

Date: July 10, 2006

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: Multiple Broadcasts on Reset

Overview

SAS requires that two BROADCAST primitives be issued if a link reset or a hard reset occurs on a link. The first occurs at the start of the reset sequence and the second occurs on completion of the reset sequence.

This creates problems in at least two cases when JBOD SAS devices are connected to a RAID implementation:

- a) when there are multiple RAIDed initiators attached on the SAS domain; and
- b) anytime discovery is initiated.

The problem is that anytime a RAID initiator does a discovery (e.g., as the result of a BROADCAST) and finds out a drive has disappeared, it will mark that drive as missing thereby causing the RAID to become exposed. If a drive reappears in a subsequent discovery (e.g., as the result of receiving a second BROADCAST) a rebuild is initiated.

In the case where a RAID initiator issues a Hard Reset as part of a recovery process, that RAID initiator knows it will receive two BROADCASTs so it could ignore the first one and only do a discovery on the second BROADCAST. However, there is no way to tell the other initiators to ignore a BROADCAST. As a result the first BROADCAST would cause any RAID initiator that does not know about the reset to start a discovery and it may find the drive being reset is missing, causing it to mark the RAID as exposed. When that RAID initiator sees the second BROADCAST it will start another discovery and will find the same drive is there but it's already too late and the damage has been done.

[This proposal recommends a new NEGOTIATED PHYSICAL LINK RATE field value that would indicate if a phy is the process of responding to a PHY CONTROL function, LINK RESET phy operation, or HARD RESET phy operation.](#)

~~This proposal recommends a new function response code to be returned for a SMP_DISCOVER function and a SMP_DISCOVER_LIST function that would indicate the phy is busy and should be checked later to determine what is attached.~~

~~There is also a RETRY_DELAY_TIMER field that contains a recommended time to wait before retrying the SMP_DISCOVER function or the SMP_DISCOVER_LIST function.~~

The recommended changes needed to incorporate this change into SAS-2 are as indicated below.

SAS-2 changes

4.7 Discover process

4.7.1 Discover process overview

Management application clients direct an SMP initiator port to request SMP functions from an SMP target port. Management application clients are located in every SAS initiator device and every self-configuring expander device. The discover process is the process of:

- a) discovering all the SAS devices and expander devices in the SAS domain (i.e., determining their device types, SAS addresses, and supported protocols); and
- b) configuring routing tables in the expander devices as needed.

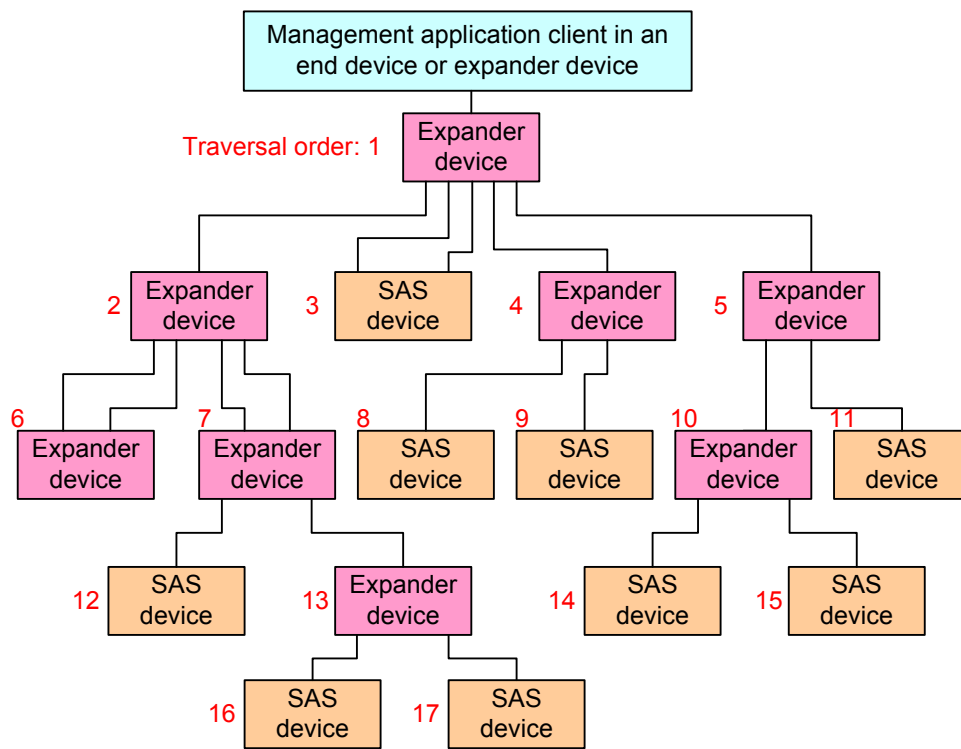
A management application client performing the discover process shall perform a level-order (i.e., breadth-first) traversal of the SAS domain. The order of traversal shall be to discover:

