

To: T10 Technical Committee
 From: Steve Johnson LSI Logic (steve.johnson@lsil.com), Brad Besmer LSI Logic
 Date: 19 April, 2006
 Subject: 06-037r2 SAS-2 SMP Lists (DISCOVER LIST)

Revision history

Revision 0 (9 January, 2006) First revision
 Revision 1 (7 April, 2006) Revised from January 9 and CC meeting discussions.
 Revision 2 (19 April, 2006).

Related documents

sas2r03 - Serial Attached SCSI 2 revision 3

Revision Overview r1 to r2

Updated the zoning and configuring bits.
 Added the routing attribute field to the descriptor.
 Updated text.

Overview

SMP DISCOVER, requires a separate SMP request and response per PHY identifier. 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. Self discovery, zoning, supervisor elections, table to table links, all using 1.1 discovery and routing procedures would add a “boat load” more of SMPs 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 and initiators necessitates the need to dramatically improved the discovery and configuration mechanisms.

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

The descriptor format provides the necessary information for each phy (up to 48 phys) in a single SMP request and response for a self configuring expander to perform self configuration of routing tables, zoning route tables, and zoning election for the requested expander.

Suggested changes

Add new SMP DISCOVER LIST function to section 10.4.3.x SMP functions of SAS-2. The DISCOVER LIST provides all the necessary data for a self configuring expander to program it’s zoning and route tables along with supervisor election information.

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> | |

Table 1 — SMP functions (FUNCTION field)

| Code | SMP function | Description | Reference |
|------------|---|---|-----------|
| 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 | 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> | Return zone information for each specified phy | |
| <u>16h</u> | <u>DISCOVER LIST</u> | <u>Return information about the specified list of phys</u> | |
| 17h - 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 | |
| <u>83h</u> | <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 | 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 |
| <u>94h</u> | <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 | | |

10.4.3.x DISCOVER LIST function

The DISCOVER LIST function returns a list of phy descriptors. This SMP function shall 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 (16h) | | | | | | | |
| 2 | Reserved | | | | | | | |
| 3 | REQUEST LENGTH (02h) | | | | | | | |
| 4 | Reserved | | | | | | | |
| 7 | Reserved | | | | | | | |
| 8 | STARTING PHY IDENTIFIER | | | | | | | |
| 9 | NUMBER OF DESCRIPTORS | | | | | | | |
| 10 | Reserved | | | | PHY FILTER | | | |
| 11 | Reserved | | | | DESCRIPTOR TYPE | | | |
| 12 | Reserved | | | | | | | |
| 15 | Reserved | | | | | | | |
| 16 | Vendor specific | | | | | | | |
| 29 | Vendor specific | | | | | | | |
| 30 | (MSB) | CRC | | | | | | (LSB) |
| 31 | (LSB) | | | | | | | |

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

The FUNCTION field shall be set to 16h.

The REQUEST LENGTH field shall be set to 02h.

The STARTING PHY IDENTIFIER field specifies the phy identifier of the first phy in the list of descriptors being requested.

The NUMBER OF DESCRIPTORS field specifies the number of descriptors requested starting with the value specified by the STARTING PHY IDENTIFIER field. If (STARTING PHY IDENTIFIER + NUMBER OF DESCRIPTORS) exceeds the NUMBER OF PHYS field reported in the REPORT GENERAL response (see 10.4.3.3) then

Editor's Note 1: [Help, How to state if STARTING PHY IDENTIFIER + NUMBER OF DESCRIPTORS leads to an invalid phy or exceeds a 1K response buffer. Also, does this belong here or on the returned value in the response?](#)

The PHY FILTER field set to 0h indicates the SMP target port shall return in the list of descriptors only phy identifiers with the ATTACHED SMP TARGET PORT (see 10.4.3.5) bit set to one. A PHY FILTER field set to 1h

indicates the SMP target port shall return in the list of descriptors only phy identifiers with the ATTACHED DEVICE TYPE (see 10.4.3.5) field set to a value other than zero. All other values are reserved.

.A DESCRIPTOR TYPE field set to 0h indicates the descriptor describes a DISCOVER response not including the CRC field (see 10.4.3.5). A DESCRIPTOR TYPE field set to 1h indicates the descriptor describes a Discover List descriptor (see Table 4 —). A DESCRIPTOR TYPE field set to Fh indicates the descriptor format is vendor specific. All other values are reserved.

The CRC field is defined in 10.4.3.2.

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 (16h) | | | | | | | |
| 2 | FUNCTION RESULT | | | | | | | |
| 3 | RESPONSE LENGTH $((n-7)/4)$ | | | | | | | |
| 4 | Reserved | | | | | | | |
| 7 | Reserved | | | | | | | |
| 8 | STARTING PHY IDENTIFIER | | | | | | | |
| 9 | NUMBER OF DESCRIPTORS | | | | | | | |
| 10 | Reserved | | | | PHY FILTER | | | |
| 11 | Reserved | | | | DESCRIPTOR TYPE | | | |
| 12 | DESCRIPTOR LENGTH | | | | | | | |
| 13 | Reserved | | | | | | | |
| 15 | Reserved | | | | | | | |
| 16 | (MSB) | | | | | | | |
| 17 | EXPANDER CHANGE COUNT | | | | | | | |
| 18 | ZONE DEVICE | ZONE ADDRESS RESOLVE DEVICE | Reserved | | | CONFIGURING | CONFIGURABLE ROUTE TABLE | |
| 19 | Reserved | | | | | | | |
| 22 | Reserved | | | | | | | |
| 23 | ACTIVE ZONE SUPERVISOR PRIORITY | | | | ZONE SUPERVISOR PRIORITY | | | |
| 24 | Reserved | | | | | | | |
| 31 | ACTIVE ZONE SUPERVISOR SAS ADDRESS | | | | | | | |

