To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 17 October 2005
Subject: 04-172r2 SAS-2 More counters

Revision history

Revision 0 (21 June 2004) First revision

Revision 1 (11 September 2005) Targeted at SAS-2 rather than SAS-1.1. Incorporated comments from July 2004 SAS protocol WG - use variable-length descriptor format rather than fixed format for SMP (making it the same as the log page format), add a mechanism to clear the peak values, define a BROADCAST for expanders to send when a counter wraps or when a peak value recorder is cleared

Revision 2 (17 October 2005) Incorporated comments from September 2005 SAS protocol WG - made the counter sources selectable in SMP but not in SSP. The number of counters is vendor-specific in both.

Related documents

sas1r09e - Serial Attached SCSI 2 revision 0

05-306r1 - SAS-2 STP connection time limits (Rob Elliott, HP) - adds ADDITIONAL DWORD fields to SMP function frame definitions

Overview

Additional standard per-phy counters or peak value registers, for both errors and non-errors, are desired to help testing and diagnosing problems, particularly in large configurations.

These counters should not be construed as required; they are all optional. For both disk drives and expanders, the number of concurrent counters is proposed as vendor-specific. For disk drives, it will be up to the drive what events each counter tracks; for expanders, software is allowed to choose.

A reasonable implementation for a disk drive would be to implement a small number of counters (e.g. 4) with vendor-specific means (e.g., new firmware loads) to configure which events are actually counted by those counters (implemented as a MUX with the various measurable events feeding each counter). If a system is suspected to be having a problem with arbitration, then devices might be loaded with firmware to count AIPs and OPEN_REJECTs to help diagnose the problem.

This is similar to the performance counters in Intel IA-64 and AMD AMD64 processors, which feature 4 counters with a 8-bit field to select the implementation-specific event counted by each.

Proposed events for both SSP target phys and expander phys:

a) Elasticity buffer overflow - not enough ALIGNS/NOTIFYs coming in, causing loss of a dword. This is important for checking ALIGN/NOTIFY insertion rates and performing clock frequency ppm tests.
b) Receive address frame error - incoming address frame has CRC or other problem (too few or too many dwords) so is ignored.
c) ERROR primitive received - phy receives an ERROR primitive (which is not counted as an invalid dword and may or may not have other effects)
d) Receive BREAK (after transmitting OPEN or during a connection)
e) Transmit BREAK - possibly differentiate the reasons: Open Timeout timer expires, Close Timeout Timer expires
f) Break Timeout timer expires. This causes the phy to assume the link is idle again (unless it chooses to start a new link reset sequence).

Proposed events for both SSP target phys and STP target phys:

a) Receive an abandon class OPEN_REJECT
b) Receive a retry class OPEN_REJECT
c) Transmit an abandon class OPEN_REJECT
d) Transmit a retry class OPEN_REJECT
e) Peak ARBITRATION WAIT TIME value used in outgoing OPENs
f) Peak PATHWAY BLOCKED COUNT value used in outgoing OPENs
g) Peak arbitration time. How long did it take to get an OPEN_ACCEPT or OPEN_REJECT?

Proposed events for SSP target phys:

a) Receive SSP frame error - SSP target port receives a bad SSP frame and generates NAK (CRC ERROR)
b) Transmit SSP frame error - SSP target port transmits an SSP frame and receives a NAK or an ACK/NAK timeout
c) Receives CREDIT_BLOCKED
d) Transmits CREDIT_BLOCKED

Proposed events for SMP target phys:

a) Receive address frame error
b) Receive SMP frame error - SMP target port receives a bad SMP frame and generates a BREAK
c) Transmit BREAK
d) Receive BREAK
e) Break Timeout. Sent a BREAK and got no reply (a very bad sign).
f) Transmit an abandon class OPEN_REJECT
g) Transmit a retry class OPEN_REJECT
h) Connection count. Can use this to tell if some phys in a wide link are not being used as much as others (or at all).
i) Peak connection time. Detect if there are any bus hogs.

Proposed events for STP target phys:

a) STP flow control buffer overflow - received too many dwords after a HOLD, causing a dword to be lost. The cable might be too long, allowing too many dwords in flight.

