T10/07-311 revision 1

Date: July 10, 2007

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: More Transport level read fixes

1 Overview

There were some issues with the read buffer offset pointers that were not successfully resolved in the 07-108 proposal. These issues are pointed out in editors notes in the current SAS revision. This proposal provides a solution to those issues. The write data path is not addressed in this proposal.

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 173.

Table 1 — ST_TTS state machine variables

State machine variable	Description
Read Data Offset	Offset into the application client's data-in buffer (i.e., the application client buffer for read data)
Balance Point Read Data Offset	Offset into the application client's data-in buffer (i.e., the application client buffer for read data) of last point at which the number of Transmission Status (ACK Received) confirmations or arguments was equal to the number of transmitted read DATA frames
Read Data Frames Transmitted	The number of Transmission Status (Frame Transmitted) confirmations received for read DATA frames
Read Data Frames ACKed	The number of Transmission Status (ACK Received) confirmations received for read DATA frames
Read Data Buffer End	One greater than the offset into the application client's data-in buffer (i.e., the application client buffer for read data) of the last location into which read data is to be placed.
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 174.

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 (see 9.2.6.3.2.3)
Data-Out	The Data-Out arguments received in the Request (Receive Data-Out) message (see 9.2.6.3.2.3)

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:

a) set the Read Data Offset state machine variable to the Data-In Application Client Buffer Offset argument;
b) set the Balance Point Read Data Offset state machine variable to the Data-In Application Client Buffer Offset argument;
c) set the Read Data Frames Transmitted state machine variable to zero;
d) set the Read Data Frames ACKed state machine variable to zero;
e) set the Read Data Buffer End state machine variable to the Data-In Application Client Buffer Offset argument plus the Data-In Request Byte Count argument; and
f) 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

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 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:

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 Read Data Buffer End state machine variable ~~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;
b) set the other fields to the same values as contained in the failed 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 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;
c) set the other fields to the same values contained in the failed XFER_RDY frame; and
d) resend a Transmit Frame (Interlocked) request to the port layer for the failed XFER_RDY frame.

Table 3 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 3 — 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 Read Data Buffer End state machine variable ~~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)

Table 4 defines messages that this state shall send to the ST_TFR state machine upon receipt of the listed confirmations and arguments.

Table 4 — 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
b) tag.

9.2.6.3.3.3.2 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.

9.2.6.3.3.3.3 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 if the Read Data Offset state machine variable ~~minus the Data-In Application Client Buffer Offset argument~~ is less than the Read Data Buffer End state machine variable ~~Data-In Request Byte Count argument~~ (i.e., there is more read data to transfer).

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).

9.2.6.3.3.3.4 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 Data-In Device Server Buffer argument 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 Data-In Destination SAS Address 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 Data-In Tag 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, set the DATA field to the information in the Device Server Buffer argument that corresponds to the read data to be transferred. The amount of data shall be:
   A) the lesser of:
      a) the Read Data Buffer End state machine variable ~~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;
      or
   B) the maximum size of the read Data information unit, if the Read Data Buffer End state machine variable minus the Read Data Offset state machine variable is equal to the maximum size of the read Data information unit.

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

a) set the CHANGING DATA POINTER bit in the frame to one;
b) set the DATA OFFSET field to the Balance Point Read Data Offset state machine variable;
c) set the Read Data Offset state machine variable to the Balance Point Read Data Offset state machine variable;
d) set the Read Data Frames Transmitted state machine variable to zero;
e) set the Read Data Frames ACKed state machine variable to zero; and
f) in the information unit, set the DATA field to the information in the Device Server Buffer argument that corresponds to the read data to be transferred. The amount of data shall be:
   A) the lesser of:
      a) the Read Data Buffer End state machine variable ~~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

B) the maximum size of the read Data information unit, if the Read Data Buffer End state machine variable minus the Read Data Offset state machine variable is equal to the maximum size of the read Data information unit;

or:

a) set the CHANGING DATA POINTER bit in the frame to one;
b) set the DATA OFFSET field to the Data-In Application Client Buffer Offset argument;
c) set the Read Data Offset state machine variable to the Data-In Application Client Buffer Offset argument;
d) set the Read Data Frames Transmitted state machine variable to zero;
e) set the Read Data Frames ACKed state machine variable to zero; and
f) in the information unit, set the DATA field to the information in the Device Server Buffer argument that corresponds to the read data to be transferred. The amount of data shall be:
   A) the lesser of:
      a) the Data-In Request Byte Count argument; and
      b) the maximum size of the read Data information unit for this Data-In request;

      or

   B) the maximum size of the read Data information unit, if the Data-In Request Byte Count argument is equal to the maximum size of the read Data information unit.