- 1) the device(s) to which the device containing the management application client is attached;
- 2) if an attached device is an expander device, every device attached to that expander device; and

- 3) for each expander device found, every device attached to that expander device.

This order is repeated until all expander devices have been traversed. If the management application client is inside an expander device, then the process shall be repeated on each expander port.

Figure 1 shows an example of level-order traversal.



Note: Assume that the phy with the lowest phy identifier in each expander device is on the top right, and the remaining phys have increasing phy identifiers assigned in a counter-clockwise direction

Figure 1 — Level-order traversal example

The management application client determines whether an expander device or SAS device is attached at each point in the traversal. For the first device (i.e., the device directly attached), this is determined from the DEVICE TYPE field in the IDENTIFY address frame information received by the phy that the management application client is using. For other devices (i.e., devices not directly attached), this is determined from ATTACHED DEVICE TYPE field the SMP DISCOVER response (see 10.4.3.5) and the SMP DISCOVER LIST response (see 10.4.3.11.3).

If an expander device is attached, the management application client shall use the SMP REPORT GENERAL function (see 10.4.3.3) to determine how many phys are in the expander device and then use the SMP DISCOVER function (see 10.4.3.5) or the SMP DISCOVER LIST function (see 10.4.3.11) to determine what is attached to each expander phy (e.g., the device type, SAS address, and supported protocol(s)).

If the expander device's CONFIGURABLE ROUTE TABLE bit is set to one in the SMP REPORT GENERAL response, its own management application client shall configure its own expander route table as described in 4.7.3 and 4.7.4.

If a self-configuring expander device's CONFIGURING bit is set to one in the SMP REPORT GENERAL response, connection requests for destination ports two or more levels beyond the self-configuring expander device may result in responses such as OPEN_REJECT (NO DESTINATION) (see 7.2.5.12). If a management application client (e.g., in an end device) detects a self-configuring expander device with the CONFIGURING bit set to one, the SAS port shall stop the I_T Nexus Loss timer for I_T nexuses involving such destination ports as long as the CONFIGURING bit is set to one (see 8.2.2.1).

If a SAS device is attached, the discover process is not required to obtain any more information about the SAS device. Additional discovery software may access that SAS device, however:

- a) if the SAS device supports an SMP target port, the management application client may use SMP functions (e.g., REPORT GENERAL and REPORT MANUFACTURER INFORMATION) to determine additional information about the SAS target device;
- b) if the SAS device supports an SSP target port, a SCSI application client may transmit SCSI commands (e.g., INQUIRY and REPORT LUNS) to determine additional information about the SCSI target device; and
- c) if the end device supports an STP target port, an ATA application client may transmit ATA commands (e.g., IDENTIFY DEVICE and IDENTIFY PACKET DEVICE) to determine additional information about the ATA device.

The result of the discover process is that the management application client has the necessary information (e.g., the device type, SAS address, and supported protocol(s)) to communicate with each SAS device and expander device in the SAS domain and each configurable expander device is configured with the necessary expander route entries to allow routing of connection requests through the SAS domain.

