

**From:** Gerry Houlder, Seagate Technology <gerry.houlder@seagate.com>  
**Subj:** Questions about SAS Interface Power Management proposal (08-015)  
**Date:** Sept. 8, 2008

---

## Overview

These questions were emailed to me from Amit Shah (Intel). These questions are intended for discussion during Sept. 8 SAS Protocol meeting.

Hi Gerry,

Attached are the agenda items for our meeting next week. Some of the initial bullets are from my comments that I had sent out couple of days back – as I want to discuss your views in detail. Let me know if you have any particular feedback on these agenda items before the meeting.

1. We support “No” Power Management (PM) within STP connection. D-SATA may support PM, but when we talk about STP, expander to drive link can go in PM state, but that does not mean that host to expander link should also go in PM State. For SAS, we should PM support only out of connection else we can open up the logic to a lot of complexity.
2. SATA 2.6 GOLD, section 6.7.3 talks about Presence detection. If we use similar methodology, we can know when a drive removal event occurred. Something to consider.
3. Less than 500msec to detect a hot plug event during PM. we would prefer 10 to 20 msec timeout.
4. D-SATA – For D-SATA, if the host transmits ComWake and if it does not get a ComWake back from drive in a specified interval time, it will re-transmit ComWake till a timeout occurs. Is this something to consider for SAS also?
5. What is the impact of this proposal to NOTIFY ENABLE SPIN-Up? Typically host or expanders would like to transmit NOTIFY ENABLE SPINUP in a periodic and staggered fashion to limit power usage surge during drive spin-up. But once the link is in power management state, what happens to SPINUP support?
  - a. There are two choices
  - b. (1) don't send NOTIFY to targets that are in power managed state [this would require target to not power manage the interface when it needs a NOTIFY in order to proceed]
  - c. (2) expander could wake up the interface, send a NOTIFY, power manage it again.
6. Once a PMREQ is received, how much time does a device / host have to respond iwth a PMACK or a PMNAK? Should we specify a limit of ay 50 DW time or something? IF we dont specify any time limit, will it not create issues?
7. How do we handle a conflict of PMREQ transmission and OAF coming in at the same time? Should we wait for PMNAK / PMACK or should we process OAF?
8. How will the management layer perform discovery process if the link is in power management state and so no BROADCAST CHANGE primitive can be received on the link?
  - a. management layer would have to wake up any sleeping links in order to send BROADCAST CHANGE primitives,
9. I think the 08-249 proposal covers this case. Receipt of OAF is treated as cancelling the PMREQ (equivalent to receiving PMNAK followed by OAF). If this isn't clear, we need to add this.
  - a. I thought that 08-249 is a parallel proposal to what you have. Is it that we have to review your proposal and consider 08-249 an addition to it?
10. Is there any SMP management function that can help bring a particular link (say a particular drive to expander link) out of power management state? Not sure if we need this functionality of not, but maybe useful

11. For 6G SAS PM, when we are coming out of PM state via ComWake exchange, we have to save DFE parameters so that we can achieve lock sooner. But still not sure about how much time it takes to lock? (cannot do it within 10us... it would need approx. 30us to lock)
12. For partial/slumber output are allowed to float in SAS.. Is this a requirement (shall) or a (may). For SATA, output is allowed to float in slumber and needs to be in common mode in partial state
13. How is PMREQ / PMREQ contention resolved for different combinations of Partial and Slumber.
14. Wide port power management -- PM is on a per phy basis...Basically I can power down 1 lane of 4 lanes in a wide port. Is there a consensus on this? There should explicit mention in the standard that there is no confusion on this.
15. Shouldn't the power management FEATURE be in the idf frame and not in SSP protocol specific mode Page. Because the mode page can be programmed by the host whereas the IDF is pre-programmed by the manufacturer..
16. Should a new OPEN REJECT primitive be created with a forced 5 or 10 ms retry time be created?
  - a. This is not required as it should be taken care by the other SAS feature related to avoidance of flooding of AWT retry
17. 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.
  - a. I would suggest any pattern that will allow 6G receivers to lock as fast as they can. If its ALIGN primitive its great. If there is any other kind of data or primitive sequence (like Train ) that can help receives lock fast, we should choose that. We need to be sensitive on ensuring there is a fast recovery time to ensure to enable extended usage of these modes like PCI-Ex.
18. 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.
19. Does SATA interface power management for STP links need to be prohibited if SAS power management for SSP links is supported?
20. Should targets be allowed to automatically change from partial to slumber power condition? I say no, because the expander has to know the target state so that appropriate initiator response to OPEN frame is made. If we made the expander response the same for both cases, perhaps this feature could be allowed. Any advantages/ disadvantage to this?
  - a. SATA is schedule to add this feature to squeeze additional power saving. So what if a SAS drive is directly connected to a initiator, we support this feature? If behind an expander, I agree there is some complication of supporting this feature.
21. How do we track the meeting minutes and associated changes to the proposal
22. Should we post part / all of this meeting minutes on T10 so that other people can comment on it?
23. Also are there other proposals that I need to review to get a good picture of what changes could come with SAS power management proposal?
24. When is the next telecom to discuss power management proposal / T10 direction on this?
25. What are other proposals and whats the current status on the direction?
26. What is the timeframe for finalizing the proposal for SAS 2.1?

Amit: Review 08-249 proposal ????

Thanks,  
Amit