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.

3 ST_ITS transport layer read data flowcharts

DATA OFFSET field = Contains the first offset location into read data buffer for the current DATA information unit
Data-In Buffer Offset = Offset into read data buffer for the last received data frame.
Data-In Buffer Size = The number of bytes to be read as requested by the application client.

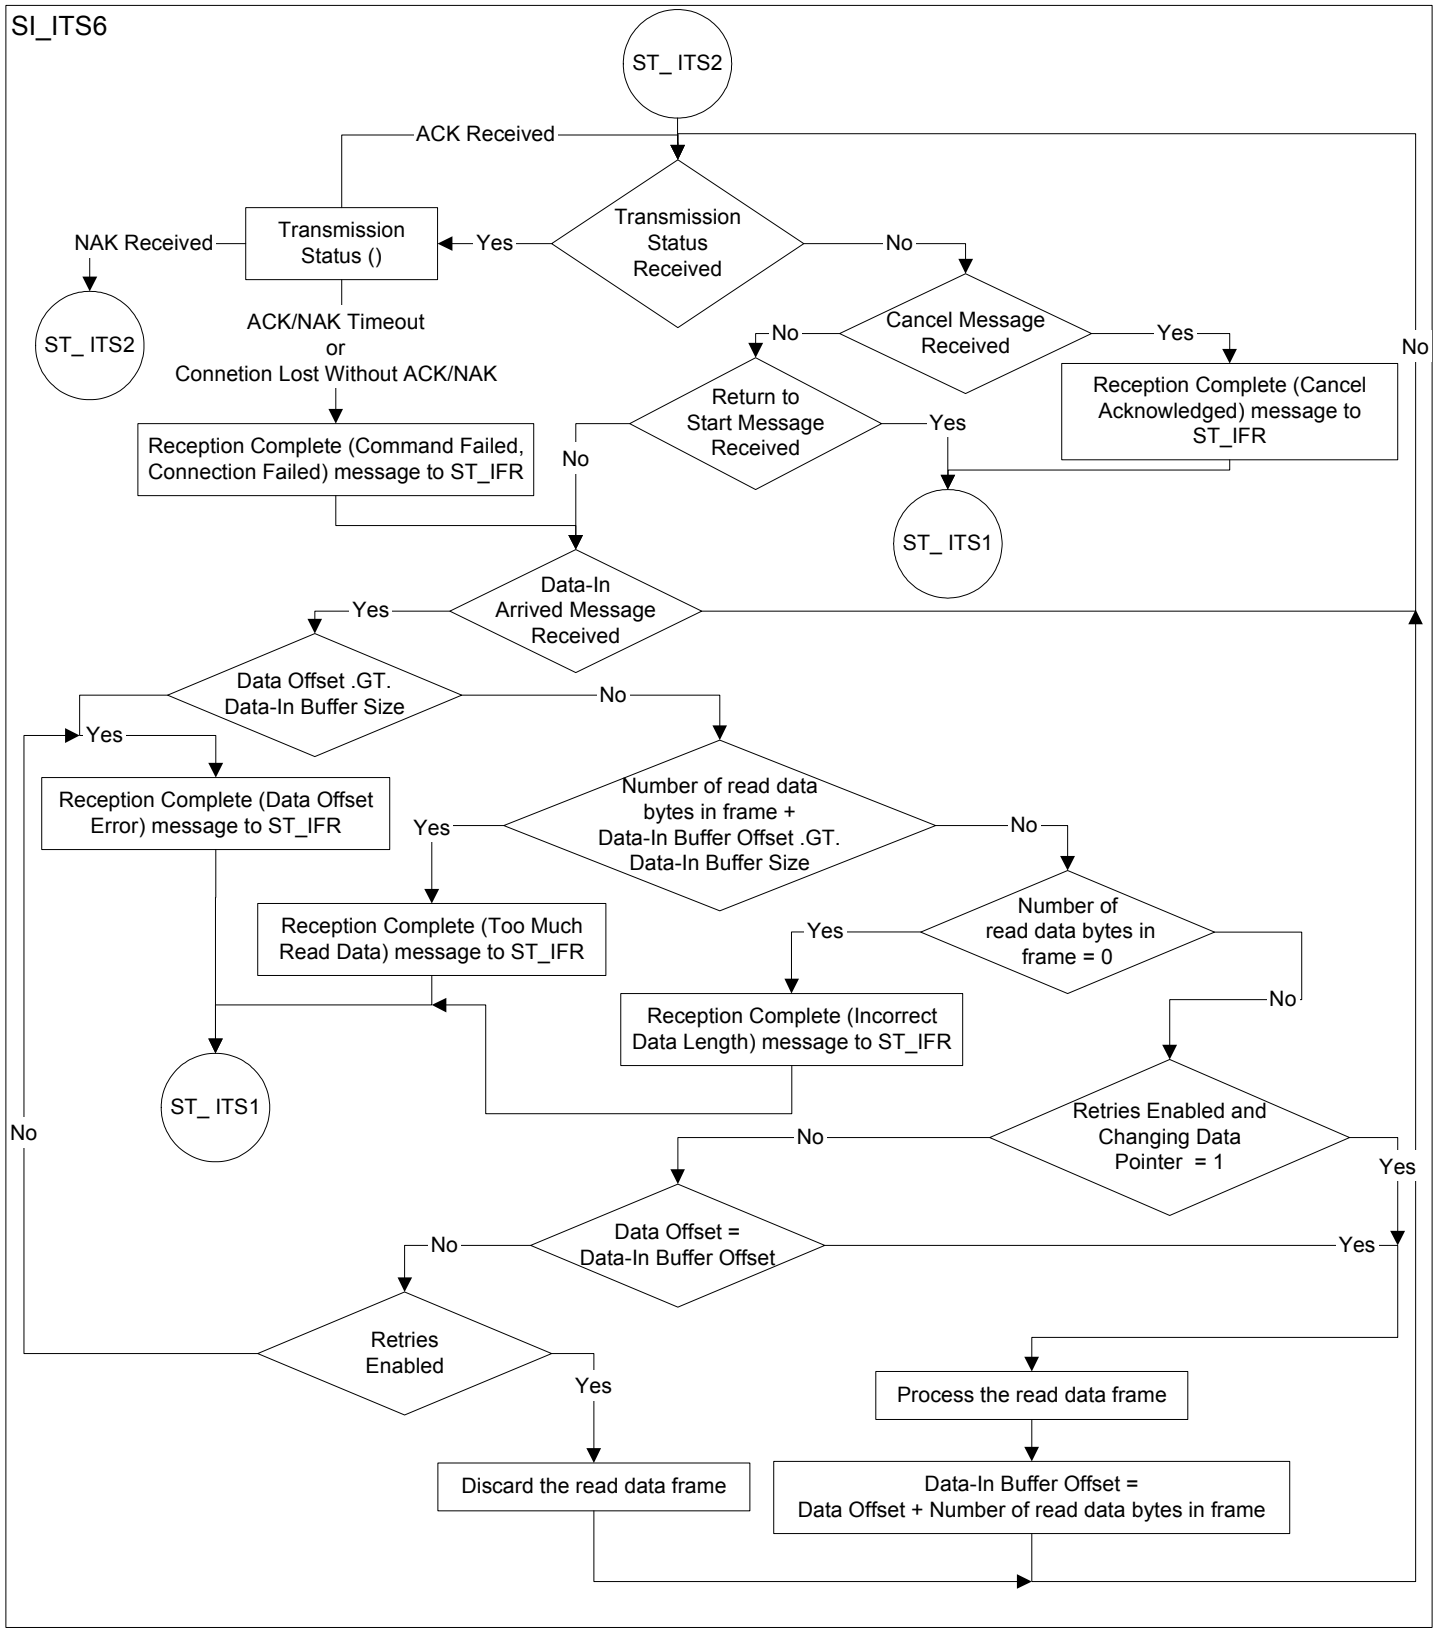

Figure 1 — Representation of transport layer (i.e., ST_ITS6) read data operation

4 ST_TTS transport layer read data flowcharts

Read Data Offset = Offset into application client read data buffer
Read Data Frames Transmitted = The number of Transmission Status (Frame Transferred) confirmations received
Read Data Frames ACKed = The number of Transmission Status (ACK Received) confirmation received.
Balance Point Read Data Offset = Offset into the application client read data buffer for last data frame that the number of frames transmitted = number ACKs received
Data-In Request Byte Count = The number of bytes requested to be transferred. Set by the device server.
Read Data Buffer End = Offset into the application client read buffer of last location into which data may be placed

