

Date: November 06, 2007

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: Initiator Write Transport Layer Write Flowcharts

1 Overview

This is a proposal that takes a detailed look at the write data path. In this revision only the initiator side write flow charts and fixes are shown. The flow charts reflect the wording as stated SAS-2 revision 11.

2 Fixes for writes in initiator transport layer description

9.2.6.2.3 ST_ITS (initiator transport server) state machine

9.2.6.2.3.1 ST_ITS state machine overview

The ST_ITS state machine performs the following functions:

- a) receives and processes messages from the ST_IFR state machine;
- b) sends messages to the ST_IFR state machine;
- c) sends request to the port layer regarding frame transmission;
- d) receives confirmations from the port layer regarding frame transmission; and
- e) receives HARD_RESET Received confirmations from the port layer.

This state machine consists of the following states:

- a) ST_ITS1:Initiator_Start state (see 9.2.6.2.3.2) (initial state);
- b) ST_ITS2:Initiator_Send_Frame state (see 9.2.6.2.3.3);
- c) ST_ITS3:Prepare_Command state (see 9.2.6.2.3.4);
- d) ST_ITS4:Prepare_Task state (see 9.2.6.2.3.5);
- e) ST_ITS5:Prepare_Data_Out state (see 9.2.6.2.3.6); and
- f) ST_ITS6:Receive_Data_In state (see 9.2.6.2.3.7).

This state machine shall start in the ST_ITS1:Initiator_Start state after power on.

If this state machine receives a HARD_RESET Received confirmation, then this state machine shall transition to the ST_ITS1:Initiator_Start state.

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

Table 1 — ST_ITS state machine variables

State machine variable	Description
Data-In Buffer Offset	Current offset in the application client's data-in buffer (i.e., the application client buffer for read data)
Data-Out Buffer Offset	Current offset in the application client's data-out buffer (i.e., the application client buffer for write data)
Previous Requested Offset	Offset in the application client's data-out buffer (i.e., the application client buffer for write data) from the last XFER_RDY frame received
Previous Write Data Length	Write data length from the last XFER_RDY frame received

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

Table 2 — ST_ITS state machine arguments

State machine argument	Description
Command	Consists of the Command arguments received in the Request (Send Command) message
Task	Consists of the arguments received in the Request (Send Task) message
Xfer_Rdy	Consists of the arguments received in the XFER_RDY Arrived message
Data-Out Buffer	The location of the application client's data-out buffer (i.e., the application client buffer for write data)
Data-Out Buffer Size	The size in bytes of the application client's data-out buffer (i.e., the application client buffer for write data)
Data-In Buffer Size	The size in bytes of the application client's data-in buffer (i.e., the application client buffer for read data)

9.2.6.2.3.2 ST_ITS1:Initiator_Start state

9.2.6.2.3.2.1 State description

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

Upon entry into this state, this state shall set the Data-In Buffer Offset state machine variable to zero.

Upon entry into this state, this state shall set the Data-Out Buffer Offset state machine variable to zero.

9.2.6.2.3.2.2 Transition ST_ITS1:Initiator_Start to ST_ITS3:Prepare_Command

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

9.2.6.2.3.2.3 Transition ST_ITS1:Initiator_Start to ST_ITS4:Prepare_Task

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

9.2.6.2.3.3 ST_ITS2:Initiator_Send_Frame state

If this state is entered from the ST_ITS3:Prepare_Command state for transmission of a COMMAND frame, then this state shall send a Transmit Frame (Interlocked) request to the port layer.

If this state is entered from the ST_ITS6:Receive_Data_In state, and the vendor-specific number of retries has not been reached for the COMMAND frame requesting a read operation, then this state shall send a Transmit Frame (Interlocked) request to the port layer.

If this state is entered from the ST_ITS4:Prepare_Task state for transmission of an TASK frame, then this state shall send a Transmit Frame (Interlocked) request to the port layer.

If this state is entered from the ST_ITS5:Prepare_Data_Out state for transmission of a write DATA frame, then this state shall send a Transmit Frame (Non-Interlocked) request to the port layer after this state has received an XFER_RDY Arrived message.

If this state is entered from the ST_ITS5:Prepare_Data_Out state for transmission of a write DATA frame and first bust is enabled, then this state shall send a Transmit Frame (Non-Interlocked) request to the port layer after this state has received a Transmission Status (Frame Transmitted) confirmation and a Transmission Status (ACK Received) confirmation for the COMMAND frame.

A Transmit Frame request shall include the COMMAND frame from the ST_ITS3:Prepare_Command state or from the ST_ITS6:Receive_Data_In state, the TASK frame from the ST_ITS4:Prepare_Task state, or the write

DATA frame from the ST_ITS5:Prepare_Data_Out state and the following arguments to be used for any OPEN address frame:

- a) initiator port bit set to one;
- 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 is Transmission Status (I_T Nexus Loss), then this state shall send a Transmission Complete (I_T Nexus Loss) message to the ST_IFR state machine. This Transmission Complete message shall include the tag as an argument.

If the confirmation is not Transmission Status (Frame Transmitted) or Transmission Status (I_T Nexus Loss) (see table 152 in 8.2.2.3.4), and the Transmit Frame request was for a COMMAND frame or a DATA frame, then this state shall send a Transmission Complete (Command Failed, Connection Failed) message to the ST_IFR state machine. The message shall include the tag.

If the confirmation is not Transmission Status (Frame Transmitted) or Transmission Status (I_T Nexus Loss) (see table 152 in 8.2.2.3.4), and the Transmit Frame request was for a TASK frame, then this state shall send a Transmission Complete (Task Failed, Connection Failed) message to the ST_IFR state machine. The message shall include the tag.

