Maxtor Corporation
500 McCarthy Boulevard
Milpitas, CA 95035 USA

To: **T10 SAS Protocol Working Group**
Contact: **Mark Evans**
Phone: 408-894-5310
Email: mark_evans@maxtor.com
Date: **23 February 2004**

Subject: **SAS-1.1, adding the Terminate Data Transfer protocol service**

Introduction

This proposal adds the new protocol service, Terminate Data Transfer, to the SSP transport layer state machines and clarifies what happens with tasks when a task management request is received by the transport layer. Terminate Data Transfer provides a mechanism for a logical unit to terminate requests that have been sent to SCSI target ports. Without the Terminate Data Transfer service, it is possible that requests for a logical unit could remain in a SCSI target port after that logical unit was reset as the result of a hard reset received on another SCSI target port. Included in this proposal are changes to the SSP initiator state machines to clarify how initiators clear tasks from ports as the result of task management requests and other clean up of the SSP transport layer state machine descriptions.

As was agreed at the working group in Chandler in January, this revision of the proposal includes the portions for clause 9 of the 03-165r5 SAS-1.1 Transport layer retries proposal integrated into a rewrite of that clause. This proposal is based on SAS1r03.

Item 1) In **Figure 106 - PL_OC (port layer overall control)** state machine in clause **8.2.2.1 PL_OC state machine overview**: add a Cancel message going from PL_OC2:Overall_Control state to the PL_PM state machines.

Item 2) In clause **8.2.2.3.7 PL_OC2:Overall_Control state frame transmission cancellations**: change the text to be as follows:

Cancel requests cause this state to cancel previous Transmit Frame requests. A Cancel request shall include the following arguments:

a) the destination SAS address; and
b) the tag.

If this state receives a Cancel request and has not sent a Tx Frame message to a PL_PM state machine for the Transmit Frame request specified by the Cancel request, then this state shall:

a) discard all Transmit Frame requests for the destination SAS address and tag; and
b) send a Transmission Status (Cancel Acknowledge) confirmation to the transport layer.

If this state receives a Cancel request and a Tx Frame message has been sent to a PL_PM state machine for the Transmit Frame request specified by the Cancel request, then this state shall send a Cancel message to

the PL_PM state machine to which the Tx Frame message was sent. The Cancel message shall include the tag.

Item 3) In **Figure 108 - PL_PM (port layer phy manager) state machine (part 2)** in clause **8.2.3.1 PL_PM state machine overview**: add a Cancel message coming from the PL_OC state machine going to the PL_PM3:Connected state.

Item 4) In clause **8.2.3.4.1 PL_PM3:Connected state description**: add the following (possibly after the paragraph that begins, “If this state receives an ACK/NAK Timeout confirmation...”).

If this state receives a Cancel message, then this state shall:

a) discard all Tx Frame requests for the destination SAS address and tag; and
b) send a Transmission Status (Cancel Acknowledge) confirmation to the transport layer.

Item 5) Replace clause 9 with the following (minor editorial changes are marked in this section with change bars, significant changes are marked with change bars and have the text in blue and underlined:

9 Transport layer

9.1 Transport layer overview

The transport layer defines frame formats. Transport layer state machines interface to the application layer and port layer and construct and parse frame contents. For SSP, the transport layer only receives frames from the port layer for which an ACK will be transmitted by the link layer.

9.2 SSP transport layer

9.2.1 SSP frame format

Table 1 defines the SSP frame format.

Table 1 — SSP frame format

Byte\Bit	7	6	5	4	3	2	1	0
0	FRAME TYPE							
1	(MSB)	HASHED DESTINATION SAS ADDRESS						
3								(LSB)
4	Reserved							
5	(MSB)	HASHED SOURCE SAS ADDRESS						
7								(LSB)
8	Reserved							
9	Reserved							
10	Reserved					RETRY DATA FRAMES	RETRANSMIT	CHANGING DATA POINTER
11	Reserved						NUMBER OF FILL BYTES	
12	Reserved							
13	Reserved							
15								
16	(MSB)	TAG						
17								(LSB)
18	(MSB)	TARGET PORT TRANSFER TAG						
19								(LSB)
20	(MSB)	DATA OFFSET						
23								(LSB)
24	INFORMATION UNIT							
m								
	Fill bytes, if needed							
n - 3	(MSB)	CRC						
n								(LSB)

Table 2 defines the frame type field, which defines the format of the information unit field.

Table 2 — FRAME TYPE **field**

Code	Name of frame	Information unit	Originator	Information unit size (bytes)	Reference
01h	DATA frame	Data	SSP initiator port or SSP target port	1 to 1024	9.2.2.4
05h	XFER_RDY frame	Transfer ready	SSP target port	12	9.2.4.3
06h	COMMAND frame	Command	SSP initiator port	28 to 284	9.2.4.1
07h	RESPONSE frame	Response	SSP target port	24 to 1024	9.2.4.5
16h	TASK frame	Task management function	SSP initiator port	28	9.2.4.2
F0h - FFh	Vendor specific				
All others	Reserved				

The HASHED DESTINATION SAS ADDRESS field contains the hashed value of the destination SAS address (see 4.2.3). See 9.2.6.2.5 and 9.2.6.3.2 for transport layer requirements on checking this field.

The HASHED SOURCE SAS ADDRESS field contains the hashed value of the source SAS address (see 4.2.3). See 9.2.6.2.5 and 9.2.6.3.2 for transport layer requirements on checking this field.

The RETRY DATA FRAMES bit is set to one for XFER_RDY frames under the conditions described in 9.2.4 and shall be set to zero for all other frame types. When set to one this bit indicates the SSP initiator port may retry write DATA frames that fail.

The RETRANSMIT bit is set to one for TASK frames, RESPONSE frames, and XFER_RDY frames under the conditions defined in 9.2.4 and shall be set to zero for all other frame types. When set to one this bit indicates the frame is a retransmission after the SSP target port failed in its previous attempt to transmit the frame.

The CHANGING DATA POINTER bit is set to one for DATA frames under the conditions described in 9.2.4 and shall be set to zero for all other frame types. When set to one this bit indicates the frame is a retransmission after the SSP target port failed in its previous attempt to transmit the frame or a subsequent frame and the data offset field of the frame may not be sequentially increased from that of the previous frame.

The NUMBER OF FILL BYTES field indicates the number of fill bytes between the information unit field and the CRC field. The number of fill bytes field shall be set to zero for all frame types except DATA frames (i.e., all other frame types are four-byte aligned).

The TAG field contains a value that allows the SSP initiator port to establish a context for commands and task management functions.

For COMMAND and TASK frames, the SSP initiator port shall set the TAG field to a value that is unique for the I_T nexus established by the connection (see 7.12). An SSP initiator port shall not reuse the same tag when transmitting COMMAND or TASK frames to different logical units in the same SSP target port. An SSP initiator port may reuse a tag when transmitting frames to different SSP target ports. The TAG field in a COMMAND frame contains the task tag defined in SAM-3. The tag field in a TASK frame does not correspond to a SAM-3 task tag, but corresponds to a SAM-3 association (see 10.2.1). The tag space used in the TAG fields is shared across COMMAND and TASK frames (e.g., if a tag is used for a COMMAND frame, it is not used for a concurrent TASK frame).

For DATA, XFER_RDY, and RESPONSE frames, the SSP target port shall set the TAG field to the tag of the command or task management function to which the frame pertains.

The TARGET PORT TRANSFER TAG field provides an additional optional method for an SSP target port to establish a write data context when it has sent more than one XFER_RDY frame and may receive data for

more than one tag. SSP target ports may set the target port transfer tag field to any value when transmitting a frame. SSP target ports that use this field should set the field in every XFER_RDY frame to a value that is unique for the L_Q portion of the I_T_L_Q nexus.

SSP initiator ports shall set the TARGET PORT TRANSFER TAG field as follows:

a) For each DATA frame that is sent in response to a XFER_RDY frame, the SSP initiator port shall set the TARGET PORT TRANSFER TAG field to the value that was in the corresponding XFER_RDY frame;
b) For each DATA frame that is sent containing first burst data (see 9.2.2.4), the SSP initiator port shall set the TARGET PORT TRANSFER TAG field to FFFFh; and
c) For frames other than DATA frames, the SSP initiator port shall set the TARGET PORT TRANSFER TAG field to FFFFh.

For DATA frames, the DATA OFFSET field is described in 9.2.2.4. For all other frame types, the DATA OFFSET field shall be ignored.

The INFORMATION UNIT field contains the information unit, the format of which is defined by the frame type field. The maximum size of the INFORMATION UNIT FIELD is 1 024 bytes, making the maximum size of the frame 1052 bytes (1024 bytes of data + 24 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the INFORMATION UNIT field so the CRC field is aligned on a four byte boundary. The number of fill bytes are indicated by the NUMBER OF FILL BYTES field. The contents of the fill bytes are vendor specific.

The CRC field contains a CRC value (see 7.5) that is computed over the entire SSP frame prior to the CRC field including the fill bytes (i.e., all data dwords between the SOF and EOF). The CRC field is checked by the link layer (see 7.16), not the transport layer.

9.2.2 Information units

9.2.2.1 COMMAND information unit

Table 3 defines the command IU. A COMMAND frame is sent by an SSP initiator port to request that a command be processed by the device server in a logical unit.

Table 3 — COMMAND information unit

Byte\Bit	7	6	5	4	3	2	1	0
0	LOGICAL UNIT NUMBER							
7								
8	Reserved							
9	ENABLE FIRST BURST	Reserved				TASK ATTRIBUTE		
10	Reserved							
11	ADDITIONAL CDB LENGTH (n dwords)						Reserved	
12	CDB							
27								
28	ADDITIONAL CDB BYTES							
27+n×4								

The LOGICAL UNIT NUMBER field contains the address of the logical unit. The structure of the LOGICAL UNIT NUMBER field shall be as defined in SAM-3. If the addressed logical unit does not exist, the task manager shall follow the rules for selection of invalid logical units defined in SAM-3.

If the ENABLE FIRST BURST bit is set to one, then first burst data shall be transferred as defined by the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see 10.2.6.1). If the ENABLE FIRST BURST bit set to zero, then the FIRST BURST SIZE field in the Disconnect-Reconnect mode page shall be ignored (i.e., there shall be no first burst data transferred for the command). Application clients shall only set the ENABLE FIRST BURST bit to one if:

a) the FIRST BURST SIZE field in the Disconnect-Reconnect mode page is non-zero or changeable; and
b) the logical unit and target port comply with SAS-1.1 or later (e.g., as reported in the standard INQUIRY data version descriptors (see SPC-3)).

The TASK ATTRIBUTE field is defined in table 4.

Table 4 — TASK ATTRIBUTE field

Code	Task attribute	Description
000b	SIMPLE	Specifies that the task is to be managed according to the rules for a simple task attribute (see SAM-3).
001b	HEAD OF QUEUE	Specifies that the task is to be managed according to the rules for a head of queue task attribute (see SAM-3).
010b	ORDERED	Specifies that the task is to be managed according to the rules for an ordered task attribute (see SAM-3).
011b	Reserved	
100b	ACA	Specifies that the task is to be managed according to the rules for an automatic contingent allegiance task attribute (see SAM-3).
101b-111b	Reserved	

The ADDITIONAL CDB LENGTH field contains the length in dwords (four bytes) of the additional cdb field.

The CDB and ADDITIONAL CDB BYTES fields together contain the CDB to be interpreted by the addressed logical unit. Any bytes between the end of the CDB and the end of the two fields shall be ignored (e.g., a six-byte CDB occupies the first six bytes of the CDB field; the remaining ten bytes are ignored; and the ADDITIONAL CDB BYTES field is not present).

The contents of the CDB are defined in the SCSI command standards (e.g., SPC-3).

9.2.2.2 TASK information unit

Table 5 defines the task management function IU. A TASK frame is sent by an SSP initiator port to request that a task management function be processed by the task manager in a logical unit.

Table 5 — TASK information unit

Byte\Bit	7	6	5	4	3	2	1	0
0	LOGICAL UNIT NUMBER							
7								
8	Reserved							
9	Reserved							
10	TASK MANAGEMENT FUNCTION							
11	Reserved							
12	(MSB) TAG OF TASK TO BE MANAGED							
13								(LSB)
14	Reserved							
27								

The LOGICAL UNIT NUMBER FIELD contains the address of the logical unit. The structure of the LOGICAL UNIT NUMBER field shall be as defined in SAM-3. If the addressed logical unit does not exist, the task manager shall return a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and its RESPONSE CODE field set to INVALID LOGICAL UNIT NUMBER.

Table 6 defines the TASK MANAGEMENT FUNCTION field.

Table 6 — Task management functions

Code	Task management function	Uses LOGICAL UNIT NUMBER field	Uses TAG OF TASK TO BE MANAGED field	Description
01h	ABORT TASK	yes	yes	The task manager shall perform the ABORT TASK task management function with L set to the value of the LOGICAL UNIT NUMBER field and Q set to the value of the TAG OF TASK TO BE MANAGED field (see SAM-3).
02h	ABORT TASK SET	yes	no	The task manager shall perform the ABORT TASK SET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
04h	CLEAR TASK SET	yes	no	The task manager shall perform the CLEAR TASK SET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
08h	LOGICAL UNIT RESET	yes	no	The task manager shall perform the LOGICAL UNIT RESET task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
20h	Reserved [a]			
40h	CLEAR ACA	yes	no	The task manager shall perform the CLEAR ACA task management function with L set to the value of the LOGICAL UNIT NUMBER field (see SAM-3).
80h	QUERY TASK	yes	yes	The task manager shall perform the QUERY TASK task management function with L set to the value of the LOGICAL UNIT NUMBER field and Q set to the value of the TAG OF TASK TO BE MANAGED field (see SAM-3).
All others	Reserved			

[a] The TARGET RESET task management function defined in SAM-3 is not supported.

If the TASK MANAGEMENT FUNCTION field contains a reserved or unsupported value, the task manager shall return a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and its RESPONSE CODE field set to TASK MANAGEMENT FUNCTION NOT SUPPORTED.

If the TASK MANAGEMENT FUNCTION field is set to ABORT TASK or QUERY TASK, the TAG OF TASK TO BE MANAGED field specifies the tag value from the COMMAND frame that contained the task to be aborted or checked. For all other task management functions, the TAG OF TASK TO BE MANAGED field shall be ignored.

9.2.2.3 XFER_RDY information unit

Table 7 defines the transfer ready IU. An XFER_RDY frame is sent by an SSP target port to request write data from the SSP initiator port.

Table 7 — XFER_RDY information unit

Byte\Bit	7	6	5	4	3	2	1	0
0	(MSB)			REQUESTED OFFSET				
3								(LSB)
4	(MSB)			WRITE DATA LENGTH				
7								(LSB)
8				Reserved				
11								

The REQUESTED OFFSET field contains the application client buffer offset of the segment of write data that the SSP initiator port may transmit to the logical unit (using DATA frames). The requested offset shall be a multiple of four (i.e., each DATA frame shall begin transferring data on a dword boundary). The REQUESTED OFFSET field shall be zero for the first XFER_RDY frame of a command unless:

a) the ENABLE FIRST BURST field in the COMMAND frame (see 9.2.4.1) was set to one; and
b) the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see 10.2.6.1.5) is not set to zero.