~~If the discover process occurs and any phy within the target expander device is in the process of a link reset sequence resulting from a PHY CONTROL function, LINK RESET phy operation, or HARD RESET phy operation processes of link reset sequence, then any SMP DISCOVER function (see 10.4.3.5) or any SMP DISCOVER LIST function (see 10.4.3.11) should shall return the NEGOTIATED PHYSICAL LINK RATE field set to RESET_IN_PROGRESS (see table 7). be terminated with a function result of PHY BUSY with the RETRY-DELAY TIMER field set to a length of time that the management application client should wait before retrying the discover process.~~

Annex L contains an example implementation of how a management application client may perform the discover process.

The discover process may be aborted prior to completion and restarted if there is an indication that it may be based on incorrect information (e.g., reception of a Broadcast (Change) or a change in the EXPANDER CHANGE COUNT field returned in an SMP response).

10.4.3.2 SMP function response frame format

An SMP response frame is sent by an SMP target port in response to an SMP request frame. Table 1 defines the SMP response frame format.

Table 1 — SMP response frame format

Byte/Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION								
2	FUNCTION RESULT								
3	RESPONSE LENGTH ((n - 7) / 4)								
4	ADDITIONAL RESPONSE BYTES								
m									
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

The SMP FRAME TYPE field is included in each frame format defined in this clause, although that field is parsed by the SMP transport layer (see 9.4). The SMP FRAME TYPE field is set to 41h.

The FUNCTION field indicates the SMP function to which this frame is a response, and is defined in table 197 in 10.4.3.1.

The FUNCTION RESULT field is defined in table 2.

Table 2 — FUNCTION RESULT field (part 1 of 3)

Code	Name	SMP function(s)	Description
00h	SMP FUNCTION ACCEPTED	All	The SMP target port supports the SMP function. The ADDITIONAL RESPONSE BYTES field contains the requested information.
01h	UNKNOWN SMP FUNCTION	Unknown	The SMP target port does not support the requested SMP function. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
02h	SMP FUNCTION FAILED	All	The SMP target port supports the SMP function, but the requested SMP function failed. The ADDITIONAL RESPONSE BYTES may be present but shall be ignored.
03h	INVALID REQUEST FRAME LENGTH	All	The SMP target port supports the SMP function, but the SMP request frame length was invalid (i.e., did not match the frame size defined for the function). The ADDITIONAL RESPONSE BYTES may be present but shall be ignored.
04h	INVALID EXPANDER CHANGE COUNT	CONFIGURE GENERAL, CONFIGURE ROUTE INFORMATION, PHY CONTROL, PHY TEST FUNCTION, CONFIGURE PHY EVENT INFORMATION	The SMP target port supports the SMP function, but the EXPECTED EXPANDER CHANGE COUNT field does not match the current expander change count.
10h	PHY DOES NOT EXIST	DISCOVER, DISCOVER LIST, REPORT PHY ERROR LOG, REPORT PHY SATA, REPORT ROUTE INFORMATION, REPORT PHY EVENT INFORMATION, CONFIGURE ROUTE INFORMATION, PHY CONTROL, PHY TEST FUNCTION, CONFIGURE PHY EVENT INFORMATION	The phy specified by the PHY IDENTIFIER field or the STARTING PHY IDENTIFIER field in the SMP request frame does not exist (e.g., the value is not within the range of zero to the value of the NUMBER OF PHYs field reported in the SMP REPORT GENERAL response). The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.