If the confirmation is Transmission Status (Frame Transmitted), and the Transmit Frame request was for a COMMAND frame not requesting a read operation, a COMMAND frame not requesting a write operation, a TASK frame, or a write DATA frame where the number of data bytes that have been transmitted equal the Data-Out Buffer Size state machine argument, 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 COMMAND frame requesting a write operation, or a write DATA frame where the number of data bytes that have been transmitted is less than the Data-Out Buffer Size state machine argument and the write data length from the previous XFER_RDY 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);
- d) Transmission Status (Connection Lost Without ACK/NAK); or
- e) XFER_RDY Arrived message.

If a XFER_RDY Arrived message is received, then the ST_ITS shall respond to the XFER_RDY frame as if a Transmission Status (ACK Received) was received.

NOTE 1 - If the number of data bytes requested to be transmitted for the Send SCSI Command protocol service request are fewer than the number of bytes in the service request, then this state may send additional Transmit Frame requests for write DATA frames for the protocol service 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 for Transmit Frame requests for previous write DATA frames sent for the I_T_L_Q nexus.

After a Transmission Status (Frame Transmitted) is received, if a confirmation of Transmission Status (NAK Received) is received, the Transmit Frame request was for a COMMAND frame, and the vendor-specific number of retries has not been reached, then this state shall send a Transmit Frame (interlocked) request to the port layer (i.e., the last COMMAND frame is retransmitted).

After a Transmission Status (Frame Transmitted) is received, if a confirmation of Transmission Status (NAK Received) is received, the Transmit Frame request was for a TASK frame, and the vendor-specific number of retries has not been reached, then this state shall send a Transmit Frame (interlocked) request to the port layer (i.e., the last TASK frame is retransmitted).

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

Table 3 — Messages sent to the ST_IFR state machine

Confirmation received from the port layer	Conditions under which confirmation was received	Message sent to ST_IFR state machine
Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK)	The Transmit Frame request was for a COMMAND frame.	Transmission Complete (Command Failed, ACK/NAK Timeout)
	The Transmit Frame request was for a TASK frame.	Transmission Complete (Task Failed, ACK/NAK Timeout)
Transmission Status (NAK Received)	The Transmit Frame request was for a COMMAND frame and the vendor-specific number of retries has been reached.	Transmission Complete (Command Failed, NAK Received)
	The Transmit Frame request was for a TASK frame and the vendor-specific number of retries has been reached.	Transmission Complete (Task Failed, NAK Received)
Transmission Status (NAK Received)	The Transmit Frame request was for a write DATA frame and: a) the RETRY DATA FRAMES bit was set to zero in the XFER_RDY frame requesting the data; or b) the RETRY DATA FRAMES bit was set to one in the XFER_RDY frame requesting the data, and the vendor-specific number of retries has been reached.	Transmission Complete (Data-Out Failed, NAK Received)
Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK)		Transmission Complete (Data-Out Failed, ACK/NAK Timeout)

After this state sends a Transmission Complete (Command Failed, ACK/NAK Timeout) this state shall continue processing messages and confirmations.

NOTE 2 - The application client may determine the command was received and is being processed by the device server and allow the command to complete. The application client may accomplish this by the use of the QUERY TASK task management request.

If this state receives a Return to Start message or a Return to Start 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. This state may also send a Cancel request to the port layer to cancel a previous Transmit Frame request.

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 (Cancel Acknowledged) message to the ST_IFR 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. This state may also send a Cancel request to the port layer to cancel a previous Transmit Frame request. The Cancel request shall include the following arguments:

- a) destination SAS address; and

b) tag.

NOTE 3 - 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.

If this state receives a Transmission Status (Cancel Acknowledged) confirmation, then this state shall send a Transmission Complete (Cancel Acknowledged) message to the ST_IFR state machine.

If this state receives an XFER_RDY Arrived message, then this state shall verify the Xfer_Rdy state machine argument as specified in table 4. If the verification fails, then this state sends the Transmission Complete message specified in table 4 to the ST_IFR state machine.

Table 4 — Transmission Complete messages for XFER_RDY frame verification failures

Message sent to ST_IFR ^a	Condition
Transmission Complete (XFER_RDY Incorrect Write Data Length)	The Write Data Length Xfer_Rdy state machine argument is zero.
	The value in the Requested Offset Xfer_Rdy state machine argument plus the Write Data Length Xfer_Rdy state machine argument is greater than the Data-Out Buffer Size state machine argument.
Transmission Complete (XFER_RDY Requested Offset Error)	First burst is disabled, this is the first XFER_RDY frame for a command, and the value in the Requested Offset Xfer_Rdy state machine argument is not set to zero.
	First burst is enabled, this is the first XFER_RDY frame for a command, and the value in the Requested Offset Xfer_Rdy state machine argument is not equal to the value indicated by the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see 10.2.7.2.5).
	Transport layer retries are disabled and the Requested Offset Xfer_Rdy state machine argument is not equal to the Previous Requested Offset state machine variable plus the Previous Write Data Length Field state machine variable.
	Transport layer retries are enabled, the Retransmit Bit Xfer_Rdy state machine argument is set to zero, and the Requested Offset Xfer_Rdy state machine argument is not equal to the Previous Requested Offset state machine variable plus the Previous Write Data Length state machine variable.
	Transport layer retries are enabled, this is not the first XFER_RDY frame for the command, the Retransmit Bit Xfer_Rdy state machine argument is set to one, and the Requested Offset Xfer_Rdy state machine argument is not equal to the Previous Requested Offset state machine variable.; and e) the Requested Offset Xfer_Rdy state machine argument is not equal to the Previous Requested Length state machine variable plus the Previous Write Data Length Field state machine variable.
^a If more than one condition is true, then this state shall send the Transmission Complete (XFER_RDY Incorrect Write Data Length) message to the ST_IFR state machine.	