The new counters wrap rather than saturate. This facilitates easier multi-initiator usage, where clearing by one initiator can confuse the others. Wrapping versions of the four existing counters (invalid dword, running disparity error, loss of dword synchronization, and phy reset problem) are also defined.

The counters are all assumed to be 32 bits.

Other possibilities not in this proposal

In a multiple initiator environment, several of the items would be more useful if recorded on a per destination address basis, such as a peak ARBITRATION WAIT TIME transmitted, a peak PATHWAY BLOCKED COUNT transmitted, peak number of RRDYs received during a connection, and a received CREDIT_BLOCKED counter. This is not proposed due to complexity.

Suggested changes

4.x Phy event information

Phys shall count the following events using saturating counters and report them in the Protocol-Specific log page (see 10.x.x.x) and/or the SMP REPORT PHY ERROR LOG function:

a) invalid dwords received;
b) dwords received with running disparity errors;
c) loss of dword synchronization; and
d) phy reset problems.

The saturating counters are up to 4 bytes wide.

Phys may also count certain events (e.g., elasticity buffer overflows) using wrapping counters and record peak values for certain events (e.g., the longest connection time), reporting them in the Protocol-Specific log page (see 10.x.x.x) and/or the SMP REPORT PHY EVENT INFORMATION function (see 10.x.x.x). The wrapping counters and peak value detectors are each 4 bytes wide.

For phys not controlled by SMP target ports, the number of additional events counted/recorded and which events to count/record is vendor-specific.

For phys controlled by SMP target ports, the number of additional events that are simultaneously counted/recorded is vendor-specific, but the the SMP CONFIGURE PHY EVENT INFORMATION function (see 10.x.x.x) allows the events to count/record to be specified.

The PHY EVENT INFORMATION SOURCE field, defined in table 1, is used in the Protocol-Specific Port log page (see 10.2.x), the REPORT PHY EVENT INFORMATION function (see 10.3.x), and the CONFIGURE PHY EVENT INFORMATION function (see 10.3.x) indicates the type of phy event information being reported in the PHY EVENT INFORMATION field.

Table 1 — PHY EVENT INFORMATION SOURCE **field** (part 1 of 3) **[all new table]**

Code	Name	Description
00h	No event	No event. The PHY EVENT INFORMATION field is not valid.
01h	Invalid dword count	Number of invalid dwords (see 3.1.100) that have been received outside of phy reset sequences (i.e., between when the SP state machine (see 6.8) sends a Phy Layer Ready (SAS) confirmation or Phy Layer Ready (SATA) confirmation and when it sends a Phy Layer Not Ready confirmation to the link layer)
02h	Running disparity error count	Number of dwords containing running disparity errors (see 6.2) that have been received outside of phy reset sequences
03h	Loss of dword synchronization count	Number of times the phy has restarted the link reset sequence because it lost dword synchronization (see 6.9)(i.e., the SP state machine transitioned from SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready to SP0:OOB_COMINIT (see 6.8))
04h	Phy reset problem count	Number of times the phy did not obtain dword synchronization during the final SAS speed negotiation window (see 6.7.4.2)
05h	Elasticity buffer overflow count	Number of times the phy's receive elasticity buffer (see 7.3) has overflowed (e.g., because it did receive a sufficient number of ALIGNs and/or NOTIFYs)
06h	Received ERROR count	Number of times the phy received an ERROR primitive
07h - 1Fh	Reserved for phy layer-based phy event information	
20h	Received address frame error count	Number of times the phy detected an invalid address frame (e.g., because of a CRC error)
21h	Received OPEN_REJECT abandon count	Number of times the phy transmitted an OPEN address frame and received an abandon-class OPEN_REJECT (see 7.2.5.11). In expander devices, forwarded OPEN_REJECTs shall not be counted.
22h	Received OPEN_REJECT retry count	Number of times the phy transmitted an OPEN address frame and received a retry-class OPEN_REJECT (see 7.2.5.11). In expander devices, forwarded OPEN_REJECTs shall not be counted.
23h	Transmitted OPEN_REJECT abandon count	Number of times the phy received an OPEN address frame and transmitted an abandon-class OPEN_REJECT (see 7.2.5.11) In expander devices, forwarded OPEN_REJECTs shall not be counted.