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

| Byte\Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
|----------------------------|------------------|-----|---|---|---|---|---|---|--|
| 32 | Vendor specific | | | | | | | | |
| 47 | | | | | | | | | |
| List of descriptors | | | | | | | | | |
| 48 | FIRST DESCRIPTOR | | | | | | | | |
| m | | | | | | | | | |
| ... | ... | | | | | | | | |
| y | LAST DESCRIPTOR | | | | | | | | |
| n - 4 | | | | | | | | | |
| n - 3 | (MSB) | CRC | | | | | | | |
| n | | | | | | | | | |

Editor's Note 2: What about PHY_VACANT status? How do we communicate that some phys will have this status returned? Should we add a status to descriptor? What if the format requested is DISCOVER response?

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

The FUNCTION field shall be set to 16h.

The FUNCTION RESULT field is defined in 10.4.3.2.

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

The STARTING PHY IDENTIFIER field specifies the phy identifier of the first phy in the list of descriptors being returned. The PHY IDENTIFIERS shall be returned in order from low to high.

Editor's Note 3: Need verbiage The starting phy may not be the starting phy depending on the filter.

The NUMBER OF DESCRIPTORS field specifies the number of descriptors returned in the list.

The PHY FILTER field is defined in the DISCOVER LIST request.

The DESCRIPTOR TYPE field is defined in the DISCOVER LIST request.

The DESCRIPTOR LENGTH field is determined by the descriptor specified in the DESCRIPTOR TYPE field.

The EXPANDER CHANGE COUNT field is defined in 10.4.3.3

The CONFIGURABLE ROUTE TABLE field is defined in 10.4.3.3

The CONFIGURING field is defined in 10.4.3.3

The ZONE ADDRESS RESOLVE DEVICE field is defined in 10.4.3.3

The ZONE DEVICE field is defined in 10.4.3.3

The ZONE SUPERVISOR PRIORITY field is defined in 10.4.3.3

The ACTIVE ZONE SUPERVISOR PRIORITY field is defined in 10.4.3.3

The ACTIVE ZONE SUPERVISOR SAS ADDRESS field is defined in 10.4.3.3

The CRC field is defined in 10.4.3.2

Table 4 defines the descriptor format.

Table 4 — DISCOVER LIST descriptor

| Byte\Bit | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|-----------------------------|----------------------|---|---|-------------------------------|------------------------|------------------------|----------------------|
| 0 | PHY IDENTIFIER | | | | | | | |
| 1 | Reserved | ATTACHED DEVICE TYPE | | | Reserved | | | |
| 2 | Reserved | | | | NEGOTIATED PHYSICAL LINK RATE | | | |
| 3 | Reserved | | | | ATTACHED SSP INITIATOR | ATTACHED STP INITIATOR | ATTACHED SMP INITIATOR | ATTACHED SATA HOST |
| 4 | ATTACHED SATA PORT SELECTOR | Reserved | | | ATTACHED SSP TARGET | ATTACHED STP TARGET | ATTACHED SMP TARGET | ATTACHED SATA DEVICE |
| 5 | Reserved | | | | | | | |
| 6 | VIRTUAL PHY | Reserved | | | ROUTING ATTRIBUTE | | | |
| 7 | ATTACHED SAS ADDRESS | | | | | | | |
| 15 | ATTACHED SAS ADDRESS | | | | | | | |
| 16 | ATTACHED PHY IDENTIFIER | | | | | | | |
| 17 | PHY CHANGE COUNT | | | | | | | |
| 18 | Reserved | | | | ZONE ADDRESS RESOLVED | ZONE GROUP PERSISTENT | ZONE PARTICIPATING | ZONING ENABLED |
| 19 | ZONE GROUP | | | | | | | |
| 20 | Reserved | | | | | | | |

The STARTING PHY IDENTIFIER field is defined in 10.4.3.5

The ATTACHED DEVICE TYPE field is defined in 10.4.3.5

The NEGOTIATED PHYSICAL LINK RATE field is defined in 10.4.3.5

The ATTACHED SATA HOST field is defined in 10.4.3.5

The ATTACHED SMP INITIATOR field is defined in 10.4.3.5

The ATTACHED STP INITIATOR field is defined in 10.4.3.5

The ATTACHED SSP INITIATOR field is defined in 10.4.3.5

The ATTACHED SATA DEVICE field is defined in 10.4.3.5

The ATTACHED SMP TARGET field is defined in 10.4.3.5

- The ATTACHED STP TARGET field is defined in 10.4.3.5
- The ATTACHED SSP TARGET field is defined in 10.4.3.5
- The ATTACHED SATA PORT SELECTOR field is defined in 10.4.3.5
- The ROUTING ATTRIBUTE field is defined in 10.4.3.5
- The VIRTUAL PHY field is defined in 10.4.3.5
- The ATTACHED SAS ADDRESS field is defined in 10.4.3.5
- The ATTACHED PHY IDENTIFIER field is defined in 10.4.3.5
- The PHY CHANGE COUNT field is defined in 10.4.3.5
- The ZONE ENABLED field is defined in 10.4.3.5
- The ZONE PARTICIPATING field is defined in 10.4.3.5
- The ZONE GROUP PERSISTENT field is defined in 10.4.3.5
- The ZONE ADDRESS RESOLVE field is defined in 10.4.3.5
- The ZONE GROUP field is defined in 10.4.3.5