If the ENABLE FIRST BURST field in the COMMAND frame (see 9.2.4.1) was set to one, then in the initial XFER_RDY frame for the command, the SSP target port shall set the REQUESTED OFFSET field to the value indicated by the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see 10.2.6.1.5). If any additional XFER_RDY frames are required, the REQUESTED OFFSET field shall be set to the value of the previous XFER_RDY frame's REQUESTED OFFSET field plus the value of the previous XFER_RDY frame's WRITE DATA LENGTH field.

The WRITE DATA LENGTH field contains the number of bytes of write data the SSP initiator port may transmit to the logical unit (using DATA frames) from the application client buffer starting at the requested offset. The SSP target port shall set the WRITE DATA LENGTH field to a value greater than or equal to 00000001h. If the value in the MAXIMUM BURST SIZE field in the Disconnect-Reconnect mode page is not zero, the SSP target port shall set the WRITE DATA LENGTH field to a value less than or equal to the value in the MAXIMUM BURST SIZE field (see 10.2.6.1.4).

If an SSP target port transmits a XFER_RDY frame containing a value in the WRITE DATA LENGTH field that is not divisible by four, the SSP target port shall not transmit any subsequent XFER_RDY frames for that command (i.e., only the last XFER_RDY for a command may request a non-dword multiple write data length).

9.2.2.4 DATA information unit

Table 8 defines the data IU. A DATA frame is sent by an SSP initiator port to deliver write data and is sent by an SSP target port to deliver read data. The maximum size of the data IU is the maximum size of any IU in an SSP frame (see 9.2.1). The minimum size of the data IU is one byte.

Table 8 — DATA information unit

Byte\Bit	7	6	5	4	3	2	1	0
0				DATA				
n-1								

The data field contains the read or write data.

An SSP initiator port shall only transmit a DATA frame:

a) in response to an XFER_RDY frame; or
b) after transmitting a COMMAND frame if the ENABLE FIRST BURST field in the COMMAND frame was set to one (see 9.2.4.1) and the FIRST BURST SIZE field in the Disconnect-Reconnect mode page is not zero (see 10.2.6.1.5).

If the value in the MAXIMUM BURST SIZE field in the Disconnect-Reconnect mode page is not zero, the maximum amount of data that is transferred at one time by an SSP target port per I_T_L_Q nexus is limited by the value in the MAXIMUM BURST SIZE field (see 10.2.6.1.4).

The DATA frame shall only contain write data for a single XFER_RDY frame.

If an SSP target port transmits a DATA frame containing a non-zero value in the NUMBER OF FILL BYTES field in the frame header (see 9.2.1), the SSP target port shall not transmit any subsequent DATA frames for that command (i.e., only the last read DATA frame for a command may have data with a length that is not a multiple of four).

An SSP initiator port may set the NUMBER OF FILL BYTES field to a non-zero value in the last DATA frame that it transmits in response to a XFER_RDY. An SSP initiator port shall set the NUMBER OF FILL BYTES field in the frame header (see 9.2.1) to zero in all other DATA frames that it transmits.

NOTE 1 - Combined with the restrictions on the WRITE DATA LENGTH field in the XFER_RDY frame (see 9.2.2.3), this ensures that only the last write DATA frame for a command may have data with a length that is not a multiple of four bytes).

An SSP initiator port shall not transmit a DATA frame for a given I_T_L_Q nexus after it has sent a TASK frame that terminates that task (e.g., an ABORT TASK).

The DATA OFFSET field in the frame header (see 9.2.1) contains the application client buffer offset as described by SAM-3. The data offset shall be a multiple of four (i.e., each DATA frame shall transfer data beginning on a dword boundary).

An SSP target port shall set the DATA OFFSET field in the initial read DATA frame for a given command to zero. If any additional read DATA frames are required for the command, then the SSP target port shall set the DATA OFFSET field to the value of the previous read DATA frame's data offset plus the previous read DATA frame's data length.

An SSP initiator port shall set the DATA OFFSET field in the initial write DATA frame for a given command to zero. If any additional write DATA frames are required for the command, then the SSP initiator port shall set the DATA OFFSET field to the value of the previous read DATA frame's data offset plus the previous read DATA frame's data length.

9.2.2.5 RESPONSE information unit

9.2.2.5.1 RESPONSE information unit overview

Table 9 defines the response IU. A RESPONSE frame is sent by an SSP target port to deliver SCSI status (e.g., GOOD or CHECK CONDITION) and sense data, or to deliver SSP-specific status (e.g., illegal frame

format). The maximum size of the RESPONSE frame is the maximum size of any IU in an SSP frame (see 9.2.1).

Table 9 — RESPONSE information unit

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved							
9								
10	Reserved						DATAPRES	
11	STATUS							
12	Reserved							
15								
16	(MSB)							
19	SENSE DATA LENGTH (n bytes)							(LSB)
20	(MSB)							
23	RESPONSE DATA LENGTH (m bytes)							(LSB)
24	RESPONSE DATA							
23+m								
24+m	SENSE DATA							
23+m+n								

The DATAPRES field indicates the format and content of the status field, sense data length field, response data length field, response data field, and sense data field. Table 10 defines the values for the DATAPRES field.

Table 10 — DATAPRES field

Code	Name	Description	Reference
00b	NO DATA	No data present	9.2.2.5.2
01b	RESPONSE_DATA	Response data present	9.2.2.5.3
10b	SENSE_DATA	Sense data present	9.2.2.5.4
11b	Reserved		

If a command completes without sense data to return, then the SSP target port shall return a RESPONSE frame with the DATAPRES field set to NO_DATA.

If a command completes with sense data to return (e.g., CHECK CONDITION status), then the SSP target port shall return a RESPONSE frame with the DATAPRES field set to SENSE_DATA.

An SSP target port shall return a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA in response to every TASK frame and in response to errors that occur while the transport layer is processing a COMMAND frame.

If the DATAPRES field is set to a reserved value, then the SSP initiator port shall discard the RESPONSE frame.

9.2.2.5.2 RESPONSE information unit NO_DATA format

If the DATAPRES field is set to NO_DATA, then:

a) the SSP target port shall set the STATUS field to the status code for a command that has ended (see SAM-3 for a list of status codes);
b) the SSP target port shall set the SENSE DATA LENGTH field and the RESPONSE DATA LENGTH field to zero;
c) the SSP initiator port shall ignore the SENSE DATA LENGTH field and the RESPONSE DATA LENGTH field; and
d) the SSP target port shall not include the SENSE DATA field or the RESPONSE DATA field.

9.2.2.5.3 RESPONSE information unit RESPONSE_DATA format

If the DATAPRES field is set to RESPONSE_DATA, then:

a) the SSP target port shall set the STATUS field and the SENSE DATA LENGTH field to zero;
b) the SSP initiator port shall ignore the STATUS field and the SENSE DATA LENGTH field;
c) the SSP target port shall not include the SENSE DATA field;
d) the SSP target port shall set the RESPONSE DATA LENGTH field to four. Other lengths are reserved for future standardization; and
e) the SSP target port shall include the RESPONSE DATA field.

The RESPONSE DATA field contains information describing protocol failures detected during processing of a request received by the SSP target port. The RESPONSE DATA field shall be present if the SSP target port detects any of the conditions described by a non-zero response code value and shall be present for a RESPONSE frame sent in response to a TASK frame. Table 11 defines the RESPONSE DATA field.

Table 11 — RESPONSE DATA field

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved							
1	Reserved							
2	Reserved							
3	RESPONSE CODE							

The RESPONSE CODE field indicates the error condition or the completion status of a task management function. See 10.2.1.5 and 10.2.1.13 for the mapping of these response codes to SCSI service responses. Table 12 defines the RESPONSE CODE field.

Table 12 — RESPONSE CODE field

Code	Description
00h	TASK MANAGEMENT FUNCTION COMPLETE [a]
02h	INVALID FRAME
04h	TASK MANAGEMENT FUNCTION NOT SUPPORTED [a]
05h	TASK MANAGEMENT FUNCTION FAILED [a]
08h	TASK MANAGEMENT FUNCTION SUCCEEDED [a]
09h	INVALID LOGICAL UNIT NUMBER [a]
All others	Reserved

[a] Only valid when responding to a TASK frame

9.2.2.5.4 RESPONSE information unit SENSE_DATA format

If the DATAPRES field is set to SENSE_DATA, then:

a) the SSP target port shall set the STATUS field to the status code for a command that has ended (see SAM-3 for a list of status codes);
b) the SSP target port shall set the RESPONSE DATA LENGTH field shall to zero;
c) the SSP initiator port shall ignore the RESPONSE DATA LENGTH field;
d) the SSP target port shall not include the RESPONSE DATA field;
e) the SSP target port shall set the SENSE DATA LENGTH field to a non-zero value indicating the number of bytes in the SENSE DATA field. The SSP target port shall not set the value in the SENSE DATA LENGTH field to be greater than 1000 (see table 2); and
f) the SSP target port shall set the SENSE DATA field to the sense data (see SAM-3).

The value in the SENSE DATA LENGTH field need not be a multiple of four. If it is not, the NUMBER OF FILL BYTES field in the SSP frame header is non-zero and fill bytes are present.

9.2.3 Sequences of SSP frames

Figure 1, figure 2, figure 3, and figure 4 show examples of the sequences of frames for single task management functions and commands. Frames may be interleaved in any order when multiple commands and/or task management functions are outstanding. Frames may be transmitted during one or more connections (e.g., the COMMAND frame could be transmitted in a connection originated by the SSP initiator port, and the DATA frames and RESPONSE frame transmitted in one or more connections originated by the SSP target port). RESPONSE frames may be returned in any order (i.e., the order in which TASK frames and COMMAND frames are sent has no effect on the order that RESPONSE frames are returned).

Figure 1 shows the sequence of SSP frames for a task management function, including the transport protocol services (see 10.2.1) invoked by the SCSI application layer.

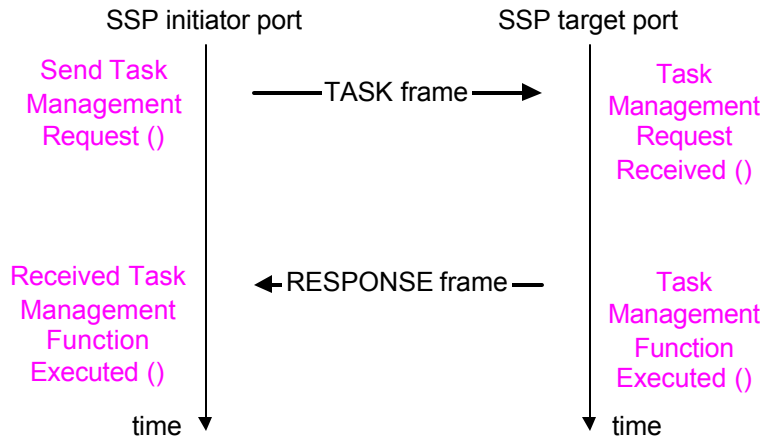

Figure 1 — Task management function sequence of SSP frames

Figure 2 shows the sequence of SSP frames for a write command, including the transport protocol services (see 10.2.1) invoked by the SCSI application layer.

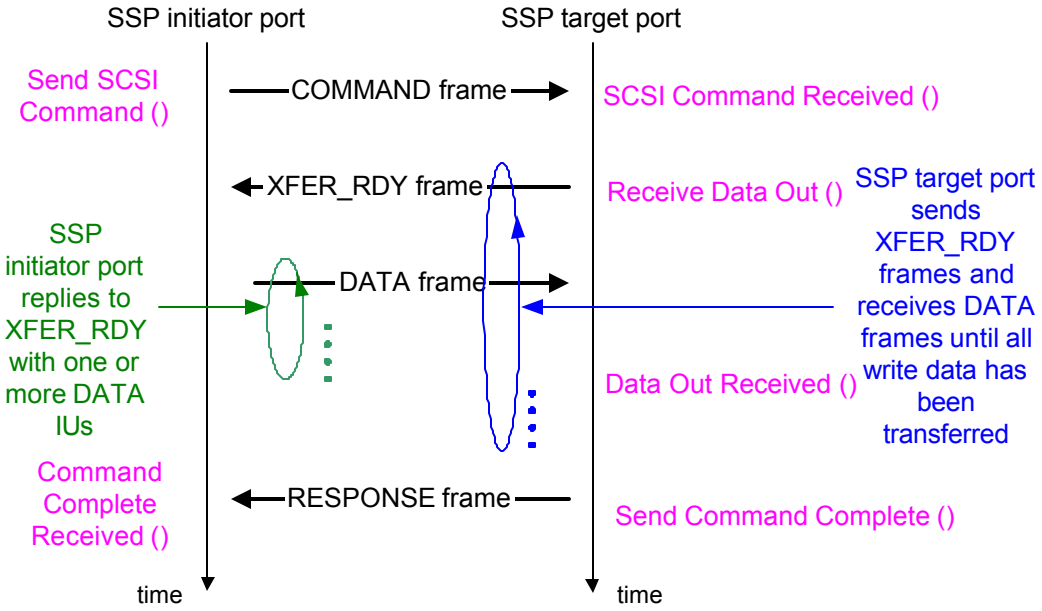

Figure 2 — Write command sequence of SSP frames

Figure 3 shows the sequence of SSP frames for a read command, including the transport protocol services (see 10.2.1) invoked by the SCSI application layer.

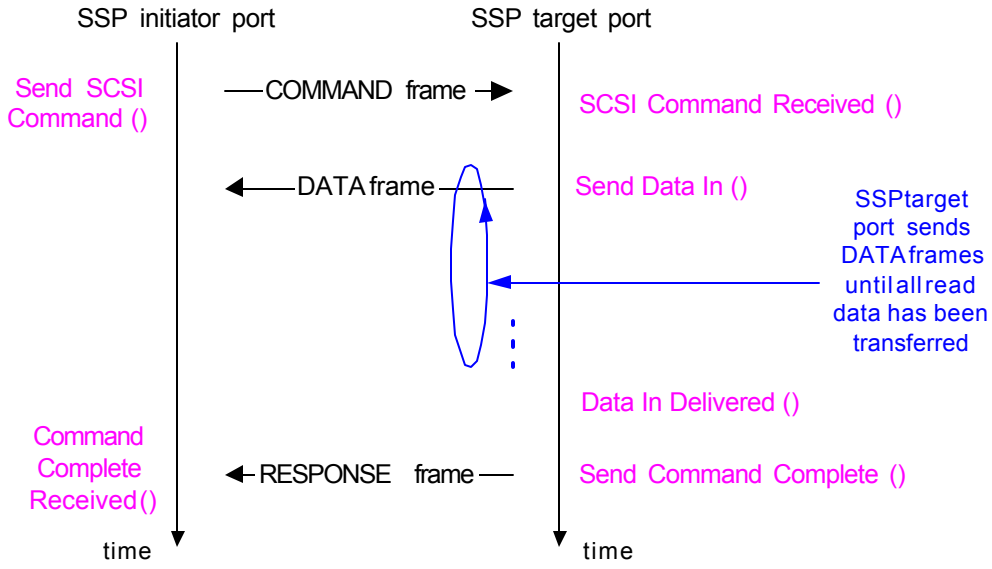

Figure 3 — Read command sequence of SSP frames

Figure 4 shows the sequence of SSP frames for a bidirectional command, including the transport protocol services (see 10.2.1) invoked by the SCSI application layer.