Table 1 — PHY EVENT INFORMATION SOURCE **field** (part 2 of 3) **[all new table]**

Code	Name	Description
24h	Transmitted OPEN_REJECT retry count	Number of times the phy received an OPEN address frame and transmitted a retry-class OPEN_REJECT (see 7.2.5.11). In expander devices, forwarded OPEN_REJECTs shall not be counted.
25h	Received AIP (WAITING ON PARTIAL) count	Number of times the phy received an AIP (WAITING ON PARTIAL) or AIP (RESERVED WAITING ON PARTIAL). In expander devices, forwarded AIPs shall be counted.
26h	Received AIP (WAITING ON CONNECTION) count	Number of times the phy received an AIP (WAITING ON CONNECTION).In expander devices, forwarded AIPs shall be counted.
27h	Received BREAK count	Number of times the phy received a BREAK that was not a response to a BREAK that it transmitted
28h	Transmitted BREAK count	Number of times the phy transmitted a BREAK that was not a response to a BREAK it received (e.g., a Close Timeout was detected by the SL state machine interfacing to the SMP target port).
29h	Break Timeout count	Number of times the phy transmitted a BREAK and did not receive a BREAK in response (e.g., as detected by the XL state machine and/or the SL state machine interfacing to the SMP target port).
2Ah	Connection count	Number of connections in which the phy was involved
2Bh	Peak transmitted pathway blocked count	Peak value of a PATHWAY BLOCKED field in an OPEN address frame transmitted by the phy. Since the maximum value is FFh, only byte 3 is used.
2Ch	Peak transmitted arbitration wait time	Peak value of an ARBITRATION WAIT TIME field in an OPEN address frame transmitted by the phy. Since the maximum value is FFFFh, only bytes 2 and 3 are used.
2Dh	Peak arbitration time	Peak time in microseconds after transmitting an OPEN address frame that the phy has waited for connection response (e.g., OPEN_ACCEPT or OPEN_REJECT).
2Eh	Peak connection time	The peak duration, in microseconds, of any connection in which the phy was involved.
2Fh - 3Fh	Reserved for SAS arbitration-related phy information	
40h	Received SSP frame error count	Number of times the phy was used in a connection involving the SSP target port, detected an invalid frame, and transmitted a NAK (CRC ERROR) (e.g., because of a CRC error)
41h	Transmitted SSP frame error count	Number of times the phy was used in a connection involving the SSP target port, transmitted a frame, and received a NAK or an ACK/NAK timeout
42h	Received CREDIT_BLOCKED count	Number of times the phy received a CREDIT_BLOCKED
43h	Transmitted CREDIT_BLOCKED count	Number of times the phy transmitted a CREDIT_BLOCKED

Table 1 — PHY EVENT INFORMATION SOURCE **field** (part 3 of 3) **[all new table]**

Code	Name	Description
44h	Transmitted SSP data dword count	Number of data dwords transmitted inside SSP frames (i.e., between SOF and EOF)
45h	Received SSP data dword count	Number of data dwords received inside SSP frames (i.e., between SOF and EOF)
46h - 4Fh	Reserved for SSP-related phy event information	
50h	Transmitted SATA data dword count	Number of data dwords transmitted inside STP or SATA frames (i.e., between SATA_SOF and SATA_EOF)
51h	Received SATA data dword count	Number of data dwords received inside STP or SATA frames (i.e., between SATA_SOF and SATA_EOF)
52h	SATA flow control buffer overflow count	Number of times the phy's STP flow control buffer (see 7.17.3) has overflowed (e.g., because it received more data dwords than allowed after transmitting HOLD during an STP connection). This count should be maintained in the phy transmitting the HOLD and receiving the data dwords, but may be maintained in the phy receiving the HOLD and transmitting the data dwords.
53h - 5Fh	Reserved for STP and SATA-related phy event information	
60h	Receive SMP frame error count	Number of times the phy was used for to access the SMP target port and the SMP target port detected an invalid frame and transmitted a BREAK (e.g., because of a CRC error).
61h - 6Fh	Reserved for STP-related phy event information	
70h - CFh	Reserved	
D0h - FFh	Vendor specific	

6.7.4.2.3 SAS speed negotiation sequence

...

