To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 27 August 2005
Subject: 05-305r0 SAS-2 Maximum SMP connection time

Revision history
Revision 0 (27 August 2005) First revision

Related documents
sas1r09e - Serial Attached SCSI revision 9e
05-306r0 SAS-2 STP connection time limits and STP/SMP I_T nexus loss

Overview
SMP initiators need to select and implement a vendor-specific timeout for SMP connections. Since there is no standard timeout defined and no way to report the response time of a particular SMP target, there's a risk that an initiator will timeout more quickly than an expander can respond. There's also a risk that expanders will be designed to take an excessive amount of time to respond to SMP requests, since the standard provides no guidance about any limit.

Proposal:

a) Define a 1.9 ms maximum time limit from the SMP request frame to the SMP response frame for the SMP target phy (transport layer state machine)
b) Define a 2 ms connection time limit for SMP, enforced by the port layer state machine (by both the target and initiator)
c) This implies a 0.1 ms maximum time limit from OPEN_ACCEPT to the SMP request frame for the SMP initiator phy. Although none of the state machines (transport, port, link) mention any delay that might occur at this time, there is some delay in real implementations. No 0.1 ms time limit is proposed for any particular state machine. The reader will have to realize that 2 ms - 1.9 ms leaves 0.1 ms.
d) Remove the SMP Frame Receive Timeout timer (which is currently vendor-specific) from the SMP initiator phy transport layer state machine. The new SMP connection time limit replaces it.
e) Start the Bus Inactivity Time Limit timer after connection establishment before the first frame is transmitted. Currently, it only starts on reception of ACK or NAK (which is SSP specific).
f) Honor the Maximum Connect Time Limit timer expiration at any time, not just on reception of ACK or NAK (which is SSP specific and doesn't honor expiration before the first frame is sent, if any).

Another possibility not currently proposed: Define different timeouts for different SMP functions. There could be a different value for each function, or functions could be classified into short (e.g. 1 ms) and long (e.g. 10 ms) categories. Some of the SMP functions for zoning batch together actions that would traditionally be done separately, reducing overall bandwidth usage but increasing latency for other traffic in the fabric. Although such functions require longer processing time, they eliminates many additional SMP function requests (and connections) and might be worth the complexity to allow a longer timeout.

Suggested changes

7.18 SMP link layer

7.18.1 SMP frame transmission and reception

Inside an SMP connection, the ~~source device~~ SMP initiator phy transmits a single SMP_REQUEST frame within 100 µs and the ~~destination device~~SMP target phy responds with a single SMP_RESPONSE frame within 1 900 µs (see 9.4).

Frames are surrounded by SOF and EOF as shown in figure 1. See 7.18.4 for error handling details.

NOTE 1 - Unlike SSP, there is no acknowledgement of SMP frames with ACK and NAK and there is no credit exchange with RRDY.

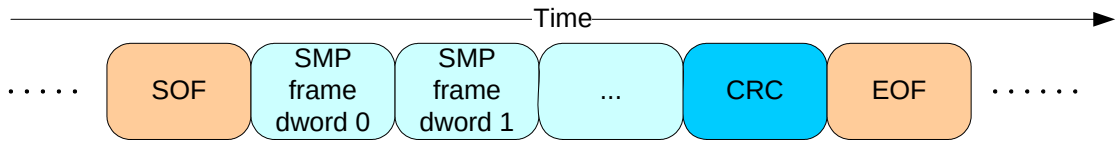

Figure 1 — SMP frame transmission

The last data dword after the SOF prior to the EOF always contains a CRC (see 7.5). The SMP link layer state machine checks that the frame is not too short and that the CRC is valid (see 7.18.4).

7.18.2 SMP flow control

By accepting an SMP connection, the ~~destination device~~SMP target phy indicates it is ready to receive one SMP_REQUEST frame.

After the ~~source device~~ SMP initiator phy transmits one SMP_REQUEST frame, it shall be ready to receive one SMP_RESPONSE frame.

7.18.3 Closing an SMP connection

After receiving the SMP_RESPONSE frame, the ~~source device~~ SMP initiator phy shall transmit a CLOSE (NORMAL) to close the connection.

After transmitting the SMP_RESPONSE frame, the ~~destination device~~SMP target phy shall reply with a CLOSE (NORMAL).

See 7.12.6 for additional details on closing connections.