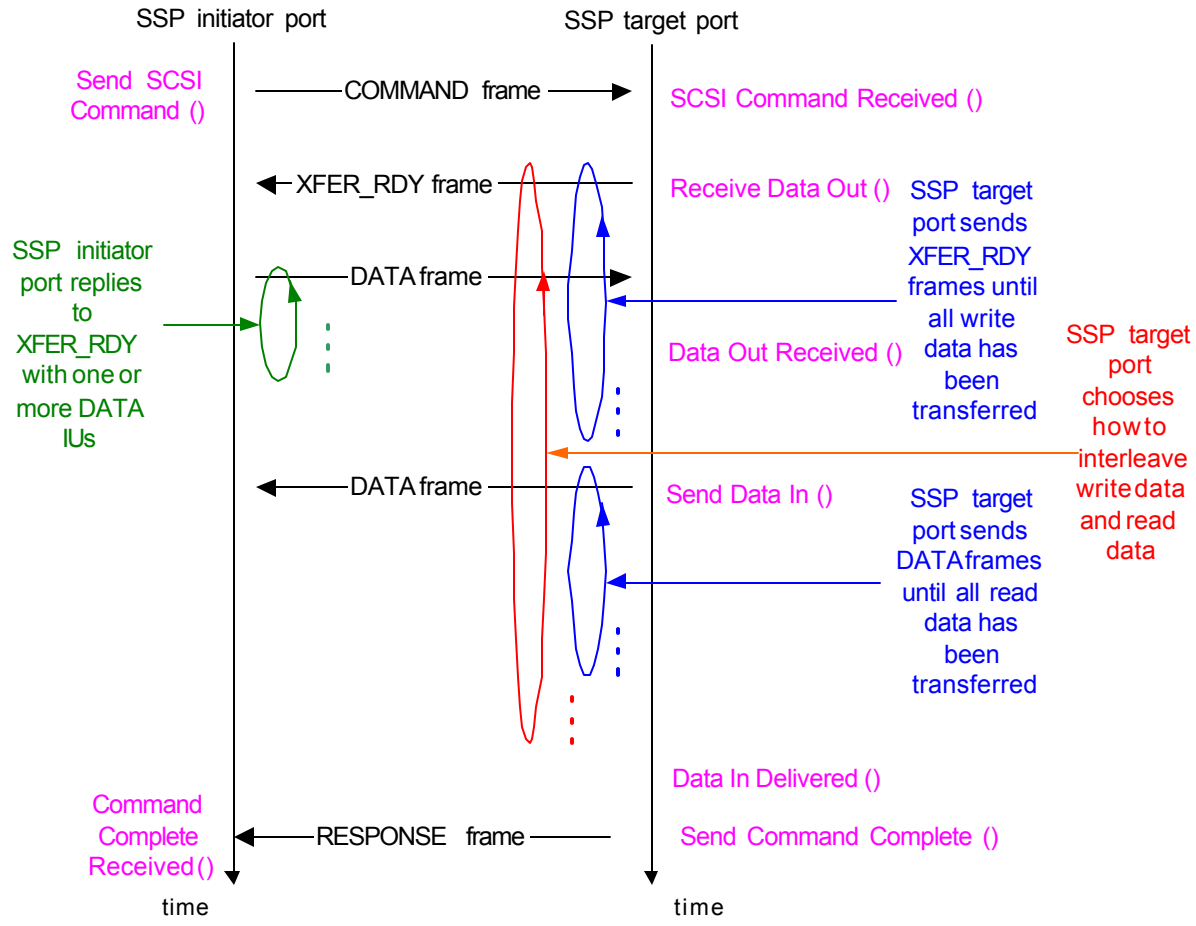

Figure 4 — Bidirectional command sequence of SSP frames

9.2.4 SSP transport layer handling of link layer errors

9.2.4.1 SSP transport layer handling of link layer errors overview

The transport layer, sometimes assisted by the application layer, handles some link layer errors (e.g., NAKs and ACK/NAK timeouts).

Link layer errors that are encountered when transmitting XFER_RDY and DATA frames are handled differently based on the TRANSPORT LAYER RETRIES bit in the Protocol Specific Logical Unit mode page (see 10.2.6.3) of the logical unit that is the source of the frame.

If the TRANSPORT LAYER RETRIES bit is set to zero, a logical unit:

a) does not support transport layer retries;
b) sets the RETRY DATA FRAMES bit to zero in each XFER_RDY frame;
c) may or may not select a different value for the TARGET PORT TRANSFER TAG field in each XFER_RDY frame than that used in the previous XFER_RDY frame for that I_T_L_Q nexus;
d) processes XFER_RDY frame link layer errors as described in 9.2.4.4.3; and
e) processes DATA frame link layer errors as described in 9.2.4.5.3.

If the TRANSPORT LAYER RETRIES bit is set to one, the logical unit:

a) supports transport layer retries;
b) supports the QUERY TASK task management function (see SAM-3);
c) sets the RETRY DATA FRAMES bit to one in each XFER_RDY frame;

d) selects a different value for the TARGET PORT TRANSFER TAG field in each XFER_RDY frame than that used in the previous XFER_RDY frame for that I_T_L_Q nexus;
e) processes XFER_RDY frame link layer errors as described in 9.2.4.4.2; and
f) processes DATA frame link layer errors as described in 9.2.4.5.2.

9.2.4.2 COMMAND frame - handling of link layer errors

If an SSP initiator port transmits a COMMAND frame and does not receive an ACK or NAK for that frame (e.g., times out or the connection is broken) it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) transmits, in a new connection, a QUERY TASK task management function to determine whether the command was received (see 10.2.2). The TAG OF TASK TO BE MANAGED field is set to the tag of the COMMAND frame.

If the SSP initiator port receives an XFER_RDY frame for the I_T_L_Q nexus of the command before the RESPONSE frame for the QUERY TASK, then the COMMAND frame was received and is being processed by the target port, and the XFER_RDY frame is valid.

If the SSP initiator port receives a RESPONSE frame for the I_T_L_Q nexus of the command before the RESPONSE frame for the QUERY TASK, then the COMMAND frame was received by the target port, the RESPONSE frame is valid, and the command processing is complete. The SSP initiator port may reuse the tag of the COMMAND frame.

If the SSP initiator port receives a RESPONSE frame for the QUERY TASK with a response code of TASK MANAGEMENT FUNCTION SUCCEEDED, then the COMMAND frame was received by the target port (i.e., ACKed) and the command is being processed.

If the SSP initiator port receives a RESPONSE frame for the QUERY TASK with a response code of TASK MANAGEMENT FUNCTION COMPLETE, then the COMMAND frame is not being processed. If neither an XFER_RDY frame nor a RESPONSE frame has been received for the I_T_L_Q nexus of the command, then the COMMAND frame was not received (i.e., NAKed or lost). The SSP initiator port may reuse the tag of the COMMAND frame.

An SSP initiator port should retransmit each COMMAND frame that does not receive an ACK at least one time.

9.2.4.3 TASK frame - handling of link layer errors

If an SSP initiator port transmits a TASK frame and does not receive an ACK or NAK for that frame (e.g., times out or the connection is broken) it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) retransmits, in a new connection, the TASK frame using the same tag (see 10.2.2) and with the RETRANSMIT bit set to one (see 10.2.2).

If the SSP initiator port receives a RESPONSE frame for the TASK frame that arrives before the ACK or NAK for the TASK frame, then the TASK frame was received (i.e., ACKed, the RESPONSE frame is valid, and the task management function is complete. The initiator port may reuse the tag of the TASK frame.

An SSP initiator port should retransmit each TASK frame that does not receive an ACK at least one time.

9.2.4.4 XFER_RDY frame - handling of link layer errors

9.2.4.4.1 XFER_RDY frame overview

If the TRANSPORT LAYER RETRIES bit is set to one in the Protocol-Specific Logical Unit mode page (see 10.x.x.x), then the SSP target port processes XFER_RDY frame link layer errors as described in 9.2.4.4.2.

If the TRANSPORT LAYER RETRIES bit is set to zero, then the SSP target port processes XFER_RDY frame link layer errors as described in 9.2.4.4.3.

9.2.4.4.2 XFER_RDY frame with transport layer retries

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), it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) 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 (see 9.2.6.3.3.6).

If an SSP target port transmits an XFER_RDY frame and receives a NAK for that frame, it retransmits the XFER_RDY frame with a different value in the TARGET PORT TRANSFER TAG field and with the RETRANSMIT bit set to one (see 9.2.6.3.3.6).

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, it should stop processing the previous XFER_RDY frame (i.e., stop sending write DATA frames) and start servicing the new XFER_RDY frame. The SSP initiator port does not send any write DATA frames for the previous XFER_RDY after sending a write DATA frame for the new XFER_RDY.

An SSP target port may reuse the value in the TARGET PORT TRANSFER TAG field from the previous XFER_RDY frame when 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.

9.2.4.4.3 XFER_RDY frame without transport layer retries

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), it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) transmits, in a new connection, a RESPONSE frame returning a CHECK CONDITION status for that command with a sense key of ABORTED COMMAND and an additional sense code of ACK/NAK TIMEOUT (see 10.2.3).

If an SSP target port transmits an XFER_RDY frame and receives a NAK for that frame, it returns a CHECK CONDITION status for that command with a sense key of ABORTED COMMAND and an additional sense code of NAK RECEIVED (see 10.2.3).

9.2.4.5 DATA frame - handling of link layer errors

9.2.4.5.1 DATA frame overview

If an SSP target port sends a read DATA frame for a logical unit that has its TRANSPORT LAYER RETRIES bit set to one in the Protocol-Specific Logical Unit mode page (see 10.x.x.x), then the SSP target port processes read DATA frame link layer errors as described in 9.2.4.5.2. If the logical unit has its TRANSPORT LAYER RETRIES bit set to zero, then the SSP target port processes read DATA frame link layer errors as described in 9.2.4.5.3.

An SSP initiator port processes write DATA frames sent in response to an XFER_RDY frame that has its RETRY DATA FRAMES bit set to one as described in 9.2.4.5.2. An SSP initiator port processes write DATA frames sent in response to an XFER_RDY frame that has its RETRY DATA FRAMES bit set to zero as described in 9.2.4.5.3.

9.2.4.5.2 DATA frame with transport layer retries

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), it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) retransmits, in a new connection, all the read DATA frames since a previous time when ACK/NAK balance occurred (see 7.16.7.3).

If an SSP target port transmits a read DATA frame and receives a NAK for that frame, it retransmits, in the same or in a new connection, all the read DATA frames since a previous time when ACK/NAK balance occurred.

If an SSP initiator port transmits a write DATA frame and does not receive an ACK or NAK for that frame (e.g., times out, or the connection is broken), it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) retransmits, in a new connection, all the write DATA frames for the previous XFER_RDY.

While processing the lack of ACK or NAK, if that SSP initiator port receives a new XFER_RDY or a RESPONSE frame for the command, it processes the XFER_RDY or RESPONSE frame and should stop sending the retransmitted write DATA frames. It does not send a write DATA frame for the previous XFER_RDY after sending a write DATA frame for the new XFER_RDY.

If an SSP initiator port transmits a write DATA frame and receives a NAK for that frame, it retransmits, in the same or in a new connection, all the write DATA frames for the previous XFER_RDY.

For both reads and writes, the first retransmitted DATA frame has its CHANGING DATA POINTER bit set to one and subsequent DATA frames have their CHANGING DATA POINTER bits set to zero.

An SSP port retransmits each DATA frame that does not receive an ACK at least one time. The number of times it retransmits each DATA frame is vendor-specific.

9.2.4.5.3 DATA frame without transport layer retries

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), it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) transmits, in a new connection, a RESPONSE frame returning a CHECK CONDITION status for that command with a sense key of ABORTED COMMAND and an additional sense code of ACK/NAK TIMEOUT (see 10.2.3).

If an SSP target port transmits a read DATA frame and receives a NAK for that frame, it returns a CHECK CONDITION status for that command with a sense key of ABORTED COMMAND and an additional sense code of NAK RECEIVED (see 10.2.3).

If an SSP initiator port transmits a write DATA frame and does not receive an ACK or NAK for that frame (e.g., times out, or the connection is broken), it closes the connection with DONE (ACK/NAK TIMEOUT) and aborts the command (see 10.2.2).

If an SSP initiator port transmits a write DATA frame and receives a NAK for that frame, it aborts the command (see 10.2.2).

9.2.4.6 RESPONSE frame - handling of link layer errors

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), it:

1) closes the connection with DONE (ACK/NAK TIMEOUT); and
2) retransmits, in a new connection, the RESPONSE frame with the RETRANSMIT bit set to one.

If an SSP target port transmits a RESPONSE frame and receives a NAK for that frame, it retransmits the RESPONSE frame at least one time with the RETRANSMIT bit set to zero (see 9.2.6.3.3.8).

An SSP target port retransmits each RESPONSE frame that does not receive an ACK at least one time. 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, it discards the extra RESPONSE frame. If it has not previously received the RESPONSE frame, it considers it to be the valid RESPONSE frame (see 10.2.2).

9.2.5 SSP transport layer error handling

9.2.5.1 SSP target port transport layer error handling summary

This clause contains a summary of how an SSP target port processes transport layer errors. This summary does not include every error case. For each instance in this clause there is a cross reference to where specific behavior is defined in this standard.

If an SSP target port receives an XFER_RDY frame or an unsupported frame type, then the SSP target port discards the frame (see 9.2.6.3.2).

If an SSP target port receives a COMMAND frame, and the frame is too short to contain a LUN field or the frame is too short to contain a CDB, then the SSP target port returns a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE set to INVALID FRAME (see 9.2.6.3.2).

If an SSP target port receives a COMMAND frame, and the ADDITIONAL CDB LENGTH field indicates that the frame should be a different length, then the SSP target port returns a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE set to INVALID FRAME (see 10.2.3).

If an SSP target port receives a TASK frame that is too short, then the SSP target port returns a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE set to INVALID FRAME (see 9.2.6.3.2).

If an SSP target port receives a COMMAND frame with a tag that is already in use, then the SSP target port may return a CHECK CONDITION status with a sense key of ABORTED COMMAND and an additional sense code of OVERLAPPED COMMANDS DETECTED (see 9.2.6.3.2).

If an SSP target port receives a TASK frame with a tag that is already in use, then the SSP target port may return a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE set to INVALID FRAME (see 9.2.6.3.2).

If an SSP target port receives a DATA frame with an unknown tag, then the SSP target port discards the frame (see 9.2.6.3.2).

If an SSP target port receives a DATA frame that does not contain first burst data and for which there is no XFER_RDY frame outstanding, then the SSP target port discards the frame (see 9.2.6.3.2).

If an SSP target port receives a TASK frame with an unknown logical unit number, then the SSP target port returns a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE set to INVALID LOGICAL UNIT (see 10.2.3).

If an SSP target port receives a COMMAND frame or TASK frame with a target port transfer tag set to a value other than FFFFh, then the SSP target port may return a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE set to INVALID FRAME (see 9.2.6.3.2).

If an SSP target port is using target port transfer tags and it receives a DATA frame with an unknown target port transfer tag, then the SSP target port discards the frame (see 9.2.6.3.3.5.1).

If an SSP target port receives a DATA frame with a data offset that was not expected, then the SSP target port discards that frame and any subsequent DATA frames received for that command and terminates the command with a CHECK CONDITION status with a sense key of ABORTED COMMAND and an additional sense code of DATA OFFSET ERROR (see 9.2.6.3.3.5.1).

If an SSP target port receives a DATA frame with more write data than expected (i.e., the length of the DATA frame extends past the end of the expected write data length), then the SSP target port discards the frame and terminates the command with a CHECK CONDITION status with a sense key of ABORTED COMMAND and an additional sense code of TOO MUCH WRITE DATA (see 9.2.6.3.3.5.1).