If the phy does not achieve dword synchronization during the final speed negotiation window, the SAS speed negotiation sequence fails. This is called a phy reset problem and may be counted and reported in the PHY RESET PROBLEM COUNT field in the SMP REPORT PHY ERROR LOG page (see 10.4.3.6) and the ~~REPORT PHY ERROR LOG~~ Protocol-Specific Port log page (see 10.2.8.1).

...

7.2.2/3 Primitive summary/encoding

Change BROADCAST (RESERVED 1) to BROADCAST (EXPANDER)

7.2.3.4 BROADCAST

BROADCASTs are used to notify all SAS ports in a domain of an event.

The versions of BROADCAST representing different reasons are defined in table 2.

Table 2 — BROADCAST primitives

Primitive	Description
BROADCAST (CHANGE)	Notification of a configuration change.
BROADCAST (RESERVED CHANGE 0)	Reserved. Processed the same as BROADCAST (CHANGE) by SAS ports (i.e, SAS initiator ports and SAS target ports).
BROADCAST (RESERVED CHANGE 1)	Reserved. Processed the same as BROADCAST (CHANGE) by SAS ports (i.e., SAS initiator ports and SAS target ports).
BROADCAST (SES)	Notification of an asynchronous event from a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-3 and SES-2) in the SAS domain.
BROADCAST (~~RESERVED 1~~ EXPANDER)	~~Reserved~~ Notification of a phy event information peak value recorder reaching its maximum value or being cleared.
BROADCAST (RESERVED 2)	Reserved.
BROADCAST (RESERVED 3)	Reserved.
BROADCAST (RESERVED 4)	Reserved.

When an expander port receives a BROADCAST it shall transmit the same BROADCAST on at least one phy in all other expander ports. BROADCAST shall only be sent outside of connections after the phy reset sequence has completed.

An expander device is not required to queue multiple identical BROADCASTs for the same expander port. If a second identical BROADCAST is requested before the first BROADCAST has been transmitted, the second BROADCAST may be ignored.

BROADCAST (CHANGE) is sent by an expander device to notify SAS initiator ports and other expander devices that a configuration change has occurred. BROADCAST (CHANGE) may also be transmitted by SAS initiator ports. BROADCAST (CHANGE) shall be ignored by SAS target ports.

BROADCAST (SES) is sent by a SAS target port to notify SAS initiator ports that an asynchronous event has occurred in an enclosure, and SSP initiator ports should poll all the logical units with peripheral device types set to 0Dh (i.e., enclosure services devices)(see SPC-3 and SES-2) in the SAS domain. BROADCAST (SES) shall be ignored by SAS target ports.

BROADCAST (EXPANDER) is sent by an expander device to notify SAS initiator ports that one or more of its phy event information peak detectors has reached its maximum value, or that the peak detectors have been cleared by an SMP CONFIGURE PHY EVENT INFORMATION function.

BROADCAST (RESERVED CHANGE 0) and BROADCAST (RESERVED CHANGE 1) shall be processed the same as BROADCAST (CHANGE) by SAS ports. ~~BROADCAST (RESERVED 1),~~ BROADCAST (RESERVED 2), BROADCAST (RESERVED 3), and BROADCAST (RESERVED 4) shall be ignored by SAS ports.

See 7.11 for details on SAS domain changes. See 10.4.3.3 for details on counting BROADCAST (CHANGE) generation in an expander device.

10.2.7 SCSI log parameters

10.2.7.1 Protocol-Specific log page

The Protocol Specific log page for SAS defined in table 3 is used to ~~report errors that have occurred on~~ return phy event information about the SAS target device's phy(s).

Table 3 — Protocol-Specific log page for SAS

Byte\Bit	7	6	5	4	3	2	1	0
0	PAGE CODE (18h)							
1	Reserved							
2	(MSB)			PAGE LENGTH (m - 3)				
3								(LSB)
	Protocol-specific log parameters							
4				First protocol-specific log parameter				
...	...							
				n^{th} protocol-specific log parameter				
m								

The PAGE CODE field shall be set to 18h.

