

To: T10 Technical Committee
 From: Jim Hatfield, Seagate (James.C.Hatfield@seagate.com),
 Date: Sept 2, 2007
 Subject: Minutes from Aug. 18, 2008 meeting: SAS Interface Power Management

Related Documents

08-015r4 - SAS-2+ Add low power transceiver options
 08-206rx - SAS-2+m low power transceiver options, phy and link states
 08-249rx - SAS 2.+ Link layer power management
 T10 Reflector msg from Gerry Houlder, dated Aug. 11 2008 3:11pm:
 "Announcement for Aug. 18 Interface power Management meeting:

Overview

On Aug. 18, 2008, a meeting was held (1pm - 5pm) to resolve open issues related to proposal 08-015r4.

Location:

Doubletree Hotel
 2050 Gateway Place
 San Jose, CA 95110

Attendees:

Name	Company
Jim Hatfield	Seagate
Dal Allan	ENDL
Dan Colegrove	Hitachi GST
Pak Seto	Intel
Kishore Karthikeyan	Intel
Kevin Marks	Dell
Mark Evans	Western Digital
Mike Fitzpatrick	Fujitsu
George Penokie	LSI
Fred Knight	NetApp
Curtis Stevens	Western Digital
Tim Symons	PMC-Sierra
Edward Chang	Samsung
Mr. Emami (was present but did not sign in)	NetApp

Proceedings:

1. Mark Evans, Fred Knight and George Penokie dominated a discussion about how to get to into the states that the proposals dealt with. This took most of the time allotted. I was unable to record the gist of the discussion due to the rapid flow and my unfamiliarity with the subject matter. Please ask Mark, Fred or George for details of the discussion.
2. Regarding the agenda question:
 Should a new OPEN REJECT primitive be created with a forced 5 or 10 ms retry time be created? This might reduce the number of unsuccessful retries (e.g., OPEN frame, OPEN REJECT (RETRY) response, repeat quickly) that could occur during slumber recovery. Disadvantage: older hosts with newer expander/target might not understand the new primitive and would treat it same as OPEN REJECT (RETRY) anyway.

The group suggested that this may not be an issue, and that a new primitive is not needed, but that a new SMP timer for OPEN REJECT may be preferable.

3. Regarding the agenda question

The proposal currently suggests ALIGN(0) and ALIGN(1) as the synchronization primitives in the PM recovery sequence. There was a suggestion in an earlier teleconference that a different pair of ALIGN primitives might be a better choice. Lets decide this.

The group suggested that the current ALIGN(0) and ALIGN(1) be used.

4. Regarding the agenda question:

The AIP (WAITING ON DEVICE) primitive has a "can only be sent one time" restriction on it and is only effective for 1 ms. Should the restriction be lifted so multiple AIPs can be sent? Decision has bearing on timeout interval for recovery from partial power condition

The group suggested that it is not a requirement to only send it one time.

5. Regarding the agenda question:

Does SATA interface power management for STP links need to be prohibited if SAS power management for SSP links is supported?

The group will consider allowing SATA power management for STP links.

6. Regarding the agenda question:

I have an issue with consequences to an expander if (1) target and expander agree on partial condition; (2) target is hot-plugged; (3) initiator tries to open the target; (4) how long of a "hot-plug timeout" can the expander tolerate before recognizing the target is gone and changing to a different response [e.g., OPEN REJECT (NO DESTINATION)]. It has been suggested that a regular hot-plug timeout interval (500 msec) be used for this but this is too long to respond with AIP (WAITING ON DEVICE) and might also be too long for OPEN REJECT (RETRY). Since the expander has two different responses (depending on whether partial or slumber is expected) perhaps there should be two different timeouts [e.g., 20 us and 20 ms respectively] that are more in tune with the responses the expander has to generate?

The group commented:

- a) There is a possible 500 msec worth of AIPs
- b) Will something else timeout first ?
- c) What about the upper levels of the system ?

7. Regarding 08-015r4:

a. figure New1

- i. contains multiple figures. please separate them
- ii. if this page is merely overview or informative material, consider deleting it.
- iii. this looks like OOB where there are no state machines, instead of the Link layer where there are state machines
- iv. see figure 154 of sas2r14c for a different type of diagram. these need to show max/min timing, vertical lines (e.g.'window labels') and there needs to be a table containing the actual numbers
- v. SAS phy x Tx/ s/B Phy A Tx
- vi. make the diagram look like it is not starting from D.C. idle
- vii. SAS phy y Rx s/b Phy B Rx

- viii. decide whether this is partial or slumber: have a specific case
 - ix. Bus Idle s/b D.C. Idle
 - b. text for figure new1
 - i. make match the figure changes
 - ii. add timer info as requested
 - c. figure new2
 - i. this is the part that overlaps with proposal 08-206
 - ii. the horizontal line extending to the right
 - 1. replace with "..."
 - iii. scrambled random data s/b dwords from link layer (ref fig 147 in SAS2r14c)
 - iv. what context is maintained in the phy, specifically in slumber ?
 - d. text for figure new2
 - i. add bullet 7): if ports are multiplexed
 - ii.
8. Regarding 08-206r0
- a. figure a
 - i. state: "SP31:PM_Recovery_Start"
 - 1. rename this state
 - 2. there needs to be a new idle state both before and after it
 - ii. note: SP15 detects the power mgmt request, and puts the device into the power mode, then comes to this state machine
 - iii. the recovery process is identical for partial and slumber - - only the timing varie
 - iv. entry from SP0 s/b entry from SP15
 - v.
9. Regarding 08-249r0:
- a. (page 7, in the list after "The SL receiver sends the following messages to the SL state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2)"
 - i. add ACK and NAK to the list
 - b. (page 10: "If a PMREQ Received message is received and power management is **enabled** (see xxx), then this state **shall** send a Transmit PMACK message to the SL transmitter. If a PMREQ Received message is received and power management is not enabled, then this state shall send a Transmit PMNAK message to the SL transmitter.")
 - i. consider changing the first 'shall' to "should"
 - ii. re: 'enabled':
 - 1. the scope is SL, not global scope.
 - 2. What is the default value ?
 - 3. In what layer does the default lie ?
 - 4. Need more qualifications for both enabled and disabled cases. (George will revise the paragraph)
 - c. (page 10: **7.14.4.2.2 Transition SL_CC0:Idle to SL_CC1:ArbSel**)
 - i. **In CC0:** need a signal going up from the cloud