If an SSP target port receives a zero length DATA frame, then the SSP target port discards the frame and terminates the command with a CHECK CONDITION status with a sense key of ABORTED COMMAND and an additional sense code of INFORMATION UNIT TOO SHORT (see 9.2.6.3.3.5.1).

9.2.5.2 SSP initiator port transport layer error processing summary

This clause contains a summary of how an SSP initiator port processes transport layer errors. This summary does not include every error case. For each instance in this clause there is a cross reference to where specific behavior is defined in this standard.

If an SSP initiator port receives a COMMAND or TASK frame or an unsupported frame type, then the SSP initiator port discards the frame (see 9.2.6.2.2).

If an SSP initiator port receives a DATA, XFER_RDY, or RESPONSE frame with an unknown TAG field value (including a tag for which it has sent a COMMAND or TASK frame but not yet received an ACK), then the SSP initiator port discards the frame (see 9.2.6.2.2). The SSP initiator port may then abort the command with that tag (see 10.2.2).

If an SSP initiator port receives an XFER_RDY frame that is not 12 bytes long, then the SSP initiator port discards the frame (see 9.2.6.2.2). The SSP initiator port may then abort the command with that tag (see 10.2.2).

If an SSP initiator port receives an XFER_RDY frame in response to a command with no write data, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2).

If an SSP initiator port receives an XFER_RDY frame requesting more write data than expected, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2).

If an SSP initiator port receives an XFER_RDY frame requesting zero bytes, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2).

If an SSP initiator port receives an XFER_RDY frame with a requested offset that was not expected, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2).

If an SSP initiator port receives a DATA frame with more read data than expected, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2). The SSP initiator port may receive a RESPONSE for the command before being able to abort the command.

If an SSP initiator port receives a DATA frame with zero bytes, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2). The SSP initiator port may receive a RESPONSE for the command before being able to abort the command.

If an SSP initiator port receives a DATA frame with a data offset that was not expected, then the SSP initiator port discards the frame (see 9.2.6.2.2) and aborts the command (see 10.2.2). The SSP initiator port may receive a RESPONSE for the command before being able to abort the command.

9.2.6 ST (transport layer for SSP ports) state machines

9.2.6.1 ST state machines overview

The ST state machines perform the following functions:

a) receive and process transport protocol service requests and transport protocol service responses from the SCSI application layer;
b) receive and process other SAS connection management requests from the application layer;
c) send transport protocol service indications and transport protocol service confirmations to the SCSI application layer;

d) send requests to the port layer to transmit frames and manage SAS connections; and
e) receive confirmations from the port layer.

The Transmission Status and Frame Received confirmations received from the port layer include the following arguments:

a) the tag;
b) the destination SAS address; and
c) the source SAS address;

These arguments are used to route the confirmations to the correct ST state machines.

9.2.6.2 ST_I (transport layer for SSP initiator ports) state machines

9.2.6.2.1 ST_I state machines overview

The ST_I state machines are as follows:

a) ST_IFR (initiator frame router) state machine (see 9.2.6.2.2);
b) ST_ISF (initiator send frame) state machine (see 9.2.6.2.3);
c) ST_IPD (initiator process data) state machine (see 9.2.6.2.4); and
d) ST_IPR (initiator process response) state machine (see 9.2.6.2.5).

update the following figure.

Figure 5 shows the ST_I state machines.

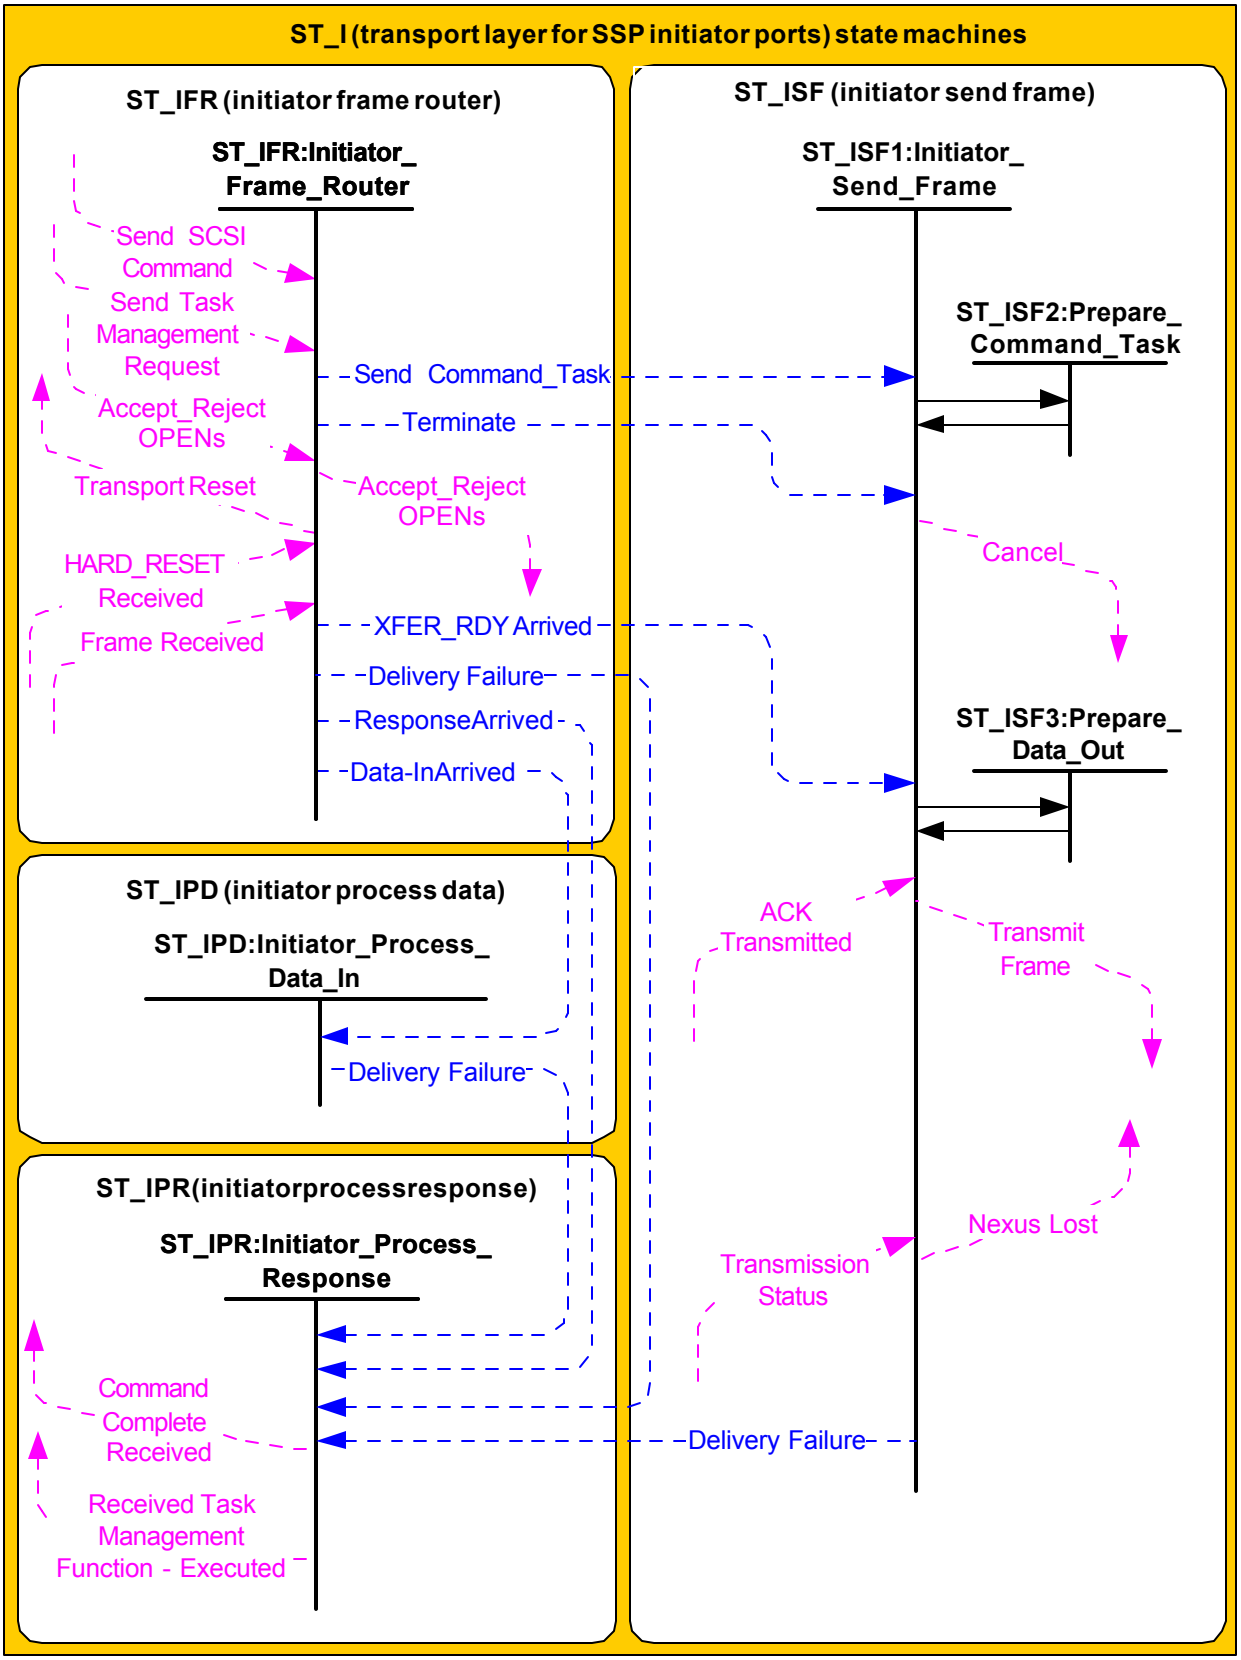

Figure 5 — ST_I (transport layer for SSP initiator ports) state machines

9.2.6.2.2 ST_IFR (initiator frame router) state machine

The ST_IFR state machine:

a) receives Send SCSI Command and Send Task Management transport protocol service requests from the SCSI application layer;
b) sends Send Command_Task, Terminate, and XFER_RDY Arrived messages to the ST_ISF state machine;
c) sends Data-In Arrived messages to the ST_IPD state machine;
d) sends Response Arrived messages to the ST_IPR state machine;
e) receives HARD_RESET Received confirmations from the port layer;
f) sends Transport Reset confirmations to the application layer;
g) receives Accept_Reject OPENs requests from the SCSI application layer; and
h) sends Accept_Reject OPENs requests to the port layer.

This state machine consists of one state.

This state machine shall start after power on.

If this state receives a Send SCSI Command transport protocol service request or a Send Task Management Request transport protocol service request, then this state shall send a Send Command_Task message to an ST_ISF state machine. The message shall include the following to be used in any OPEN address frames required to service the request:

a) connection rate;
b) initiator connection tag; and
c) destination SAS address.

If the request is a Send SCSI Command transport protocol service request, then the message shall also include the following to be used in any SSP frame for the request:

a) logical unit number;
b) tag;
c) task attribute;
d) additional CDB length;
e) CDB; and
f) additional CDB bytes.

If the request is for a data-out command and first burst is enabled (see 9.2.2.1), then the message shall also include the number of bytes for the first burst size for the logical unit.

If the request is a Send Task Management Request transport protocol service request, then the message shall include the following to be used in the TASK frame:

a) logical unit number;
b) tag;
c) task management function;
d) tag of task to be managed; and
e) retransmit bit.

If the request is a Send SCSI Command transport protocol service request, and the request causes one or more outstanding tasks in the port to be cleared or aborted (e.g., PREEMPT), then this state shall send a Terminate message to any ISF state machine that is processing one of the affected tasks.

If the request is a Send Task Management Request transport protocol service request, and the request causes one or more outstanding tasks in the port to be cleared or aborted, then this state shall send a Terminate message to any ISF state machine that is processing one of the affected tasks.

If this state machine receives a Frame Received (ACK/NAK Balanced) or Frame Received (ACK/NAK Not Balanced) confirmation, then this state shall check the frame type in the received frame. If the confirmation was Frame Received (ACK/NAK Balanced) and the frame type is not XFER_RDY, RESPONSE, or DATA, then this state machine shall discard the frame. If the confirmation was Frame Received (ACK/NAK Not Balanced) and the frame type is not DATA, then this state machine shall discard the frame.

If the frame type is correct relative to the confirmation, then this state may check that the hashed source SAS address matches the SAS address of the SAS port transmitting the frame and the hashed destination SAS address in the frame matches the SAS address of the SAS port receiving the frame based on the connection. If this state checks these SAS addresses and they do not match, then this state machine shall discard the frame.

If the frame type is correct relative to the confirmation, then this state shall check the tag. If the tag does not specify a valid state machine, then this state shall discard the frame and may send a vendor-specific confirmation to the SCSI application layer to abort the command using that tag.

If the frame type is XFER_RDY and the tag is for a task with no write data, then this state shall;

a) discard the frame;
b) send a Delivery Failure (XFER_RDY Not Expected) message to the ST_IPR state machine; and
c) if there is an ST_ISF state machine for the tag, send a Terminate message to that state machine.

If the frame type is XFER_RDY and the items checked in the frame correct, then this state shall send an XFER_RDY Arrived message to the ST_ISF1:Send_Frame state specified by the tag. The message shall include the content of the frame as an argument.

If the frame type is RESPONSE, then this state shall check the tag. If this state has received a RESPONSE frame for this I_T_L_Q nexus, then this state shall discard the frame.

If the frame type is RESPONSE, the items checked in the frame are correct, and this state has not received a RESPONSE frame for this I_T_L_Q nexus, then this state shall send a Response Arrived message to an ST_IPR state machine. The message shall include the content of the frame as an argument.

If the frame type is DATA and the items checked in the frame are correct, then this state shall send a Data-In Arrived message to an ST_IPR state machine. The message shall include the content of the frame as an argument.

If this state machine receives an Accept_Reject OPENs (Accept SSP) or Accept_Reject OPENs (Reject SSP) request, then this state shall send an Accept_Reject OPENs request along with the received argument to the port layer.

If this state machine receives a HARD_RESET Received confirmation, then this state shall send a Transport Reset event notification to the SCSI application layer and send a Cancel message to any ISF state machine processing a task.

9.2.6.2.3 ST_ISF (initiator send frame) state machine

9.2.6.2.3.1 ST_ISF state machine overview

The ST_ISF state machine:

a) receives Send Command_Task and XFER_RDY Arrived messages from the ST_IFR state machine;
b) constructs COMMAND, TASK, and data-out DATA frames;
c) sends Transmit Frame requests to the port layer;
d) receives Transmission Status and ACK Transmitted confirmations from the port layer;
e) receives Terminate messages from the ST_IFR state machine;
f) sends Cancel requests to the port layer; and
g) communicates to the ST_IPR state machine regarding service delivery subsystem failures.

This state machine consists of the following states:

a) ST_ISF1:Send_Frame (see 9.2.6.2.3.2)(initial state);
b) ST_ISF2:Prepare_Command_Task (see 9.2.6.2.3.3); and
c) ST_ISF3:Prepare_Data_Out (see 9.2.6.2.3.4).

This state machine shall be started when a when a Send Command_Task message is received from the ST_IFR state machine.