Table 2 — FUNCTION RESULT field (part 2 of 3)

Code	Name	SMP function(s)	Description
11h	INDEX DOES NOT EXIST	REPORT ROUTE INFORMATION, CONFIGURE ROUTE INFORMATION	The phy specified by the PHY IDENTIFIER field in the SMP request frame does not have the table routing attribute (see 4.6.7.1), or the expander route index specified by the EXPANDER ROUTE INDEX field does not exist (i.e., the value is not in the range of 0000h to the value of the EXPANDER ROUTE INDEXES field in the SMP REPORT GENERAL response). The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
12h	PHY DOES NOT SUPPORT SATA	REPORT PHY SATA and PHY CONTROL (TRANSMIT SATA PORT SELECTION SIGNAL)	The phy specified by the PHY IDENTIFIER field in the SMP request frame is not part of an STP target port. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
13h	UNKNOWN PHY OPERATION	PHY CONTROL	The operation specified by the PHY OPERATION field in the SMP request frame is unknown. The SMP function had no affect. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
14h	UNKNOWN PHY TEST FUNCTION	PHY TEST FUNCTION	The operation specified by the PHY TEST FUNCTION field in the SMP request frame is unknown. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
15h	PHY TEST FUNCTION IN PROGRESS	PHY TEST FUNCTION	The specified phy is already performing a phy test function. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
16h	PHY VACANT	DISCOVER, DISCOVER LIST, 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 SMP REPORT GENERAL response. The ADDITIONAL RESPONSE BYTES field may be present but shall be ignored.
17h	PHY EVENT INFORMATION SOURCE NOT SUPPORTED	CONFIGURE PHY EVENT INFORMATION	The phy event information source specified by a PHY EVENT INFORMATION SOURCE field is not supported.

Table 2 — FUNCTION RESULT field (part 3 of 3)

Code	Name	SMP function(s)	Description
20h	SMP ZONE VIOLATION	CONFIGURE GENERAL, ZONED BROADCAST, PHY CONTROL, PHY TEST FUNCTION, CONFIGURE PHY EVENT INFORMATION	The SMP target port supports the function, but zoning is enabled and the SMP initiator port does not have access to a necessary zone group according to the zone permission table (see 4.8.3.2).
xxh	PHY BUSY	DISCOVER, DISCOVER LIST	The SMP target port processing the SMP request frame temporarily does not have access to a phy (e.g., because of a link reset sequence). The RETRY DELAY TIMER field may indicate how long the management application client should wait before restarting the discover process.
All others	Reserved		

The RESPONSE LENGTH field indicates the number of dwords that follow, not including the CRC field. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h sometimes indicates a non-zero number of dwords; this is defined in the function description.

The ADDITIONAL RESPONSE BYTES field definition depends on the SMP function requested. The maximum size of the ADDITIONAL RESPONSE BYTES field is 1 024 bytes, making the maximum size of the frame 1 032 bytes (i.e., 1 024 bytes of data + 4 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the ADDITIONAL RESPONSE BYTES field so the CRC field is aligned on a four byte boundary. The contents of the fill bytes are vendor specific.

The CRC field is included in each response frame format defined in this clause, although that field is defined by the SMP transport layer (see 9.4.1) and parsed by the SMP link layer (see 7.18).

10.4.3.5 DISCOVER function

The DISCOVER function returns information about 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.

NOTE 1 - The DISCOVER LIST function (see 10.4.3.11) returns information about one or more phys.

Table 3 defines the request format.

Table 3 — 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	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5	Reserved							
6	Reserved							
7	Reserved							
8	Reserved							IGNORE ZONE GROUP
9	PHY IDENTIFIER							
10	Reserved							
11	Reserved							
12	(MSB)	CRC						(LSB)
15	Reserved							

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.

An IGNORE ZONE GROUP bit set to one specifies that the SMP target port shall return information about the specified phy (i.e., the phy specified by the PHY IDENTIFIER field) regardless of the zone permission table.

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 a function result of PHY VACANT in the response frame.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for which the information is being requested.

The CRC field is defined in 10.4.3.1.

Table 4 defines the response format.

