

To: T10 Technical Committee
From: Steve Johnson LSI Logic (steve.johnson@lsil.com)
Date: 21 June 2006
Subject: 06-078r3 SAS-2 Expander Route Table (REPORT EXPANDER ROUTE TABLE)

Revision history

Revision 0 (19 April 2006) First revision
Revision 1 (1 May 2006) Incorporated feedback from April 20 Denver meeting.
Revision 2 (26 May 2006) Incorporated feedback from May Denver meeting.
Revision 3 (21 June 2006) Incorporated various feedback from emails and June 20 Meeting. Accepted for inclusion in SAS-2.

Related documents

sas2r04 - Serial Attached SCSI 2 revision 4.
06-214.
06-213.

Revision Overview r1 to r3

Added CONFIGURING bit and STARTING PHY IDENTIFIER.
Removed EXPANDER ROUTE INDEXES changes and created new field in REPORT GENERAL to specify the total Number of route table entries for the expander.
Changed name to PHY IDENTIFIER BIT MAP
Moved fields around.
Expander change count starts and wraps to non zero value.
Reworked zone group definition
Added the ZONE GROUP VALID bit to the descriptor.

Overview

In SAS 1.1 each table routing phy has it's own routing table. The tables are required to be programmed on a phy basis in a very specific order to facilitate several initiators reading and writing them at the same time. These rules do not make sense for self configuring expanders.

Some disadvantages of this method are:

- 1) Expander add and removal (depending on location) can cause the tables to have to be completely rebuilt.
- 2) To program and build the tables complete discovery must be performed to determine routing phy attributes.
- 3) All unused entries are required to be zeroed by initiators
- 4) The exact number of route entries for a given phy route table can only be determined by a process of trial an error requiring each initiator to send the SMP REPORT ROUTE INFORMATION function or the SMP CONFIGURE ROUTE INFORMATION until a status of INDEX DOSE NOT EXIST is returned.

Another way to implement SAS routing tables is to view the routing tables as one single table for all the phys of the expander, where there is a single list of SAS Addresses and phys that the addresses are routed to. This table may be sparse and addresses are not programmed in any particular order.

Some advantages of this method are:

- 1) The table does not have to be in any specific order.
- 2) Adding and removing SAS ADDRESSES is independent of order.
- 3) Allows for partial discoveries.
- 4) Table does not have to be rebuilt or reordered when expanders (or any device types) are added or removed.
- 5) No zeroing of table entries.

- 6) Matches most hardware implementations.
- 7) The exact table size can be reported.

How self configuring expanders manage their route table or “tables” is vendor specific; A method to report the table that more closely matches the actual implementations is needed.

Suggested changes

Add the SMP REPORT EXPANDER ROUTE TABLE function to section 10.4.3.x SMP functions of the SAS-2
The REPORT EXPANDER ROUTE TABLE function is being proposed to report the current expander routing table.

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	
<u>03h</u>	<u>REPORT ZONE PERMISSION</u>	<u>Return zone permission table entries</u>	
04h - 0Fh	Reserved for general SMP input functions		
10h	DISCOVER	Return information about the specified phy	10.4.3.5
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	
13h	REPORT ROUTE INFORMATION	Return route table information for the specified phy	10.4.3.8
14h	REPORT PHY EVENT INFORMATION	Return phy event information for the specified phy	10.4.3.9
<u>15h</u>	<u>REPORT ZONE ROUTE TABLE</u>	<u>Return zone information for each specified phy</u>	
<u>16h</u>	<u>DISCOVER LIST</u>	<u>Return information about the specified list of phys</u>	
<u>17h</u>	<u>REPORT EXPANDER ROUTE TABLE</u>	<u>Return expander route table information</u>	
<u>18h</u> - 1Fh	Reserved for phy-based SMP input functions		
20h - 3Fh	Reserved for SMP input functions		
40h - 7Fh	Vendor specific		
80h	CONFIGURE GENERAL	Configure the device	10.4.3.10
81h	Reserved for a general SMP output function		
82h	WRITE GPIO REGISTER	See SFF-8485	

Table 1 — SMP functions (FUNCTION field)

Code	SMP function	Description	Reference
83h	<u>CONFIGURE ZONE PERMISSION</u>	<u>Change zone permission table information</u>	
84h	Reserved for general SMP output functions		
85h	ZONED BROADCAST	Transmit the specified BROADCAST on the expander ports in the specified zone group(s)	10.4.3.11
86h - 8Fh	Reserved for general SMP output functions		
90h	CONFIGURE ROUTE INFORMATION	Change route table information <u>for the specified phy</u>	10.4.3.11
91h	PHY CONTROL	Request actions by the specified phy	10.4.3.12
92h	PHY TEST FUNCTION	Request a test function by the specified phy	10.4.3.13
93h	CONFIGURE PHY EVENT INFORMATION	Configure phy event information for the specified phy	10.4.3.14
94h	<u>CONFIGURE PHY ZONE</u>	<u>Change phy entries within a zone route table</u>	
95h - 9Fh	Reserved for phy-based SMP output functions		
A0h - BFh	Reserved for SMP output functions		
C0h - FFh	Vendor specific		

T

10.4.3.x REPORT EXPANDER ROUTE TABLE function

The REPORT EXPANDER ROUTE TABLE function reports a list of SAS ADDRESSES of table routed phys. The list may be in any order. Self configuring expander devices shall support this function.