9.2.6.2.3.2 ST_ISF1:Send_Frame state

9.2.6.2.3.2.1 State description

If this state receives an XFER_RDY Arrived message, then this state shall check the length of the information unit. If the length of the information unit is not 12 bytes, then this state shall:

a) discard the frame;
b) if the information unit contains fewer than 12 bytes, then send a Delivery Failure (XFER_RDY Information Unit Too Short) message to the ST_IPR state machine;
c) if the information unit contains more than 12 bytes, then send a Delivery Failure (XFER_RDY Information Unit Too Long) message to the ST_IPR state machine; and
d) terminate after sending a Delivery Failure message to the ST_IPR state machine.

If this state receives an XFER_RDY Arrived message, and the write data length is zero or exceeds the amount of data remaining to be transferred for the data-out command, then this state shall:

a) send a Delivery Failure (XFER_RDY Incorrect Write Data Length) message to the ST_IPR state machine; and
b) terminate after sending the message.

If this state machine receives an XFER_RDY Arrived message and the requested offset is not expected, then this state shall:

a) send a Delivery Failure (XFER_RDY Requested Offset Error) message to the ST_IPR state machine; and
b) terminate after sending the message.

If this state is entered from the ST_ISF2:Prepare_Command_Task state, then this state shall send a Transmit Frame (Interlocked) request to the port layer.

If this state is entered from the ST_ISF3:Prepare_Data_Out state, then this state shall send a Transmit Frame (Non-Interlocked) request to the port layer.

A Transmit Frame request shall include the SSP frame and the following to be used for any OPEN address frame:

a) the initiator port bit set to one;
b) protocol set to SSP;
c) connection rate;
d) initiator connection tag;
e) destination SAS address; and
f) source SAS address set to the SAS address of the SSP initiator port.

After sending a Transmit Frame request to the port layer this state shall wait for a Transmission Status confirmation (see table xx [was 92]).

If this state receives a confirmation that is not Transmission Status (Frame Transmitted), and the Transmit Frame request was for a COMMAND or TASK frame (i.e., not for a DATA frame), then this state shall send a Delivery Failure (Service Delivery Subsystem Failure) message to the ST_IPR state machine.

If this state receives a confirmation that is not Transmission Status (Frame Transmitted), the Transmit Frame request was for a data frame, and the RETRY DATA FRAMES bit was set to zero in the XFER_RDY frame for the data, then this state shall send a Delivery Failure (Service Delivery Subsystem Failure) message to the ST_IPR state machine.

If this state receives a confirmation that is not Transmission Status (Frame Transmitted), the Transmit Frame request was for a DATA frame, the RETRY DATA FRAMES bit was set to one in the XFER_RDY frame for the data, and the state machine has exceeded its vendor-specific number of retries then this state shall send a Delivery Failure (Service Delivery Subsystem Failure) message to the ST_IPR state machine.

The Delivery Failure message shall include:

a) any argument received with the Transmission Status confirmation; and
b) I_T_L_Q nexus information (i.e., destination SAS address and tag);

If the confirmation is Transmission Status (Frame Transmitted) and the Transmit Frame request was for a COMMAND frame, TASK frame, or DATA frame where the number of data bytes that have been transmitted for the Send SCSI Command or Send Task Management transport protocol service request equal the number of bytes in the service request, then this state shall wait to receive one of the following:

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).

NOTE 2 - If the number of data bytes that have been transmitted for the Send SCSI Command or Send Task Management transport protocol service request are fewer than the number of bytes in the service request, then this state may send additional Transmit Frame requests for 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 DATA frames sent for the I_T_L_Q nexus.

If the confirmation is Transmission Status (ACK Received) and the Transmit Frame request was for a COMMAND frame, TASK frame, or DATA frame where the number of data bytes that have been transmitted for the Send SCSI Command or Send Task Management transport protocol service request equal the number of bytes in the service request, then this state shall terminate.

If this state receives a Transmission Status (NAK Received), and the Transmit Frame request was not for a DATA frame, then this state shall send a Delivery Failure (Service Delivery Subsystem Failure - NAK Received) message to the ST_IPR state machine.

If this state receives a Transmission Status (ACK/NAK Timeout) or a Transmission Status (Connection Lost Without ACK/NAK), and the Transmit Frame request was not for a DATA frame, then this state shall send a Delivery Failure (Service Delivery Subsystem Failure - Connection Failed) message to the ST_IPR state machine.

If this state receives a Transmission Status (NAK Received), the Transmit Frame request was for a DATA frame, and the RETRY DATA FRAMES bit was set to zero in the XFER_RDY frame for the data, then this state shall send a Delivery Failure (Service Delivery Subsystem Failure - NAK Received) message to the ST_IPR state machine.

If this state receives a Transmission Status (ACK/NAK Timeout) or a Transmission Status (Connection Lost Without ACK/NAK), the Transmit Frame request was for a DATA frame, and the RETRY DATA FRAMES bit was set to zero in the XFER_RDY frame for the data, then this state shall send a Delivery Failure (Service Delivery Subsystem Failure - Connection Failed) message to the ST_IPR state machine.

If this state receives a Transmission Status (NAK Received), the Transmit Frame request was for a DATA frame, the RETRY DATA FRAMES bit was set to one in the XFER_RDY frame for the data, and the state machine has exceeded its vendor-specific number of retries then this state shall send a Delivery Failure (Service Delivery Subsystem Failure - NAK Received) message to the ST_IPR state machine.

If this state receives a Transmission Status (ACK/NAK Timeout) or a Transmission Status (Connection Lost Without ACK/NAK), the Transmit Frame request was for a DATA frame, the RETRY DATA FRAMES bit was set to one in the XFER_RDY frame for the data, and the state machine has exceeded its vendor-specific number of retries then this state shall send a Delivery Failure (Service Delivery Subsystem Failure - Connection Failed) message to the ST_IPR state machine.

The Delivery Failure message shall include:

a) any argument received with the Transmission Status confirmation; and
b) I_T_L_x nexus information (i.e., the destination SAS address and tag).

After sending a Delivery Failure message to the ST_IPR state machine, this state machine shall terminate.

If this state receives a Terminate message from the ST_IFR state machine, and this state has received confirmations for all Transmit Frame requests sent to the port layer, then this state machine shall terminate.

If this state receives a Terminate message from the ST_IFR state machine, 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. A Cancel request shall include the following arguments:

a) the destination SAS address; and
b) the tag.

This state machine shall terminate upon receipt of a Transmission Status (Cancel Acknowledge) confirmation.

9.2.6.2.3.2.2 Transition ST_ISF1:Send_Frame to ST_ISF2:Prepare_Command_Task

This transition shall occur after this state receives a Send Command_Task message.

9.2.6.2.3.2.3 Transition ST_ISF1:Send_Frame to ST_ISF3:Prepare_Data_Out

This transition shall occur after this state receives:

a) a Transmission Status (ACK Received) confirmation for a COMMAND frame for a data-out operation if first burst is enabled;
b) an XFER_RDY Arrived message followed by an ACK Transmitted confirmation;
c) a Transmission Status (Frame Transmitted) confirmation for a Transmit Frame (Non-Interlocked) request if the number of data bytes that has been transmitted for the request is less than the first burst size or the write data length specified in the XFER_RDY; or
d) a Transmission Status (Frame Transmitted) confirmation and a confirmation other than Transmission Status (ACK Received) for a DATA frame for which a Delivery Failure message was not sent to the ST_IPR state machine (i.e., in order to retry transmitting the frame).

9.2.6.2.3.3 ST_ISF2:Prepare_Command_Task state

9.2.6.2.3.3.1 State description

This state shall construct either a COMMAND or TASK frame.

If the frame to be constructed is a COMMAND frame, then this state shall include the following values received from the Send Command_Task message in the frame:

a) logical unit number;
b) tag;
c) task attribute;
d) additional CDB length;
e) CDB; and
f) additional CDB bytes.

If the frame to be constructed is a TASK frame, then this state shall include the following values received from the Send Command_Task message in the frame:

a) logical unit number;
b) tag;
c) task management function;
d) tag of task to be managed; and
e) retransmit bit.

This state shall generate and include the following values in either a COMMAND or TASK frame:

a) frame type;
b) hashed destination SAS address;
c) hashed source SAS address; and
d) number of fill bytes.

9.2.6.2.3.3.2 Transition ST_ISF2:Prepare_Command_Task to ST_ISF1:Send_Frame

This transition shall occur after this state constructs a COMMAND or TASK frame.

9.2.6.2.3.4 ST_ISF3:Prepare_Data_Out state

9.2.6.2.3.4.1 State description

This state shall construct a DATA frame.

This state shall include the following values in the frame, received either from the Send Command_Task message (i.e., if the first burst is not zero) or included in an XFER_RDY Arrived message:

a) tag;
b) target port transfer tag;
c) data offset; and
d) data.

This state shall generate and include the following values in the frame:

a) frame type;
b) hashed destination SAS address;
c) hashed source SAS address;
d) retransmit bit set to zero;
e) number of fill bytes; and
f) fill bytes.

If this is the first DATA frame constructed by this state, then this state shall set the CHANGING DATA POINTER bit to zero.

If this state is entered after the ST_ISF1:Send_Frame state receives a Transmission Status (Frame Transmitted) confirmation for a DATA frame and that state has only received confirmations of Transmission Status (Frame Transmitted) and Transmission Status (ACK Received), then this state shall set the CHANGING DATA POINTER bit to zero and shall set the DATA OFFSET field to the value in the DATA OFFSET field in the previous DATA frame plus the number of bytes in the previous DATA information unit.

If this state is entered after the ST_ISF1:Send_Frame state receives a Transmission Status (Frame Transmitted) confirmation and a confirmation other than Transmission Status (ACK Received) for a DATA frame for which a Delivery Failure message was not sent to the ST_IPR state machine (i.e., in order to retry transmitting the frame), then this state shall set the CHANGING DATA POINTER bit in the frame to one and shall set the DATA OFFSET field to a data offset value associated with a previous ACK/NAK balance.

9.2.6.2.3.4.2 Transition ST_ISF3:Prepare_Data_Out to ST_ISF1:Send_Frame

This transition shall occur after this state constructs a DATA frame.

9.2.6.2.4 ST_IPD (initiator process data) state machine

The ST_IPD state machine receives and processes a message from the ST_IFR state machine containing a DATA frame.

This state machine consists of one state.

This state machine shall be started when a Data-In Arrived message is received.

This state shall check the length and data offset of the DATA information unit.

If the length of the information unit is zero, then this state shall send a Delivery Failure (DATA Incorrect Read Data Length) message to the ST_IPR state machine.

If the length of the information unit exceeds the amount of data remaining to be transferred for the data-in command, then this state shall send a Delivery Failure (DATA Too Much Read Data) message to the ST_IPR state machine.

If the data offset is not an expected offset (i.e., the changing data pointer bit is set to one and the data offset is not set to a data offset associated with a previous ACK/NAK balance, or the changing data pointer bit is set to zero and the data offset is not set to the previous data offset plus the previous data size), then this state shall send a Delivery Failure (DATA Offset Error) message to the ST_IPR state machine.

If the DATA information unit is valid, this state shall process the data-in data.

This state machine shall terminate after sending a message or processing the data-in data.

9.2.6.2.5 ST_IPR (initiator process response) state machine

The ST_IPR state machine receives:

a) a message from the ST_IFR state machine containing a RESPONSE frame;
b) a message from the ST_IFR state machine containing a a service delivery subsystem failure;
c) a message from the ST_IPD state machine containing a a service delivery subsystem failure; or
d) a message from the ST_ISF state machine containing a service delivery subsystem failure.

This state machine processes the RESPONSE frame or the service delivery subsystem failure and sends a transport protocol service confirmation to the SCSI application layer.

This state machine consists of one state.

This state machine shall be started when a Response Arrived message is received or a Delivery Failure message is received.

If a Delivery Failure message is received, this state shall send a Command Complete Received or Received Task Management Function – Executed confirmation to the SCSI application layer with the Service Response argument set as defined in table 13.

Table 13 — Delivery Failure to Command Complete Received mapping

Delivery Failure argument	Command Complete Received (Service Response)
XFER_RDY Information Unit Too Short	Service Delivery or Target Failure - XFER_RDY Information Unit Too Short
XFER_RDY Information Unit Too Long	Service Delivery or Target Failure - XFER_RDY Information Unit Too Long
XFER_RDY Incorrect Write Data Length	Service Delivery or Target Failure - XFER_RDY Incorrect Write Data Length
XFER_RDY Requested Offset Error	Service Delivery or Target Failure - XFER_RDY Requested Offset Error
XFER_RDY Not Expected	Service Delivery or Target Failure - XFER_RDY Not Expected
DATA Incorrect Data Length	Service Delivery or Target Failure - DATA Incorrect Data Length
DATA Too Much Read Data	Service Delivery or Target Failure - DATA Too Much Read Data
DATA Offset Error	Service Delivery or Target Failure - DATA Offset Error
Service Delivery Subsystem Failure - NAK Received	Service Delivery or Target Failure - NAK Received
Service Delivery Subsystem Failure - Connection Failed	Service Delivery or Target Failure - Connection Failed

If a Response Arrived message is received, this state shall check the length of the RESPONSE information unit.

If the length of the information unit is correct, then this state shall send one of the following transport protocol service confirmations (based on the type of task for which the response was received) to the SCSI application layer:

a) a Command Complete Received with the Service Response argument set to Task Complete; or
b) a Command Complete Received with the Service Response argument set to Linked Command Complete; or
c) a Received Task Management Function – Executed with the Service Response argument set to Function Complete.

The confirmation shall also include a Retransmit argument indicating the state of the retransmit bit.

If the length of the information unit is not correct, then this state shall send one of the following transport protocol service confirmations (based on the type of task for which the response was received) to the SCSI application layer:

a) a Command Complete Received with the Service Response argument set to Service Delivery or Target Failure; or
b) a Received Task Management Function – Executed with the Service Response argument set to Service Delivery or Target Failure.

The confirmation shall include the tag.

This state machine shall terminate after sending a confirmation.

9.2.6.3 ST_T (transport layer for SSP target ports) state machines

9.2.6.3.1 ST_T state machines overview

The ST_T state machines are as follows:

a) ST_TFR (target frame router) state machine (see 9.2.6.3.2); and
b) ST_TTS (target transport server) state machine (see 9.2.6.3.3).

If implemented, this state machine shall maintain the timers listed in table 14.

Table 14 — ST_T state machine timers

Timer	Initial value
Initiator Response Timeout	The value in the INITIATOR RESPONSE TIMEOUT field in the Protocol-Specific Port mode page (see 10.2.6.2).

Figure 6 shows the ST_T state machines.

