To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@intel.com)
Date: 12 January 2007
Subject: 06-371r1 SAS2: ST_TTS retransmitted DATA frames

Revision history
Revision 0 (20 July 2006) First revision
Revision 1 (12 January 2007) Incorporate feedback from September & November 2006 SAS WG.

Related documents
SAS-2-r07: Serial Attached SCSI - 2 (SAS-2) revision 07
06-467r1: SAS-2: Initiator handling of retransmit read DATA frames
06-470r2 SAS-2: Transport layer read data flowchart

Overview

This proposal originally addressed both transport-layer retries as it applies to the ST_ITS state machine receipt of transmitted DATA frames during transfers-in as well as the ST-TTS state machine selection of the data offset for retry DATA frames transmitted from an SSP target port to an SSP initiator port. The first part, affecting the ST_ITS state machine, was moved into a new proposal, 06-467. Only the portion affecting the ST-TTS state machine remains in this proposal.

This proposal 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 9.2.6.3.3 as follows:

9.2.6.3.3 ST_TTS (target transport server) state machine

9.2.6.3.3.1 ST_TTS state machine overview

The ST_TTS state machine performs the following functions:

a) receives and processes messages from the ST_TFR state machine;
b) sends messages to the ST_TFR state machine;
c) communicates with the port layer using requests and confirmations regarding frame transmission; and
d) receives HARD_RESET Received confirmations from the port layer.

This state machine consists of the following states:

a) ST_TTS1:Target_Start (see 9.2.6.3.3.2) (initial state);
b) ST_TTS2:Target_Send_Frame (see 9.2.6.3.3.3);
c) ST_TTS3:Prepare_Data_In (see 9.2.6.3.3.4);
d) ST_TTS4:Prepare_Xfer_Rdy (see 9.2.6.3.3.5);
e) ST_TTS5:Receive_Data_Out (see 9.2.6.3.3.6); and
f) ST_TTS6:Prepare_Response (see 9.2.6.3.3.7).

This state machine shall start in the ST_TTS1:Target_Start state after power on.

If this state machine receives a HARD_RESET Received confirmation, then this state machine shall transition to the ST_TTS1:Target_Start state.

The state machine shall maintain the state machine variables defined in table 1.

Table 1 — ST_TTS state machine variables

State machine variable	Description
Read Data Offset	Offset into the application client buffer for read data
Balance Point Read Data Offset	Offset into the device server buffer for read data of last point at which the number of Transmission Status (ACK Received) confirmations or arguments equals the number of transmited read DATA frames ~~DATA frame to have received an ACK Transmitted confirmation~~
Frames Transmitted Count	The number of Transmission STatus (Frame Transmitted) confirmations received
Ack Received Count	The number of Transmission Status (ACK Received) confirmations received.
Requested Write Data Offset	Device server requested offset in the application client buffer for write data
Requested Write Data Length	Amount of write data requested by the Device server from the application client buffer

This state machine shall maintain the state machine arguments defined in table 2.

Table 2 — ST_TTS state machine arguments

State machine argument	Description
Data-In	The Data-In arguments received in the Request (Send Data-In) message
Data-Out	The Data-Out arguments received in the Request (Receive Data-Out) message

9.2.6.3.3.2 ST_TTS1:Target_Start state

9.2.6.3.3.2.1 State description

This state is the initial state of the ST_TTS state machine.

Upon entry into this state, this state shall set the Read Data Offset state machine variable to the Data-In application client buffer offset argument ~~zero~~.

Upon entry into this state, this state shall set the Balance Point Read Data Offset state machine variable to the Data-In application client buffer offset argument ~~zero~~.

Upon entry into this state, this state shall set the Frames Transmitted Count state machine variable to zero.

Upon entry into this state, this state shall set the Requested Write Data Offset state machine variable to zero.

If this state was entered without an Enable First Burst argument, then the Requested Write Data Length state machine variable shall be set to the Request Byte Count Data-Out state machine argument.

If this state was entered with an Enable First Burst argument, then the Requested Write Data Length state machine variable shall be set to the First Burst Size argument.

9.2.6.3.3.2.2 Transition ST_TTS1:Target_Start to ST_TTS3:Prepare_Data_In

This transition shall occur after this state receives a Request (Send Data-In) message.

9.2.6.3.3.2.3 Transition ST_TTS1:Target_Start to ST_TTS4:Prepare_Xfer_Rdy

If this state was entered without an Enable First Burst argument, then this transition shall occur after a Request (Receive Data-Out) message is received.