Table 4 — 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 (0Fh)							
4	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5								
6	RETRY DELAY TIMER field Reserved							
7								
8	Reserved							
9	PHY IDENTIFIER							
10	Reserved							
11								
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								
24	ATTACHED SAS ADDRESS							
31								
32	ATTACHED PHY IDENTIFIER							
33	Reserved							ATTACHED BREAK_REPLY CAPABLE
34	Reserved							
39								
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							

Table 4 — DISCOVER response (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0	
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								
50	Vendor specific								
51	Vendor specific								
52	ATTACHED DEVICE NAME								
59	ATTACHED DEVICE NAME								
60	Reserved				ZONE ADDRESS RESOLVED	ZONE GROUP PERSISTENT	ZONE PARTICIPATING	ZONING ENABLED	
61	Reserved								
62	Reserved								
63	ZONE GROUP								
64	(MSB)	CRC							
67							(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 0Fh. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h indicates that there are 12 dwords before the CRC field.

The EXPANDER CHANGE COUNT field is defined in the SMP REPORT GENERAL response (see 10.4.3.3).

The RETRY DELAY TIMER field is defined in table 5.

Table 5 — ~~RETRY DELAY TIMER field~~

Function result	Retry delay timer code	Description
<u>PHY-BUSY</u>	<u>0000h</u>	<u>No addition information</u>
	<u>0001h–FFEFh</u>	<u>The number of 100 milliseconds increments the management application client should wait before retrying the discover process.</u>
	<u>FFF0h–FFFDh</u>	<u>Reserved</u>
	<u>FFFEh</u>	<u>The management application client should handle this as it would a PHY-VACANT function result.</u>
	<u>FFFFh</u>	<u>Reserved</u>
<u>All others</u>	<u>0000h–FFFFh</u>	<u>Reserved</u>

The PHY IDENTIFIER field indicates the phy for which physical configuration link information is being returned.

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence and is defined in table 6.

Table 6 — ATTACHED DEVICE TYPE field

Code	Description
000b	No device attached
001b	End device
010b	Edge expander device
011b	Fanout expander device
All others	Reserved

The ATTACHED DEVICE TYPE field shall only be set to a value other than 000b after:

- a) the identification sequence is complete if a SAS device or expander device is attached; or
- b) the initial Register - Device to Host FIS has been received if a SATA phy is attached.

The NEGOTIATED PHYSICAL LINK RATE field is defined in table 7 and indicates the physical link rate negotiated during the link reset sequence and other conditions relating to the phy. The negotiated physical link rate may be less than the programmed minimum physical link rate or greater than the programmed maximum physical link rate if the programmed physical link rates have been changed since the last link reset sequence.

Table 7 — NEGOTIATED PHYSICAL LINK RATE field

Code	Name	Description
0h	UNKNOWN	Phy is enabled; unknown physical link rate. ^a
1h	DISABLED	Phy is disabled.
2h	PHY_RESET_PROBLEM	Phy is enabled; the phy obtained dword synchronization for at least one physical link rate during the SAS speed negotiation sequence (see 6.7.4.2), but the SAS speed negotiation sequence failed (i.e., the last speed negotiation window, using a physical link rate expected to succeed, failed). These failures may be logged in the SMP REPORT PHY ERROR LOG function (see 10.4.3.6), the SMP REPORT PHY EVENT INFORMATION function (see 10.4.3.9), and/or the Protocol-Specific Port log page (see 10.2.8.1).
3h	SPINUP_HOLD	Phy is enabled; detected a SATA device and entered the SATA spinup hold state. The LINK RESET and HARD RESET operations in the SMP PHY CONTROL function (see 10.4.3.16) may be used to release the phy. This field shall be updated to this value at SATA spinup hold time (see 6.8.7 and 6.10)(i.e., after the COMSAS Detect Timeout timer expires during the SATA OOB sequence) if SATA spinup hold is supported.
4h	PORT_SELECTOR	Phy is enabled; detected a SATA port selector. The physical link rate has not been negotiated since the last time the phy's SP state machine entered the SP0:OOB_COMINIT state. The SATA spinup hold state has not been entered since the last time the phy's SP state machine entered the SP0:OOB_COMINIT state. The value in this field may change to 3h, 8h, or 9h if attached to the active phy of the SATA port selector. Presence of a SATA port selector is indicated by the ATTACHED SATA PORT SELECTOR bit.
5h	RESET_IN_PROGRESS	<u>Phy is enabled and the expander phy is performing a PHY CONTROL function, LINK RESET phy operation, or HARD RESET phy operation. The value in this field shall change when this field is updated to any other value. This value shall only occur if the specified phy contained a value of 8h or 9h in this field when the PHY CONTROL function, LINK RESET phy operation, or HARD RESET phy operation is processed.</u>
8h	G1	Phy is enabled; 1,5 Gbps physical link rate. This field shall be updated to this value after the speed negotiation sequence completes.
9h	G2	Phy is enabled; 3,0 Gbps physical link rate. This field shall be updated to this value after the speed negotiation sequence completes.
All others	Reserved.	
^a This code may be used by an application client in its local data structures to indicate an unknown negotiated physical link rate (e.g., before the discover process has queried the phy).		

