T10/07-397 revision 0

Date: September 13, 2007

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: Indeterminate response lenght to a SMP REPORT GENERAL function

1 Overview

As a result of proposal 05-306r2 (SAS-2 STP connection time limits and STP/SMP I_T nexus loss) a REQUEST LENGTH field was added to all the SMP function and a RESPONSE LENGTH field was added to all the SMP responses (even though neither of those has anything to do with STP or I_T nexus loss).

This change created a minefield for SAS 1.1 and SAS 2 compatibility by changing the SAS 1.1 SMP requests and responses from fixed structures to variable length structures (to understand the magnitude of this change consider what would happen if we changed any of the existed fixed length SCSI CDBs to a variable length CDB). Also, on SCSI CDBs that have parameters lists that are returned there is an allocation length specified which tells the target the maximum amount of data that can be sent. That is there to allow parameters lists to become longer in future generations of standards without impacting past implementations. There was no allocation length like field added in the SMP functions with the length additions so there will forever be having a problem with response length changes.

The only thing that keeps this from being a total disaster is that for all except two of the SMP functions the new REQUEST LENGTH field had to contain a non-zero value for SAS-2 compliance and all the new RESPONSE LENGTH fields have to contain non-zero values if the SMP request contained a non-zero value in the REQUEST LENGTH field. This works except that there is a good chance that a SAS 1.1 SMP device may fail a SAS 2 SMP function as the a reserved field contains a value. But the SAS-2 device knowing that this could happen would have to adjust to sending SAS 1.1 SMP functions. If it were not for the two SMP functions that have the same response length for both SAS 1.1 and SAS-2 then all this would be manageable (if not pretty).

The two SMP function that have the REQUEST LENGTH field set to zero in both SAS 1.1 and the current version of SAS-2 are the REPORT GENERAL function and the REPORT MANUFACTURER INFORMATION function. Of those REPORT MANUFACTURER INFORMATION function has no difference in the length of the response length so it should work (as long as the SAS 1.1 initiator ignores the value in the new RESPONSE LENGTH field).

The real problem is that the REPORT GENERAL function which has different lengths for SAS 1.1 (i.e., 32 bytes) and SAS-2 (i.e., 72 bytes). The problem occurs when a SAS 1.1 device issues a REPORT GENERAL function to a SAS-2 SMP device. The SAS-2 SMP device is required to deliver 72 bytes. That can cause the SAS 1.1 initiator to choke as it is only expecting 32 bytes.

This proposal addresses this issue by adding a bit to the REPORT GENERAL function that specifies if this initiator is requesting a response length of 32 bytes or 72 bytes. It also includes a bit in the in the REPORT GENERAL response to specify if the SMP device supports the short response length or the long response length.

The bit in the REPORT GENERAL function allow SAS-2 SMP devices to know the length of the response data. The bit in the REPORT GENERAL response allows a SAS-2 initiator to know if it is talking to an SMP device that support the long or short SMP response.

Of this to work without the possibility of any errors occurring is that a SAS-2 initiator would have to first issue a REPORT GENERAL function with the bit set to short. If the response contains the I support long response indication then it can send a REPORT GENERAL function with the bit set to long. If the response contains the I don't support long indication them it will have to use the SAS 1.1 SMP function formats for all SMP functions to that SMP device.

2 Proposed SAS-2 changes

2.0.0.1 REPORT GENERAL function

The REPORT GENERAL function returns general information about the SAS device (e.g., a SAS device contained in an expander device). This SMP function shall be implemented by all management device servers.

Table 1 defines the request format.

Table 1 — REPORT GENERAL request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (00h)							
2	REQUEST LONG	Reserved						
3	REQUEST LENGTH (00h)							
4	(MSB)	CRC						
7								(LSB)

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

The FUNCTION field shall be set to 00h.

The REQUEST LONG bit set to one specifies that the management device server shall return a non-zero value in the REPORT GENERAL response RESPONSE LENGTH field. The REQUEST LONG bit set to zero specifies that the management device server shall return a zero in the REPORT GENERAL response RESPONSE LENGTH field.

The REQUEST LENGTH field shall be set to 00h.

The CRC field is defined in 10.4.3.1.

Table 2 defines the response format.