After this state verifies an XFER_RDY frame, it shall:

- a) set the Data-Out Buffer Offset state machine variable to the Requested Offset Xfer_Rdy state machine argument;
- b) set the Previous Requested Offset state machine variable to the Requested Offset Xfer_Rdy state machine argument; and
- c) set the Previous Write Data Length state machine variable to the Requested ~~Offset~~ [write data length](#) Xfer_Rdy state machine argument.

9.2.6.2.3.3.3 Transition ST_ITS2:Initiator_Send_Frame to ST_ITS1:Initiator_Start

This transition shall occur after:

- a) this state has sent one of the following to the ST_IFR state machine:
 - A) a Transmission Complete (Command Failed, NAK Received) message;
 - B) a Transmission Complete (Task Failed, ACK/NAK Timeout) message;
 - C) a Transmission Complete (Task Failed, NAK Received) message;
 - D) a Transmission Complete (Command Failed, ACK/NAK Timeout) message and the command was for a non-data operation;
 - E) a Transmission Complete (Data-Out Failed, NAK Received) message;
 - F) a Transmission Complete (Data-Out Failed, ACK/NAK Timeout) message;
 - G) a Transmission Complete (XFER_RDY Incorrect Write Data Length) message;
 - H) a Transmission Complete (XFER_RDY Requested Offset Error) message; or
 - I) a Transmission Complete (Cancel Acknowledged) message;
- or
- b) this state has received a Return To Start message or Return To Start argument, and has received:
 - A) confirmations for all Transmit Frame requests sent to the port layer; or
 - B) a Transmission Status (Cancel Acknowledged) confirmation.

9.2.6.2.3.3.4 Transition ST_ITS2:Initiator_Send_Frame to ST_ITS5:Prepare_Data_Out

If first burst is enabled, this transition shall occur and include the First Burst argument after this state receives:

- a) a Transmission Status (Frame Transmitted) confirmation followed by a Transmission Status (ACK Received) for a COMMAND frame requesting a write operation; or
- b) a Transmission Status (Frame Transmitted) confirmation for a Transmit Frame (Non-interlocked) request if the Data-Out Buffer Offset state machine variable is less than the first burst size.

This transition shall occur after this state receives:

- a) an XFER_RDY Arrived message; or
- b) a Transmission Status (Frame Transmitted) confirmation for a Transmit Frame (Non-interlocked) request if the Data-Out Buffer Offset state machine variable is less than the Requested Offset Xfer_Rdy state machine argument plus the Write Data Length Xfer_Rdy state machine argument.

NOTE 4 - This transition occurs even if this state has not received a Transmission Status (ACK Received) for the write DATA frame.

This transition shall include a Retry argument and occur after:

- a) this state receives one of the following confirmations or arguments for a write DATA frame:
 - A) Transmission Status (NAK Received);
 - B) Transmission Status (ACK/NAK Timeout); or
 - C) Transmission Status (Connection Lost without ACK/NAK);
- b) the RETRY DATA FRAMES bit is set to one in the XFER_RDY frame for the write operation;
- c) the Data-Out Buffer Offset state machine variable is set to the Requested Offset Xfer_Rdy state machine argument;
- d) [all write DATA frames that have received a Transmission Status \(Frame Transmitted\) confirmation have received a Transmission Status confirmation](#); and
- e) the vendor-specific number of retries, if any, for the write DATA frame has not been reached.

9.2.6.2.3.3.5 Transition ST_ITS2:Initiator_Send_Frame to ST_ITS6:Process_Data_In

This transition shall occur after this state receives a Transmission Status (Frame Transmitted) confirmation for a COMMAND frame for a command requesting a read operation.

NOTE 5 - This transition occurs even if this state has not received a Transmission Status (ACK Received) for the COMMAND frame.

9.2.6.2.3.4 ST_ITS3:Prepare_Command state

9.2.6.2.3.4.1 State description

This state shall construct a COMMAND frame using the Command arguments:

- a) FRAME TYPE field set to 06h (i.e., COMMAND frame);
- b) HASHED DESTINATION SAS ADDRESS field set to the hashed value of the Destination SAS Address Commands argument;
- c) HASHED SOURCE SAS ADDRESS field set to the hashed value of the SSP initiator port's SAS address;
- d) RETRY DATA FRAMES bit set to zero;
- e) RETRANSMIT bit set to zero;
- f) CHANGING DATA POINTER bit set to zero;
- g) NUMBER OF FILL BYTES field set zero.
- h) TAG field set to the Tag Command argument;
- i) TARGET PORT TRANSFER TAG field set to FFFFh;
- j) DATA OFFSET field set to zero;
- k) in the information unit, LOGICAL UNIT NUMBER field set to the Logical Unit Number Command argument;
- l) in the information unit, ENABLE FIRST BURST bit set to the Enable First Burst Command argument;
- m) in the information unit, TASK PRIORITY field set to the Task Priority Command argument;
- n) in the information unit, TASK ATTRIBUTE field set to the Task Attribute Command argument;
- o) in the information unit, ADDITIONAL CDB LENGTH field set to the Additional CDB Length Command argument;
- p) in the information unit, CDB field set to the CDB Command argument;
- q) in the information unit, ADDITIONAL CDB BYTES field set to the Additional CDB Bytes Command argument, if any; and
- r) no fill bytes.

