

To: T10 Technical Committee  
From: Steve Johnson, Brad Besmer, LSI Logic (sjohnson@lsil.com)  
Date: 9 January, 2006  
Subject: 06-037r0 SAS-2 SMP Lists

### **Revision history**

Revision 0 (8 January, 2006) First revision

### **Related documents**

sas2r01 - Serial Attached SCSI 2 revision 0

### **Overview**

#### **SMP List**

SMP DISCOVER, REPORT PHY ERROR LOG, REPORT PHY SATA, REPORT ROUTE INFORMATION CONFIGURE ROUTE INFORMATION, PHY CONTROL, PHY TEST, and other newly proposed SMP functions require a separate SMP request and response per PHY identifier and ROUTE entry. In typical SAS-1.1 topologies the overhead of discovering and configuring the topology can become a significant hindrance to active I/O and fail-over scenarios. Zoning, self discovery, supervisor elections, table to table links, all using 1.1 discovery and routing procedures adds a "boat load" of overhead to the SMP storm. The discussion (or requirement) of very large SAS-2 topologies containing 1000's of end devices, dozens of self configuring expanders necessitates the need to dramatically improved the discovery and configuration.

Part of the solution is to reduce the number of SMP request and responses by combining multiple highly used SMP operations into a single request or response.

Combining SMP lists with an elected supervisor based configuration scheme will dramatically improve the discovery and configuration process for SAS-2 solutions.

#### **Expander Table Routing**

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. 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) Specific table sizes cannot be determined requiring a trial and error process of programming.
- 4) Entries are required to be zeroed by initiators.

Another way to handle SAS routing is to view the routing tables as a single table for the entire expander, where there is a list of SAS Addresses and PHYs that the addresses are routed to. Some advantages of this method are:

- 1) The table does not have to be in any specific order.
- 2) Add and remove of SAS ADDRESSES is independent of order.
- 3) Allows for partial discoveries where an initiator does not have to look at all the PHYs to talk to down stream devices (Important for new supervisor based configuration).
- 4) Table does not have to be rebuilt or reorder when expanders are added or removed.
- 5) No zeroing of table entries.
- 6) Matches most hardware implementations.
- 7) The exact table size can be reported.

**Suggested changes**

Add the following new SMP functions to section 10.4.3.x SMP functions of the SAS-2.

- 1) DISCOVER LIST
- 2) PROGRAM PHY ROUTED SAS ADDRESSES
- 3) REPORT ROUTED SAS ADDRESSES

The SMP DISCOVER LIST function returns a specified list of SMP DISCOVER responses for the specified range of phys. In addition it can optionally return a list of attached expanders. The list of attached expanders is used during the supervisor election process to expedite the work required to elect a supervisor.

The PROGRAM PHY ROUTED SAS ADDRESSES function is being proposed to provide a new mechanism to program route tables that uses lists of SAS ADDRESSES and PHYs to improve efficiency. It provides for a list of addresses to be assigned or removed to a given list of PHYs.

The REPORT PHY ROUTED SAS ADDRESSES function is being proposed to report the current routing table.

---

---

Editor's Note 1: Should PROGRAM and REPORT ROUTED SAS ADDRESSES be a separate proposal since it is a large change? What is the group preference?

---

---

table 191.

**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 - 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	10.4.3.8
<a href="#">14h</a>	<a href="#">DISCOVER LIST</a>	<a href="#">Return information about the specified list of phys</a>	<a href="#">10.4.3.x</a>
15h - 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	10.4.3.9
81h	Reserved for a general SMP output function		
82h	WRITE GPIO REGISTER	See SFF-8485	
83h - 8Fh	Reserved for general SMP output functions		
90h	CONFIGURE ROUTE INFORMATION	Change route table information	10.4.3.10
91h	PHY CONTROL	Request actions by the specified phy	10.4.3.11
92h	PHY TEST FUNCTION	Request a test function by the specified phy	10.4.3.12
<a href="#">93h</a>	<a href="#">PROGRAM PHY ROUTED SAS ADDRESSES</a>	<a href="#">Assign and delete routed SAS ADDRESSES to and from PHY route tables</a>	<a href="#">10.4.3.x</a>
<a href="#">94h</a>	<a href="#">REPORT PHY ROUTED SAS ADDRESSES</a>	<a href="#">Report the list of routed SAS ADDRESSES for the given set of PHY route tables</a>	<a href="#">10.4.3.x</a>
95h - 9Fh	Reserved for phy-based SMP output functions		
A0h - BFh	Reserved for SMP output functions		
C0h - FFh	Vendor specific		

**10.4.3.x DISCOVER LIST function**

The DISCOVER LIST function returns a list of DISCOVER SMP response data. This SMP function should be implemented by all SMP target ports.