Table 2 — REPORT GENERAL response (part 1 of 3)

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (00h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (10h)							
4	(MSB)	EXPANDER CHANGE COUNT						
5								(LSB)
6	(MSB)	EXPANDER ROUTE INDEXES						
7								(LSB)
8	LONG RESPONSE	Reserved						
9	NUMBER OF PHYS							

Table 2 — REPORT GENERAL response (part 2 of 3)

Byte\Bit	7	6	5	4	3	2	1	0
10	TABLE TO TABLE SUPPORTED	Reserved				CONFIGURES OTHERS	CONFIGURING	EXTERNALLY CONFIGURABLE ROUTE TABLE
11	Reserved							
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20	Reserved							
29								
30	(MSB) STP BUS INACTIVITY TIME LIMIT							
31								(LSB)
32	(MSB) STP MAXIMUM CONNECT TIME LIMIT							
33								(LSB)
34	(MSB) STP SMP I_T NEXUS LOSS TIME							
35								(LSB)
36	NUMBER OF ZONE GROUPS		Reserved		ZONE LOCKED	PHYSICAL PRESENCE SUPPORTED	PHYSICAL PRESENCE ASSERTED	ZONING SUPPORTED	ZONING ENABLED
37	Reserved							
38	(MSB) MAXIMUM NUMBER OF ROUTED SAS ADDRESSES							
39								(LSB)
40	ACTIVE ZONE MANAGER SAS ADDRESS							
47								
48	(MSB) ZONE LOCK INACTIVITY TIME LIMIT							
49								(LSB)
50	Reserved							
51								
52	Reserved							
53	FIRST ENCLOSURE CONNECTOR ELEMENT INDEX							
54	NUMBER OF ENCLOSURE CONNECTOR ELEMENT INDEXES							
55	Reserved							
56	REDUCED FUNCTIONALITY	Reserved						

Table 2 — REPORT GENERAL response (part 3 of 3)

Byte\Bit	7	6	5	4	3	2	1	0
57	TIME TO REDUCED FUNCTIONALITY							
58	INITIAL TIME TO REDUCED FUNCTIONALITY							
59	MAXIMUM REDUCED FUNCTIONALITY TIME							
60	(MSB)	LAST SELF-CONFIGURATION STATUS DESCRIPTOR INDEX						
61								(LSB)
62	(MSB)	MAXIMUM NUMBER OF STORED SELF-CONFIGURATION STATUS DESCRIPTORS						
63								(LSB)
64	(MSB)	LAST PHY EVENT INFORMATION DESCRIPTOR INDEX						
65								(LSB)
66	(MSB)	MAXIMUM NUMBER OF STORED PHY EVENT INFORMATION DESCRIPTORS						
67								(LSB)
68	(MSB)	CRC						
71								(LSB)

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

The FUNCTION field shall be set to 00h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 10h. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h indicates that there are 6 dwords before the CRC field.

The EXPANDER CHANGE COUNT field counts the number of Broadcast (Change)s originated by an expander device (see 7.11). Management device servers in expander devices shall support this field. Management device servers in other device types (e.g., end devices) shall set this field to 0000h. This field shall be set to at least 0001h at power on. If the expander device has originated Broadcast (Change) for any reason described in 7.11 since transmitting a REPORT GENERAL response, it shall increment this field at least once from the value in the previous REPORT GENERAL response. It shall not increment this field when forwarding a Broadcast (Change). This field shall wrap to at least 0001h after the maximum value (i.e., FFFFh) has been reached.

NOTE 1 - Application clients that use the EXPANDER CHANGE COUNT field should read it often enough to ensure that it does not increment a multiple of 65 536 times between reading the field.

NOTE 2 - Management device servers in expander devices compliant with previous versions of this standard may return an EXPANDER CHANGE COUNT field set to 0000h.

NOTE 3 - The originated Broadcast (Change) count is also reported in the REPORT BROADCAST response (see 10.4.3.8).

The EXPANDER ROUTE INDEXES field indicates the maximum number of expander route indexes per phy for the expander device (see 4.6.7.3). Management device servers in externally configurable expander devices containing phy-based expander route tables shall support this field. Management device servers in other device types (e.g., end devices, externally configurable expander devices with expander-based expander route tables, and self-configuring expander devices) shall set the EXPANDER ROUTE INDEXES field to zero. Not

all phys in an externally configurable expander device are required to support the maximum number indicated by this field.

The LONG RESPONSE bit set to one indicates that the management device server supports returning non-zero values in the SMP responses RESPONSE LENGTH field when the REQUEST LENGTH field is sent to an non-zero value or a REQUEST LONG bit is set to one. The LONG RESPONSE bit set to zero indicates that the management device server returns a value of zero in the RESPONSE LENGTH field for the following SMP functions:

a) REPORT GENERAL function;
b) REPORT MANUFACTURER INFORMATION function (see x.x.x);
c) READ GPIO REGISTER function (See SFF-8485);
d) DISCOVER function (see x.x.x);
e) REPORT PHY ERROR LOG function (see x.x.x);
f) REPORT PHY SATA function (see x.x.x);
g) REPORT ROUTE INFORMATION function (see x.x.x);
h) WRITE GPIO REGISTER (See SFF-8485) function (see x.x.x);
i) CONFIGURE ROUTE INFORMATION function (see x.x.x);
j) PHY CONTROL function (see x.x.x); and
k) PHY TEST FUNCTION function (see x.x.x).

The NUMBER OF PHYS field indicates the number of phys in the device, including any virtual phys and any vacant phys.