9.2.6.3.3.2.4 Transition ST_TTS1:Target_Start to ST_TTS5:Receive_Data_Out

If this state was entered with an Enable First Burst argument, then this transition shall occur after a Request (Receive Data-Out) message is received.

9.2.6.3.3.2.5 Transition ST_TTS1:Target_Start to ST_TTS7:Prepare_Response

This transition shall occur after this state receives a Request (Send Transport Response) message.

The transition shall include the Transport Response arguments.

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 3 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 is Transmission Status (Frame Transmitted) and the Transmit Frame request was for a read DATA frame, then this state shall increment the Frames Transmitted Count state machine variable by one, and set the Read Data Offset state machine variable to the current the value of the current Read Data Offset state machine variable plus the number of read data bytes in the transmitted read DATA frame.

If the confirmation or argument is Transmission Status (ACK Received) and the Transmit Frame request was for a read DATA frame, then this state shall increment the ACK Received Count state machine variable by one.

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 (i.e., Frames Transmitted Count state machine variable equals the ACK Received Count state machine variable); 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 1 - 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.~~

When the Frames Transmitted Count state machine variable equals the ACK Received Count state machine variable and the Transmit Frame request was for a read DATA frame, this state may set the Balance Point Read Data Offset state machine variable to the current Read Data Offset state machine variable.

Editor's Note 1: "may" is proposed instead of "shall" to give the ST-TTS the option of using an arbitrary ACK/NAK balance point. The effect on the flow chart proposed in 06-470r2 is to make execution of the box in the lower-right corner of figure two (i.e., "Balance Point Read Data Offset = Point Read Data Offset") optional.

~~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 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.

Table 3 defines the messages that this state shall send to the ST_TFR state machine upon receipt of the listed confirmations and arguments, based on the conditions under which each confirmation or argument was received.

Table 3 — Messages sent to the ST_TFR state machine

Confirmation received from the port layer or argument received from ST_TTS3	Conditions under which confirmation was received	Message sent to the ST_TFR state machine
Transmission Status (ACK Received)	The Transmit Frame request was for an XFER_RDY frame.	Transmission Complete (Xfer_Rdy Delivered) with a Target Port Transfer Tag argument
	Transmit Frame request was for a RESPONSE frame	Transmission Complete (Response Delivered)
	The Transmit Frame request was for a read DATA frame and: a) the Read Data Offset state machine variable is equal to the Data-In Request Byte Count argument; and b) the Read Data Offset state machine variable is equal to the Balance Point Read Data Offset state machine variable.	Transmission Complete (Data-In Delivered)
Transmission Status (NAK Received), Transmission Status (ACK/NAK Timeout), or Transmission Status (Connection Lost Without ACK/NAK)	The Transmit Frame request was for a RESPONSE frame and the vendor-specific number of retries has been reached.	Transmission Complete (Response Failed)
Transmission Status (NAK Received)	The Transmit Frame request was for an XFER_RDY frame and: a) if transport layer retries are disabled; or b) if transport layer retries are enabled and the vendor-specific number of retries has been reached.	Transmission Complete (Xfer_Rdy Failed, NAK Received)
Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK)		Transmission Complete (Xfer_Rdy Failed, Connection Failed)
Transmission Status (NAK Received)	The Transmit Frame request was for a read DATA frame and: a) if transport layer retries are disabled; or b) if transport layer retries are enabled and the vendor-specific number of retries has been reached.	Transmission Complete (Data-In Failed, NAK Received)
Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK)		Transmission Complete (Data-In Failed, Connection Failed)

If this state receives a Cancel message or a Cancel argument and this state has received confirmations for all Transmit Frame requests sent to the port layer, then this state shall send a Transmission Complete (Data Transfer Terminated) message to the ST_TFR state machine.

If this state receives a Cancel message or a Cancel argument and this state has not received confirmations for all Transmit Frame requests sent to the port layer, then this state shall send a Cancel request to the port layer to cancel previous Transmit Frame requests. The Cancel request shall include the following arguments:

a) destination SAS address; and
b) tag.

Upon receipt of a Transmission Status (Cancel Acknowledged) confirmation or argument this state shall send a Transmission Complete (Data Transfer Terminated) message to the ST_TFR state machine.

A Transmission Complete message to the ST_TFR state machine shall include the following arguments:

a) destination SAS address; and
b) tag.

9.2.6.3.3.3.1 Transition ST_TTS2:Target_Send_Frame to ST_TTS1:Target_Start

This transition shall occur after this state sends a Transmission Complete message other than Transmission Complete (Xfer_Rdy Delivered) to the ST_TFR state machine.