9.2.6.2.3.4.2 Transition ST_ITS3:Prepare_Command to ST_ITS2:Initiator_Send_Frame

This transition shall occur after this state:

- a) constructs a COMMAND frame;
- b) receives a Cancel message; or
- c) receives a Return To Start message.

This transition shall include the:

- a) COMMAND frame as an argument;
- b) if a Cancel message was received, then a Cancel argument; or
- c) if a Return To Start message was received, then a Return To Start argument.

9.2.6.2.3.5 ST_ITS4:Prepare_Task state

9.2.6.2.3.5.1 State description

This state shall construct a TASK frame using the Task arguments:

- a) FRAME TYPE field set to 16h (i.e., TASK frame);
- b) HASHED DESTINATION SAS ADDRESS field set to the hashed value of the Destination SAS Address Task argument;
- c) HASHED SOURCE SAS ADDRESS field set to the hashed value of the SSP initiator port's SAS address;
- d) RETRY DATA FRAMES bit set to zero;
- e) RETRANSMIT bit set to the Retransmit Bit Task argument;
- f) CHANGING DATA POINTER bit set to zero;
- g) NUMBER OF FILL BYTES field set zero.
- h) TAG field set to the Tag Task argument;
- i) TARGET PORT TRANSFER TAG field set to FFFFh;
- j) DATA OFFSET field set to zero;
- k) in the information unit, LOGICAL UNIT NUMBER field set to the Logical Unit Number Task argument;

- l) in the information unit, TASK MANAGEMENT FUNCTION field set to the Task Management Function Task argument;
- m) in the information unit, TAG OF TASK TO BE MANAGED field set to the Tag Task argument of task to be managed and
- n) no fill bytes.

9.2.6.2.3.5.2 Transition ST_ITS4:Prepare_Task to ST_ITS2:Initiator_Send_Frame

This transition shall occur after this state:

- a) constructs a TASK frame;
- b) receives a Cancel message; or
- c) receives a Return To Start message.

This transition shall include the:

- a) TASK frame as an argument;
- b) if a Cancel message was received, then a Cancel argument; or
- c) if a Return To Start message was received, then a Return To Start argument.

9.2.6.2.3.6 ST_ITS5:Prepare_Data_Out state

9.2.6.2.3.6.1 State description

This state shall construct a write DATA frame using the following Xfer_Rdy state machine arguments and Command state machine arguments:

- 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 Command argument;
- c) HASHED SOURCE SAS ADDRESS field set to the hashed value of the SSP initiator port's SAS address;
- d) RETRY DATA FRAMES bit set to zero;
- e) RETRANSMIT bit set to zero;
- f) CHANGING DATA POINTER bit set as specified in this subclause;
- g) NUMBER OF FILL BYTES field set to the number of fill bytes, based on the length of the specified data;
- h) TAG field set to the Tag Command argument;
- i) TARGET PORT TRANSFER TAG field set to FFFFh if this state received a First Burst argument or the Target Port Transfer Tag Xfer_Rdy argument if this state did not receive a First Burst argument;
- j) DATA OFFSET field set to the Data-Out Buffer Offset state machine variable;
- ~~k) in the information unit, DATA field set to the information that starts at the location in the Data-Out Buffer state machine argument pointed to by the Data-Out Buffer Offset state machine variable and shall contain the amount of data indicated by the Write Data Length Xfer_Rdy argument; and~~
- l) in the information unit, set the DATA field to the information that starts at the location in the Data-Out Buffer state machine argument pointed to by the Data-Out Buffer Offset state machine variable. If the number of bytes remaining to be transferred as defined by the following calculation:
bytes remaining to be transferred = Write Data Length Xfer_Rdy state machine argument - (Data-Out Buffer Offset state machine argument - Requested Offset Xfer_Rdy state machine argument)
is equal to the maximum size of the write Data information unit, then the amount of data shall be the maximum size of the write Data information unit. Otherwise, the amount of data shall be the lesser of:
A) the bytes remaining to be transferred; and
B) the maximum size of the Write information unit;
and
- m) fill bytes, if any.

If this state is entered without a Retry argument, then this state shall set the CHANGING DATA POINTER bit to zero.

If this state is entered with a Retry argument, then this state shall set the CHANGING DATA POINTER bit to one.

After constructing the write DATA frame, this state shall set the Data-Out Buffer Offset state machine variable to the value of the DATA OFFSET field plus the number of bytes in the DATA field in the write Data information unit.

9.2.6.2.3.6.2 Transition ST_ITS5:Prepare_Data_Out to ST_ITS2:Initiator_Send_Frame

This transition shall occur after this state:

- a) constructs a write DATA frame;
- b) receives a Cancel message; or
- c) receives a Return To Start message.

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

- a) write DATA frame as an argument;
- b) if a Cancel message was received, then a Cancel argument; or
- c) if a Return To Start message was received, then a Return To Start argument.

2 ST_ITS transport layer write data flowcharts

Data-Out Buffer Offset = Offset into the application client write buffer.
 Data-In Buffer Offset = Offset into the application client read buffer.