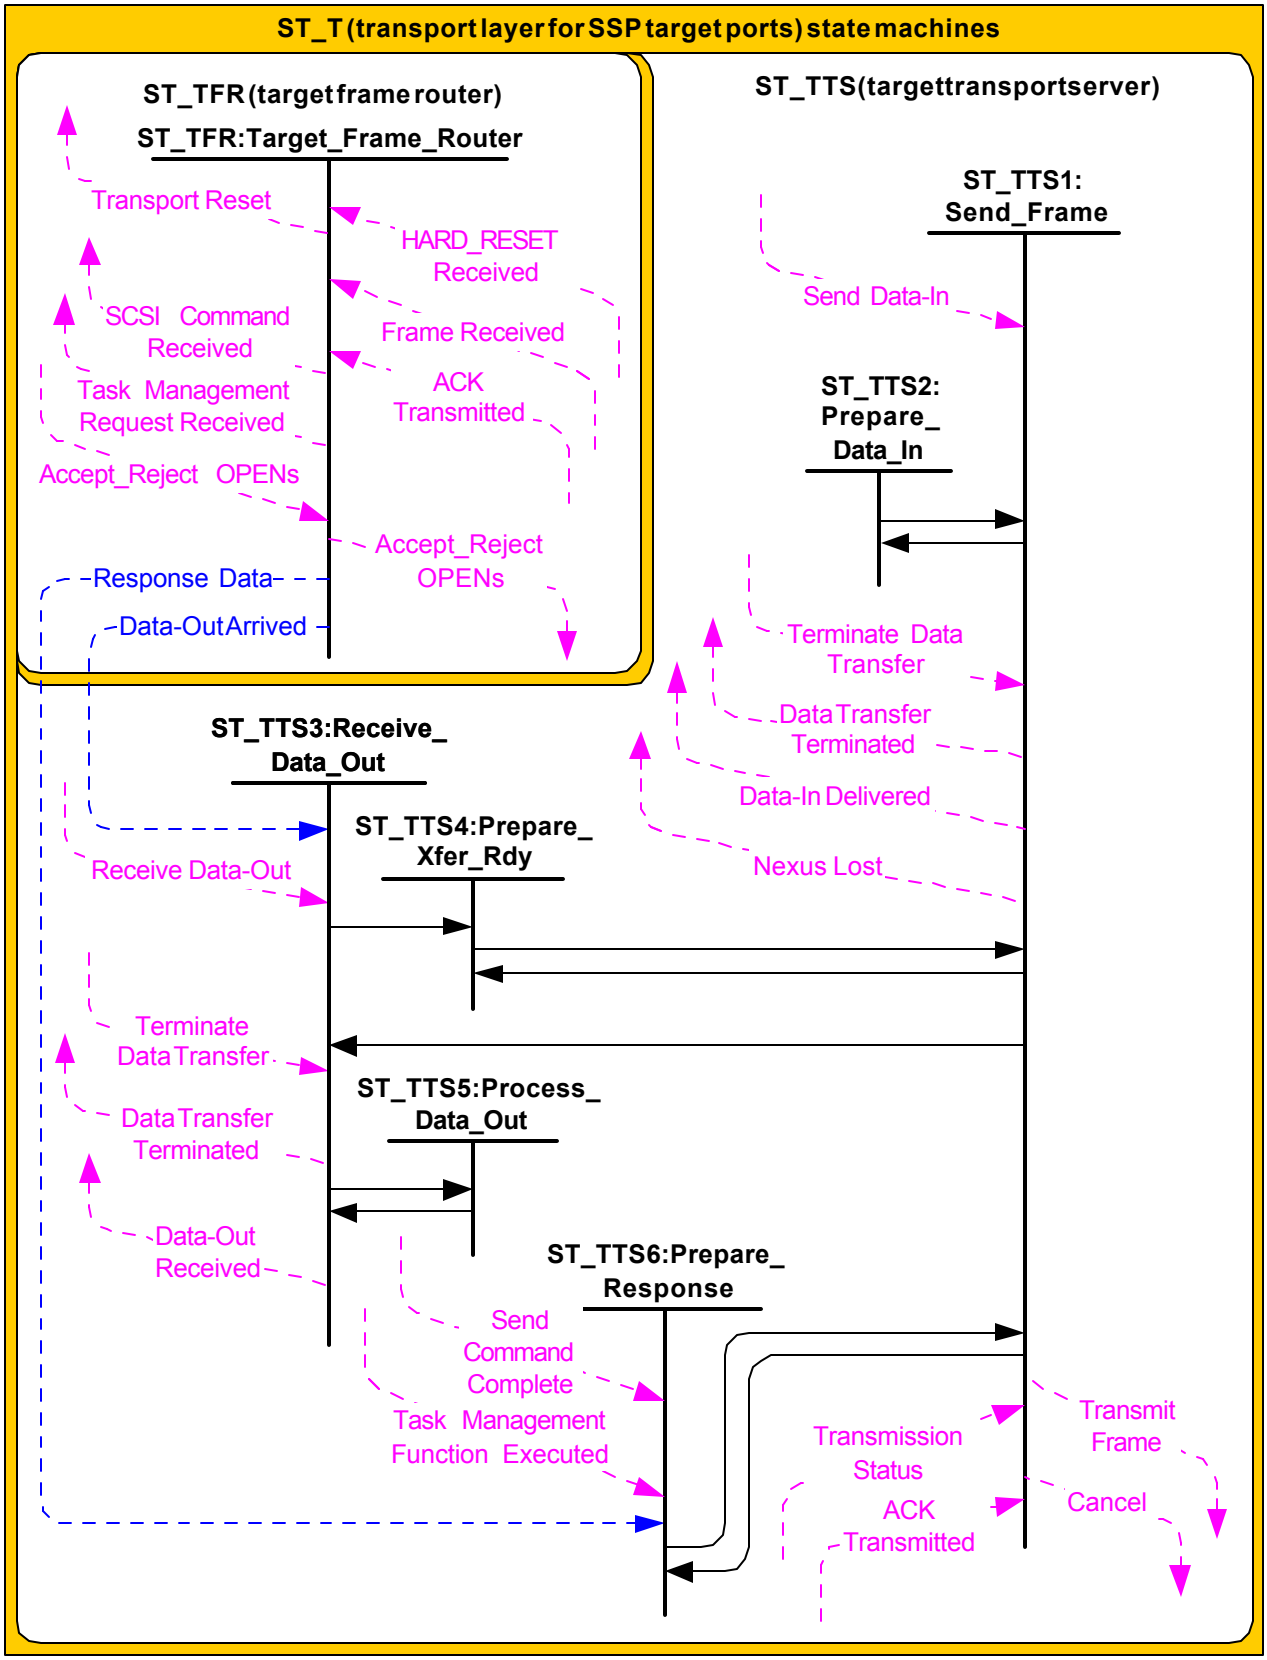

Figure 6 — ST_T (transport layer for SSP target ports) state machines

9.2.6.3.2 ST_TFR (target frame router) state machine

The ST_TFR state machine:

The ST_TFR state machine:

a) receives confirmations from the port layer;
b) receives transport protocol service requests from the SCSI application layer;
c) sends transport protocol service indications to the SCSI application layer;
d) sends messages to the ST_TTS state machine
e) receives Accept_Reject OPENs requests from the application layer; and
f) sends Accept_Reject OPENs requests to the port layer.

This state machine consists of one state.

This state machine shall be started after power on.

If this state machine receives an Accept_Reject OPENs (Accept SSP) or Accept_Reject OPENs (Reject SSP) request, then this state machine shall send a corresponding Accept_Reject OPENs request to the port layer.

If this state machine receives a Frame Received (ACK/NAK Balanced) or Frame Received (ACK/NAK Not Balanced) confirmation, then this state machine shall check the frame type in the received frame (see table 2). If the frame type is not COMMAND, TASK, or DATA, then this state machine shall discard the frame. If the confirmation was Frame Received (ACK/NAK Not Balanced) and the frame type is not DATA, then this state machine shall discard the frame.

This state machine may check that reserved fields in the frame are zero. If any reserved fields are not zero, then this state machine may send a Response Data (Invalid Frame) message to the ST_TTS6:Prepare_Response state including the logical unit number and tag.

NOTE 3 - The check of reserved fields described above does not apply to the reserved fields within the CDB in a COMMAND frame. Checking of reserved fields in a CDB is described in SAM-3.

If the frame type is correct relative to the confirmation, then this state machine may check that the hashed source SAS address matches the SAS address of the SAS port transmitting the frame and the hashed destination SAS address in the frame matches the SAS address of the SAS port receiving the frame based on the connection. If this state machine checks these SAS addresses and they do not match, then this state machine shall discard the frame.

If the frame type is COMMAND, then this state machine shall check the length of the information unit. If the length of the information unit is not correct (see 9.2.2.1), then this state machine shall send a Response Data (Invalid Frame) message to the ST_TTS6:Prepare_Response state in the ST_TTS state machine. The message shall include the logical unit number and tag.

If the frame type is TASK, then this state machine shall check the length of the information unit. If the length of the information unit is not correct (see 9.2.2.2), then this state machine shall send a Response Data (Invalid Frame) message to the ST_TTS6:Prepare_Response state in the ST_TTS state machine. The message shall include the logical unit number and tag.

If the frame type is COMMAND or TASK, then this state machine may check if the tag conflicts with an existing tag (i.e., an existing command or task management function). If this state machine checks the tag and the tag conflicts, then this state machine shall send a Response Data (Invalid Frame) message to the ST_TTS6:Prepare_Response state in the ST_TTS state machine. The message shall include the logical unit number and tag.

If the frame type is COMMAND or TASK, then this state machine may check the target port transfer tag. If target port transfer tag is set to a value other than FFFFh, then this state machine may send a Response Data (Invalid Frame) message to the ST_TTS6:Prepare_Response in the ST_TTS state machine. The message shall include the logical unit number and tag.

If the frame type is TASK, then this state machine shall check the logical unit number. If there is no logical unit at the specified logical unit number, then this state machine shall send a Response Data (Invalid Logical Unit Number) message to the ST_TTS6:Prepare_Response state in the ST_TTS state machine. The message shall include the logical unit number and tag.

If the frame type is COMMAND or TASK and the items checked in the frame are correct, then this state machine shall wait to receive an ACK Transmitted confirmation.

If the frame type is COMMAND, the items checked in the frame are correct, and this state machine receives an ACK Transmitted confirmation, then this state shall send a SCSI Command Received transport protocol service indication to the SCSI application layer. The indication shall include:

a) the source SAS address;
b) the tag;
c) the logical unit number;
d) the task attribute;
e) the CDB; and
f) any additional CDB bytes.

If the frame type is TASK, the items checked in the frame are correct, and this state machine receives an ACK Transmitted confirmation, then this state machine shall send a Task Management Request Received transport protocol service indication to the SCSI application layer. The indication shall include:

a) the source SAS address;
b) the tag;
c) the logical unit number;
d) the task management function; and
e) the tag of the task to be managed.

If the frame type is DATA, and the tag does not match a tag for an outstanding data-out command, then this state machine shall discard the frame.

If the frame type is DATA, and the tag matches a tag for an outstanding data-out command without first burst data for which no XFER_RDY frame is outstanding, then this state machine shall discard the frame.

If the frame type is DATA, then this state machine may check the target port transfer tag. If this state machine checks the target port transfer tag and it does not specify a valid state machine, then this state shall discard the frame.

If the frame type is DATA and the items checked in the frame are correct, then this state machine shall send a Data-Out Arrived message to the ST_TTS3:Receive_Data_Out state in the ST_TTS state machine. The message shall include the content of the SAS frame.

If this state machine receives a HARD_RESET Received confirmation, then this state machine shall send a Transport Reset event notification to the SCSI application layer.

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 transport protocol service requests and responses from the SCSI application layer;
b) processes and sends transport protocol service confirmations to the SCSI application layer; and
c) communicates with the port layer via requests and confirmations regarding frame transmission.

This state machine consists of the following states:

a) ST_TTS1:Send_Frame (see 9.2.6.3.3.2);
b) ST_TTS2:Prepare_Data_In (see 9.2.6.3.3.3);
c) ST_TTS3:Receive_Data_Out (see 9.2.6.3.3.4);
d) ST_TTS4:Prepare_Xfer_Rdy (see 9.2.6.3.3.5);
e) ST_TTS6:Process_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 be started in the ST_TTS1:Send_Frame state when this state receives a Send Data-In transport protocol service request from the SCSI application layer.

This state machine shall be started in the ST_TTS3:Receive_Data_Out state if this state machine is not already running and:

a) this state receives a Receive Data-Out transport protocol service request; or
b) this state receives a Data-Out Arrived message and first burst is enabled.

This state machine shall be started in the ST_TTS6:Prepare_Response state when:

a) this state receives a Task Management Function Executed transport protocol service response;
b) this state receives a Send Command Complete transport protocol service response.
c) this state machine is not already running and this state receives a Response Data message.

If transport layer retries are enabled, this state machine shall retain the data offset for the last DATA frame transmitted for which ACK/NAK Balance was achieved (i.e., when the number of DATA frames sent matches the number of ACK Received confirmations received) for use as the restart point in case of a retry.

9.2.6.3.3.2 ST_TTS1:Send_Frame state

9.2.6.3.3.2.1 State description

If this state is entered from the ST_TTS2:Prepare_Data_In state for transmission of a 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.

NOTE 4 - The interlocked frame transmission rules ensure that wide ports do not send an XFER_RDY or RESPONSE frame for an I_T_L_Q nexus on a phy until all the ACKs have been transmitted for write DATA frames for the nexus on a different phy. In a narrow port, the link layer ensures that there is ACK/NAK balance before transmitting an interlocked frame.

A Transmit Frame request from this state shall include the SSP frame and the following to be used for any OPEN address frame:

a) the initiator port bit set to zero;
b) protocol set to SSP;
c) connection rate;
d) initiator connection tag;
e) destination SAS address; and
f) source SAS address set to the SAS address of the SSP target port.

After sending a Transmit Frame request this state shall wait to receive a Transmission Status confirmation (see table xx [was 92]).

If the confirmation is Transmission Status (I_T Nexus Loss), this state shall send a Nexus Lost confirmation to the SCSI application layer. This state machine shall terminate after sending a Nexus Lost confirmation.

If the confirmation is Transmission Status (Frame Transmitted) and the Transmit Frame request was for an XFER_RDY frame, a RESPONSE frame, or a DATA frame where the number of bytes that have been transmitted equal the number of bytes in the Send Data-In transport protocol service request, then the state machine 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).

NOTE 5 - If the number of data bytes that have been transmitted for the Send Data-In transport protocol service request are fewer than the number of bytes in the service request, then this state transitions to the ST_TTS2:Prepare_Data_In state to construct the additional 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 for Transmit Frame requests for previous DATA frames sent for the I_T_L_Q nexus.

If the confirmation is Transmission Status (ACK Received) and the Transmit Frame request was for a RESPONSE frame, then this state shall send a Data-In Delivered transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY SUCCESSFUL.

If the confirmation is Transmission Status (ACK Received), the Transmit Frame request was for a DATA frame where the number of data bytes that have been transmitted for the Send Data-In transport protocol service request equal the number of bytes in the service request, and this state has received a Transmission Status (ACK Received) confirmation for each of the DATA frames transmitted for the Send Data-In transport protocol service request, then this state shall send a Data-In Delivered transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY SUCCESSFUL.

If the confirmation is Transmission Status (NAK Received), Transmission Status (ACK/NAK Timeout), or Transmission Status (Connection Lost Without ACK/NAK), and the Transmit Frame request was for a RESPONSE frame generated by a Response Data message from the ST_TFR:Target_Frame_Router state machine (e.g., the RESPONSE frame is to report that the received information unit was not the correct length), and the vendor specific number of retries has been reached, then this state machine shall terminate.

If the confirmation is Transmission Status (NAK Received), and the Transmit Frame request was for an XFER_RDY frame or a RESPONSE frame generated by a protocol service response from the SCSI application layer (e.g., a Send Command Complete), and the vendor specific number of retries has been reached, then this state machine shall send a Data-In Delivered transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - NAK RECEIVED.

