From: Steven Fairchild, HP (steve.Fairchild@hp.com)
Date: 6 January 2006
Subject: SAS BROADCAST (SCSI ASYNCHRONOUS EVENT)

In a multi-initiator and/or a dual port environment, a hard reset sequence to a target device will flush the outstanding commands for all initiators on all target ports. This will result in commands timing out on initiators that are not aware of the hard reset sequence. So that the initiators in the domain(s) are notified when a hard reset sequence is performed, this proposal defines a new BROADCAST (SCSI ASYNCHRONOUS EVENT) primitive that the target device would use to indicate that a hard reset sequence has completed on one or both ports. In the case of a wide target port, it would be broadcast on all links of a wide port.

When an initiator received the BROADCAST (SCSI ASYNCHRONOUS EVENT) primitive, it would take one of two actions;

a) If directly attached to the target device, then the initiator would determine if it generated the hard reset sequence. If the initiator did not generate the hard reset sequence then the source would be from another port (i.e. on a dual port target device).
b) If attached to the target device through an expander infrastructure, then the initiator would determine the target device that was the source of the BROADCAST (SCSI ASYNCHRONOUS EVENT) primitive. The initiator may choose to query the expander infrastructure to identify the source (i.e. a new SMP command to query for BROADCAST counts) or may choose to use its command queue to determine if any of the devices with outstanding commands still have the commands pending (i.e. issue TEST UNIT READY to each device in its queue to see if a UNIT ATTENTION is returned, or use QUERY TASK for on or more command tasks in the queue).

In either case, the initiator should notify the upper layers of the driver stack of any outstanding commands that have been aborted. This avoids excessive delays waiting for commands to timeout.

There is a mechanism that exists, the TAS bit on the CONTROL mode page (0Ah, SPC-3) that attempts to address this problem. However it requires the target device to make a connection to each initiator notifying it that commands are being flushed from the queue. This presents a problem if the cause of the hard reset sequence is a misbehaving target device, then it may not be able to comply with the requested behavior.

Requiring a target device to generate a BROADCAST (SCSI ASYNCHRONOUS EVENT) primitive when it completes its HARD LINK RESET seems more likely to succeed in the potential failing situations.

Requested changes;

Redefine BROADCAST (RESERVED x) to BROADCAST (SCSI ASYNCHRONOUS EVENT) in table 79, with a description of;

BROADCAST (SCSI ASYNCHRONOUS EVENT)	Notification from a target port that a SCSI asynchronous event has occurred (e.g. hard reset sequence).

Add the following text to section 7.2.5.4 BROADCAST

BROADCAST (SCSI ASYNCHRONOUS EVENT) is sent by a SAS target port to notify SAS initiator ports that a hard reset sequence has been processed by this target port or another target port in the SAS target device.

> Editor's note: Should the BROADCAST (SCSI ASYNCHRONOUS EVENT) be consolidated with BROADCAST (SES) since BROADCAST (SES) is effectively just a subset of SCSI asynchronous events.

Define a new SMP function that returns the BROADCAST counters for all phys.

10.4.3.x REPORT BROADCAST COUNTS function

The REPORT BROADCAST COUNTS function returns the broadcast primitives received counts for the specified phy. This SMP function should be implemented by all SMP target ports.

The expander device is not required to increment the counts contained in the REPORT BROADCAST COUNTS response again unless a REPORT BROADCAST COUNTS response is transmitted. The counts contained in the REPORT BROADCAST COUNTS response shall not be incremented when forwarding a BROADCAST primitive from another expander device. The counts contained in the REPORT BROADCAST COUNTS response shall wrap to zero after the maximum value (i.e., FFh) has been reached.

NOTE xx - Application clients that use the REPORT BROADCAST COUNTS response should request it often enough to ensure that the counts contained in the REPORT BROADCAST COUNTS response do not increment a multiple of 256 times between requests.

Table xxx defines the request format.

Table xxx – REPORT BROADCAST COUNTS request

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (xxh)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	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 xxh.

The REQUEST LENGTH field shall be set to 02h.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for the broadcast counters being requested.

The CRC field is defined in 10.4.3.1

Table yyy defines the response format.