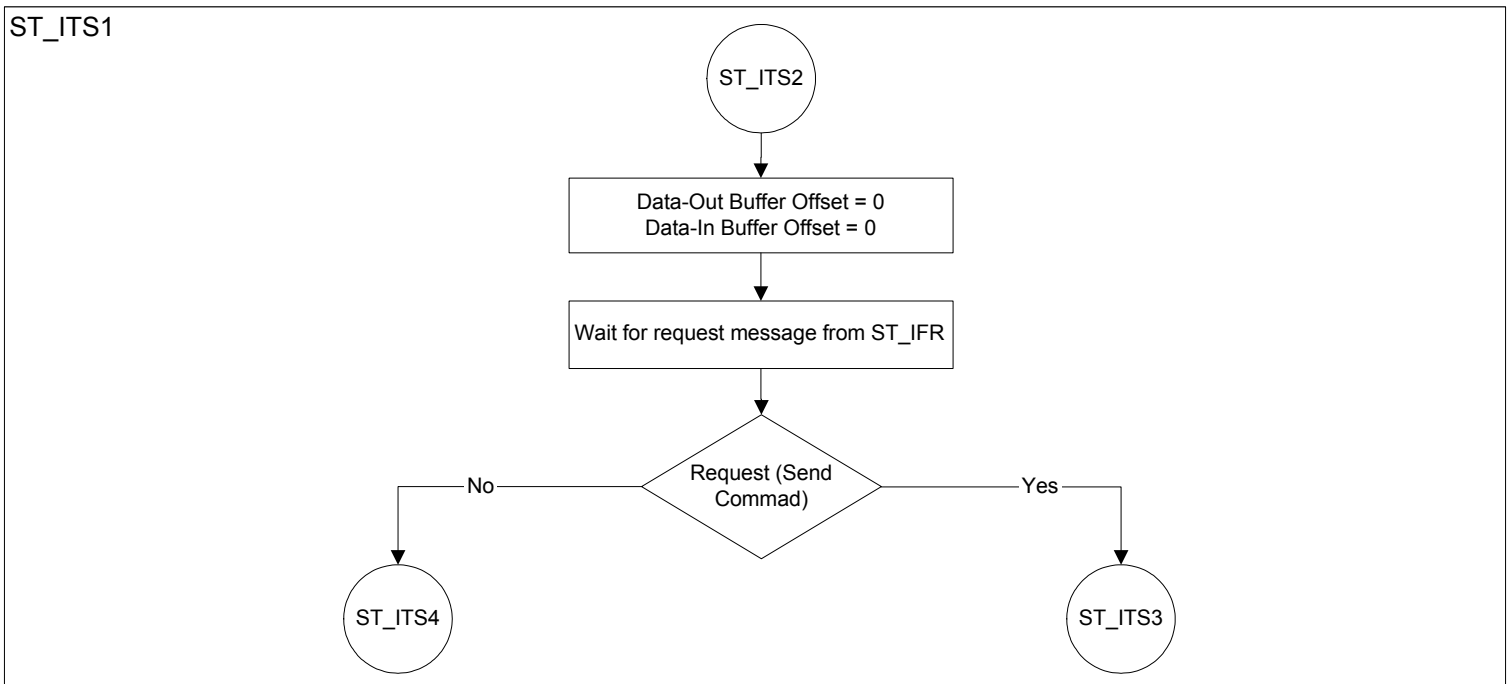


Figure 1 — Representation of transport layer (i.e., ST_ITS1) write data operation

Data-Out Buffer Size = Size of the application client write buffer.
 Data-Out Buffer Offset = Offset into the application client write buffer.
 Requested Offset = The value in the current Xfer_Rdy's Requested Offset.
 Write Data Length = The value of the current Xfer_Rdy's Write Data Length.

