

Date: February 5, 2003
To: T10 Committee Membership
From: Steve Johnson, LSI Logic
Subject: PHY CHANGE COUNT, SMP Discover response

Overview of issue

Currently, there are no SMP commands that the initiator can use to determine which PHY or PHY's caused a CHANGE broadcast if the PHY link rate has not changed. Initiators should be able to determine which command sets were cleared due to link resets and hard resets. SAS expanders have the opportunity to provide a simple and deterministic method using SMP functions to help initiators quickly and efficiently determine this.

Link Reset can occur on loss of dword sync and other causes. This will cause an expander to broadcast CHANGE. For SATA devices, any OOB is treated as a device reset and all commands are lost. With some SATA devices it means only one command is lost. When queuing is available, the entire queue is lost. Using current SMP functions there is no way for an initiator to determine what caused the CHANGE. The PHY CHANGE COUNT will essentially provide the SCSI unit attention behavior for SATA devices.

With SAS devices, a link reset is treated as a temporary loss of signal and does not clear the command set. A HARD RESET clears the command set for all initiators and the SAS devices should raise unit attention for all initiators. Only issuing a SCSI command and getting a unit attention returned will exactly determine this. Currently, an initiator will have to send a SCSI command (TUR) to each SAS device connected to the expander to determine which drive or drives caused the CHANGE and lost it's command set. The PHY CHANGE COUNT would help initiators quickly determine which drives to send a TUR to.

Again, using current SMP functions there is no means to determine which PHY or PHY's caused the CHANGE broadcast if the PHY link rate has not changed. The REPORT GENERAL has the EXPANDER CHANGE COUNT, the initiator can use this to know which expander broadcasted the change. There are no fields in the DISCOVER response and the REPORT PHY SATA that help. The REPORT ERROR LOG response does have the "LOSS OF DWORD SYNCHRONIZATION COUNT" but I don't think that was the intent of this request and it does not cover all the cases.

Changes required to DISCOVER Response

Add a PHY CHANGE COUNT field to byte 38 of the SMP DISCOVER response.

The PHY CHANGE COUNT field counts the number of BROADCAST (CHANGE)s originated by an expander PHY. Expander devices shall support this field. Other device types shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it transmits a BROADCAST (CHANGE) for one of the following reasons:

- a) after the expander phy has lost dword synchronization; or
- b) after the link reset sequence completes for the phy.

The expander device need not increment this field again unless a DISCOVER response is transmitted. The PHY CHANGE COUNT field shall wrap to zero after the maximum value (i.e., FFh) has been reached.

NOTE xx - Application clients that use the PHY CHANGE COUNT field should read it often enough to ensure that it does not increment a multiple of 256 times between reading the field.