

To: T10 SAS Protocol Working Group

From: Brian Day

Subject: SAS 2 : 08-425 TRAIN_DONE race condition fix

Revision History

Revision 0 - Initial draft

Revision 1 - Just added an editor comment that the state machine figure also needs to be updated to reflect the argument being passed between the SP29 and SP30 states.

Related Documents

sas2r14f - Serial Attached SCSI - 2 Draft revision 14f

Overview

As originally indicated on the T10 reflector by Bill Martin of Emulex (posted 10/27/2008), there is a race condition relative to when devices transition from SP29:SAS_Train to SP30:SAS_TrainingDone.

From the reflector email:

“There is an issue of devices not coming out of the train done sequence properly in the following case:

Device A and Device B begin transmitting the TRAIN pattern.

Device A locks on the TRAIN pattern and begins transmitting TRAIN_DONE

Device A completes sending 4 TRAIN_DONE patterns and continues sending TRAIN_DONE patterns waiting for the TRAIN_DONE primitive to be received

Device A sends TRAIN_DONE primitive sequence and one of the dword of the TRAIN_DONE pattern

Device B completes TRAIN and sends TRAIN_DONE

Device A detects TRAIN_DONE before completing the current TRAIN_DONE pattern and exits the TRAIN SNW (not sending any more TRAIN_DONE primitives

Device B never detects TRAIN_DONE primitive from device A”

This then results in Device A continuing to SP15:SAS_PHY_Ready, sending its IDENTIFY frame, timing out after 1ms, and going back to COMINIT.

Since there is some expectation that there will be devices in production that use the “current” training protocol with the existing race condition, this proposal adds a new receiver requirement that both fixes the problem above, and increases interoperability with these early production devices.

Proposed Changes

6.8.4.12 SP29:SAS_Train state

6.8.4.12.1 State description

Upon entry into this state, the phy shall:

- a) initialize and start the MTT timer;
- b) initialize and start the TLT timer;
- c) send a Start Training message to the SP receiver; and
- d) send a Start DWS message to the SP_DWS state machine.

This state shall repeatedly send Transmit TRAIN Pattern messages to the SP transmitter.

Each time this state receives a DWS Lost message, this state shall send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization.

If the MTT timer expires, then this state shall send an Abort Training message to the SP receiver.

A phy reset problem occurs if:

- a) the MTT timer expires; and
- b) the Commonly Supported Settings state machine variable does not contain additional commonly supported settings.

6.8.4.12.2 Transition SP29:SAS_Train to SP0:OOB_COMINIT

This transition shall occur after receiving a COMINIT Detected message.

Before the transition, this state shall set the ResetStatus state machine variable to UNKNOWN.

6.8.4.12.3 Transition SP29:SAS_Train to SP1:OOB_AwaitCOMX

This transition shall occur if a phy reset problem occurs.

Before the transition, this state shall set the ResetStatus state machine variable to PHY_RESET_PROBLEM.

6.8.4.12.4 Transition SP29:SAS_Train to SP28:SAS_TrainSetup

This transition shall occur if:

- a) the MTT timer expires; and
- b) the Commonly Supported Settings state machine variable contains additional commonly supported settings.

6.8.4.12.5 Transition SP29:SAS_Train to SP30:SAS_TrainingDone

This transition shall occur if:

- a) the TLT timer has not expired;
- b) this state receives a Training Completed message; ~~and~~
- c) dword synchronization is acquired; ~~and~~
- d) this state receives a TRAIN Received message or TRAIN_DONE Received message.

If a TRAIN_DONE Received message was received, then the transition shall include a TRAIN_DONE Received argument.

Editor's Note 0: Figure 168 needs to add the argument being passed between the SP29 and SP30 states.

6.8.4.13 SP30:SAS_TrainingDone state

6.8.4.13.1 State description

This state shall repeatedly send Transmit TRAIN_DONE Pattern messages to the SP transmitter.

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.

This state waits for the MTT timer to expire or a TRAIN_DONE Received message from the receiver.

This state shall send a Start SL_IR Receiver confirmation to the link layer when a TRAIN_DONE Received message is received.

A phy reset problem occurs if:

- a) TRAIN_DONE Received message is not received before the MTT timer expires; and
- b) the Commonly Supported Settings state machine variable does not contain additional commonly supported settings.

6.8.4.13.2 Transition SP30:SAS_TrainingDone to SP0:OOB_COMINIT

This transition shall occur after receiving:

- a) a DWS Lost message if this state does not send a Start DWS message; or
- b) a COMINIT Detected message.

Before the transition, this state shall set the ResetStatus state machine variable to UNKNOWN.

6.8.4.13.3 Transition SP30:SAS_TrainingDone to SP1:OOB_AwaitCOMX

This transition shall occur if a phy reset problem occurs.

Before the transition, this state shall set the ResetStatus state machine variable to PHY_RESET_PROBLEM.

6.8.4.13.4 Transition SP30:SAS_TrainingDone to SP28:SAS_TrainSetup

This transition shall occur if:

- a) the MTT timer expires; and
- b) the Commonly Supported Settings state machine variable contains additional commonly supported settings.

6.8.4.13.5 Transition SP30:SAS_TrainingDone to SP15:SAS_PHY_Ready

This transition shall occur if this state receives: at least four TRAIN_DONE Pattern Transmitted messages; ~~and~~ and:

- a) a TRAIN_DONE Received message ~~before the MTT timer expires;~~ or
- b) this state was entered with a TRAIN_DONE Received argument. ~~;~~