Table 2 — REPORT EXPANDER ROUTE TABLE request

Byte/Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (17h)								
2	Reserved								
3	REQUEST LENGTH (06h)								
4	Reserved								
7	Reserved								
8	MAXIMUM NUMBER OF DESCRIPTORS								
9	Reserved								
10	STARTING EXPANDER ROUTE INDEX								
15	Reserved								
16	Reserved								
18	Reserved								
19	STARTING PHY IDENTIFIER								
20	Reserved								
27	Reserved								
28	(MSB)	CRC							
31							(LSB)		

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

The FUNCTION field shall be set to 17h.

The REQUEST LENGTH field shall be set to 06h.

The MAXIMUM NUMBER OF DESCRIPTORS field specifies the maximum number of descriptors contained in the REPORT EXPANDER ROUTE TABLE descriptor list.

The STARTING EXPANDER ROUTE INDEX field specifies the first route table index to be returned in the REPORT EXPANDER ROUTE TABLE descriptor list.

The STARTING PHY IDENTIFIER field is the first phy identifier of the PHY IDENTIFIER BIT MAP returned in the REPORT EXPANDER ROUTE TABLE descriptor (see table 4). The value of this field shall be aligned to a boundary of 48 (e.g. 0, 48, 96) and not exceed the value of the NUMBER OF PHYS field reported in the REPORT GENERAL response (see 10.4.3.3).

Table 3 — REPORT EXPANDER ROUTE TABLE response

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (17h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH ((n-7) / 4)							
4	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5								
6	(MSB)	EXPANDER ROUTE TABLE CHANGE COUNT						(LSB)
7								
8	Reserved						CONFIGURING	Reserved
9	Reserved							
10								
11	NUMBER OF DESCRIPTORS							
12								
13	FIRST EXPANDER ROUTE INDEX							
14								
15	LAST EXPANDER ROUTE INDEX							
16								
18	Reserved							
19	STARTING PHY IDENTIFIER							
20								
31	Reserved							
REPORT EXPANDER ROUTE TABLE descriptor list								
32	REPORT EXPANDER ROUTE TABLE descriptor (first)							
47								
...	...							
n - 20	REPORT EXPANDER ROUTE TABLE descriptor (last)							
n - 4								
n - 3	(MSB)	CRC						(LSB)
n								

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

The FUNCTION field shall be set to 17h.

The RESPONSE LENGTH field shall be set to ((n - 7) / 4).

The EXPANDER CHANGE COUNT field is defined in 10.4.3.3.

The EXPANDER ROUTE TABLE CHANGE COUNT field counts the number of times the routing table has been modified by the self configuring expander device. Self configuring expander devices shall support this field. This field shall be set to at least 0001h at power on. If the self configuring expander device modified the expander route table since responding to a previous REPORT EXPANDER ROUTE TABLE request, it shall increment this field at least once from the value in the previous REPORT EXPANDER ROUTE TABLE response. 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 ROUTE TABLE 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.

The CONFIGURING bit is defined in 10.4.3.3.

The NUMBER OF DESCRIPTORS field specifies the number of REPORT EXPANDER ROUTE TABLE descriptors contained in the REPORT EXPANDER ROUTE TABLE descriptor list.

The FIRST EXPANDER ROUTE INDEX field specifies the route table index of the first REPORT EXPANDER ROUTE TABLE descriptor reported in the REPORT EXPANDER ROUTE TABLE descriptor list.

The LAST EXPANDER ROUTE INDEX field specifies the route table index of the last REPORT EXPANDER ROUTE TABLE descriptor reported in the REPORT EXPANDER ROUTE TABLE descriptor list. The value of this field plus one may be used in the REPORT EXPANDER ROUTE TABLE request as the STARTING EXPANDER ROUTE INDEX.

The STARTING PHY IDENTIFIER field specified in the request. Table 4.

Table 4 defines the REPORT EXPANDER ROUTE TABLE descriptor format.

Table 4 — REPORT EXPANDER ROUTE TABLE descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	(STARTING PHY IDENTIFIER + 47)							(STARTING PHY IDENTIFIER + 39)
	PHY IDENTIFIER BIT MAP							
5	(STARTING PHY IDENTIFIER + 7)							(STARTING PHY IDENTIFIER)
10	ZONE GROUP VALID	Reserved						
6	Reserved							
7	ZONE GROUP							
8	ROUTED SAS ADDRESS							
15								

The PHY IDENTIFIER BIT MAP field indicates the phy identifiers to which connection requests to the routed SAS ADDRESS are forwarded. This field is a bit map where each bit position set to one indicates a corresponding phy identifier (i.e., bit zero of byte 5 indicates the phy indicated by the STARTING PHY IDENTIFIER). A bit position set to zero indicates the ROUTED SAS ADDRESS is not associated to the corresponding phy identifier. A Bit positions that indicate phys beyond the value of the NUMBER OF PHYS field reported in the REPORT GENERAL response (see 10.4.3.3) shall be set to zero.

The ZONE GROUP VALID bit indicates if the value in ZONE GROUP field is valid or not. If the ZONE GROUP VALID bit is set to zero then the ZONE GROUP field shall be ignored. If the ZONE GROUP VALID bit is set to one then the value in the ZONE GROUP field is a valid zone group associated with the ROUTED SAS ADDRESS.

The ZONE GROUP field is defined in 4.8.3.1.

Editor's Note 1: Add to zone group "This field shall be set to zero if ZONING ENABLED bit in the REPORT GENERAL response is set to zero". see proposal 06-213. Expand definition of zone group to include address resolve.

The ROUTED SAS ADDRESS field contains the SAS address in the expander route entry (see 4.6.7.3).