03-034r3 SAS Expander internal devices

To: T10 Technical Committee
From: Steven Fairchild, HP(steve.fairchild@hp.com)
Date: 21 January 2003
Subject: 03-034r3 SAS Expander internal devices

Revision history

Revision 0 (17 December 2002) First revision Revision 1 (10 January 2003) changed to minimize changes to existing specification Revision 2 (17 January 2003) added feedback from WG Revision 3 (21 January 2003) fixed typos

Related documents

sas-r03 - Serial Attached SCSI revision 3

Overview

There is confusion on how expanders need to handle internal initiator and target ports with different protocols with respect to addressability.

This proposal is to clarify that internal "end devices" should be treated the same as "end devices" connected to external phys.

In the specification, expanders are described as potentially containing internal initiator and target ports with different protocols supported. These internal ports reside at the same SAS address as the expander function.

Instead of this approach, expanders should only report the SMP bit (initiator and/or target) at their address. An internal port/device of the expander should have a different SAS address than the expander with the appropriate protocol bits set.

The expander should then increase the NUMBER OF PHYS field in the REPORT GENERAL to include each internal port as a phy.

As an example, an edge expander (SAS address, X) with 16 external phys and an internal device (SAS address, Y) that supported the SSP protocol would report the following information;

IDENTIFY FRAME device type - edge expander, 010b protocol bits - SMP target set, all others cleared SAS address - X

REPORT GENERAL number of phys - 17

DISCOVER -

for phy identifier - 0 - 15, information about external attached devices

...

for phy identifier - 16, information about internal device attached device type - end device, 001b routing attribute - direct routing, 0000b negotiated link rate - (phy is enabled, unknown link rate) or (link rate associated with expander phy responding to SMP request) protocol bits - attached SSP target, all others cleared attached SAS address - Y SAS address - X programmed minimum physical link rate - based on expander characteristics programmed maximum physical link rate - based on expander characteristics hardware minimum physical link rate - based on expander characteristics hardware maximum physical link rate - based on expander characteristics

1

partial pathway timeout value - based on expander characteristics

An initiator attached to the expander would be able to distinguish an internal port from a device attached by the discover response. A phy representing an internal port would have the INTERNAL PHY bit set in the discover response.

Because of the increased phy count due to internal ports consuming phy identifiers, the current limit of 64 phy identifiers within an edge device set should be increased to128.

Suggested changes to SAS

Change section 4.1.8.2 text to:

"The number of devices attached to an edge expander device set shall not exceed 64 128.

Change section 4.1.10.

"No more than one fanout expander device shall be included in a SAS domain. The fanout expander device may be attached to up to 64 128 edge expander device sets, initiator ports, or target ports. Each edge expander devices set shall contain no more than 64 128 physical links to edge expander devices, initiator devices, or target devices. Each edge expander device set shall not be attached to more than one fanout expander device. An edge expander device set may be attached to one other edge expander device set if that is the only other edge expander device set in the domain and there are no fanout expander devices in the domain."

Change figure 14, all text references from 64 to 128.

Change figure 15, all text references from 64 to 128.

Change figure 16, all text references from 64 to 128.

Change section 4.2.6 text to:

"Phy identifiers shall be greater than or equal to 00h and less than 40 80h.

Change section 4.4.2 text to:

"If the port is part of an expander device, the expander function and other expander ports in the expander device shall not be affected by hard reset. If the expander device is also a SCSI device and/or an ATA device-i.e., it contains internal SCSI or ATA initiator ports or target ports), the SCSI and/or ATA device roles are reset. An internal port of an expander device that supports STP, SSP or SATA shall not be affected by a hard reset to the expander device.

An internal port of an expander device that supports STP, SSP or SATA shall be reset when an SMP PHY CONTROL function with a phy operation of HARD RESET and phy identifier for the internal port is executed by the expander which contains the internal port."

Change section 4.6.1 list to:

"An expander device may contain the following:

- a) additional internal SMP target ports;
- b) internal SMP initiator port(s);
- c) internal SSP target port(s);
- d) internal SSP initiator port(s);
- e) internal STP target port(s); and
- f) internal STP initiator port(s); and-
- g) internal SATA device port(s)."

