transport layer	
80h - 9Fh	Reserved for status reported by the transport layer
Status reported by the application layer	
A0h	SMP response frame is too short
A1h	SMP response frame contains field(s) with unsupported values
A2h	SMP response frame contains results inconsistent with other SMP response frames (e.g., the DISCOVER response ATTACHED SAS ADDRESS field does not contain the SAS address the expander device expected)
A3h	This is the SAS address of a self-configuring expander device that returned a REPORT GENERAL response with the CONFIGURING bit set to one. Accesses to SAS addresses two levels beyond this expander device may not succeed until that expander device completes configuration. This is not necessarily an error.
A0h - BFh	Reserved for status reported by the application layer
Other status	
C0h - DFh	Reserved
E0h - FFh	Vendor-specific

Editor's Note 7: Although not likely, expander devices can send multiple SMP requests at a time - up to one per expander phy. So, that many SAS addresses could be returned with code 01h. The same SAS address could even be reported multiple times, but with different phy identifiers.

A FINAL bit set to one indicates the expander device has given up trying to access the SMP target port with the indicated SAS address (e.g., all attempts at retrying, perhaps through different phys, have failed) and the descriptor contains the final status for that SAS address. A FINAL bit set to zero indicates the expander device has not given up trying to access the SMP target port with the indicated SAS address.

A VITAL bit set to one indicates the specified SAS address was being accessed to configure the expander device's own expander route table. A VITAL bit set to zero indicates the specified SAS address was being accessed for other purposes (e.g., performing the discover/configuration process through a subtractive port, looking for configurable expander devices and configuring their expander route tables).

[Editor's Note 8: there are probably better names than VITAL and FINAL](#)

The PHY IDENTIFIER field indicates the phy (see 4.2.7) that was used to request a connection with the SMP target port with the indicated SAS address.

The SAS ADDRESS field indicates the SAS address of the SMP target port to which the expander device established a connection or attempted to establish a connection.

The CRC field is defined in 10.4.3.2.