To: T10 Technical Committee
From: Steve Finch (steve.finch@st.com)
Date: June 5, 2006
Subject: 06-273r0 SAS2: Bus Inactivity Timeout Timer Is Broken

Revision history
Revision 0 (5 June 2005) First revision

Related documents
sas2r04 - Serial Attached SCSI - 2 (SAS-2) revision 4
05-305r9 - SAS-2 Maximum SMP connection time

Overview

When sas2r00 was created base upon sas1 and approved changes, I believe we accidentally broke the Bus Inactivity Timeout Timer.

In SAS1, the Bus Inactivity Timeout timer used to start at the first reception of ACK/NAK. The phy wouldn't start the timer if it did not send a frame (because it never receives an ACK or NAK).

I gather the theory was that if this phy initiate a connection then it would have a frame to send and the ACK or NAK would have started the Bus Inactivity Timeout timer the first time after the frame was sent.

But if a phy was the recipient of the connection open request, then the other side must have had something to send and it would have started timers to catch any issues. So the recipient's side would only start the time if it had something to send.

Proposal 05-305r0 made a change. Now the Bus Inactivity Timeout timer simply starts at the beginning of the connection, and both phys can decide the connection has been inactive too long and CLOSE the connection.

The problem is that, as the text now reads, the timer is stopped and initialized when the PL_PM state machine receives a Tx Frame message. It is not restarted.

Take the sequence that both sides of the connection, after the connection is established, send one frame and each receives an ACK. As it stands now, there is no timer running. Yes, someone should be closing the connection, but the PL_OC has a lot of "mays" and "shoulds" associated with the sending of a Close Connection message.

What should happen is a Bus Inactivity Timeout.

I believe this problem came in when 05-305r0 changed the paragraph in section 8.2.3.4.1 that used to read:

"If a Bus Inactivity Time Limit timer has been created and this state receives an ACK Received or NAK Received confirmation, then this state shall start the Bus Inactivity Time Limit timer. If the Bus Inactivity Time Limit timer expires before this state receives a Tx Frame message, then this state shall send a Close Connection request to the link layer."

was edited and became:

"If a Bus Inactivity Time Limit timer expires then:

a) if the connection is SSP and there is no Tx Frame request outstanding (i.e., this state is not waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer; and
b) if the connection is SSP and there is a Tx Frame request outstanding (i.e., this state is waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer after receiving an ACK Received or NAK Received confirmation."

It interesting to note that case b) above can not happen because the timer was stopped when the frame was sent. Case a) can happen only if this phy never sends a frame during the current connection.

Proposal

I think it the original first sentence should be re-instated with some minor modifications so that the final text in this area should read:

"If a Bus Inactivity Time Limit timer has been created, this state receives an ACK Received or NAK Received confirmation and there are no additional Tx Frame requests outstanding, then this state shall start the Bus Inactivity Time Limit timer.

If a Bus Inactivity Time Limit timer expires then:

a) if the connection is SSP and there is no Tx Frame request outstanding (i.e., this state is not waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer; and
b) if the connection is SSP and there is a Tx Frame request outstanding (i.e., this state is waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer after receiving an ACK Received or NAK Received confirmation."

With this change, the Bus Inactivity Timer would be initialized and started on entry to this state. It will be stopped and re-initialized when a frame is transmitted. It would be re-started when upon receipt of an ACK or NAK and there are no other outstanding requests outstanding, i.e., the bus is inactive.