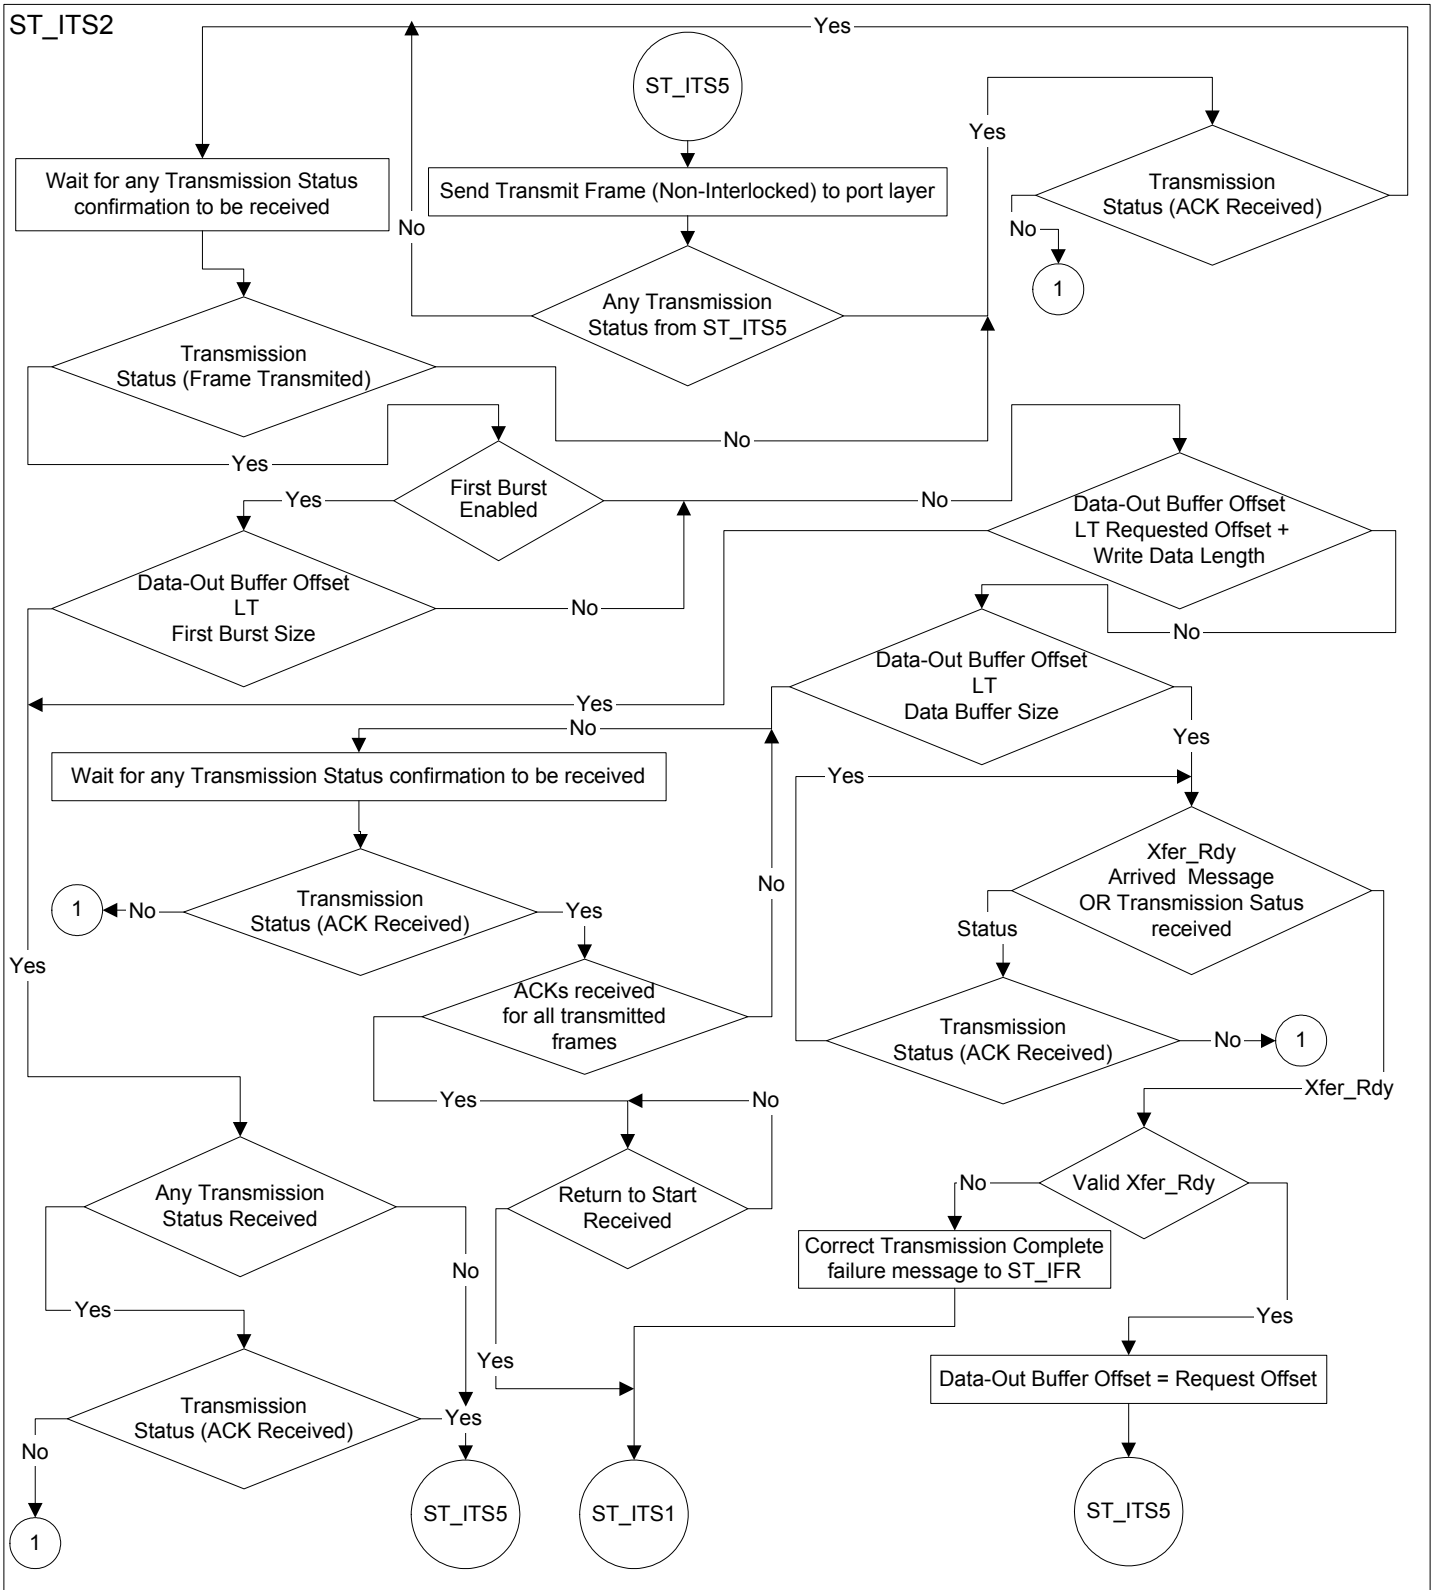


Figure 2 — Representation of transport layer (i.e., ST_ITS2) write data operation (part 1 of 2)

