To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@intel.com)
Date: 11 August 2006
Subject: 06-371r0 SAS2: Transport-layer handling of extraneous retransmit DATA frames

Revision history

Revision 0 (20 July 2006) First revision

Related documents

SAS-2-r05a - Serial Attached SCSI - 2 (SAS-2) revision 05a

Overview

While SAS-2 defines clearly how the transport layer handles (ST_ITS) receipt of retransmitted DATA frames when previous data frames have been lost, it does not specify how the transport layer handles receipt of retransmitted data frames when the original data frames have been received and processed without error (i.e., the ACK was lost). This proposal adds text to the ST_ITS6:Receive_Data_In state to define how this is handled, and allows the ST_ITS to either discard the previously received read DATA frame and replace it with the retry DATA frame, or to discard the retry read DATA frame if the read DATA frame it would replace was received without error.

This proposal also addresses an inconsistency in the transport layer definition of the ACK/NAK balance point where the target (ST_TTS) state machine may choose to start transmitting retry read DATA frames.

Subclause 9.2.4.5.2 reads as follows:

> **9.2.4.5.2 DATA frame with transport layer retries enabled**
>
> If an SSP target port transmits a read DATA frame and receives a NAK for that frame, then the read DATA frame was not received. The SSP target port retransmits, in the same or in a new connection, all the read DATA frames since ***a*** previous time when ACK/NAK balance occurred (see 9.2.6.3.3.4).
>
> If an SSP target port transmits a read DATA frame and does not receive an ACK or NAK for that frame (e.g., times out, or the connection is broken):
>
> 1) the SSP_TF state machine closes the connection with DONE (ACK/NAK TIMEOUT) (see 7.16.8.6.5); and
> 2) the ST_TTS state machine retransmits, in a new connection, all the read DATA frames since a previous time when ACK/NAK balance occurred (see 9.2.6.3.3.4).

This indicates the ACK/NAK balance point chosen does not have to be the most recent ACK/NAK balance point. However the text describing the retransmission of frames in the ST_TTS (target transport server) state machine requires setting of the Balance Point Read Data Offset state machine variable so as to select only the most recent ACK/NAK balance point for retransmitting read DATA frames. This proposal provides text that makes it possible for the ST_TTS state machine to choose an earlier ACK/NAK balance point.

Suggested changes

Modify subclause 9.2.6.2.3.7.1 as follows:

9.2.6.2.3.7 ST_ITS6:Receive_Data_In state

9.2.6.2.3.7.1 State description

If this state receives a Data-In Arrived message, then this state shall verify the values in the read DATA frame received with the message as defined in table 1.

If the verification fails, then this state sends the Reception Complete message specified in table 1 to the ST_IFR state machine.

If:

a) transport layer retries are enabled;
b) the CHANGING DATA POINTER bit is set to zero;

Table 1 — Reception Complete messages for read DATA frame verification failures

Message sent to ST_IFR [a]	Condition
Reception Complete (Data Offset Error)	Transport layer retries are disabled, and the DATA OFFSET field in the read DATA frame is not equal Data-In Buffer Offset state machine variable.
	The DATA OFFSET field in the read DATA frame is greater than the Data-In Buffer Size state machine argument.
Reception Complete (Too Much Read Data)	The number of bytes in the DATA field in the read DATA information unit plus the Data-In Buffer Offset state machine variable is greater than the Data-In Buffer Size state machine argument.
Reception Complete (Incorrect Data Length)	The number of bytes in the DATA field in the read DATA information unit is zero.

[a] If more than one condition is true, then this state shall select which message to send to the ST_IFR state machine using the following order:
1) Reception Complete (Data Offset Error);
2) Reception Complete (Too Much Read Data); or
3) Reception Complete (Incorrect Data Length).

c) the DATA OFFSET field is not set to the Data-In Buffer Offset state machine variable;
d) the DATA OFFSET field is less than the Data-In Buffer Size state machine argument; and
e) the DATA OFFSET field plus the length of the Data information unit is less than or equal to the Data-In Buffer Size state machine argument,

then this state should discard all Data-In Arrived messages until a read DATA frame is received in which the CHANGING DATA POINTER bit is set to one. This state shall resume processing additional Data-In Arrived messages when it receives a Data-In Arrived message with the CHANGING DATA POINTER bit set to one.

