

T10/06-470 revision 0

Date: October 27, 2006

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: Transport layer read data flowchart

1 Overview

What follows is a flowchart that represents the transport layers handling of read data operations at the SCSI target that are handled in the ST_TTS1:Target_Start state, ST_TTS3:Prepare_Data_In state, and ST_TTS2:Target_Send_Frame state.

It is not intention of this proposal nor is this proposal requesting the inclusion of these flowcharts into the SAS-2 standard as these flowcharts do not represent all the operations handled by the transport layer. However, it is intended that any future proposal that would impact the area of the transport layer covered by these flowcharts will be required to use them to show how any proposed change would effect the these flowcharts. The source code of the flowcharts is Visio and will be available for modification.

In the process of converting the words that are currently in SAS-2 there were no errors found, however, some of the wording may be somewhat unclear. Also, it was discovered that one of the conditions described in ST_TTS3 can never happen and another is already checked in ST_TTS2 so this proposal is recommending that they be deleted. See below the flowcharts for the suggested changes.

2 ST_TTS transport layer read data flowcharts

Read Data Offset = Offset into read data buffer
 Balance Point Read Data Offset = Offset into read data buffer for last data frame that was ACKed
 Data-In Request Byte Count = The number of bytes requested to be transferred. Set by the device server.

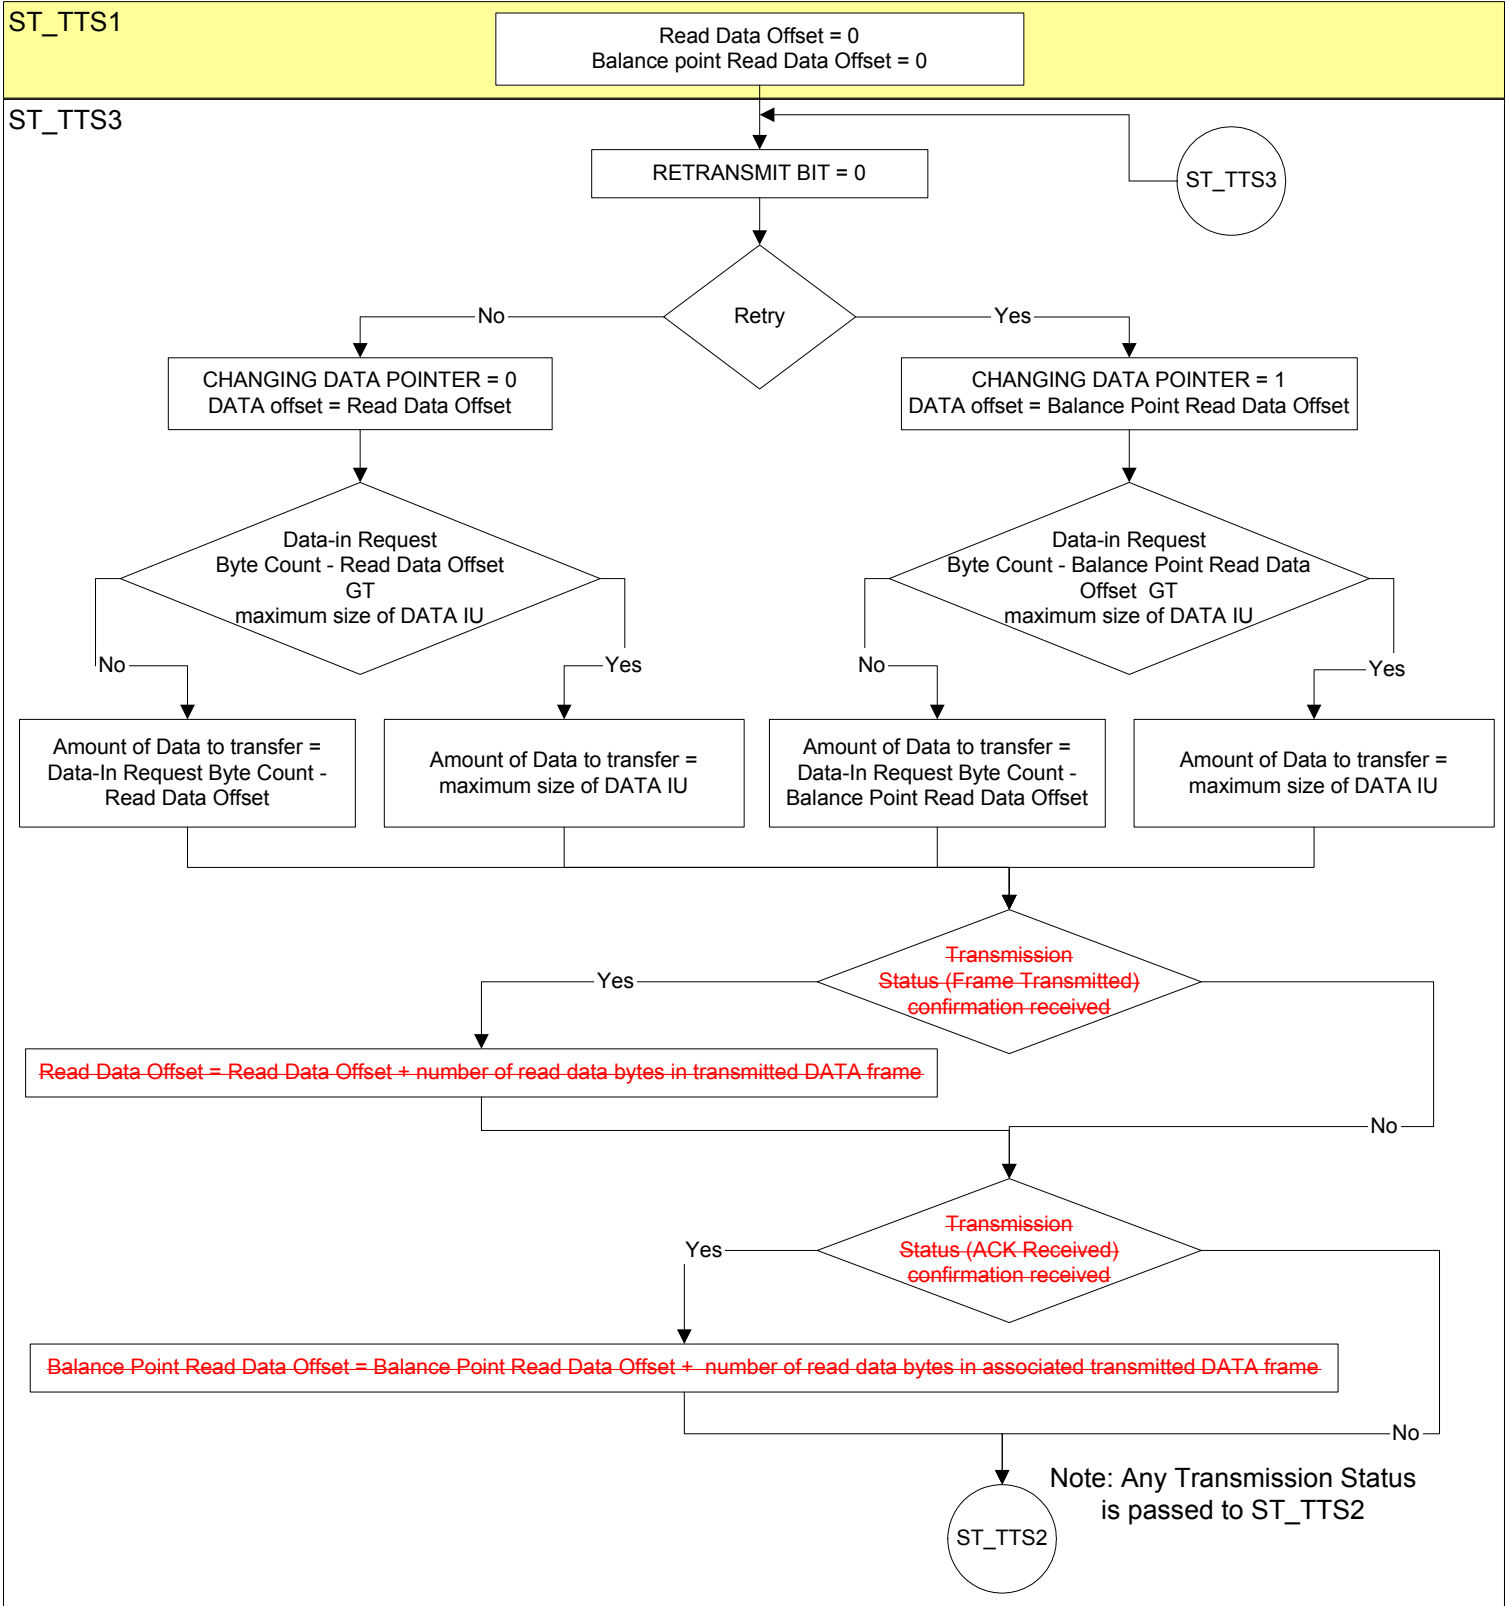


Figure 1 — Representation of transport layer (i.e., ST_TTS1 and ST_TTS3) read data operation

