

To: T10 Technical Committee
From: William Martin (bill.martin@emulex.com)
Date: 6 December 2007
Subject: 08-010r0 SAS-2 Enabling and disabling Transport Layer Retries

Problem statement:

With the introduction of multiplexing, there was an error introduced in the interaction between the SP and SL state machines. In the diagrams in 4.4.1 as well as in the state descriptions, the message to start the SL state machines is originated as soon as speed negotiation sequence is completed; however, this would cause MUX0/MUX1 to be sent at the same time as either the IDENTIFY frame or the HARD_RESET. It is not clear that the multiplexing of the transmitter signal in figure 31 resolves this issue. I would propose the following changes to the SAS2 standard to clarify the expected operation.

4.4.1 Reset overview

Figure 40 illustrates the reset terminology used in this standard:

- a) link reset sequence;
- b) phy reset sequence (see 6.7);
- c) SATA OOB sequence (see 6.7.2.1);
- d) SATA speed negotiation sequence (see 6.7.2.2);
- e) SAS OOB sequence (see 6.7.4.1);
- f) SAS speed negotiation sequence (see 6.7.4.2);
- g) hard reset sequence (see 7.9); and
- h) identification sequence (see 7.9).

Modify diagram 40 for the SAS diagram, there is a discrepancy between the SAS diagram and the SAS (with hard reset) diagrams. The first of these has SP enters SAS_PHY_Ready state after the Multiplexing sequence; however, since this sequence is part of the Phy reset sequence, the SAS (with hard reset) diagram would imply this state transition after the Multiplexing sequence has been transmitted.

The phy reset sequences, including the OOB sequence, ~~and~~ speed negotiation sequences, **and Multiplexing sequence** are implemented by the SP state machine and are described in 6.7 and 6.8. The hard reset sequence and identification sequence are implemented by the SL_IR state machine and are described in 7.9.

6.8.2 SP transmitter and receiver

The SP transmitter transmits OOB signals and dwords on the physical link based on messages from the SP state machine (see 6.8).

The SP transmitter receives the following messages from the SP state machine:

- a) Transmit COMINIT;
- b) Transmit COMSAS;
- c) Transmit COMWAKE;
- d) Transmit SATA Port Selection Signal;
- e) Transmit D10.2;
- f) Set Rate with a Physical Link Rate argument, and an SSC On or an SSC Off argument;
- g) Transmit ALIGN with an argument indicating the specific type (e.g., Transmit ALIGN (0));
- h) Transmit Phy Capabilities Bits;
- i) Transmit TRAIN Pattern;
- j) Transmit TRAIN_DONE Pattern; and
- k) Transmit MUX Sequence.

When not otherwise instructed, the SP transmitter transmits D.C. idle.

Upon receiving a Transmit MUX Sequence message, the SP transmitter transmits:

- 1) MUX (LOGICAL LINK 0);
- 2) MUX (LOGICAL LINK 1);
- 3) MUX (LOGICAL LINK 0);
- 4) MUX (LOGICAL LINK 1);
- 5) MUX (LOGICAL LINK 0); and
- 6) MUX (LOGICAL LINK 1).

The SP transmitter shall complete any physical link rate change requested with the Set Rate message within RCDT (see table 87 in 6.7.4.2).

The SP transmitter sends the following messages to the SP state machine:

- a) COMINIT Transmitted;
- b) COMSAS Transmitted;
- c) COMWAKE Transmitted;
- d) SATA Port Selection Signal Transmitted;
- e) TRAIN_DONE Pattern Transmitted; **and**
- f) Phy Capabilities Bits Transmitted; **and**
- g) **MUX Sequence Transmitted.**

6.8.4.1 SAS speed negotiation states overview

Figure 143 shows the SAS speed negotiation states, in which the phy has detected that it is attached to a SAS phy or expander phy rather than a SATA phy, and performs the SAS speed negotiation sequence. These states are indicated by state names with a prefix of SAS.