Table 8 describes the ATTACHED SATA PORT SELECTOR bit and the ATTACHED SATA DEVICE bit.

Table 8 — ATTACHED SATA PORT SELECTOR and ATTACHED SATA DEVICE bits

ATTACHED SATA PORT SELECTOR bit value ^{a b}	ATTACHED SATA DEVICE bit value ^{c d}	Description
0	0	Neither a SATA port selector nor a SATA device is attached and ready on the selected phy.
0	1	The attached phy is a SATA device phy. No SATA port selector is present (i.e., the SP state machine did not detect COMWAKE in response to the initial COMINIT, but sequenced through the normal (non-SATA port selector) SATA device OOB sequence).
1	0	The attached phy is a SATA port selector host phy, and either: a) the attached phy is the inactive host phy, or b) the attached phy is the active host phy and a SATA device is either not present or not ready behind the SATA port selector (i.e., the SP state machine detected COMWAKE while waiting for COMINIT).
1	1	The attached phy is a SATA port selector's active host phy and a SATA device is present behind the SATA port selector (i.e., the SP state machine detected COMWAKE while waiting for COMINIT, timed out waiting for COMSAS, and exchanged COMWAKE with an attached SATA device).
<p>^a The ATTACHED SATA PORT SELECTOR bit is invalid if the NEGOTIATED PHYSICAL LINK RATE field is set to UNKNOWN (i.e., 0h), or DISABLED (i.e., 1h), <u>or RESET IN PROGRESS (i.e., 5h).</u></p> <p>^b Whenever the ATTACHED SATA PORT SELECTOR bit changes, the phy shall originate a Broadcast (Change)(see 7.11).</p> <p>^c For the purposes of the ATTACHED SATA DEVICE bit, the SATA port selector is not considered a SATA device.</p> <p>^d The ATTACHED SATA DEVICE bit shall be updated at SATA spin-up hold time (see 6.8.7 and 6.10).</p>		

10.4.3.11 DISCOVER LIST function

10.4.3.11.1 DISCOVER LIST function overview

The DISCOVER LIST function returns information about the device (i.e., some fields from the REPORT GENERAL response (see 10.4.3.3)) and one or more phys (i.e., some fields from the DISCOVER response (see 10.4.3.5)). This SMP function shall be implemented by all SMP target ports. This function is intended to provide the necessary information in a single SMP response for a self-configuring expander device to perform the discover process and configure its own expander routing table.

10.4.3.11.2 DISCOVER LIST request

Table 9 defines the request format.

Table 9 — 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 (06h)							
4	Reserved							
7	Reserved							
9	STARTING PHY IDENTIFIER							
9	MAXIMUM NUMBER OF DESCRIPTORS							
10	IGNORE ZONE GROUP	Reserved			PHY FILTER			
11	Reserved				DESCRIPTOR TYPE			
12	Reserved							
15	Reserved							
16	Vendor-specific							
27	Vendor-specific							
28	(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 06h.

The STARTING PHY IDENTIFIER field specifies the phy identifier of the first phy for which the information is being requested.

The MAXIMUM NUMBER OF DESCRIPTORS field specifies the maximum number of complete DISCOVER LIST descriptors that the management device server shall return.

