

To: T10 Technical Committee
 From: Bob Sheffield (Robert.L.Sheffield@intel.com)
 Date: 27 April 2007
 Subject: 07-166r1: SAS-2 Clarify scope of retransmitted XFER_RDY & RESPONSE

Revision history

Revision 0 (3 April 2007) First revision

Revision 1 (27 April 2007) Incorporate feedback from April 07 SAS WG. Require retransmitted XFER_RDY to match previous XFER_RDY (requested offset, length, etc.), and extend to include RESPONSE frames.

Related documents

SAS-2-r09 - Serial Attached SCSI - 2 (SAS-2) revision 09

Overview

It is the intent but not clearly stated in the subclause describing retransmission of XFER_RDY frames that an SSP target port is only permitted to retransmit an XFER_RDY frame that corresponds to the immediately preceding XFER_RDY frame that did not receive an ACK (i.e., either received a NAK or the connection was closed with an ACK/NAK timeout before receiving an ACK or NAK for that frame). This proposal makes it clear that the retransmitted XFER_RDY frame is identical to the previous one that didn't receive an ACK or NAK, except that the RETRANSMIT bit is set to one and it has a different value in the TARGET PORT TRANSFER TAG field. The same rule is applied to retransmitted RESPONSE frames (rev-1 of the proposal).

Suggested changes

Modify subclause 9.2.4.4.2 as follows:

9.2.4.4.2 XFER_RDY frame with transport layer retries enabled

If an SSP target port transmits an XFER_RDY frame and receives a NAK for that frame, the SSP target port retransmits, in the same or a new connection, the XFER_RDY frame with a different value in the TARGET PORT TRANSFER TAG field, ~~and~~ with the RETRANSMIT bit set to one, and with the other fields set to the same values as the original XFER_RDY frame (see 9.2.6.3.3.3) (~~see 9.2.6.3.3.5).~~

If an SSP target port transmits an XFER_RDY 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 SSP target port retransmits, in a new connection, the XFER_RDY frame with a different value in the TARGET PORT TRANSFER TAG field, ~~and~~ with the RETRANSMIT bit set to one, and with the other fields set to the values of the corresponding fields in the original XFER_RDY frame (see 9.2.6.3.3.3) (~~see 9.2.6.3.3.5).~~

For a given I_T_L_Q nexus, an SSP target port shall not retransmit an XFER_RDY frame except one that has the same requested offset as the immediately preceding XFER_RDY that did not receive an ACK (i.e., the XFER_RDY frame received a NAK or the connection was closed with ACK/NAK timeout before the XFER_RDY received an ACK or NAK).

If an SSP initiator port receives a new XFER_RDY frame with the RETRANSMIT bit set to one while processing the previous XFER_RDY frame for that I_T_L_Q nexus, the ST_ITS state machine stops processing the previous XFER_RDY frame (i.e., stops transmitting write DATA frames) and starts servicing the new XFER_RDY frame (see 9.2.6.2.3). The ST_ITS state machine does not transmit any write DATA frames for the previous XFER_RDY frame after transmitting a write DATA frame for the new XFER_RDY frame.

The SSP target port may reuse the value in the TARGET PORT TRANSFER TAG field from the previous XFER_RDY frame after it receives a write DATA frame for the new XFER_RDY frame.

An SSP target port retransmits each XFER_RDY frame that does not receive an ACK at least one time.

Modify subclause 9.2.4.6 as follows:**9.2.4.6 RESPONSE frame - handling of link layer errors**

If an SSP target port transmits a RESPONSE frame and receives a NAK for that frame, the SSP target port retransmits, in the same or a new connection, the RESPONSE frame at least one time with the RETRANSMIT bit set to one, and with the other fields set to the same values as the original RESPONSE frame (see 9.2.6.3.3.3).

If an SSP target port transmits a RESPONSE 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 SSP target port retransmits, in a new connection, the RESPONSE frame with the RETRANSMIT bit set to one, and with the other fields set to the same values as the original RESPONSE frame (see 9.2.6.3.3.3).

The ST_TTS state machine retransmits each RESPONSE frame that does not receive an ACK at least one time (see 9.2.6.3.3). The number of times it retransmits each RESPONSE frame is vendor-specific.

If an SSP initiator port receives a RESPONSE frame with a RETRANSMIT bit set to one, and it has previously received a RESPONSE frame for the same I_T_L_Q nexus, the ST_IFR state machine discards the extra RESPONSE frame (see 9.2.6.3.2). If the ST_IFR state machine has not previously received a RESPONSE frame for the I_T_L_Q nexus, then it considers the RESPONSE frame to be the valid RESPONSE frame.

Modify subclause 9.2.3.3.3 as follows:**9.2.6.3.3.3 ST_TTS2:Target_Send_Frame state****9.2.6.3.3.3.1 State description**

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 171 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:

- a) increment the Read Data Frames Transmitted state machine variable by one; and
- b) set the Read Data Offset state machine variable to the current Read Data Offset state machine variable plus the number of read data bytes transmitted in the DATA frame associated with the Transmission Status (Frame Transmitted) confirmation.

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

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

- a) Transmission Status (ACK Received) confirmations or arguments for each outstanding read DATA frame (i.e., Read Data Frames Transmitted state machine variable equals the Read Data Frames ACKed 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 Read Data Frames Transmitted state machine variable equals the Read Data Frames ACKed state machine variable and the Transmit Frame request was for a read DATA frame, this state shall:

- a) not modify the Balance Point Read Data Offset state machine variable (i.e., the balance point remains at the last point at which balance occurred); or
- b) set the Balance Point Read Data Offset state machine variable to the current Read Data Offset state machine variable.

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) set the other fields to the same values as contained in the previous RESPONSE frame; and
- c) 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) set the other fields to the same values as contained in the previous XFER_RDY frame; and
- d) resend a Transmit Frame (Interlocked) request to the port layer for the failed XFER_RDY frame.

Table 171 defines 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 171 — Messages sent to the ST_TFR state machine

Confirmation received from the port layer or argument received from ST_TTS3:Prepare_Data_In	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 minus the Data-In Application Client Buffer Offset argument 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	Transmission Complete (Xfer_Rdy Failed, NAK Received)
Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK)	b) if transport layer retries are enabled and the vendor-specific number of retries has been reached.	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	Transmission Complete (Data-In Failed, NAK Received)
Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK)	b) if transport layer retries are enabled and the vendor-specific number of retries has been reached.	Transmission Complete (Data-In Failed, Connection Failed)

Table 172 defines messages that this state shall send to the ST_TFR state machine upon receipt of the listed confirmations and arguments.

Table 172 — Additional messages sent to the ST_TFR state machine

Confirmation received from the port layer or argument received from ST_TTS3:Prepare_Data_In	Message sent to the ST_TFR state machine
Transmission Status (Bad Destination)	Transmission Complete (Connection Failed)
Transmission Status (Connection Rate Not Supported)	Transmission Complete (Connection Failed)
Transmission Status (Protocol Not Supported)	Transmission Complete (Connection Failed)
Transmission Status (Reserved Abandon 1)	Transmission Complete (Connection Failed)
Transmission Status (Reserved Abandon 2)	Transmission Complete (Connection Failed)
Transmission Status (Reserved Abandon 3)	Transmission Complete (Connection Failed)
Transmission Status (STP Resources Busy)	Transmission Complete (Connection Failed)
Transmission Status (Wrong Destination)	Transmission Complete (Connection Failed)
Transmission Status (Zone Violation)	Transmission Complete (Connection Failed)
Transmission Status (Break Received)	Transmission Complete (Data Transfer Terminated)

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
- tag.