

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com) and Tim Symons, PMC-Sierra (tim_symons@pmc-sierra.com)
Date: 29 March 2006
Subject: 06-176r0 SAS-2 Add phy zone information to SMP DISCOVER

Revision history

Revision 0 (29 March 2006) First revision

Related documents

sas2r03 - Serial Attached SCSI - 2 (SAS-2) revision 3
06-019r5 SAS-2 zoning (Tim Symons and Heng Liao, PMC-Sierra)(the source for 06-176)
06-166 SAS-2 Restrict access to SMP ZONED BROADCAST function (Rob Elliott, HP)(this defines zone group 3)

Overview

SAS-2 revision 3 defines zone phy information for expander phys, but does not define a way to retrieve that information.

The fields should be included in the SMP DISCOVER response.

Also (as previously proposed in 06-019r5) an IGNORE ZONE GROUP bit is added to the DISCOVER request. If the bit is 0, only the phys to which the SMP initiator port has access are reported as attached at all. If the bit is 1, all phys are reported. This ensures that legacy software performing the discover process doesn't try to establish connections to SAS addresses to which it does not have permission and that will return OPEN_REJECT (ZONE VIOLATION). That should be handled correctly, but results in lots of worthless SMP traffic.

If only that were done, any SMP initiator port that chooses to set the IGNORE ZONE GROUP bit to one could learn information about zone groups to which it has no access rights. To close this hole, only SMP initiator ports with access to zone group 2 (in sas2r03) or zone group 3 (defined by 06-166) are allowed to use that bit. Any SMP initiator port trusted to claim where a BROADCAST came from by sending ZONED BROADCAST is thus allowed to discover information about the constituents of the other zones (provided it has access to the zone group of the expander device itself).

Editor's Note 1: ZG 2 provides strong write access. ZG 3 provides weak write access. Perhaps ZG 4 is appropriate for read-only access, rather than overloading 3? How many access levels are appropriate? Is the range of reserved ZGs (3 to 7) large enough?

Editor's Note 2: This proposal allows either ZG 2 or ZG 3 to provide access. Is there any reason why access to ZG 2 should not imply access to ZG 3 as well (for this or any other similar function)?

Suggested changes

10.4.3.2 SMP function response frame format

...

The FUNCTION RESULT field is defined in table 1.

Table 1 — FUNCTION RESULT field

Code	Name	SMP function(s)	Description
04h	SMP ZONE VIOLATION	TBD	The SMP target port supports the function, but the application zone permission bit is set to zero (e.g., the ZP[s, 2] bit is set to zero).
16h	PHY VACANT	DISCOVER, REPORT PHY ERROR LOG, REPORT PHY SATA, REPORT ROUTE INFORMATION, REPORT PHY EVENT INFORMATION, CONFIGURE ROUTE INFORMATION, PHY CONTROL, CONFIGURE PHY EVENT INFORMATION	The SMP target port processing the SMP request frame does not have access to the phy (e.g., because of zoning or vendor-specific reasons), although the value is within the range of zero to the value of the NUMBER OF PHYs field reported in the REPORT GENERAL function. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
...

Editor's Note 3: 06-166 adds ZONED BROADCAST to the list of functions for code 04h and changes ZG 2 to 3 in the e.g.

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								
7	Reserved								
8	Reserved							IGNORE ZONE GROUP	
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.

If the SMP initiator port has a zone group with access to zone group 2 or zone group 3 and the IGNORE ZONE GROUP bit is set to one, the SMP target port shall return information about the specified phy regardless of the zone permission table. If the SMP initiator port has a zone group without access to zone group 2 or zone group 3 and the IGNORE ZONE GROUP bit is set to one, the SMP target port shall return an SMP function response of PHY VACANT.

An IGNORE ZONE GROUP bit set to zero specifies that the SMP target port shall:

- a) if the SMP initiator port has access to the specified phy based on the zone permission table, return the requested information; and
- b) if the SMP initiator port does not have access to the specified phy, return an SMP function response of PHY VACANT.

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			
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							

Table 3 — DISCOVER response (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0	
50	Vendor specific								
51									
52	ATTACHED DEVICE NAME								
59									
60	Reserved			ZONE ADDRESS RESOLVED	ZONE GROUP PERSISTENT	ZONE PARTICIPATING	ZONE FIELDS VALID		
61	Reserved								
62	Reserved								
63	ZONE GROUP								
6064	(MSB)	CRC							
6367							(LSB)		

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

The FUNCTION field shall be set to 10h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to ~~0Eh~~0Fh. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h specifies that there are 12 dwords before the CRC field.

...

[The ZONE ADDRESS RESOLVED bit contains the value of the ZONE ADDRESS RESOLVED bit in the zone phy information \(see 4.8.3.1\).](#)

[The ZONE GROUP PERSISTENT bit contains the value of the ZONE GROUP PERSISTENT bit in the zone phy information \(see 4.8.3.1\).](#)

[The ZONE PARTICIPATING bit contains the value of the ZONE PARTICIPATING bit in the zone phy information \(see 4.8.3.1\).](#)

[A ZONE FIELDS VALID bit set to one indicates that the ZONE ADDRESS RESOLVED bit, the ZONE GROUP PERSISTENT bit, the ZONE PARTICIPATING bit, and the ZONE GROUP field are valid. A ZONE FIELDS VALID bit set to zero indicates that those fields are reserved.](#)

Editor's Note 4: We are not sure about the need for a ZONE FIELDS VALID bit. We don't want non-zoning software to assume zone group 0 applies if all these bits are 0. Relying on ZONING ENABLED bit in the REPORT GENERAL response would eliminate confusion, but that's a bit far away from here. If zoning status could ever change between when the REPORT GENERAL and DISCOVER functions are requested, it'd be better to keep a valid bit here.

[The ZONE GROUP field contains the value of the ZONE GROUP field in the zone phy information \(see 4.8.3.1\). Zone group values between 128 and 255, inclusive, are reserved.](#)

The CRC field is defined in 10.4.3.2.