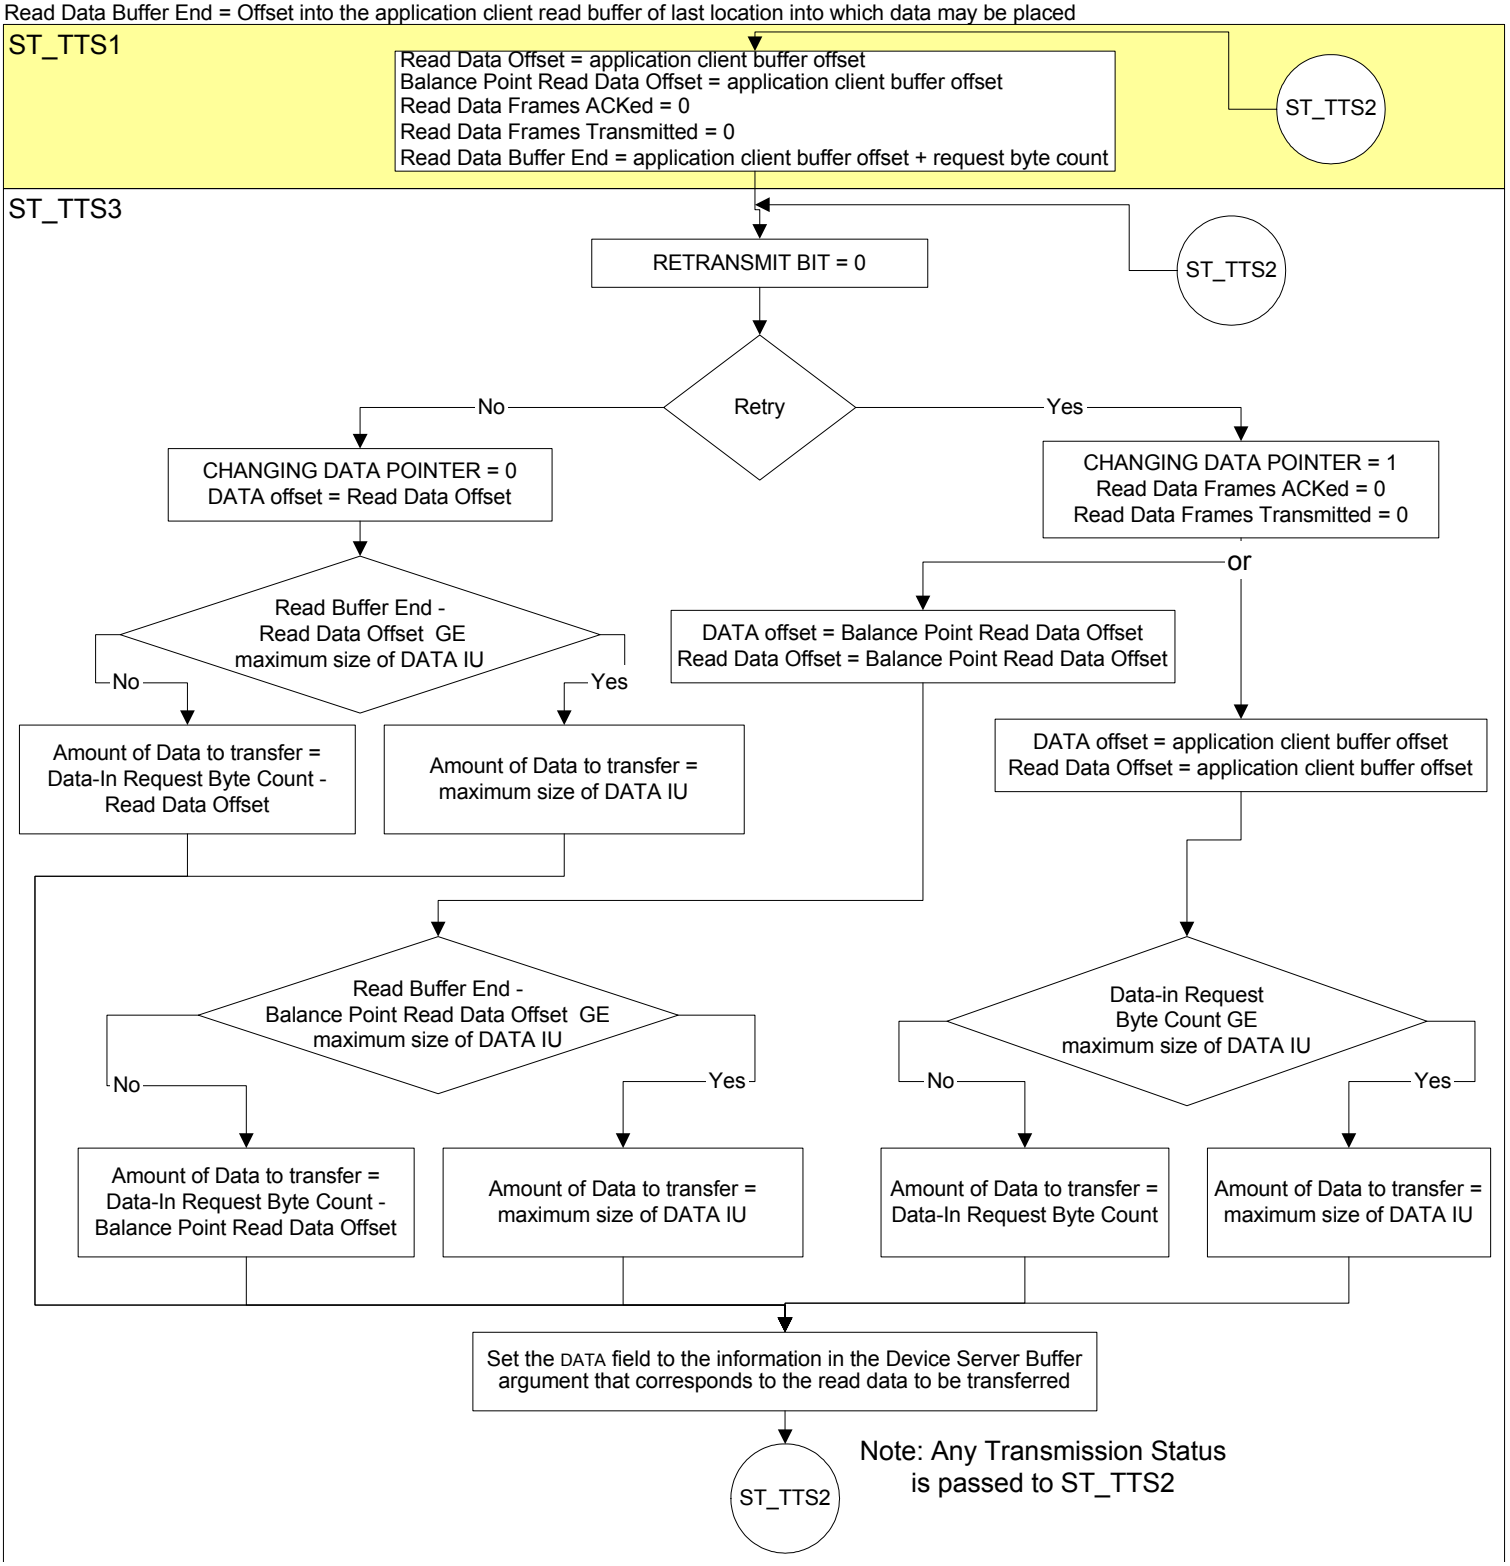

Figure 2 — Representation of transport layer (i.e., ST_TTS1 and ST_TTS3) read data operation