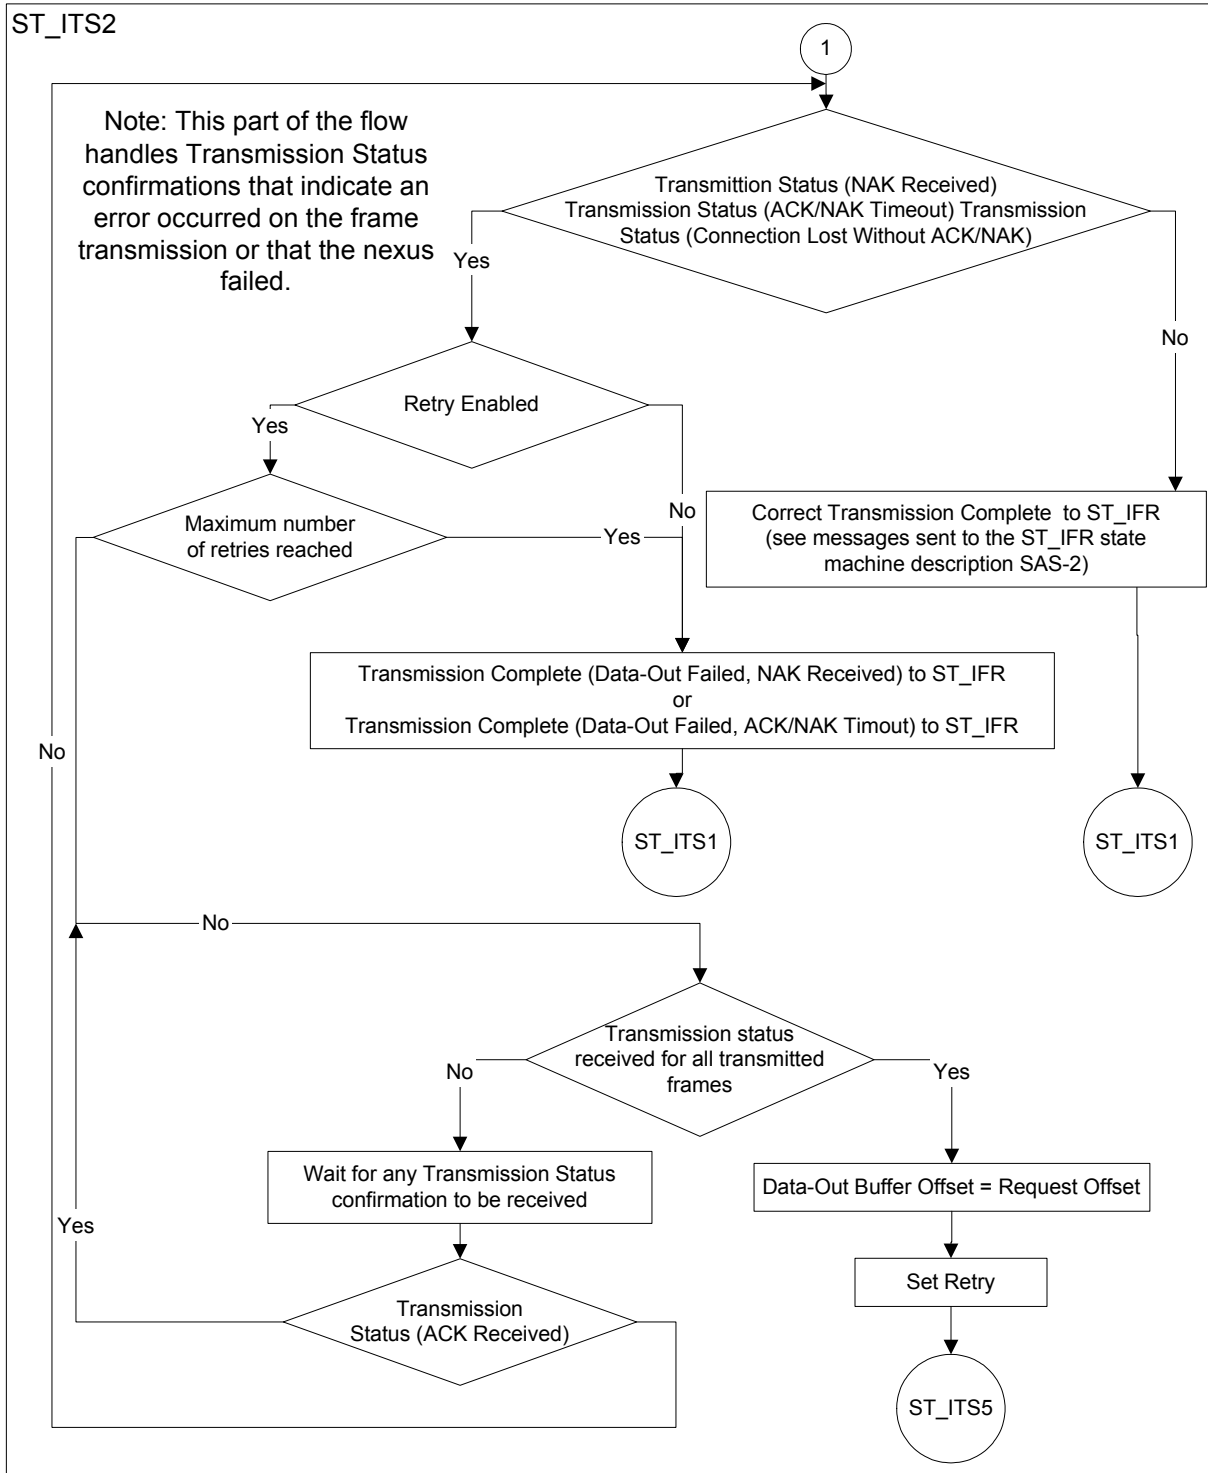


Figure 3 — Representation of transport layer (i.e., ST_ITS2) write data operation (part 2 of 2)

Data-Out Buffer Offset = Offset into the application client write buffer.
 Requested Offset = The value in the current Xfer_Rdy's Requested Offset.
 Write Data Length = The value of the current Xfer_Rdy's Write Data Length.