7.18.4 SMP (link layer for SMP phys) state machines

7.18.4.1 SMP state machines overview

The SMP state machines control the flow of dwords on the physical link during an SMP connection. The SMP state machines are as follows:

a) SMP_IP (link layer for SMP initiator phys) state machine (see 7.18.4.3); and
b) SMP_TP (link layer for SMP target phys) state machine (see 7.18.4.4).

7.18.4.2 SMP transmitter and receiver

The SMP transmitter receives the following messages from the SMP state machines specifying dwords and frames to transmit:

a) Transmit Idle Dword; and
b) Transmit Frame.

The SMP transmitter sends the following messages to the SMP state machines based on dwords that have been transmitted:

a) Frame Transmitted.

When there is no outstanding message specifying a dword to transmit, the SMP transmitter shall transmit idle dwords.

The SMP receiver sends the following messages to the SMP state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2):

a) SOF Received;
b) Data Dword Received;
c) EOF Received;

d) ERROR Received; and
e) Invalid Dword Received.

The SMP receiver shall ignore all other dwords.

7.18.4.3 SMP_IP (link layer for SMP initiator phys) state machine

7.18.4.3.1 SMP_IP state machine overview

The SMP_IP state machine's function is to transmit an SMP request frame and then receive the corresponding response frame. This state machine consists of the following states:

a) SMP_IP1:Idle (see 7.18.4.3.2)(initial state);
b) SMP_IP2:Transmit_Frame (see 7.18.4.3.3); and
c) SMP_IP3:Receive_Frame (see 7.18.4.3.4).

This state machine shall start in the SMP_IP1:Idle state on receipt of an Enable Disable SMP (Enable) message from the SL state machines (see 7.14).

The SMP_IP state machine shall terminate after receiving an Enable Disable SMP (Disable) message from the SL state machines.

Figure 2 shows the SMP_IP state machine.