9.2.6.3.3.3.2 Transition ST_TTS2:Target_Send_Frame to ST_TTS3:Prepare_Data_In

This transition shall occur after this state receives a Transmission Status (Frame Transmitted) confirmation for a read DATA frame and Read Data Offset state machine variable is less than the Request Byte Count Data-In argument.

If transport layer retries are enabled and the vendor-specific number of retries, if any, for the read DATA frame has not been reached, this transition shall occur and include a Retry argument after this state receives one of the following confirmations for a read DATA frame:

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

9.2.6.3.3.3.3 Transition ST_TTS2:Target_Send_Frame to ST_TTS5:Receive_Data_Out

This transition shall occur after this state sends a Transmission Complete (Xfer_Rdy Delivered) message to the ST_TFR state machine.

9.2.6.3.3.4 ST_TTS3:Prepare_Data_In state

9.2.6.3.3.4.1 State description

This state retrieves the data from the Device Server Buffer and constructs a read DATA frame.

This state shall construct a read DATA frame using the Data-In arguments as follows:

a) FRAME TYPE field set to 01h (i.e., DATA frame);
b) HASHED DESTINATION SAS ADDRESS field set to the hashed value of the Destination SAS Address Data-In argument;
c) HASHED SOURCE SAS ADDRESS field set to the hashed value of the SSP target port's SAS address;
d) RETRY DATA FRAMES bit set to zero;
e) RETRANSMIT bit set to zero;
f) CHANGING DATA POINTER set as specified in this subclause;
g) NUMBER OF FILL BYTES field set to the number of fill bytes needed for the specified read data;
h) TAG field set to the Tag Data-In argument;
i) TARGET PORT TRANSFER TAG field set to zero;
j) DATA OFFSET field set as specified in this subclause;
k) in the information unit, DATA field set as specified in this subclause; and
l) fill bytes, if required.

If this state is entered without a Retry argument then this state shall:

a) set the CHANGING DATA POINTER bit set to zero;

b) set the DATA OFFSET field to the Read Data Offset state machine variable; and
c) in the information unit, DATA field set to the information that starts at the location in the specified device server buffer pointed to by the Read Data Offset state machine variable and shall contain the amount of data that is the lesser of:
   A) the Data-In Request Byte Count argument minus the Read Data Offset state machine variable; and
   B) the maximum size of the read Data information unit for this Data-In request.

If this state is entered with a Retry argument then this state shall:

a) set the CHANGING DATA POINTER bit in the frame to one;
b) set the Read Data Offset state machine variable to the Balance Point Read Data Offset state machine variable;
c) set the DATA OFFSET field to Balance Point Read Data Offset state machine variable; and
d) in the information unit, DATA field set to the information that starts at the location in the specified device server buffer pointed to by the Balance Point Read Data Offset state machine variable and shall contain the amount of data that is the lesser of:
   A) the Data-In Request Byte Count argument minus the Balance Point Read Data Offset state machine variable; and
   B) the maximum size of the read Data information unit for this Data-In request;

   or

a) set the CHANGING DATA POINTER bit in the frame to one;
b) set the Read Data Offset state machine variable to the Data-In application client buffer offset argument;
c) set the DATA OFFSET field to the Data-In application client buffer offset argument; and
d) in the information unit, DATA field set to the information that starts at the offset location the Data-In application client buffer offset argument in the specified device server buffer and shall contain the amount of data that is the lesser of:
   A) the Data-In Request Byte Count argument; and
   B) the maximum size of the DATA information unit for this Data-In request.

Editor's Note 2: With the "may" suggested in changes noted in subclause 9.2.6.3.3.3 (see Editor's Note 1:), what follows the "or" here may not be necessary. It's equivalent to the ST_TTS2 state *always* choosing not to update the Balance Point Read Data Offset state machine variable.

~~If a confirmation of Transmission Status (Frame Transmitted) is received, then this state shall set the Read Data Offset state machine variable to the current read data offset plus the number of read data bytes in the transmitted read DATA frame.~~

Editor's Note 3: Was there a reason the update of the Read Data Offset state machine variable was originally placed AFTER the earlier comparisons with the Balance Point Read Data Offset?

~~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.~~

9.2.6.3.3.4.2 Transition ST_TTS3:Prepare_Data_In to ST_TTS2:Target_Send_Frame

This transition shall occur after this state:

a) constructs a read DATA frame; or
b) receives a Cancel message.

This transition shall include the received Transmission Status, if any, as an argument and the:

a) read DATA frame as an argument; or
b) if a Cancel message was received, then a Cancel argument.

The remainder of subclause 9.2.6.3.3 remains unchanged