Date: 27 January 2006
To: T10 Technical Committee
From: Tim Symons
Subject: This is how SAS self-configuring expanders are already defined?

Referenced Document
SAS-2 revision 2 (http://www.t10.org/ftp/t10/drafts/sas2/sas2r02.pdf)
T10/06-087r0 - How do SAS self-configuring expanders self configure?

Overview
Proposal 06-087r0 could incorrectly imply that self-configuring zoning expander devices are not properly defined in SAS1.1 and SAS2r02.

It is unfortunate that the references in 06-087r0 were edited such that the context of each statement was indeterminate. By providing a more extensive reference, it is shown that self configuring expanders are explicitly defined in SAS1.1 and SAS2r02, and that their characteristics of self configuring follow a defined set of rules.

Self-configuring expander devices provide robustness against the case of a single point failure in SAS domains. For zoning applications self-configuring expanders provide a level of independence appropriate to device group isolation in shared infrastructures. Self-configuring expanders are available in many different products. They have been available for a significant period of time, and have been successfully used in many interoperability SAS plug-fest events.

Further searches on the definition of the elements required of a self configuring expander device such as the management application client yields more definition of self-configuring expander devices.

References made in 06-087r0 are expanded to show their context, and the relevance of each reference to SAS1.1 and SAS-2 systems including self configuring expander device functions.

==========
Reference SAS2r02, Section 3.1, Definitions

3.1.50 discover process: The algorithm used by a management application client to configure the SAS domain. See 4.7.
==========
Reference SAS2r02, Section 4.1.5, Last Paragraph

“Expander devices with expander phys with the table routing attribute contain an expander route table. The expander route table may be configurable. An expander device with a configurable route table depends on a management application client within the SAS domain to use the discover process (see 4.7) to configure the expander route table. An expander device with expander phys with the table routing attribute that does not have a configurable route table shall be self-configuring, **and shall contain a management application client and SMP initiator port to perform the discover process to configure its own expander route table**.”

>This is the paragraph referenced in 06-087r0. that unfortunately was abbreviated and appeared out of context. The text clearly states that a self configuring expander device shall contain a management application client and SMP initiator port to perform the discover process to configure its own expander route table. The discovery process is described in 4.7.

==========

Reference SAS2r02,

4.6.5 Broadcast primitive processor (BPP)

…

The BPP receives broadcast primitive requests from each expander phy and requests transmission of those requests on all expander ports except the expander port from which the broadcast primitive request was received.

In a self-configuring expander device (see 4.1.5), the BPP requests transmission of a BROADCAST (CHANGE) when it completes configuration (see 10.4.3.3).

>This is the paragraph referenced in 06-087r0. that unfortunately was abbreviated and appeared out of context. When placed with the associated text it is clear that this event relates to the completion of the discovery process.

==========

Reference SAS2r02,

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.**

> This section not only states that an self configuring expander device contains a management client, but it also states the discover process to be used.

==========

Reference SAS2r02, Section 4.7.3 Discover process optimization

…

If the management application client supports the discover process optimization, then the management application client should provide a vendor-specific method for initiating a check of the resulting expander route tables. The check should be performed under the following situations:

a) when an OPEN_REJECT (NO DESTINATION) is received for a connection request to a SAS address that is expected to be accepted;
b) when a discover process has been completed;
c) when another SMP initiator port is discovered in the SAS domain; or
d) when a self-configuring expander device is discovered in the SAS domain.

If the management application client detects an inconsistency in the expander route tables when the discover process optimization is enabled (e.g., detects entries that appear to have been filled

in by a discover process with optimization disabled), then the management application client shall report an error in a vendor-specific manner and shall disable the discover process optimization. The management application client should then re-initiate a discover process with the discover process optimization disabled.

> Describes discovery process optimization for all applications which includes self-configuring expander devices.
==========
Reference SAS2r02, Section 7.11 SAS domain changes
…
Expander devices shall transmit BROADCAST (CHANGE) for the following reasons:

a) after an expander phy's SP state machine transitions from the SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready state to the SP0:OOB_COMINIT state (see 6.8);

NOTE 29 - This occurs when the expander phy is reset or disabled with the SMP PHY CONTROL function DISABLE, LINK RESET, HARD RESET, or TRANSMIT SATA PORT SELECTION SIGNAL phy operations (see 10.4.3.12) as well as when dword synchronization is unexpectedly lost;

b) after a virtual phy has been disabled with the SMP PHY CONTROL function DISABLE phy operation or started processing a reset requested by the LINK RESET or HARD RESET phy operations (see 10.4.3.12);
c) after an expander phy's SP state machine reaches the SATA spinup hold state (see 6.8.7and 6.10);
d) after an expander phy's SP state machine receives a COMWAKE Detected message in states SP0:OOB_COMINIT, SP1:OOB_AwaitCOMX, SP3:OOB_AwaitCOMINIT_Sent, or SP4:OOB_COMSAS if the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero prior to receiving the COMWAKE detected message (see 6.8.3 and table 201 in 10.4.3.5);
e) after an expander phy's SP state machine transitions from the SP1:OOB_AwaitCOMX state to the SP0:OOB_COMINIT state if the value of the ATTACHED SATA PORT SELECTOR bit was one in the DISCOVER response upon entry to SP1:OOB_AwaitCOMX, and if no COMWAKE detected message was received while in SP1:OOB_AwaitCOMX before the transition to SP0:OOB_COMINIT (see 6.8.3.3.2);
f) after the link reset sequence completes (see 7.9);
g) after a virtual phy has been enabled or completed processing a reset requested by the SMP PHY CONTROL function LINK RESET or HARD RESET phy operations (see 10.4.3.12);
h) after a self-configuring expander device has completed configuration and has changed its
CONFIGURING bit from one to zero in the SMP REPORT GENERAL function (see 10.4.3.3);
i) after an STP/SATA bridge receives an initial Register - Device to host FIS (see 9.3.1); and
j) after an expander phy receives BROADCAST (CHANGE).

> In context the definition describes many condition, including the definition for self-configuring expander devices.
==========
Reference SAS2r02, Section 10.4.3.3 REPORT GENERAL function
…
A CONFIGURING bit set to one indicates that a self-configuring expander device has not completed configuring its expander route table. A CONFIGURING bit set to zero indicates that configuration is complete and the expander device is ready for connection requests. Changes in this bit from one to zero result in a BROADCAST (CHANGE) being originated. SMP target ports in self-configuring expander devices shall support this bit. SMP target ports in configurable expander devices and in other device types shall set the CONFIGURING bit to zero.

The CONFIGURABLE ROUTE TABLE bit indicates whether the expander device has an expander route table that is required to be configured with the SMP CONFIGURE ROUTE INFORMATION function (see 4.6.7.3). An expander device with a configurable route table shall set the CONFIGURABLE ROUTE TABLE bit to one. An expander device without a configurable route table or a device with any other device type shall set the CONFIGURABLE ROUTE TABLE bit to zero.

> definition of bit fields in report general supporting self-configuring expander devices.
==========