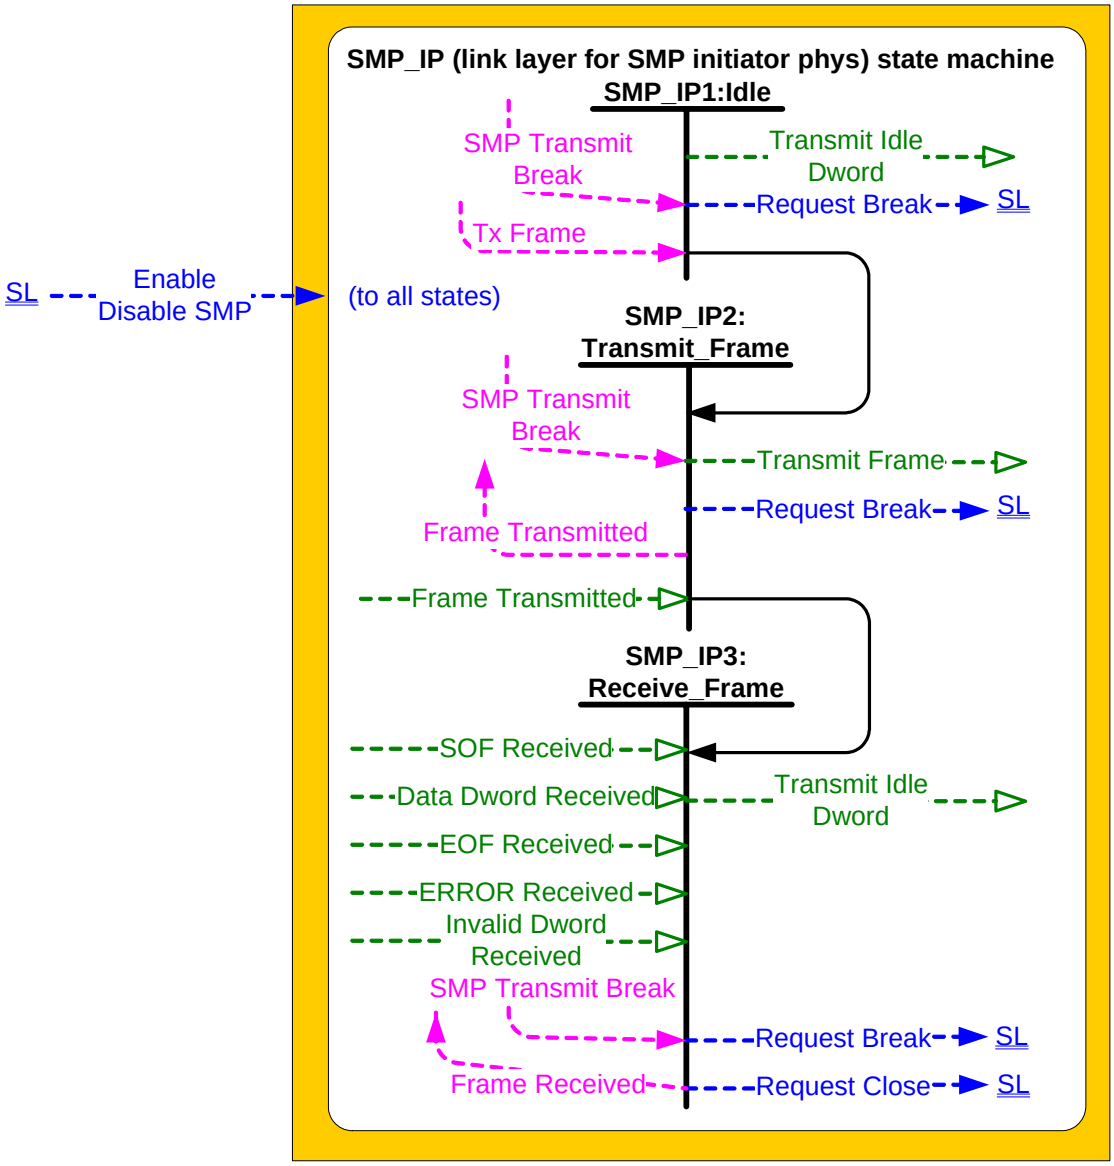

Figure 2 — SMP_IP (link layer for SMP initiator phys) state machine

7.18.4.3.2 SMP_IP1:Idle state

7.18.4.3.2.1 State description

This state is the initial state.

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SMP transmitter.

If an SMP Transmit Break request is received, this state shall send a Request Break message to the SL state machines (see 7.14).

7.18.4.3.2.2 Transition SMP_IP1:Idle to SMP_IP2:Transmit_Frame

This transition shall occur after a Tx Frame request is received.

7.18.4.3.3 SMP_IP2:Transmit_Frame state

7.18.4.3.3.1 State description

This state shall send a Transmit Frame message to the SMP transmitter.

If an SMP Transmit Break request is received, this state shall send a Request Break message to the SL state machines (see 7.14) and terminate.

After the Frame Transmitted message is received, this state shall send a Frame Transmitted confirmation to the port layer.

7.18.4.3.3.2 Transition SMP_IP2:Transmit_Frame to SMP_IP3:Receive_Frame

This transition shall occur after sending a Frame Transmitted confirmation to the port layer.

7.18.4.3.4 SMP_IP3:Receive_Frame state

This state checks the SMP response frame and determines if the SMP response frame was successfully received (e.g., no CRC error).

If this state receives a subsequent SOF Received message after receiving an SOF Received message but before receiving an EOF Received message (i.e., SOF, data dwords, SOF, data dwords, and EOF instead of SOF, data dwords, EOF, SOF, data dwords, and EOF), then this state shall discard the Data Dword Received messages received before the subsequent SOF Received message.

This state shall discard the frame, send a Frame Received (SMP Failure) confirmation to the port layer, send a Request Break message to the SL state machines, and terminate the state machine if:

a) this state receives more than 258 Data Dword Received messages after an SOF Received message and before an EOF Received message; or
b) this state receives fewer than 2 Data Dword Received messages after an SOF Received message and before an EOF Received message.

If this state receives an Invalid Dword Received message or an ERROR Received message after an SOF Received message and before an EOF Received message, then this state machine shall:

a) ignore the invalid dword or ERROR; or
b) discard the frame, send a Frame Received (SMP Failure) confirmation to the port layer, send a Request Break message to the SL state machines, and terminate the state machine.

If the SMP response frame is received with a CRC error, this state shall discard the frame, send a Frame Received (SMP Failure) confirmation to the port layer, send a Request Break message to the SL state machines, and terminate the state machine.

If the SMP response frame is received with no CRC error and the SMP response frame is valid, this state shall:

a) send a Frame Received confirmation to the port layer; and
b) send a Request Close message to the SL state machines (see 7.14).

If an SMP Transmit Break request is received, this state shall send a Request Break message to the SL state machines and this state machine shall terminate.

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SMP transmitter.