The IGNORE ZONE GROUP bit is defined in the SMP DISCOVER request (see 10.4.3.5).

The PHY FILTER field is defined in table 10 and specifies a filter limiting which phys that the management device server shall return in the DISCOVER LIST descriptor list in the DISCOVER response.

Table 10 — PHY FILTER field

Code	Description
0h	All phys.
1h	Phys with the ATTACHED DEVICE TYPE field (see 10.4.3.5) set to 010b or 011b (i.e., attached to expander devices).
2h	Phys with the ATTACHED DEVICE TYPE field (see 10.4.3.5) set to a value other than zero (i.e., attached to end devices or expander devices).
All others	Reserved

The DESCRIPTOR TYPE field is defined in table 11 and specifies the DISCOVER LIST descriptor format and length.

Table 11 — DESCRIPTOR TYPE field

Code	DISCOVER LIST descriptor format	Descriptor length
0h	DISCOVER response defined in table 4 (see 10.4.3.5) not including the CRC field	The length of the DISCOVER response not including the CRC field
1h	SHORT FORMAT descriptor defined in table 227 (see 10.4.3.11.4)	24 bytes
All others	Reserved	

The CRC field is defined in 10.4.3.1.

10.4.3.11.3 DISCOVER LIST response

Table 12 defines the response format.

Table 12 — 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	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5								
6	RETRY DELAY TIMER field Reserved							
7								
8	STARTING PHY IDENTIFIER							
9	NUMBER OF DESCRIPTORS							
10	Reserved				PHY FILTER			
11	Reserved				DESCRIPTOR TYPE			

Table 12 — DISCOVER LIST response (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
12	DESCRIPTOR LENGTH							
13	Reserved							
15	Reserved							
16	ZONING SUPPORTED	ZONING ENABLED	Reserved			CONFIGURING	CONFIGURABLE ROUTE TABLE	
17	Reserved							
31	Reserved							
32	Vendor specific							
47	Reserved							
DISCOVER LIST descriptor list								
48	DISCOVER LIST descriptor (first)(see table 11 in 10.4.3.11.1, and table 4 in 10.4.3.5 or table 227 in 10.4.3.11.4)							
	...							
	DISCOVER LIST descriptor (last)(see table 11 in 10.4.3.11.1, and table 4 in 10.4.3.5 or table 227 in 10.4.3.11.4)							
n - 4								
n - 3	(MSB)							
n	CRC (LSB)							

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 contains the number of dwords that follow not including the CRC field.

~~The RETRY DELAY TIMER field is defined in table 5.~~

The EXPANDER CHANGE COUNT field is defined in the SMP REPORT GENERAL response (see 10.4.3.3).

The STARTING PHY IDENTIFIER field indicates the phy identifier of the first phy in the DISCOVER LIST descriptor list.

10.4.3.16 PHY CONTROL function

The PHY CONTROL function requests actions by the specified phy. This SMP function may be implemented by any SMP target port. In zoning expander devices, if zoning is enabled then this function shall only be processed from SMP initiator ports that have access to zone group 2 or the zone group of the specified phy (see 4.8.3.2).

Table 13 defines the request format.

Table 13 — PHY CONTROL request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (91h)							
2	Reserved							
3	REQUEST LENGTH (09h)							
4	(MSB)	EXPECTED EXPANDER CHANGE COUNT						(LSB)
5	Reserved							
6	Reserved							
8	Reserved							
9	PHY IDENTIFIER							
10	PHY OPERATION							
11	Reserved							UPDATE PARTIAL PATHWAY TIMEOUT VALUE
12	Reserved							
31	Reserved							
32	PROGRAMMED MINIMUM PHYSICAL LINK RATE				Reserved			
33	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				Reserved			
34	Reserved							
35	Reserved							
36	Reserved				PARTIAL PATHWAY TIMEOUT VALUE			
37	Reserved							
39	Reserved							
40	(MSB)	CRC						(LSB)
43	Reserved							

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

The FUNCTION field shall be set to 91h.