The PAGE LENGTH field shall be set to the total length in bytes of the log parameters.

Table 4 defines the format for ~~a~~the SAS log parameter. The SAS log parameter is a list parameter (i.e., not a data counter) and only has cumulative (i.e., not threshold) values (see SPC-3).

Table 4 — Protocol-Specific log parameter format for SAS

Byte\Bit	7	6	5	4	3	2	1	0
0	(MSB)			PARAMETER CODE (relative target port identifier)				
1								(LSB)

Table 4 — Protocol-Specific log parameter format for SAS

Byte\Bit	7	6	5	4	3	2	1	0
2	DU	DS	TSD	ETC	TMC		LBIN	LP
3	PARAMETER LENGTH (y - 3)							
4	Reserved				PROTOCOL IDENTIFIER (6h)			
5	Reserved							
6								
7	NUMBER OF PHYS							
SAS phy log descriptors								
8	First SAS phy log descriptor							
	...							
	Last SAS phy log descriptor							
y								

The PARAMETER CODE field contains the relative target port identifier (see SPC-3) of the SSP target port that this log parameter describes.

Table 5 defines the values for the log parameter control bits for this log parameter.

Table 5 — Parameter control bits for SAS log parameters

~~Bit~~ Field	Value for LOG SENSE	Value for LOG SELECT	Description
DU	0	0 or 1	~~The value is provided by the device server.~~ The DU bit is not defined for list parameters, so shall be set to zero when read with the LOG SENSE command and shall be ignored when written with the LOG SELECT command.
DS	0	0 or 1	The device server shall support~~s~~ saving of the log parameter.
TSD	0	0 or 1	The device server ~~manages saving of the parameter~~ shall support implicitly saving the log parameter at vendor specific intervals.
ETC	0	0 or 1	~~No threshold comparison is made on this value.~~ The ETC bit is not defined for list parameters, so shall be set to zero when read with the LOG SENSE command and shall be ignored when written with the LOG SELECT command.
TMC	~~any~~ 00b	any	~~This field is ignored when the ETC bit is 0.~~ The TMC field is not defined for list parameters, so shall be set to 00b when read with the LOG SENSE command and shall be ignored when written with the LOG SELECT command.
LBIN	1	1	The log parameter is in binary format.
LP	1	1	The log parameter is a list parameter.

The PARAMETER LENGTH field is set to the length of the log parameter minus three.

The PROTOCOL IDENTIFIER field is set to 6h.

The NUMBER OF PHYS field contains the number of SAS phy log descriptors that follow.

Table 6 defines the SAS phy log descriptor. ~~Each SAS phy log descriptor is the same length~~.

Table 6 — SAS phy log descriptor (part 1 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved							
1	PHY IDENTIFIER							
2	Reserved							
3	SAS PHY LOG DESCRIPTOR LENGTH (m - 3)							
4	Reserved	ATTACHED DEVICE TYPE			Reserved			
5	Reserved				NEGOTIATED PHYSICAL LINK RATE			
6	Reserved				ATTACHED SSP INITIATOR PORT	ATTACHED STP INITIATOR PORT	ATTACHED SMP INITIATOR PORT	Reserved
7	Reserved				ATTACHED SSP TARGET PORT	ATTACHED STP TARGET PORT	ATTACHED SMP TARGET PORT	Reserved
8	SAS ADDRESS							
15								
16	ATTACHED SAS ADDRESS							
23								
24	ATTACHED PHY IDENTIFIER							
25	Reserved							
31								

Table 6 — SAS phy log descriptor (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
32	(MSB)			INVALID DWORD COUNT				
35								(LSB)
36	(MSB)			RUNNING DISPARITY ERROR COUNT				
39								(LSB)
40	(MSB)			LOSS OF DWORD SYNCHRONIZATION COUNT				
43								(LSB)
44	(MSB)			PHY RESET PROBLEM COUNT				
47								(LSB)
48				Reserved				
50								
51	NUMBER OF PHY EVENT DESCRIPTORS							
52				Phy event descriptor(s)				
m								

The SAS PHY LOG DESCRIPTOR LENGTH field indicates the number of bytes that follow in the SAS phy log descriptor. A SAS PHY LOG DESCRIPTOR LENGTH field set to zero indicates there are 44 additional bytes.