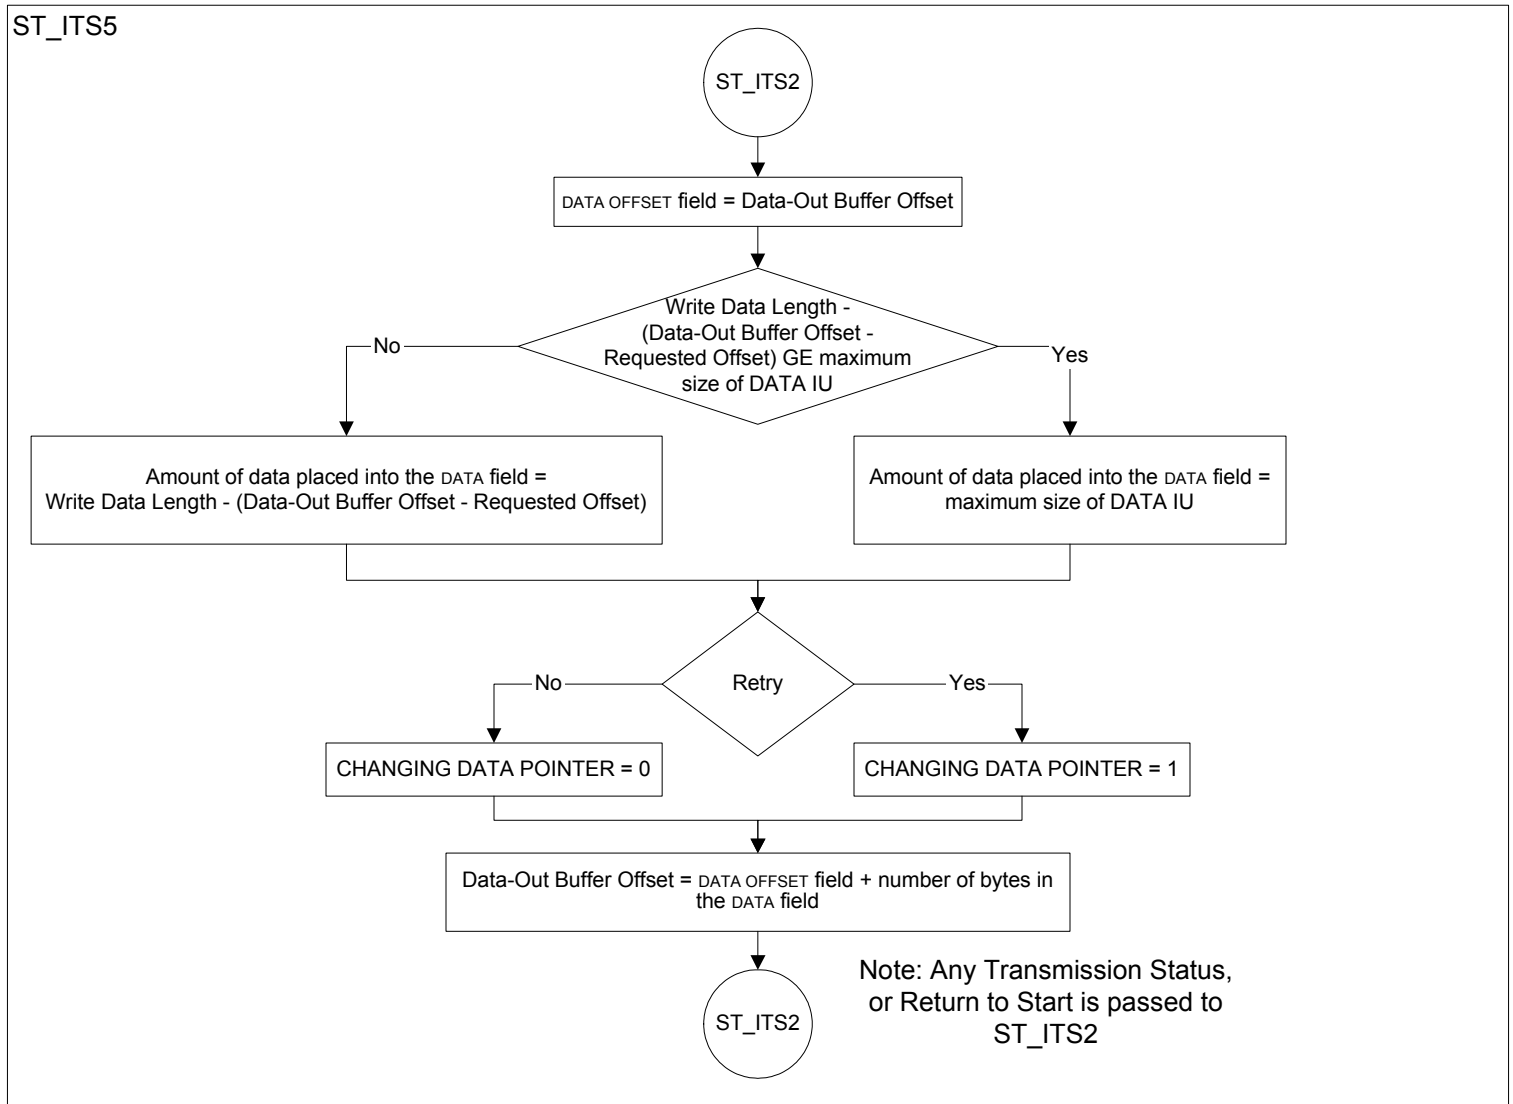


Figure 4 — Representation of transport layer (i.e., ST_ITS5) write data operation