The REQUEST LENGTH field shall be set to 09h. For compatibility with previous versions of this standard, a REQUEST LENGTH field set to 00h specifies that there are 9 dwords before the CRC field.

The EXPECTED EXPANDER CHANGE COUNT field is defined in the SMP CONFIGURE GENERAL request (see 10.4.3.13).

The PHY IDENTIFIER field specifies the phy (see 4.2.7) to which the SMP PHY CONTROL request applies.

Table 14 defines the PHY OPERATION field.

Table 14 — PHY OPERATION field (part 1 of 2)

Code	Operation	Description
00h	NOP	No operation.
01h	LINK RESET	<p>If the specified phy is not a virtual phy, perform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the specified phy is a virtual phy, perform an internal reset and enable the specified phy. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device.</p> <p>Any affiliation (see 7.17.5) shall continue to be present. The phy shall bypass the SATA spinup hold state, if implemented (see 6.8.3.9).</p> <p>The SMP response shall be returned without waiting for the link reset to complete.</p> <p>During the link reset sequence the receipt of any SMP DISCOVER function (see 10.4.3.5) or any SMP DISCOVER LIST function (see 10.4.3.11) is processed as defined in 4.7.1</p>
02h	HARD RESET	<p>If the specified phy is not a virtual phy, perform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the attached phy is a SAS phy or an expander phy, the link reset sequence shall include a hard reset sequence (see 4.4.2). If the attached phy is a SATA phy, the phy shall bypass the SATA spinup hold state. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device.</p> <p>If the specified phy is a virtual phy, perform an internal reset and enable the specified phy.</p> <p>Any affiliation (see 7.17.5) shall be cleared.</p> <p>The SMP response shall be returned without waiting for the hard reset to complete.</p> <p>During the link reset sequence with hard reset the receipt of any SMP DISCOVER function (see 10.4.3.5) or any SMP DISCOVER LIST function (see 10.4.3.11) is processed as defined in 4.7.1</p>
03h	DISABLE	Disable the specified phy (i.e., stop transmitting valid dwords and receiving dwords on the specified phy). The LINK RESET and HARD RESET operations may be used to enable the phy. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device.
04h	Reserved	
05h	CLEAR ERROR LOG	Clear the error log counters reported in the REPORT PHY ERROR LOG function (see 10.4.3.6) for the specified phy.

Table 14 — PHY OPERATION field (part 2 of 2)

Code	Operation	Description
06h	CLEAR AFFILIATION	Clear an affiliation (see 7.17.5) from the STP initiator port with the same SAS address as the SMP initiator port that opened this SMP connection. If there is no such affiliation, the SMP target port shall return a function result of SMP FUNCTION FAILED in the response frame.
07h	TRANSMIT SATA PORT SELECTION SIGNAL	<p>This function shall only be supported by phys in an expander device.</p> <p>If the expander phy incorporates an STP/SATA bridge and supports SATA port selectors, the phy shall transmit the SATA port selection signal (see 6.6) which causes the SATA port selector to select the attached phy as the active host phy and make its other host phy inactive. See 7.11 for Broadcast (Change) requirements related to this phy operation in an expander device.</p> <p>Any affiliation (see 7.17.5) shall be cleared.</p> <p>If the expander phy does not support SATA port selectors, then the SMP target port shall return a function result of PHY DOES NOT SUPPORT SATA.</p> <p>If the expander phy supports SATA port selectors but is attached to a SAS phy or an expander phy, the SMP target port shall return a function result of SMP FUNCTION FAILED.</p>
08h	CLEAR STP I_T NEXUS LOSS	The STP I_T NEXUS LOSS OCCURRED bit in the REPORT PHY SATA function (see 10.4.3.7) shall be set to zero.
All others	Reserved	

If the PHY IDENTIFIER field specifies the phy which is being used for the SMP connection and a phy operation of LINK RESET, HARD RESET, or DISABLE is requested, the SMP target port shall not perform the requested operation and shall return a function result of SMP FUNCTION FAILED in the response frame.