Table 2 defines the request format.

**Table 2 — DISCOVER LIST request**

Byte/Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (17h)								
2	Reserved								
3	REQUEST LENGTH (02h)								
4	Reserved								
7	Reserved								
8	STARTING PHY IDENTIFIER								
9	NUMBER OF PHY IDENTIFIERS								
10	Reserved							REPORT SMP TARGETS ONLY	
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.

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

The NUMBER OF PHY IDENTIFIER field specifies the number of DISCOVER responses (see 4.2.7) in the list.

The REPORT SMP TARGETS ONLY bit set to one specifies that only DISCOVER response data with the ATTACHED SMP TARGET PORT bit set to one shall be returned.

The CRC field is defined in 10.4.3.1.

**Table 3 — DISCOVER LIST response (part 1 of 2)**

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (17h)							
2	FUNCTION RESULT							

**Table 3 — DISCOVER LIST response (part 2 of 2)**

Byte/Bit	7	6	5	4	3	2	1	0	
3	RESPONSE LENGTH ((n-h)/4)								
4	Reserved								
8	Reserved								
9	STARTING PHY IDENTIFIER								
10	NUMBER OF PHY IDENTIFIERS								
11	Reserved							REPORT SMP TARGETS ONLY	
12	Reserved								
15	Reserved								
16	FIRST DISCOVER RESPONSE								
m									
...	...								
y	LAST DISCOVER RESPONSE								
n - 4									
n - 3	(MSB)	CRC							
n								(LSB)	

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

The FUNCTION field shall be set to 17h.

The FUNCTION RESULT field is defined in 10.4.3.x.

The RESPONSE LENGTH field shall be set to 0000?h.

The STARTING PHY IDENTIFIER field indicates the first phy of the list for which DISCOVER response data is being returned.

The NUMBER OF PHY IDENTIFIERS indicates the number of DISCOVER responses being returned. The PHY IDENTIFIERS shall be in consecutive order if the REPORT SMP TARGET ONLY bit is set to zero, otherwise only phys with ATTACHED SMP TARGET PORT bit set to one shall be returned.

**10.4.3.x PROGRAM PHY ROUTED SAS ADDRESSES function**

The PROGRAM PHY ROUTED SAS ADDRESSES function assigns or removes a list of SAS ADDRESSES to one or more expander table route phys. The list can be in any order. Other SMP target ports shall not support this SMP function.

Table 4 defines the request format.

**Table 4 — PROGRAM PHY ROUTED SAS ADDRESSES request**

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (93h)								
2	Reserved								
3	REQUEST LENGTH $((n - 7h) / 4)$								
4	Reserved								
7									
8	NUMBER ROUTED SAS ADDRESSES								
9									
10	Reserved							REMOVE ROUTED ADDRESSES	
11	NUMBER OF PHY IDENTIFIER								
12	FIRST PHY IDENTIFIER								
...	...								
27									
28	Reserved								
31									
32	FIRST ROUTED SAS ADDRESS								
...	...								
n - 11	LAST ROUTED SAS ADDRESS								
n - 4									
n - 3	(MSB)	CRC							
n								(LSB)	

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

The FUNCTION field shall be set to 93h.

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

The NUMBER ROUTED SAS ADDRESSES field specifies the number of routed SAS ADDRESSES.

The REMOVE ROUTED ADDRESSES bit when set to one specifies that the SAS ADDRESSES shall be removed from the route table. otherwise, the list of ROUTED SAS ADDRESSES shall be added to the route table.

The NUMBER PHY IDENTIFIERS field specifies the number of PHY IDENTIFIERS.

---



---

Editor's Note 2: Should the PHY IDENTIFIERS be described as a bit field from 0 - 255 instead of a list of 16? This would specify all possible PHYs. However, 16 covers most uses, since we really only need to describe the phys that are part of a wide port.

---



---

Table 5 defines the response format.

**Table 5 — PROGRAM ROUTED SAS ADDRESSES response**

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (93h)								
2	FUNCTION RESULT								
3	RESPONSE LENGTH (00h)								
4	(MSB)	CRC							
7								(LSB)	

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

The FUNCTION field shall be set to 93h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 00h.

The CRC field is defined in 10.4.3.2.

#### 10.4.3.x REPORT PHY ROUTED SAS ADDRESSES function

The REPORT PHY ROUTED SAS ADDRESSES function assigns or removes a list of SAS ADDRESSES to one or more expander table route phys. The list can be in any order. Other SMP target ports shall not support this SMP function.

---



---

Editor's Note 3: When the entire list does fit in one SMP transaction how do we indicate where to start from. One way is to have an index and a table write count to know where to start in the table and to know that the table has changed since the last time an initiator has read it.

---



---