Read Data Frames Transmitted = The number of Transmission Status (Frame Transferred) confirmations received
Read Data Frames ACKed = The number of Transmission Status (ACK Received) confirmation received.
Balance Point Read Data Offset = Offset into the application client read data buffer for last data frame that the number of frames transmitted = number ACKs received
Read Data Buffer End = Offset into the application client read data buffer of last location into which data may be placed

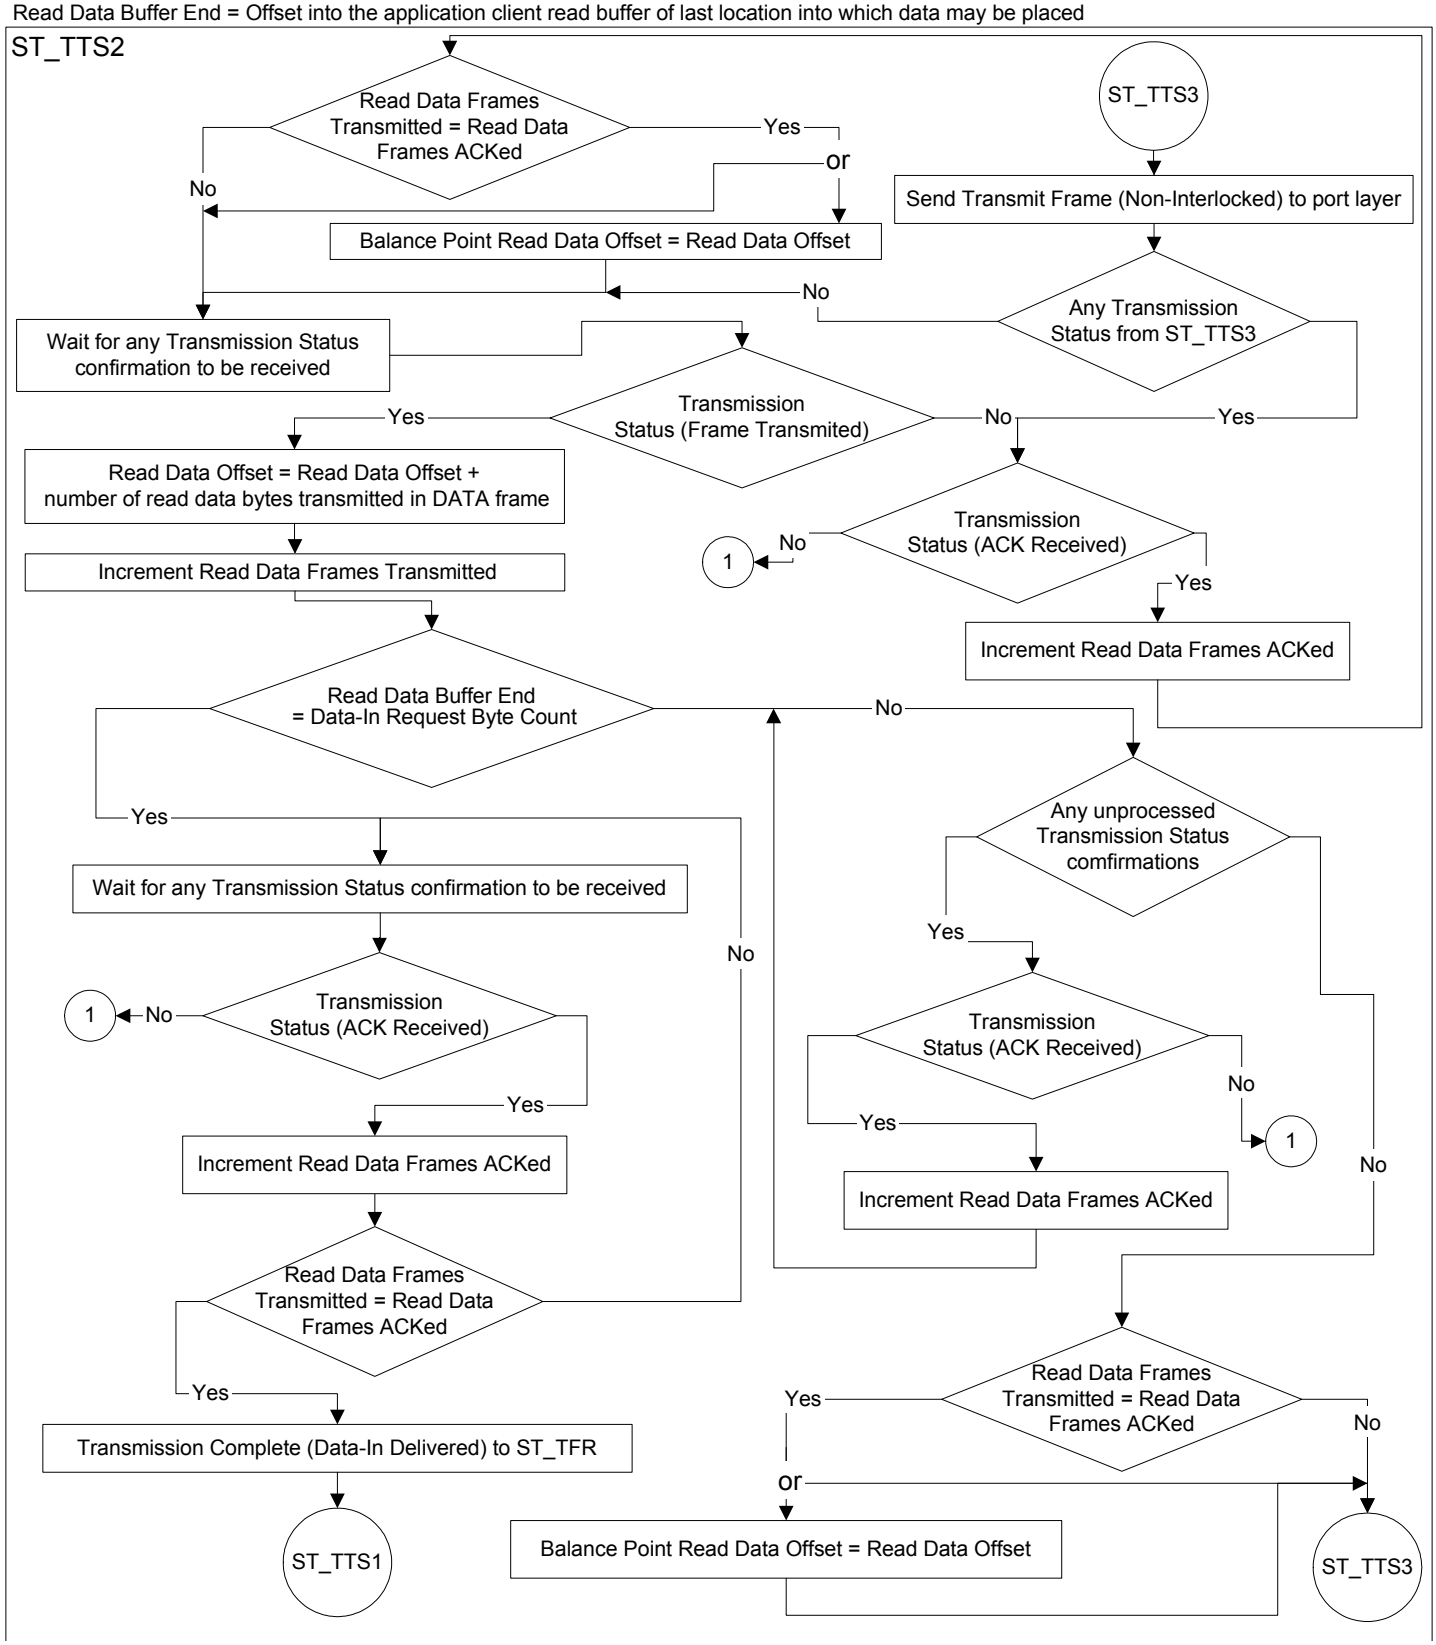

Figure 3 — Representation of transport layer (i.e., ST_TTS2) read data operation (part 1 or 2)