Table yyy – REPORT BROADCAST COUNTS response

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (xxh)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (07h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	Reserved	BROADCAST RESERVED 4 COUNT VALID	BROADCAST RESERVED 3 COUNT VALID	BROADCAST RESERVED 2 COUNT VALID	BROADCAST SCSI ASYNC EVENT COUNT VALID	BROADCAST SES COUNT VALID	BROADCAST CHANGE RESERVED 1 COUNT VALID	BROADCAST CHANGE RESERVED 0 COUNT VALID

Byte/Bit	7	6	5	4	3	2	1	0
13	Reserved							
15								
16	BROADCAST CHANGE RESERVED 0 COUNT							
17	BROADCAST CHANGE RESERVED 1 COUNT							
18	BROADCAST SES COUNT							
19	BROADCAST SCSI ASYNCHRONOUS EVENT COUNT							
20	BROADCAST RESERVED 2 COUNT							
21	BROADCAST RESERVED 3 COUNT							
22	BROADCAST RESERVED 4 COUNT							
23	Reserved							
31								
32	(MSB) CRC							
35								(LSB)

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

The FUNCTION field shall be set to xxh.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 07h.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for which the broadcast counters is being returned.

A BROADCAST CHANGE RESERVED 0 COUNT VALID field set to one indicates that the BROADCAST CHANGE RESERVED 0 COUNT field is supported. A BROADCAST CHANGE RESERVED 0 COUNT VALID field set to zero indicates that the BROADCAST CHANGE RESERVED 0 COUNT field is not supported.

A BROADCAST CHANGE RESERVED 1 COUNT VALID field set to one indicates that the BROADCAST CHANGE RESERVED 1 COUNT field is supported. A BROADCAST CHANGE RESERVED 1 COUNT VALID field set to zero indicates that the BROADCAST CHANGE RESERVED 1 COUNT field is not supported.

A BROADCAST SES COUNT VALID field set to one indicates that the BROADCAST SES COUNT field is supported. A BROADCAST SES COUNT VALID field set to zero indicates that the BROADCAST SES COUNT field is not supported.

A BROADCAST SCSI ASYNC EVENT COUNT VALID field set to one indicates that the BROADCAST SCSI ASYNCHRONOUS EVENT COUNT field is supported. A BROADCAST SCSI ASYNCHRONOUS EVENT COUNT VALID field set to zero indicates that the BROADCAST SCSI ASYNCHRONOUS EVENT COUNT field is not supported.

A BROADCAST RESERVED 2 COUNT VALID field set to one indicates that the BROADCAST RESERVED 2 COUNT field is supported. A BROADCAST RESERVED 2 COUNT VALID field set to zero indicates that the BROADCAST RESERVED 2 COUNT field is not supported.

A BROADCAST RESERVED 3 COUNT VALID field set to one indicates that the BROADCAST RESERVED 3 COUNT field is supported. A BROADCAST RESERVED 3 COUNT VALID field set to zero indicates that the BROADCAST RESERVED 3 COUNT field is not supported.

A BROADCAST RESERVED 3 COUNT VALID field set to one indicates that the BROADCAST RESERVED 3 COUNT field is supported. A BROADCAST RESERVED 3 COUNT VALID field set to zero indicates that the BROADCAST RESERVED 3 COUNT field is not supported.

The BROADCAST CHANGE RESERVED 0 COUNT field counts the number of BROADCAST (RESERVED CHANGE 0)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (RESERVED CHANGE 0) from an attached end device. The expander device shall not

increment the count for a BROADCAST (RESERVED CHANGE 0) forwarded by another expander device.

The BROADCAST CHANGE RESERVED 1 COUNT field counts the number of BROADCAST (RESERVED CHANGE 1)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (RESERVED CHANGE 1) from an attached end device. The expander device shall not increment the count for a BROADCAST (RESERVED CHANGE 1) forwarded by another expander device.

The BROADCAST SES COUNT field counts the number of BROADCAST (SES)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (SES) from an attached end device. The expander device shall not increment the count for a BROADCAST (SES) forwarded by another expander device.

The BROADCAST SCSI ASYNCHRONOUS EVENT COUNT field counts the number of BROADCAST (SCSI ASYNCHRONOUS EVENT)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (SCSI ASYNCHRONOUS EVENT) from an attached end device. The expander device shall not increment the count for a BROADCAST (SCSI ASYNCHRONOUS EVENT) forwarded by another expander device.

The BROADCAST RESERVED 2 COUNT field counts the number of BROADCAST (RESERVED 2)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (RESERVED 2) from an attached end device. The expander device shall not increment the count for a BROADCAST (RESERVED 2) forwarded by another expander device.

The BROADCAST RESERVED 3 COUNT field counts the number of BROADCAST (RESERVED 3)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (RESERVED 3) from an attached end device. The expander device shall not increment the count for a BROADCAST (RESERVED 3) forwarded by another expander device.

The BROADCAST RESERVED 4 COUNT field counts the number of BROADCAST (RESERVED 4)s received from a SAS target port attached to the phy specified by the PHY IDENTIFIER. Expander devices shall support this field. Other devices shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it receives a BROADCAST (RESERVED 4) from an attached end device. The expander device shall not increment the count for a BROADCAST (RESERVED 4) forwarded by another expander device.

The CRC field is defined in 10.4.3.1