Read Data Offset = Offset into read data buffer

Balance Point Read Data Offset = Offset into read data buffer for last data frame that was ACKed

Data-In Request Byte Count = The number of bytes requested to be transferred. Set by the device server.

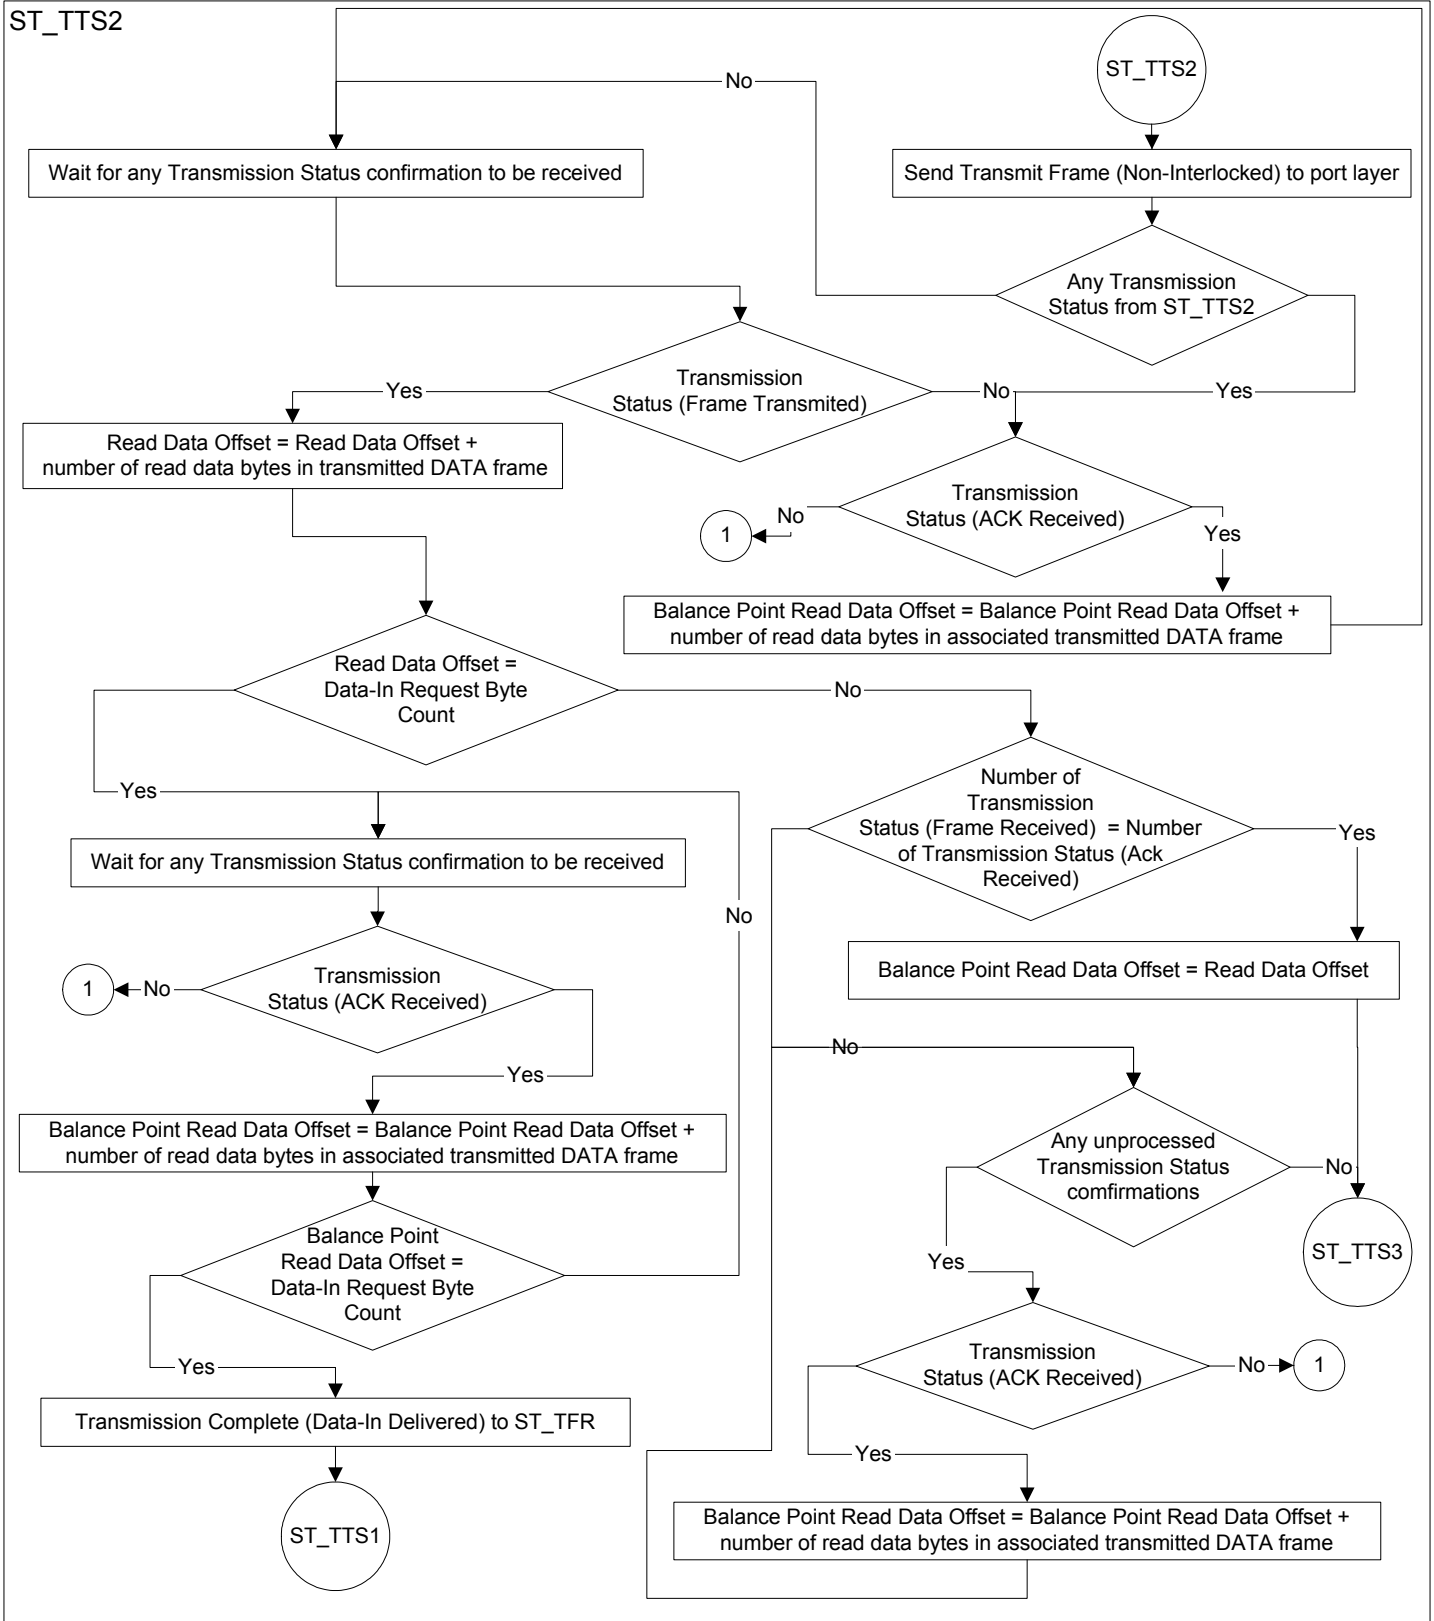


Figure 2 — Representation of transport layer (i.e., ST_TTS2) read data operation (part 1 or 2)