If:

a) transport layer retries are enabled;
b) the CHANGING DATA POINTER bit is set to one;
c) the value of DATA OFFSET field is set to the value of the Data-In Buffer Offset state machine variable;
d) the value of DATA OFFSET field is less than the the value of the Data-In Buffer Size state machine argument; and
e) the value of DATA OFFSET field plus the length of the Data information unit is less than or equal to the value of the Data-In Buffer Size state machine argument;

then this state may either:

a) discard Data-In Arrived messages where data in the INFORMATION UNIT field matches (i.e., based on the DATA OFFSET field and the number of bytes in the DATA field) data received and processed in previously received DATA frames for which this state has transmitted an ACK; or
b) discard previously received DATA frames where data in the INFORMATION UNIT field matches (i.e., based on the DATA OFFSET field and the number of bytes in the DATA field) data received and processed in previously received DATA frames for which this state has transmitted an ACK;

and this state shall forward the undiscarded DATA frames to the application client.

If the verification succeeds or after this state resumes processing Data-In Arrived messages, then this state shall process the data received in the read DATA frame and set the Data-In Buffer Offset state machine variable to the DATA OFFSET field plus the length of the Data information unit.

If this state receives Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK), then this state shall send a Reception Complete (Command Failed, Connection Failed) to the ST_IFR state machine.

After this state sends a Reception Complete (Command Failed, Connection Failed) message, this state shall continue processing messages and confirmations.

NOTE 1 - The application client may determine the command was received and is being processed by the device server and allow the command to complete.

If this state receives a Cancel message, then this state shall send a Reception Complete (Cancel Acknowledged) message to the ST_IFR state machine. The Reception Complete message shall include the tag as an argument.

NOTE 2 - The Cancel message results from a vendor-specific request from the SCSI application layer after the SCSI application layer has used a task management function to determine that the SAS target port did not receive the COMMAND frame.

Editor's Note 1: The remaining subclauses in 9.2.6.2.3.7 remain unchanged.

Modify subclause 9.2.6.3.3.3 as follows:

9.2.6.3.3.3 ST_TTS2:Target_Send_Frame state

If this state is entered from the ST_TTS3:Prepare_Data_In state for transmission of a read DATA frame, then this state shall send a Transmit Frame (Non-Interlocked) request to the port layer.

If this state is entered from the ST_TTS4:Prepare_Xfer_Rdy state for transmission of an XFER_RDY frame, then this state shall send a Transmit Frame (Interlocked) request to the port layer.

If this state is entered from the ST_TTS6:Prepare_Response state for transmission of a RESPONSE frame, then this state shall send a Transmit Frame (Interlocked) request to the port layer.

All Transmit Frame requests from this state shall include the read DATA frame from the ST_TTS3:Prepare_Data_In state, the XFER_RDY frame from the ST_TTS4:Prepare_Xfer_Rdy state, or the RESPONSE frame from the ST_TTS6:Prepare_Response state and the following arguments to be used for any OPEN address frame:

a) initiator port bit set to zero;
b) protocol set to SSP;
c) Connection Rate argument;
d) Initiator Connection Tag argument;
e) Destination SAS Address argument; and
f) Source SAS Address argument.

After sending a Transmit Frame request this state shall wait to receive a Transmission Status confirmation.

If the confirmation or argument is Transmission Status (I_T Nexus Loss), then this state shall send a Transmission Complete (I_T Nexus Loss) message to the ST_TFR state machine. The Transmission Complete message shall include the tag as an argument.

If the confirmation or argument is not Transmission Status (Frame Transmitted) or Transmission Status (I_T Nexus Loss), then this state shall send the Transmission Complete message defined in table 149 to the ST_TFR state machine. The message shall include the following arguments:

a) tag; and
b) arguments received with the Transmission Status confirmation.

If the confirmation is Transmission Status (Frame Transmitted) and the Transmit Frame request was for:

a) an XFER_RDY frame; or
b) a RESPONSE frame,

then this state shall wait to receive one of the following confirmations:

a) Transmission Status (ACK Received);
b) Transmission Status (NAK Received);

c) Transmission Status (ACK/NAK Timeout); or
d) Transmission Status (Connection Lost Without ACK/NAK).

If the confirmation or argument is Transmission Status (Frame Transmitted), the Transmit Frame request was for a read DATA frame, and the Read Data Offset state machine variable is equal to the Request Byte Count Data-In argument, then this state shall wait to receive:

a) Transmission Status (ACK Received) confirmations or arguments for each outstanding read DATA frame; or
b) one of the following:
   A) Transmission Status (NAK Received);
   B) Transmission Status (ACK/NAK Timeout); or
   C) Transmission Status (Connection Lost Without ACK/NAK).