NOTE 1 Logical units compliant with SAS and SAS-1.1 only support a 48 byte SAS phy log descriptor.

The PHY IDENTIFIER field, ATTACHED DEVICE TYPE field, NEGOTIATED PHYSICAL LINK RATE field, ATTACHED SSP INITIATOR PORT bit, ATTACHED STP INITIATOR PORT bit, ATTACHED SMP INITIATOR PORT bit, ATTACHED SSP TARGET PORT bit, ATTACHED STP TARGET PORT bit, ATTACHED SMP TARGET PORT bit, SAS ADDRESS field, ATTACHED SAS ADDRESS field, and attaCHED PHY IDENTIFIER field are defined in the SMP DISCOVER function (see 10.4.3.5).

The INVALID DWORD COUNT field, RUNNING DISPARITY ERROR COUNT field, LOSS OF DWORD SYNCHRONIZATION field, and PHY RESET PROBLEM COUNT field are each defined in the SMP REPORT PHY ERROR LOG response data (see 10.4.3.6).

For the INVALID DWORD COUNT field, RUNNING DISPARITY ERROR COUNT field, LOSS OF DWORD SYNCHRONIZATION COUNT field, and PHY RESET PROBLEM COUNT field, the phy may maintain any size counter but should maintain a 32-bit counter. If it reaches its maximum value, the counter shall stop and the device server shall set the field to FFFFFFFFh in the SAS phy log descriptor.

The NUMBER OF PHY EVENT DESCRIPTORS field indicates how many phy event (see 4.xx) descriptors follow.

Each phy event descriptor is 8 bytes long and follows the format defined for the SMP REPORT PHY EVENT INFORMATION function in table 12 (see 10.4.3.x).

10.4.3 SMP functions

10.4.3.1 SMP function request frame format

...

The FUNCTION field specifies which SMP function is being requested and is defined in table 7. If the value in the FUNCTION field is not supported by the SMP target port, it shall return a function result of UNKNOWN SMP FUNCTION as described in table 148.

Table 7 — SMP functions

Code	SMP function	Description	~~Request frame size (in bytes)~~	~~Response frame size (in bytes)~~	Reference
00h	REPORT GENERAL	Return general information about the device	~~8~~	~~32~~	10.4.3.3
01h	REPORT MANUFACTURER INFORMATION	Return vendor and product identification	~~8~~	~~64~~	10.4.3.4
02h - 0Fh	Reserved for general SMP input functions				
10h	DISCOVER	Return information about the specified phy	~~16~~	~~56~~	10.4.3.5
11h	REPORT PHY ERROR LOG	Return error logging information about the specified phy	~~16~~	~~32~~	10.4.3.6
12h	REPORT PHY SATA	Return information about a phy currently attached to a SATA device	~~16~~	~~60~~	10.4.3.7
13h	REPORT ROUTE INFORMATION	Return route table information	~~16~~	~~44~~	10.4.3.8
14h	REPORT PHY EVENT INFORMATION	Return phy event information for the specified phy	~~16~~	~~variable~~	10.4.3.xx
~~14h~~ 15h - 1Fh	Reserved for phy-based SMP input functions				
20h - 3Fh	Reserved for SMP input functions				
40h - 7Fh	Vendor specific				
80h - 8Fh	Reserved for general SMP output functions				
90h	CONFIGURE ROUTE INFORMATION	Change route table information	~~44~~	~~8~~	10.4.3.9
91h	PHY CONTROL	Request actions by the specified phy	~~44~~	~~8~~	10.4.3.10
92h	CONFIGURE PHY EVENT INFORMATION	Configure phy event information for the specified phy	~~16~~	~~variable~~	10.4.3.xx
92h - 9Fh	Reserved for phy-based SMP output functions				
A0h - BFh	Reserved for SMP output functions				
C0h - FFh	Vendor specific				

...

10.4.3.5 DISCOVER

...

The NEGOTIATED PHYSICAL LINK RATE field is defined in table 8 and indicates the physical link rate negotiated during the link reset sequence. 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 8 — 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.x), and/or the Protocol-Specific Port log page (see).
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 10.4.3.10) 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.
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).