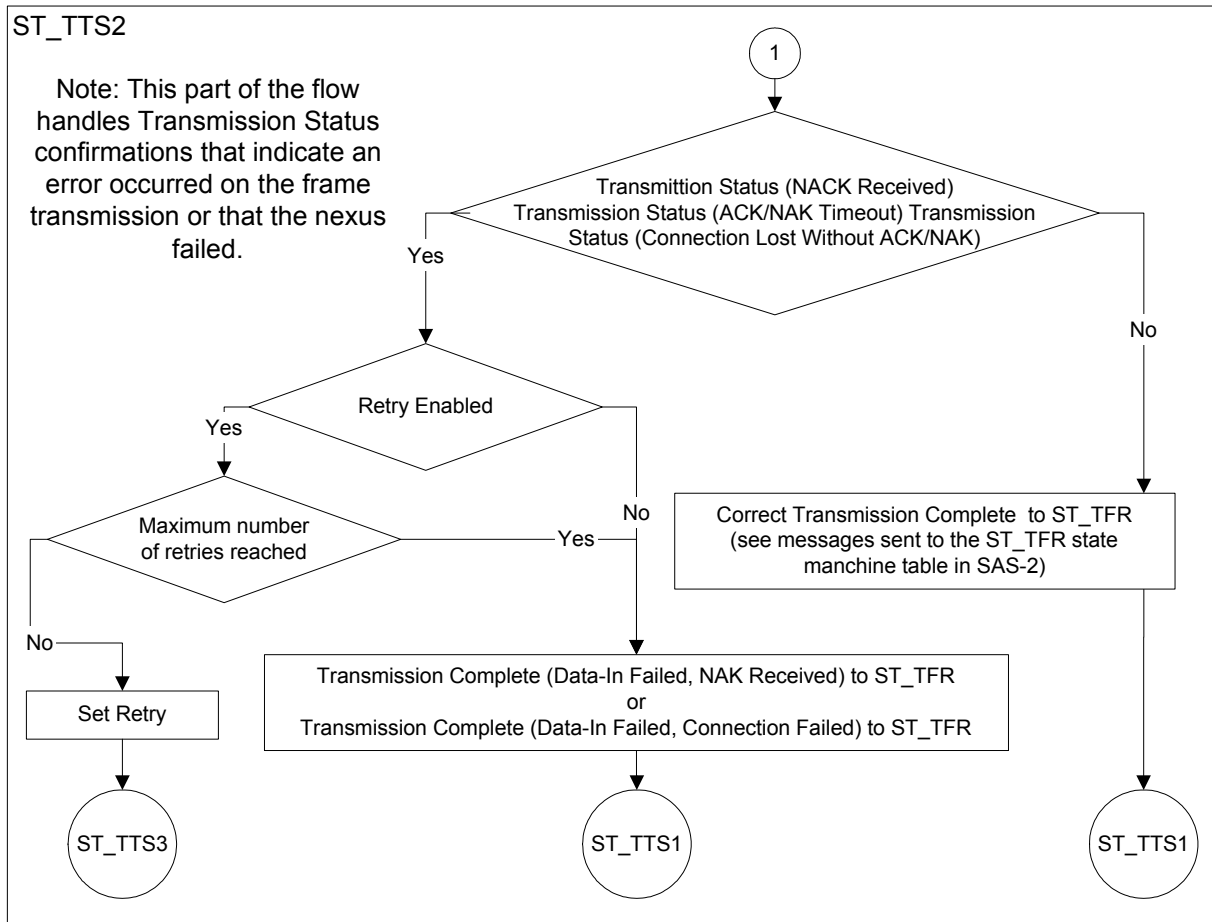


Figure 3 — Representation of transport layer (i.e., ST_TTS2) read data operation (part 2 or 2)

3 Recommended changes to SAS-2

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 DATA information unit.

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 DATA OFFSET field to Balance Point 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 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 DATA information unit.

~~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.~~

~~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.~~

9.2.6.3.3.3 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 ([Frame Transmitted](#)) 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 1 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 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 arguments](#), 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 number of Transmission Status \(Frame Transmitted\) confirmations for Transmit Frame \(Non-Interlocked\) requests is not equal to the number of Transmission Status \(ACK Received\) confirmations and arguments, then](#) if a Transmission Status (ACK Received) confirmation or [argument](#) 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 [or argument](#).

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.

T10/06-470 revision 0

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 1 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 1 — Messages sent to the ST_TFR state machine

Confirmation received from the port layer <u>or argument received from ST_TTS3</u>	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.

T10/06-470 revision 0

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.