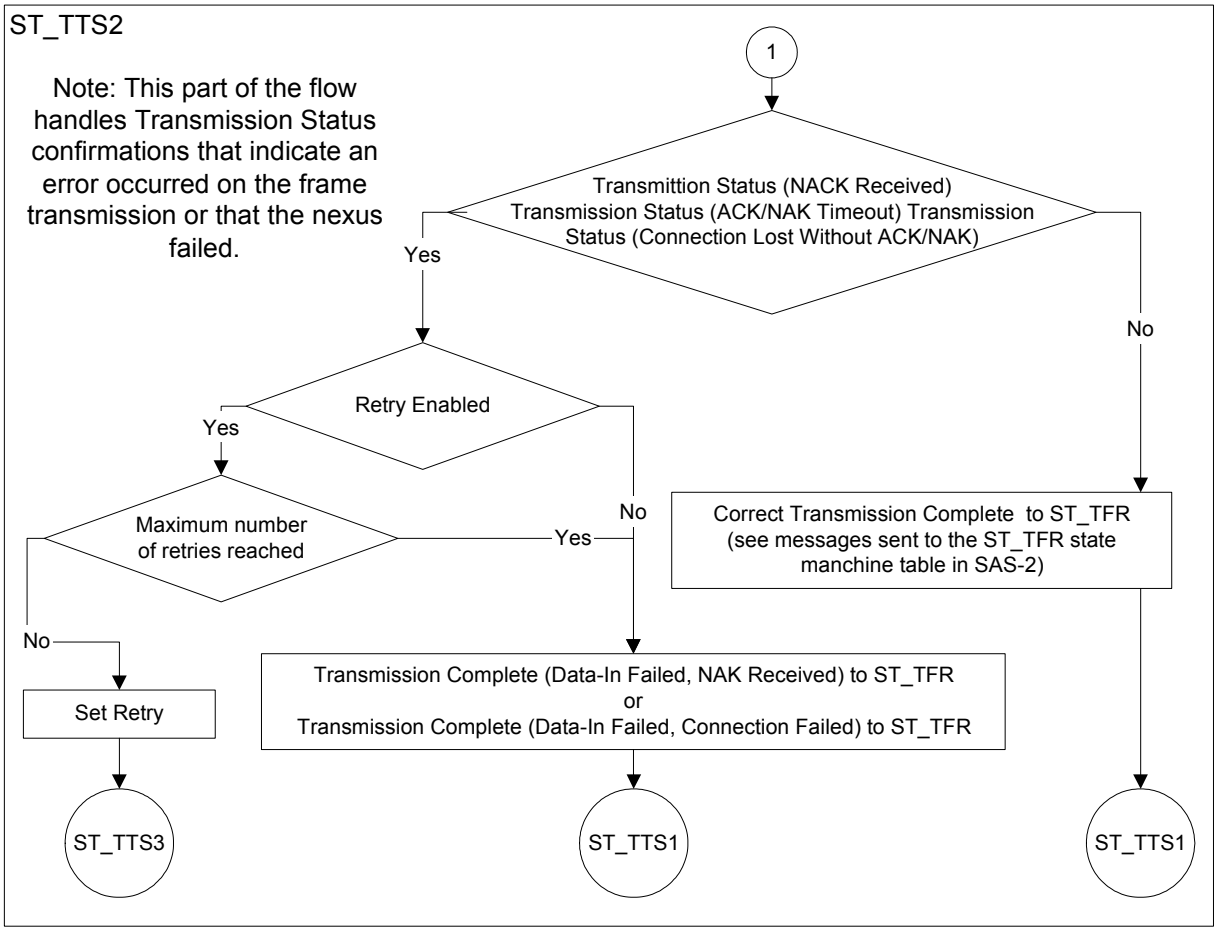

Figure 4 — Representation of transport layer (i.e., ST_TTS2) read data operation (part 2 or 2)