To: T10 Technical Committee
From: Steven Fairchild, HP(steve.fairchild@hp.com)
Date: 10 January 2003
Subject: 03-034r1 SAS Expander internal devices

Revision history
Revision 0 (17 December 2002) First revision
Revision 1 (10 January 2002) changed to minimize changes to existing specification

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 protocol 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 - X
    SAS address - Y
    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
    partial pathway timeout value - based on expander characteristics

An initiator attached to the expander would be unable to logically distinguish an internal port of the expander from a device attached to an external phy during an SMP operation.

To make the association of internal ports to specific expanders some other means would need to be employed, like enclosure services, which is beyond the scope of the SAS specification. However, this seems appropriate because the enclosure services will need to be relied upon in any case to identify the physical location of all devices within the SAS topology.

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 device 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 the STP, SSP or SATA protocol shall not be affected by a hard reset to the expander device.
An internal port of an expander device that supports the STP, SSP or SATA protocol 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 target 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 expander device's SAS address. If there is more than one internal SMP port, one internal SSP port, and one internal STP port, 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 protocols shall identify itself as an expander device supporting only the SMP target protocol. The expander device may also identify itself as supporting the SMP initiator protocol."

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 the SSP, STP or SATA protocols."

Discussion

By treating internal devices as phys the number of real devices that may be attached to an expander device set (currently 64) is reduced. To allow the number of real devices to be sufficiently large, the range of phy identifiers should be increased from 64 to 128.

This change could just be limited to the number of devices in an edge expander device set, or could also define the maximum number of phys for a fanout expander (from 64 to 128).

This would increase the total number of potential devices in the topology from 4096 to 16384.

Conversly, the total number of potential devices could be held constant and the number of phys for a fanout expander could be reduced (from 64 to 32).

The last possibility is to leave the number of fanout phys the same, so the total number of potential devices in the topology would increase to 8192.

So the options would be;

A) don't change the number of phys in either the edge expander device set or fanout expander.

B) change the number of phys in the edge expander device set to 128 only.

C) change the number of phys in the edge expander device set and fanout expander device to 128 each.

D) change the number of phys in the edge expander device set to 128 and the fanout expander device to 32.

C is documented by this proposal because it has the least impact on the specification.