Change section 4.6.2 text to:

"One internal SMP port, one internal SSP port, and one internal STP port may each share the expanderdevice's SAS address. If there is more than one internal SMP port, one internal SSP port, and one internal STPport, the additional ports shall include SAS addresses different from that of the expander device. One internal SMP port is associated with the device's SAS address. Any additional internal ports of the expander device shall have SAS addresses different from that of the expander device, shall have a unique phy identifier within the expander device and shall be identified as an end device type." Change figure 30, all text references from 64 to 128.

Change section 7.7.2 text to:

"The DEVICE TYPE field indicates the type of device containing the phy, and is defined in table 73. A device that is capable of being both an end device and an expander device shall only identify itself as an expander device in this field

An expander device with internal ports supporting SSP, STP or SATA shall identify itself as an expander device supporting only an SMP target port. The expander device may also identify itself as supporting an SMP initiator port."

Change section 10.3.1.2 text to:

"If a fanout expander device supports an expander route table, then the number of expander route indexes shall be 64 128.

The NUMBER OF PHYS field contains the number of phys in the device, including any internal ports supporting SSP, STP or SATA."

Change table 135 to:

Bit Byte	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (10h)								
2		FUNCTION RESULT							
3	Reserved								
4	lanorod								
7									
8	Reserved								
9	PHY IDENTIFIER								
10	Ignored								
11	Reserved								
12	Ignored ATTACHED DEVICE TYPE				ROUTING ATTRIBUTE				
13	Reserved				NEGOTIATED PHYSICAL LINK RATE				
14	Reserved				ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	Reserved	
15	Reserved				ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET	ATTACHED SATA TARGET	
16									
23	ATTACHED SAS ADDRESS								
24									
31									
32	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE				
33	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE				
34	Vandar aposifia								
35									
36	INTERNAL Reserved				PARTIAL PATHWAY TIMEOUT VALUE				
37		Descried							
39									
40	(MSB)								
43						(LSB)			

Table 1 — DISCOVER response

03-034r3 SAS Expander internal devices

Change section 10.3.1.4 to:

"The PARTIAL PATHWAY TIMEOUT VALUE field indicates the partial pathway timeout value set by the PHY CONTROL function (see 10.3.1.9). The default value for PARTIAL PATHWAY TIMEOUT VALUE shall be 7 μ s. The INTERNAL PHY field indicates that the phy identifier references a link to an internal port. The INTERNAL PHY bit set to one indicates the attached device is internal to the expander device. The INTERNAL PHY bit set to zero indicates the attached device is external to the expander device. The INTERNAL PHY bit set to zero indicates the attached in 9.4.3."

Change table 156 to:

Code	Operation	Description		
00h	NOP	No operation.		
01h	LINK RESET	If the specified phy is not identified as an internal phy, pPPerform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the specified phy is identified as an internal phy, perform an internal reset on the link and enable the specified phy. Any active affiliation (see 9.3.2) shall continue to be active. The phy shall bypass the SATA spinup hold state.		
02h	HARD RESET	If the specified phy is not identified as an internal phy, pPerform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the specified phy is identified as an internal phy, perform an internal reset on the link and enable the specified phy. If the attached phy is a SAS phy and is not identified as an internal phy, the link reset sequence shall include a hard reset sequence (see 4.4.2). Any active affiliation (see 9.3.2) shall be cleared. The phy shall bypass the SATA spinup hold state.		
03h	DISABLE	Disable the specified phy (i.e., stop transmitting and receiving on the specified phy). The LINK RESET and HARD RESET operations may be used to enable the phy.		
04h	NEA LOOPBACK	Set the specified phy to near-end analog loopback mode (see 7.10) if it is not disabled and it is currently operating at a valid physical link rate.		
05h	CLEAR ERROR LOG	Clear the error log counters (see 10.3.1.5) for the specified phy.		
06h	CLEAR AFFILIATION	Clear an active affiliation (see 9.3.2) 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 target port shall return a FUNCTION RESULT of SMP FUNCTION FAILED in the response frame.		
All others	Reserved.			