NOTE 3 - If the number of data bytes that have been transmitted for a Request (Send Data-In) message are fewer than the Data-In Request Byte Count argument, then this state transitions to the ST_TTS3:Prepare_Data_In state to construct the additional read DATA frames for the request before receiving a Transmission Status (ACK Received), Transmission Status (NAK Received), Transmission Status (ACK/NAK Timeout), or Transmission Status (Connection Lost Without ACK/NAK) confirmation.

When the number of Transmission Status (Frame Transmitted) confirmations for Transmit Frame (Non-Interlocked) requests equals the number of Transmission Status (ACK Received) confirmations and the Transmit Frame request was for a read DATA frame, ~~this state shall set the Balance Point Read Data Offset state machine variable to the current Read Data Offset state machine variable.~~ if the Balance Point Read Data Offset state machine variable is set to zero or is equal to the Request Byte Count Data-In argument, then this state shall set the Balance Point Read Data Offset state machine variable to the current Read Data Offset state machine variable. If the Balance Point Read Data Offset state machine variable is not set to zero and is not equal to the Request Byte Count Data-In argument, then this state shall either:

a) set the Balance Point Read Data Offset state machine variable to the current Read Data Offset state machine variable; or
b) not modify the Balance Point Read Data Offset state machine variable.

~~If a Transmission Status (ACK Received) confirmation is received, and the Transmit Frame request was for a read DATA frame, then this state shall set the Balance Point Read Data Offset state machine variable to the current balance point read data offset plus the number of read data bytes transmitted in the read DATA frame associated with Transmission Status (ACK Received) confirmation.~~

If:

a) a Transmission Status (ACK Received) confirmation is received;
b) the Transmit Frame request was for a read DATA frame; and
c) the current balance point read data offset is zero;

then this state shall set the Balance Point Read Data Offset state machine variable to the number of read data bytes transmitted in the read DATA frame associated with Transmission Status (ACK Received) confirmation.

If:

a) a Transmission Status (ACK Received) confirmation is received;
b) the Transmit Frame request was for a read DATA frame; and
c) the current balance point read data offset plus the number of read data bytes transmitted in the read DATA frame associated with Transmission Status (ACK Received) confirmation is equal to the request byte count data-in;

then this state shall set the Balance Point Read Data Offset state machine variable to the request byte count data-in.

If:

a) a Transmission Status (ACK Received) confirmation is received;
b) the Transmit Frame request was for a read DATA frame;
c) the current balance point read data offset is not zero; and

d) the current balance point read data offset plus the number of read data bytes transmitted in the read DATA frame associated with Transmission Status (ACK Received) confirmation is not equal to the request byte count data-in;

then this state shall either:

a) set the Balance Point Read Data Offset state machine variable to the current balance point read data offset plus the number of read data bytes transmitted in the read DATA frame associated with Transmission Status (ACK Received) confirmation; or
b) not modify the Balance Point Read Data Offset state machine variable.

If transport layer retries are enabled, the Transmit Frame request was for a RESPONSE frame, the vendor-specific number of retries has not been reached, and this state receives one of the following confirmations:

a) Transmission Status (NAK Received);
b) Transmission Status (ACK/NAK Timeout); or
c) Transmission Status (Connection Lost Without ACK/NAK),

then this state shall:

a) set the RETRANSMIT bit to one; and
b) resend a Transmit Frame (Interlocked) request to the port layer for the failed RESPONSE frame.

If transport layer retries are enabled, the Transmit Frame request was for a XFER_RDY frame, the vendor-specific number of retries has not been reached, and this state receives one of the following confirmations:

a) Transmission Status (NAK Received);
b) Transmission Status (ACK/NAK Timeout); or
c) Transmission Status (Connection Lost Without ACK/NAK),

then this state shall:

a) set the RETRANSMIT bit to one;
b) set the value in the TARGET PORT TRANSFER TAG field to a value that is different than the target port transfer tag in the previous XFER_RDY frame associated with the Data-out arguments and is different than any other target port transfer tag currently in use. If write data is received for a subsequent XFER_RDY frame for a command, then all target port transfer tags used for previous XFER_RDY frames for the command are no longer in use; and
c) resend a Transmit Frame (Interlocked) request to the port layer for the failed XFER_RDY frame.

Editor's Note 2: The remainder of subclause in 9.2.6.3.3.3 is unchanged.