If the confirmation is Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK), and the Transmit Frame request was for an XFER_RDY frame or a RESPONSE frame generated by a protocol service response from the SCSI application layer, and the vendor specific number of retries has been reached, then this state machine shall send a Data-In Delivered transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - CONNECTION FAILED.

If the confirmation is Transmission Status (NAK Received), and the Transmit Frame request was for a DATA frame, and the vendor specific number of retries has been reached, then this state shall send a Data-In Delivered transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - NAK RECEIVED.

If the confirmation is Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK), the Transmit Frame request was for a DATA frame, and the vendor specific number of retries has been reached, then this state shall send a Data-In Delivered transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - CONNECTION FAILED.

A Data-In Delivered transport protocol service confirmation to the SCSI application layer shall include the following:

a) any argument received from the port layer (e.g., Transmission Status (ACK Received) or Service Delivery Subsystem Failure); and
b) I_T_L_x nexus information (i.e., destination SAS address and tag).

This state machine shall terminate after sending a Data-In Delivered protocol service confirmation.

If this state receives a Terminate Data Transfer request and this state has received confirmations for all Transmit Frame requests sent to the port layer, then this state shall send a Data Transfer Terminated confirmation to the application layer. After sending the confirmation, this state machine shall terminate.

If this state receives a Terminate Data Transfer request 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. A Cancel request shall include the following arguments:

a) the destination SAS address; and
b) the tag.

Upon receipt of a Transmission Status (Cancel Acknowledge) confirmation this state shall:

a) send a Transmission Status (Cancel Acknowledge) confirmation; and
b) terminate.

9.2.6.3.3.2.2 Transition ST_TTS1:Send_Frame to ST_TTS2:Prepare_Data_In

This transition shall occur after this state receives:

a) a Send Data-In protocol service request from the SCSI application layer;
b) a Transmission Status (Frame Transmitted) confirmation for a DATA frame and the number of bytes moved for the Send Data-In transport protocol service request is less than the Request Byte Count; or
c) a Transmission Status (Frame Transmitted) confirmation and a confirmation other than Transmission Status (ACK Received) for a DATA frame for which a Data-In Delivered transport protocol service confirmation was not sent to the SCSI application layer (i.e., in order to retry transmitting the frame).

9.2.6.3.3.2.3 Transition ST_TTS1:Send_Frame to ST_TTS3:Receive_Data_Out

This transition shall occur after this state receives a Transmission Status (Frame Transmitted) confirmation followed by a Transmission Status (ACK Received) confirmation for an XFER_RDY frame.

9.2.6.3.3.2.4 Transition ST_TTS1:Send_Frame to ST_TTS4:Prepare_Xfer_Rdy

This transition shall occur after this state receives a Transmission Status (Frame Transmitted) confirmation for an XFER_RDY frame followed by a Transmission Status (NAK Received), Transmission Status (ACK/NAK Timeout), or Transmission Status (Connection Lost Without ACK/NAK) confirmation for the frame, and the vendor specific number of retries has not been reached.

9.2.6.3.3.2.5 Transition ST_TTS1:Send_Frame to ST_TTS6:Prepare_Response

This transition shall occur after this state receives a Transmission Status (Frame Transmitted) confirmation and a confirmation other than Transmission Status (ACK Received) for a RESPONSE frame generated by a protocol service response from the SCSI application layer (e.g., a Send Command Complete) and for which a Data-In Delivered transport protocol service confirmation was not sent to the SCSI application layer (i.e., in order to retry transmitting the frame).

9.2.6.3.3.3 ST_TTS2:Prepare_Data_In state

9.2.6.3.3.3.1 State description

This state fetches the data from the Device Server Buffer and constructs a DATA frame. This state shall use the tag received from the SCSI application layer to construct the frame.

This state shall generate the following values when constructing the frame:

a) frame type;
b) hashed destination SAS address;
c) hashed source SAS address;
d) retransmit bit set to zero;
e) number of fill bytes;
f) fill bytes;
g) data offset; and
h) data.

If this is the first DATA frame constructed by this state, then this state shall set the CHANGING DATA POINTER bit and the DATA OFFSET field in the DATA frame to zero.

If this state is entered from the ST_TTS1:Send_Frame state after that state receives a Transmission Status (Frame Transmitted) confirmation and that state has only received confirmations of Transmission Status (Frame Transmitted) and Transmission Status (ACK Received), then this state shall set the CHANGING DATA POINTER bit to zero and shall set the value in the DATA OFFSET field to the value of the DATA OFFSET field in the previous DATA frame plus the number of bytes in the previous DATA information unit.

If this state is entered from the ST_TTS1:Send_Frame state after that state receives a Transmission Status (Frame Transmitted) confirmation and a confirmation other than Transmission Status (ACK Received) for which a Delivery Failure message was not sent to the ST_IPR state machine (i.e., to retry transmitting a frame), then this state shall set the CHANGING DATA POINTER bit in the frame to one and shall set the DATA OFFSET field to a data offset value associated with a previous ACK/NAK balance.

9.2.6.3.3.3.2 Transition ST_TTS2:Prepare_Data_In to ST_TTS1:Send_Frame

This transition shall occur after this state constructs a DATA frame.

9.2.6.3.3.4 ST_TTS3:Receive_Data_Out state

9.2.6.3.3.4.1 State description

If a Receive Data-Out transport protocol service request caused this state machine to be started, then the request includes the following to be used in any OPEN address frames required to service the request:

a) connection rate;
b) initiator connection tag;
c) destination SAS address;
d) logical unit number;
e) tag;
f) device server buffer (e.g., starting logical block address); and
g) request byte count (e.g., transfer length).

If a Data-Out Arrived message caused this state machine to be started and a Receive Data-Out transport protocol service request has not been received (i.e., first burst is enabled), then this state shall wait to process the Data-Out Arrived message until this state receives a Receive Data-Out transport protocol service request. The data received in the Data-Out Arrived message shall be saved in a first burst buffer until this state receives a Receive Data-Out transport protocol service request.

After this state receives a Receive Data-Out transport protocol service request and a Data-Out Arrived message, then this state shall verify the received DATA frame as follows:

1) check the target port transfer tag value in the DATA frame if the target port transfer tag is being used. If the value is incorrect, then this state shall discard the frame and the state machine shall terminate;
2) check the data offset. If the data offset was not expected (i.e., the CHANGING DATA POINTER bit is set to 1 and the value in the DATA OFFSET field is not set to a data offset associated with a previous ACK/NAK balance, or the CHANGING DATA POINTER bit is set to zero and the value in the DATA OFFSET field is not set to the value in the DATA OFFSET field plus the number of bytes in the previous DATA information unit), then this state shall send a Data-Out Received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - DATA OFFSET ERROR. This confirmation shall include the tag. This state machine shall terminate after sending the confirmation;
3) check the length of the data. If an XFER_RDY frame was sent for the data (i.e., it is not first burst data) and the length of the data exceeds that specified by the XFER_RDY frame that requested the data, then this state shall send a Data-Out Received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - TOO MUCH WRITE DATA. This confirmation shall include the tag. This state machine shall terminate after sending the confirmation; and
4) check the length of the data. If the length of the data is zero, then this state shall send a Data-Out Received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - INFORMATION UNIT TOO SHORT. This confirmation shall include the tag. This state machine shall terminate after sending the confirmation.

If this state is entered from the ST_TTS1:Send_Frame state after that state received a Transmission Status (Frame Transmitted) confirmation followed by a Transmission Status (ACK Received) confirmation for an XFER_RDY frame, then this state shall wait for a Data-Out Arrived message.

If this state is entered from the ST_TTS1:Send_Frame state after that state received a Transmission Status (Frame Transmitted) confirmation followed by a Transmission Status (NAK Received) confirmation for an

XFER_RDY frame (i.e., the vendor specific number of retries has been reached), then this state shall send a Data-Out received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - NAK RECEIVED.

If this state is entered from the ST_TTS1:Send_Frame state after that state received a Transmission Status (Frame Transmitted) confirmation followed by a Transmission Status (ACK/NAK Timeout) or Transmission Status (Connection Lost Without ACK/NAK) confirmation for an XFER_RDY frame (i.e., the vendor specific number of retries has been reached), then this state shall send a Data-Out received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - CONNECTION FAILED.

If the Initiator Response Timeout timer is implemented, this state shall initialize and start the Initiator Response Timeout timer after any of the following occur:

a) a Data-Out Arrived message is received;
b) this state is entered from the ST_TTS1:Send_Frame state; or
c) this state is entered from the ST_TTS6:Process_Data_Out state.

If the Initiator Response Timeout timer expires this state shall send a Data-Out Received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY FAILURE - INITIATOR RESPONSE TIMEOUT. This confirmation shall include the tag. This state machine shall terminate after sending the confirmation.

If this state is entered from the ST_TTS6:Process_Data_Out state and the number of bytes moved for the Receive Data-Out transport protocol service request is less than the Request Byte Count, then this state shall wait for a Data-Out Arrived message.

If this state is entered from the ST_TTS6:Process_Data_Out state and number of bytes moved for the Receive Data-Out transport protocol service request equals the Request Byte Count, then this state shall send a Data-Out Received transport protocol service confirmation to the SCSI application layer with the Delivery Result argument set to DELIVERY SUCCESSFUL. This confirmation shall include the tag. If this state has no more bytes in its first burst buffer, then this state machine shall terminate after sending the confirmation. If this state has more bytes to move in its first burst buffer, then this state machine shall wait for a Receive Data-Out transport protocol service request.

If this state machine receives a Terminate Data Transfer protocol service request from the SCSI application layer, then this state machine shall terminate.

9.2.6.3.3.4.2 Transition ST_TTS3:Receive_Data_Out to ST_TTS4:Prepare_Xfer_Rdy

This transition shall occur after this state receives:

a) a Receive Data-Out transport protocol service request and first burst is not enabled; or
b) a Receive Data-Out transport protocol service request, first burst is enabled, all of the first burst data has been processed, and the first burst data did not satisfy the requested byte count.

9.2.6.3.3.4.3 Transition ST_TTS3:Receive_Data_Out to ST_TTS6:Process_Data_Out

This transition shall occur after this state:

a) receives a Receive Data-Out transport protocol service request; and
b) receives and verifies a Data-Out Arrived message (see 9.2.6.3.3.4.1).

9.2.6.3.3.5 ST_TTS4:Prepare_Xfer_Rdy state

9.2.6.3.3.5.1 State description

This state shall construct an XFER_RDY frame. This state shall use the following values received from the SCSI application layer to construct the frame:

a) logical unit number;
b) tag;
c) target port transfer tag;

d) requested offset; and
e) write data length.

If first burst is enabled, this state shall adjust the write data length to reflect the amount of first burst data.

This state shall generate the following values when constructing the frame:

a) frame type;
b) hashed destination SAS address;
c) hashed source SAS address;
d) retransmit bit set to zero;
e) retry data frames bit set to the value of the TRANSPORT LAYER RETRIES bit in the Protocol-Specific Logical Unit mode page; and
f) number of fill bytes.

It this state was entered from the ST_TTS1:Send_Frame state (i.e., the transmission of the XFER_RDY frame previously constructed by this state was unsuccessful and retries are enabled), then this state shall construct a new XFER_RDY frame using the values from the previous frame except:

a) the RETRANSMIT bit shall be set to one; and
b) the value in the TARGET PORT TRANSFER TAG field shall be set to a different value than the previous XFER_RDY frame. The new TARGET PORT TRANSFER TAG field value shall not conflict with any other target port transfer tag currently in use.

9.2.6.3.3.5.2 Transition ST_TTS4:Prepare_Xfer_Rdy to ST_TTS1:Send_Frame

This transition shall occur after this state constructs an XFER_RDY frame.

9.2.6.3.3.6 ST_TTS6:Process_Data_Out state

9.2.6.3.3.6.1 State description

This state shall process the data received in the Data-Out Arrived message using the Device Server Buffer (e.g., logical block address) to which the data is to be transferred.

9.2.6.3.3.6.2 Transition ST_TTS6:Process_Data_Out to ST_TTS3:Receive_Data_Out

This transition shall occur after this state has processed the data received in a Data-Out message.

9.2.6.3.3.7 ST_TTS6:Prepare_Response state

9.2.6.3.3.7.1 State description

This state shall construct a RESPONSE frame when this state machine receives:

a) a Response Data message;
b) a Task Management Function Executed transport protocol service response; or
c) a Send Command Complete transport protocol service response.

If this state was entered as the result of receiving a Response Data message, this state shall use the logical unit number and tag received in the message and shall construct the frame as described in table 15.

Table 15 — Response Data argument to RESPONSE frame content mapping

Response Data argument	RESPONSE frame
Information Unit Too Short	The DATAPRES field shall be set to SENSE_DATA, the STATUS field shall be set to CHECK CONDITION and the additional sense code shall be set to INFORMATION UNIT TOO SHORT.
Information Unit Too Long	The DATAPRES field shall be set to SENSE_DATA, the STATUS field shall be set to CHECK CONDITION and the additional sense code shall be set to INFORMATION UNIT TOO LONG.
Invalid Frame	The DATAPRES field shall be set to RESPONSE_DATA and the RESPONSE CODE field shall be set to INVALID FRAME.
Invalid Logical Unit Number	The DATAPRES field shall be set to RESPONSE_DATA and the RESPONSE CODE field shall be set to INVALID LOGICAL UNIT NUMBER.

If this state was entered as a result of receiving a Task Management Function Executed transport protocol service response or a Send Command Complete transport protocol service response, this state shall use the following values received from the SCSI application layer to construct the frame:

a) logical unit number;
b) tag;
c) status;
d) response data; and
e) sense data.

This state shall generate the following values when constructing the frame:

a) frame type;
b) hashed destination SAS address;
c) hashed source SAS address;
d) retransmit bit set to zero;
e) number of fill bytes;
f) fill bytes;
g) data present;
h) sense data length; and
i) response data length.

If this state was entered as the result of the ST_TTS1:Send_Frame state receiving something other than a Transmission Status (Frame Transmitted) confirmation followed by a Transmission Status (ACK Received) confirmation for a RESPONSE frame (i.e., the frame transmission was unsuccessful and the vendor specific number of retries has not been reached), then this state shall generate a new RESPONSE frame using all of the values for the previous RESPONSE frame except that the RETRANSMIT bit shall be set to one.

9.2.6.3.3.7.2 Transition ST_TTS6:Prepare_Response to ST_TTS1:Send_Frame

This transition shall occur after this state constructs a RESPONSE frame.

9.3 STP transport layer

9.3.1 Initial FIS

A SATA device phy transmits a Register - Device to Host FIS after completing the link reset sequence. The expander device shall update a set of shadow registers with these contents and shall not deliver them to any STP initiator port. The STP initiator ports may read the shadow register contents using the SMP REPORT PHY SATA function (see 10.4.3.7).

9.3.2 BIST Activate FIS

STP initiator ports and STP target ports shall not generate BIST Activate FISes and shall process any BIST Activate FISes received as frames having invalid FIS types (i.e., have the link layer generate SATA_R_ERR in response).

9.3.3 TT (transport layer for STP ports) state machines

The STP transport layer uses the transport layer state machines defined in SATA, modified to communicate with the port layer rather than directly with the link layer. These modifications are not described in this standard.

9.4 SMP transport layer

9.4.1 SMP transport layer overview

Table 16 defines the SMP frame format.

Table 16 — SMP frame format

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE							
1	Frame-type dependent bytes							
	Fill bytes, if needed							
n - 3	(MSB)	CRC						
n								(LSB)

Table 17 defines the SMP FRAME TYPE field, which defines the format of the frame-type dependent bytes.