...

10.4.3.10 PHY CONTROL

...

Table 9 defines the PHY OPERATION field.

Table 9 — PHY OPERATION **field** (part 1 of 2)

Code	Operation	Description
00h	NOP	No operation.
01h	LINK RESET	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. 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). The SMP response shall be returned without waiting for the link reset to complete.
02h	HARD RESET	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. If the specified phy is a virtual phy, perform an internal reset and enable the specified phy. Any affiliation (see 7.17.5) shall be cleared. The SMP response shall be returned without waiting for the hard reset to complete.
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 9 — 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	This function shall only be supported by phys in an expander device. 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. Any affiliation (see 7.17.5) shall be cleared. 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. 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.
All others	Reserved	

...

Editor's Note 1: Beginning all-new section (no underlines shown)

10.4.3.x REPORT PHY EVENT INFORMATION function

The REPORT PHY EVENT INFORMATION function returns phy event information (see 4.x) about the specified phy. This SMP function may implemented by any SMP target port.

Table 6 defines the request format.

Table 10 — REPORT PHY EVENT INFORMATION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (14h)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	Reserved							
5	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB) CRC							
15								(LSB)

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

The FUNCTION field shall be set to 14h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 2).

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for which information shall be reported.

The CRC field is defined in 10.4.3.1.

Table 7 defines the response format.

Table 11 — REPORT PHY EVENT INFORMATION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (14h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
4	Reserved							
5	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
14								
15	NUMBER OF PHY EVENT DESCRIPTORS							
16	Phy event descriptor(s)							
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 14h.

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 PHY IDENTIFIER field indicates the phy (see 4.2.7) for which information is being reported.

The NUMBER OF PHY EVENT DESCRIPTORS field indicates how many phy event descriptors follow.

Each phy event descriptor is 8 bytes long and follows the format defined in table 5..

Table 12 — Phy event descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved							
2								
3	PHY EVENT INFORMATION SOURCE							
4	(MSB)	PHY EVENT INFORMATION						
7								(LSB)

The PHY EVENT INFORMATION SOURCE field, defined in table 1 in 4.xx, indicates the type of phy event information being reported in the PHY EVENT INFORMATION field.

The PHY EVENT INFORMATION field contains the value (i.e., the count or peak value detected) of the event indicated by the PHY EVENT INFORMATION SOURCE field.

The CRC field is defined in 10.4.3.2.

10.4.3.xx CONFIGURE PHY EVENT INFORMATION function

The CONFIGURE PHY EVENT INFORMATION function configures phy event information (see 4.x) about the specified phy. This SMP function may implemented by any SMP target port.

Table 6 defines the request format.

Table 13 — CONFIGURE PHY EVENT INFORMATION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (92h)							
2	Reserved							
3	REQUEST LENGTH							
4	Reserved							CLEAR PEAKS
5	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11	NUMBER OF PHY EVENT INFORMATION SOURCES							
12	PHY EVENT INFORMATION SOURCE (first)							
...	...							
n - 4	PHY EVENT INFORMATION SOURCE (last)							
n - 3	(MSB) CRC							
n								(LSB)

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

The FUNCTION field shall be set to 92h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field.

A CLEAR PEAKS field set to one specifies that any event information recorders recording peak values shall be set to zero. A CLEAR PEAKS field set to zero specifies no change to the event information recorders containing peak values.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for which information shall be reported.

The NUMBER OF PHY EVENT INFORMATION SOURCES field specifies the number of phy event information sources, and shall be set to the same value as the NUMBER OF PHY EVENT DESCRIPTORS field in the SMP REPORT PHY EVENT INFORMATION function (see 10.4.3.x).

Each PHY EVENT INFORMATION SOURCE field, defined in table 1 in 4.xx, specifies the type of event that shall be recorded by the corresponding event information recorder.

The CRC field is defined in 10.4.3.1.

Table 7 defines the response format.

Table 14 — CONFIGURE PHY EVENT INFORMATION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (92h)							
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 92h.

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 (i.e., 0).

The CRC field is defined in 10.4.3.2.

Editor's Note 2: end of all-new section