[In Figure 143, add Transmit MUX Sequence as an output of SP15 and MuX Sequence Transmitted as an input to SP15.](#)

6.8.4.9 SP15:SAS_PHY_Ready state

6.8.4.9.1 State description

This state waits for a COMINIT Detected message, a DWS Lost message, or a DWS Reset message.

Upon entry into this state, the phy shall:

- a) if multiplexing is enabled (see table 92 in 6.7.4.2.3.3),
 - 1) send a Transmit MUX Sequence message to the SP transmitter; **and**
 - 2) **when MUX Sequence Transmitted is received from the SP transmitter, then send a Phy Layer Ready (SAS) confirmation to the link layer to indicate that the physical link has been brought up successfully in SAS mode;**
- b) **if multiplexing is not enabled, then** send a Phy Layer Ready (SAS) confirmation to the link layer to indicate that the physical link has been brought up successfully in SAS mode;
- c) if the SP transmitter is transmitting at 1.5 Gbps, set the ResetStatus state machine variable to G1;
- d) if the SP transmitter is transmitting at 3 Gbps, set the ResetStatus state machine variable to G2; and
- e) if the SP transmitter is transmitting at 6 Gbps, set the ResetStatus state machine variable to G3.

While in this state dwords from the link layer are transmitted at the negotiated physical link rate at the rate established in the previous SNW.

If multiplexing is disabled, each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

NOTE 32 - If multiplexing is enabled and this state receives a DWS Lost message, this state does not send a Start DWS message and the state machine transitions to SP0:OOB_COMINIT.

(There are no changes proposed in these clauses. They are included for reference only)

7.9.4.3.2 SL_IR_TIR1:Idle state

7.9.4.3.2.1 State description

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL_IR transmitter.

7.9.4.3.2.2 Transition SL_IR_TIR1:Idle to SL_IR_TIR2:Transmit_Identify

This transition shall occur after both:

- a) a Phy Layer Ready (SAS) confirmation is received; and
- b) a Transmit IDENTIFY Address Frame request is received.

7.9.4.3.2.3 Transition SL_IR_TIR1:Idle to SL_IR_TIR3:Transmit_Hard_Reset

This transition shall occur after both:

- a) a Phy Layer Ready (SAS) confirmation is received; and
- b) a Transmit HARD_RESET request is received.

7.9.4.3.3 SL_IR_TIR2:Transmit_Identify state

7.9.4.3.3.1 State description

Upon entry into this state, this state shall send either one or three Transmit IDENTIFY Address Frame messages to the SL_IR transmitter.

NOTE 49 - Phys compliant with previous versions of this standard only transmitted one Transmit IDENTIFY Address Frame message.

After this state receives an IDENTIFY Address Frame Transmitted message in response to its first Transmit IDENTIFY Address Frame message, this state shall send an Identify Transmitted message to the SL_IR_IRC state machine.

7.9.4.3.3.2 Transition SL_IR_TIR2:Transmit_Identify to SL_IR_TIR4:Completed

If this state sends one Transmit IDENTIFY Address Frame message, this transition shall occur after sending an Identify Transmitted message to the SL_IR_IRC state machine.

If this state sends three Transmit IDENTIFY Address Frame messages, this transition shall occur after receiving three Identify Transmitted messages.

7.9.4.3.4 SL_IR_TIR3:Transmit_Hard_Reset state

7.9.4.3.4.1 State description

Upon entry into this state, this state shall send a Transmit HARD_RESET message to the SL_IR transmitter.

After this state receives a HARD_RESET Transmitted message, this state shall send a HARD_RESET Transmitted confirmation to the management application layer.

7.9.4.3.4.2 Transition SL_IR_TIR3:Transmit_Hard_Reset to SL_IR_TIR4:Completed

This transition shall occur after sending a HARD_RESET Transmitted confirmation to the management application layer.