Table 17 — SMP FRAME TYPE **field**

Code	Name	Frame type	Originator	Reference
40h	SMP_REQUEST	SMP function request	SMP initiator port	9.4.2
41h	SMP_RESPONSE	SMP function response	SMP target port	9.4.3
All others	Reserved.			

The SMP target port in an expander device shall support the SMP_REQUEST and SMP_RESPONSE frames. Other SMP target ports may support these frames.

Fill bytes shall be included after the frame-type dependent bytes so the crc field is aligned on a four byte boundary. The contents of the fill bytes are vendor specific.

The CRC field contains a CRC value (see 7.5) that is computed over the entire SMP frame prior to the CRC field, and shall begin on a four-byte boundary. The CRC field is checked by the SMP link layer (see 7.18).

9.4.2 SMP_REQUEST frame

The SMP_REQUEST frame is sent by an SMP initiator port to request an SMP function be performed by a management device server. Table 18 defines the SMP_REQUEST frame format.

Table 18 — SMP_REQUEST frame format

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	REQUEST BYTES							
	Fill bytes, if needed							
n - 3	(MSB)	CRC						
n								(LSB)

The SMP FRAME TYPE field shall be set to 40h indicating this is an SMP_REQUEST frame.

The REQUEST BYTES field definition and length is based on the SMP function (see 10.4.3.1). The maximum size of the REQUEST BYTES field is 1024 bytes, making the maximum size of the frame 1032 bytes (1024 bytes of data + 4 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the additional request bytes field so the CRC field is aligned on a four byte boundary. The contents of the fill bytes are vendor specific.

The CRC field is defined in 9.4.1.

9.4.3 SMP_RESPONSE frame

The SMP_RESPONSE frame is sent by an SMP target port in response to an SMP_REQUEST frame. Table 19 defines the SMP_RESPONSE frame format.

Table 19 — SMP_RESPONSE frame format

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	RESPONSE BYTES							
	Fill bytes, if needed							
n - 3	(MSB)	CRC						
n								(LSB)

The SMP FRAME TYPE field shall be set to 41h indicating this is an SMP_RESPONSE frame.

The RESPONSE BYTES field definition and length is based on the SMP function (see 10.4.3.2). The maximum size of the RESPONSE BYTES field is 1024 bytes, making the maximum size of the frame 1032 bytes (1024 bytes of data + 4 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the additional response bytes field so the CRC field is aligned on a four byte boundary. The contents of the fill bytes are vendor specific.

The CRC field is defined in 9.4.1.

9.4.4 Sequence of SMP frames

Inside an SMP connection, the source phy transmits a single SMP_REQUEST frame and the destination phy replies with a single SMP_RESPONSE frame.

Figure 7 shows the sequence of SMP frames.

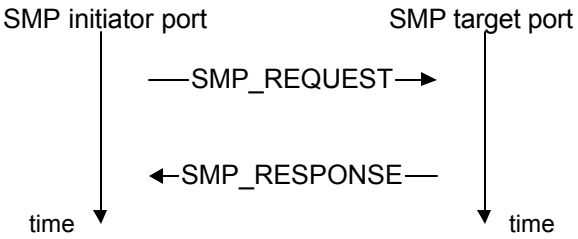

Figure 7 — Sequence of SMP frames

9.4.5 MT (transport layer for SMP ports) state machines

9.4.5.1 SMP transport layer state machines overview

The SMP transport layer contains state machines that process requests from the management application layer and returns confirmations to the management application layer. The SMP transport state machines are as follows:

a) MT_IP (transport layer for SMP initiator ports) state machine (see 9.4.5.2); and
b) MT_TP (transport layer for SMP target ports) state machine (see 9.4.5.3).

9.4.5.2 MT_IP (transport layer for SMP initiator ports) state machine

9.4.5.2.1 MT_IP state machine overview

The MT_IP state machine processes requests from the management application layer. These management requests are sent to the port layer and the resulting SMP frame or error condition is sent to the management application layer as a confirmation.

This state machine consists of the following states:

a) MT_IP1:Idle (see 9.4.5.2.2)(initial state);
b) MT_IP2:Send (see 9.4.5.2.3); and
c) MT_IP3:Receive (see 9.4.5.2.4).

This state machine shall maintain the timers listed in table 20.

Table 20 — MT_IP timers

Timer	Initial value
SMP Frame Receive Timeout timer	Vendor specific

Figure 8 describes the MT_IP state machine.

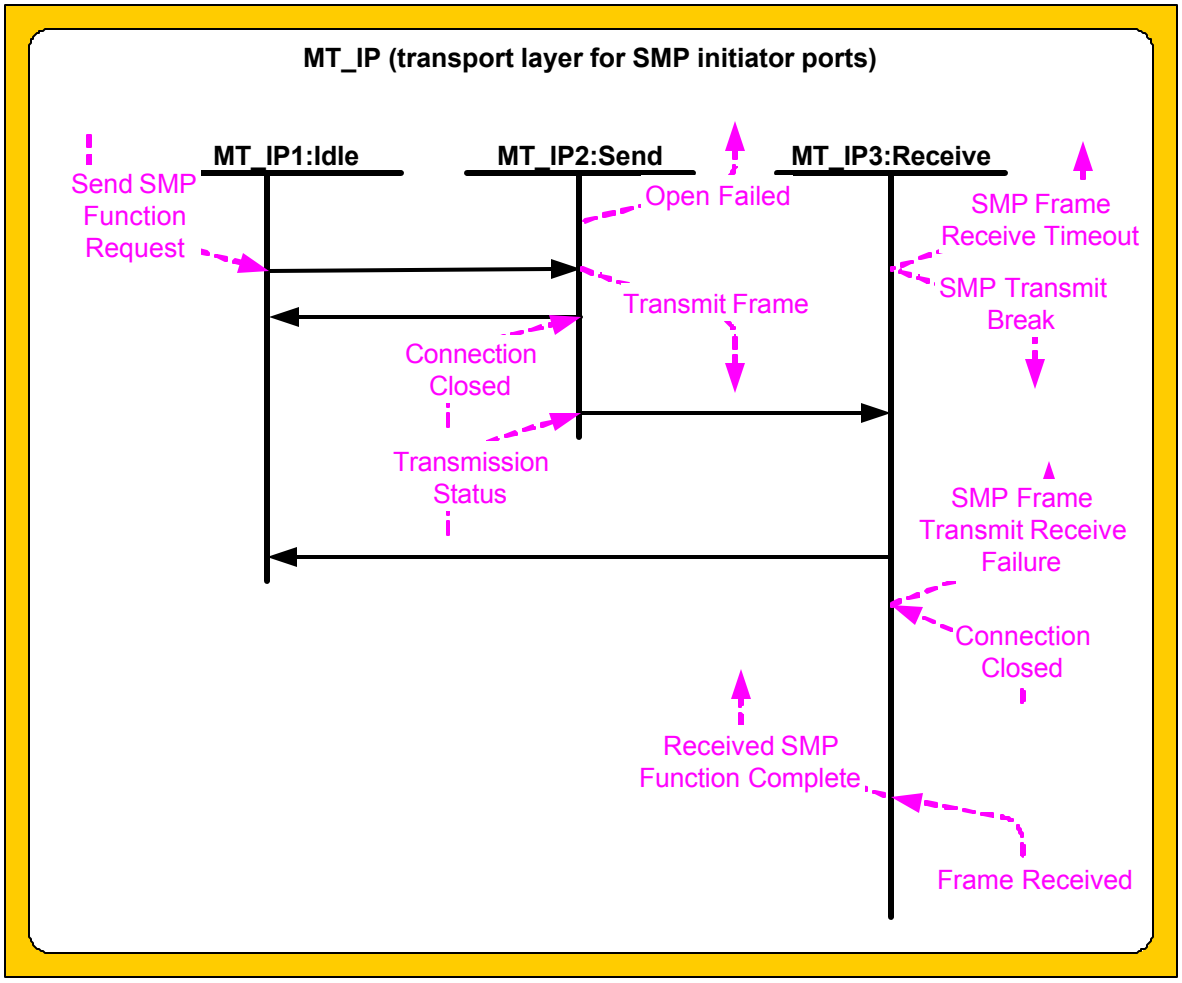

Figure 8 — MT_IP (transport layer for SMP initiator ports) state machine

9.4.5.2.2 MT_IP1:Idle state

9.4.5.2.2.1 State description

This state waits for a Send SMP Function Request request, which includes the following arguments:

a) connection rate;
b) destination SAS address; and
c) request bytes.

9.4.5.2.2.2 Transition MT_IP1:Idle to MT_IP2:Send

This transition shall occur after a Send SMP Function Request request is received. This transition shall include the following arguments:

a) connection rate;
b) destination SAS address; and
c) request bytes.

9.4.5.2.3 MT_IP2:Send state

9.4.5.2.3.1 State description

This state constructs an SMP_REQUEST frame using the following arguments received with the transition into this state:

a) request bytes;

and sends a Transmit Frame request to the port layer. The Transmit Frame request shall contain the following arguments:

a) protocol set to SMP;
b) connection rate;
c) destination SAS address;
d) request bytes;
e) source SAS address set to the SAS address of the SMP initiator port;
f) initiator port bit set to one; and
g) initiator connection tag set to FFFFh.

9.4.5.2.3.2 Transition MT_IP2:Send to MT_IP1:Idle

This transition shall occur after receiving either a Connection Closed confirmation or a Transmission Status confirmation other than a Transmission Status (Frame Transmitted) confirmation, and after sending an Open Failed confirmation to the management application layer.

9.4.5.2.3.3 Transition MT_IP2:Send to MT_IP3:Receive

This transition shall occur after receiving a Transmission Status (Frame Transmitted) confirmation.

9.4.5.2.4 MT_IP3:Receive state

9.4.5.2.4.1 State description

This state waits for a confirmation from the port layer that either an SMP frame has been received or a failure occurred.

Upon entry into this state, this state shall initialize and start the SMP Frame Receive Timeout timer.

If a Frame Received confirmation is received and the SMP frame type is equal to 41h, this state shall send a Received SMP Function Complete confirmation to the management application layer. If the SMP frame type is not equal to 41h, this state shall send an SMP Transmit Break request to the port layer.

If a Connection Closed or Frame Received (SMP Failure) confirmation is received, this state shall send an SMP Frame Transmit Receive Failure confirmation to the management application layer.

If the SMP Frame Receive Timeout timer expires before a Received SMP Function Complete confirmation is received, this state shall send an SMP Frame Receive Timeout confirmation to the management application layer and send an SMP Transmit Break request to the port layer.

9.4.5.2.4.2 Transition MT_IP3:Receive to MT_IP1:Idle

This transition shall occur after one of the following:

a) sending a Received SMP Function Complete confirmation;
b) sending an SMP Frame Transmit Receive Failure confirmation; or
c) sending an SMP Transmit Break request.

9.4.5.3 MT_TP (transport layer for SMP target ports) state machine

9.4.5.3.1 MT_TP state machine overview

The MT_TP state machine informs the management application layer of the receipt of an SMP frame. Confirmation of the receipt of an SMP frame is sent to the management application layer. The management application layer creates the corresponding SMP_RESPONSE frame and this state sends it to the port layer.

This state machine consists of the following states:

a) MT_TP1:Idle (see 9.4.5.3.2)(initial state); and
b) MT_TP2:Respond (see 9.4.5.3.3).

Figure 9 describes the MT_TP state machine.

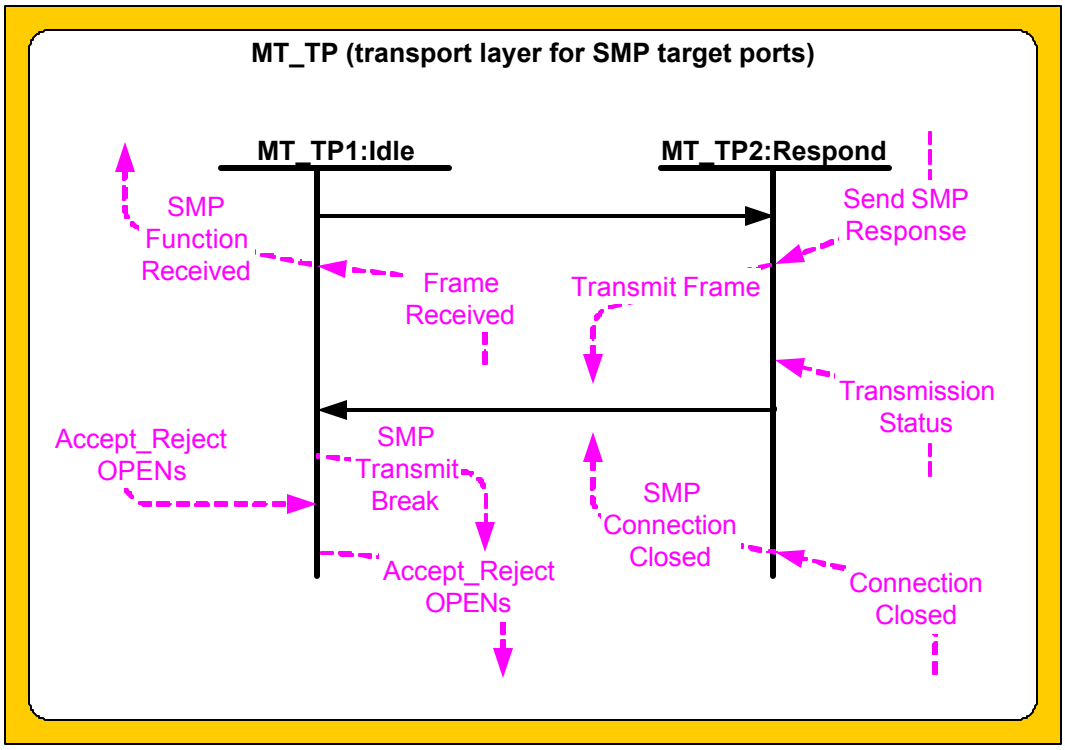

Figure 9 — MT_TP (transport layer for SMP target ports) state machine

9.4.5.3.2 MT_TP1:Idle state

9.4.5.3.2.1 State description

This state waits for a Frame Received confirmation. If the SMP frame type is not equal to 40h, this state shall discard the frame and send a SMP Transmit Break request to the port layer. Otherwise, this state shall send an SMP Function Received confirmation to the management application layer.

If an Accept_Reject OPENs (Accept SMP) or Accept_Reject OPENs (Reject SMP) request is received, this state shall send an Accept_Reject OPENs request with the same arguments to the port layer.

9.4.5.3.2.2 Transition MT_TP1:Idle to MT_TP2:Respond

This transition shall occur after sending an SMP Function Received confirmation.

9.4.5.3.3 MT_TP2:Respond state

9.4.5.3.3.1 State description

This state waits for a Send SMP Response request, which includes the following arguments:

a) response bytes.

After receiving a Send SMP Response request, this state shall construct an SMP_RESPONSE frame using the arguments from the Send SMP Response request and send a Transmit Frame request to the port layer.

If this state receives a Connection Closed confirmation, this state shall send an SMP Connection Closed confirmation to the management application layer.

9.4.5.3.3.2 Transition MT_TP2:Respond to MT_TP1:Idle

This transition shall occur after one of the following:

a) receiving a Transmission Status (Frame Transmitted) confirmation; or
b) sending an SMP Connection Closed confirmation.