

Accredited Standards Committee*
InterNational Committee for Information Technology Standards (INCITS)

Doc. No.: T10/ 05-195r0

Date: May 16, 2005

Reply to: Heng Liao

To: T10 Membership

From: Heng Liao

Subject: Minutes of T10 SAS2 Zoning Conference Call -- May 13, 2005

Attendees:

Pat Thaler
Heng Liao
Bill Martin
Bob Sheffield
Ed D'Avignon
George Penokie
Gerry Houlder
Greg Elkins
Haluk Aytak
Jason Williams
Ken Hirata
M. Tim Jones
Phillip Roberts
Ron Roberts
Steve Grorshe
Steve Johnson
Stillman Gates
Tom Grieff
Ralph Weber
Gil Romo

Meeting Notes:

Heng Liao and Tom Grief jointly presented 05-186r-1.

pp.1

Heng started with opening remarks about the objective of the presentation. The goal is to address the issues raised in May SAS WG meeting.

pp.2

Heng went through the list of key questions from last WG meeting including:

How to handle group reassignment based on device SAS address
Synchronized update of permission table
Clarification on impact of zone management update on current connection
and current open arbitration.

Is per Group Expander CHANGE counter needed to support zoning.

The committee agreed that these are important issues to look at, and the list did not overlook other important issue from last WG meeting
pp.3, 4

Heng proceed to describe the group assignment methods and the issue of group reassignment in case of a device go up and down, or move location.

The committee agreed that the group assignment based on Expander Phy (Port) is reasonable. They also agree that there need to be a need to handle group reassignment in the conditions identified:

- 1) The committee agreed that when a new device is added, the expander shall assign it to default group 0.
- 2) The committee agree that:

A) when a SAS device is removed from an expander Phy, and then reattached to the same expander Phy, the Expander shall detect the same SAS address has been attached based on IDENTIFY SAS address, and automatically restore the the previous PHY ZONE configuration of the expander Phy.

A question was raised about SATA device reassignment on the same expander Phy. Since SATA device does not have IDENTIFY frame, or even a WWN. After much discussion, the committee reaches the following consensus:

B) For SATA device, after SMP link reset or hardreset to the Expander Phy, causing the SATA device to lose and regain Phy_RDY. The expander shall automatically restore the PHY ZONE configuration of the expander Phy.

C) For SATA device, if the Expander Phy lost link to a SATA device, and later on regain link PHY_RDY not due to SMP link/hard resets conditions, the Expander Phy shall be set to default group 0 waiting for supervisor to determine whether the Phy should be assigned to the previous group, or a new group. Outside to the scope of SAS 2 specification, a supervisor (if it is STP initiator capable) could send ATA command the SATA device to obtain its WWN, or the ASCII model and serial number (and run the ASCII string through a hash function to generate a unique identifier), if the WWN or the hash identifier matches the previous identifier associated with the expander Phy, the supervisor could

send SMP commands to reassign the PHY ZONE configuration of the expander Phy.

D) If a native STP device is reattached, it shall be treated the same as any other SAS device.

3) The committee agreed that when a device is moved from one expander Phy to a different Expander Phy (of the same expander or a different expander within the topology), the Expander shall always assign it to group 0. The supervisor may reassign it to other groups after it detects the device has been attached, but this supervisor function is outside the scope of SAS 2 specification. Everyone agrees that the expander shall NOT restore the PHY ZONE configuration automatically without supervisor intervention in the case a device location move.

pp5, 6, 7

Heng describes the issue of synchronized updates among multiple supervisors. The PHY ZONE configuration updates do not need to be synchronized, as they are atomic. The PERMISSION TABLE updates need to be synchronized because each update may take multiple SMP commands, and the table needs to be propagated to multiple expanders.

Heng then introduced the idea of electing a “supervising” expander as the update proxy to synchronize operations among multiple expanders, and the process of electing a supervising expander. The idea was agreed by the committee.

Question was raised on corner cases about “supervising expander” failure in the middle of an update procedure. pp.7 provides the answer to this problem by requiring the supervisor to always redownload the permission table using the latest supervising expander address. everytime it detects an topology change that involves expanders. This ensures all expanders in the topology has a clean and consistent table from then on. The committee agrees it solves the corner cases problem identified.

pp8,9

Heng clarifies the impact of PHY ZONE changes and PERMISSION update on the data traffic and new open arbitration during the update period. The group agree with the conclusion that the new zone configuration may not take full effect until the end of the update cycle, this should be the expected behavior at the system level. Therefore, the expander shall not stop the ECM arbitration or tear down existing connection during the update cycle.

pp10

Tom use an example to show when the disruption free steps for zone management to accomplish adding a device and the permission for the new device to access members of an existing group.

pp11

Heng presented the change counter impact on Per PHY change counter, and per expander change counter. Two conclusions are reached.

1) PHY change counter do not need to be increased. The current PHY counter should be PHY, the SMP reporting of PHY change counter need to be masked by permission table based on SGID.

2) Expander change counter need to be maintain based on per source group basis. The approach suggested by Heng (by summing up PHY change counter to calculate expander change counter) may not work properly in the case of a PHY group ID reassigned to a different group ID. But the group agree that per source group change counter may be maintained in expander firmware.

pp12- 38

Heng ran through the modifications to the SMP frames. Some fields are rearranged for consistency. GENERATION CODE is removed from all commands.

REPORTE GENERAL function now returns the SUPERVISING EXPANDER ADDRESS.

The SUPERVISE THIS bit in SMP CONFIGURE PERMISSION TABLE is now changed to PROPAGETE UPDATE with different meaning to support supervising expander operation.

REPORT ZONE ROUTE TABLE now includes the ATTACHED DEVICE TYPE field.

Comments made by the group:

1) There is still one place where GENERATION CODE has not been removed.

2) Need to expander SAS 1.1 section 4.6.7.3 expander route table description to include the new fields required by expander zone route table. And incorporate the topology discovery process description based on SAS 1.1 section 4.7.

3) SOURCE CHECK should be removed from this proposal because it is not inherently part of zoning. The group agree that the SOURCE CHECK fields in SMP frames can be made reserved in this proposal to ensure the separate proposal for source check can later on can use the field positions reserved.

Summary:

In the end, the group agreed that provided the suggested changes listed above were made to the proposal, we have reached consensus on the architecture basis. The actual language in the proposal needs to be reworked significantly to fit in the T10 format and style. And further more, there is still work to be done in terms of describing the details of the affected areas of the specification, such as modification to state-machines. New issues may be uncovered as we make progress on filling in the details.

Action:

Heng Liao to revise the document to incorporate the suggested changes.

The next SAS zoning conference call is arranged for June 16, 1:00pm Pacific Time. Pat Thaler and Heng Liao to send out meeting notice to T10 mail list with conference call and WebEx information. The suggested agenda is for Heng to go over the revised document.