7.18.4.4 SMP_TP (link layer for SMP target phys) state machine

7.18.4.4.1 SMP_TP state machine overview

The SMP_TP state machine's function is to receive an SMP request frame and then transmit the corresponding SMP response frame. The SMP_TP state machine consists of the following states:

a) SMP_TP1:Receive_Frame (see 7.18.4.4.2)(initial state); and

b) SMP_TP2:Transmit_Frame (see 7.18.4.4.3).

This state machine shall start in the SMP_TP1:Receive_Frame state after receiving an Enable Disable SMP (Enable) message from the SL state machines (see 7.14).

The SMP_TP state machine shall terminate after receiving an Enable Disable SMP (Disable) message from the SL state machines.

Figure 3 shows the SMP_TP state machine.

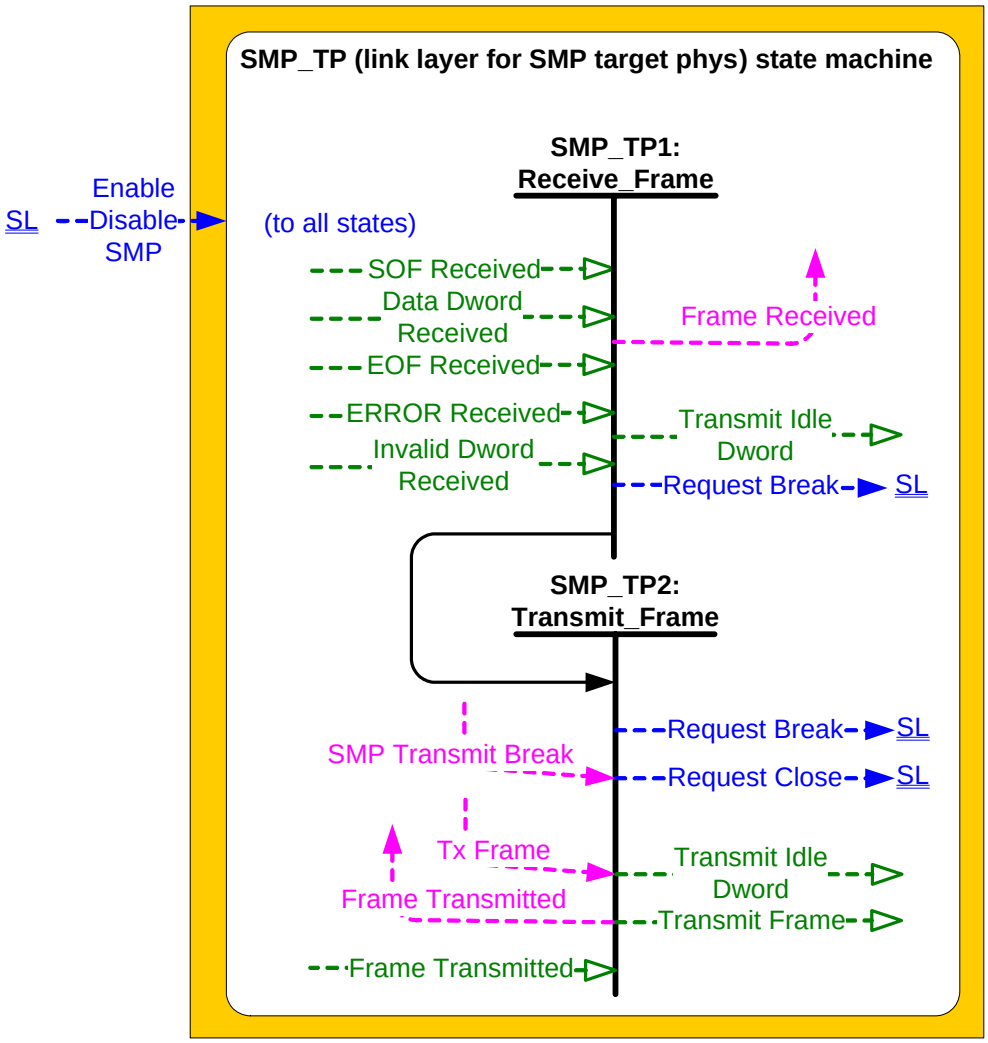

Figure 3 — SMP_TP (link layer for SMP target phys) state machine

7.18.4.4.2 SMP_TP1:Receive_Frame state

7.18.4.4.2.1 State description

This state waits for an SMP frame and determines if the SMP frame was successfully received (e.g., no CRC error).

If this state receives a subsequent SOF Received message after receiving an SOF Received message but before receiving an EOF Received message (i.e., SOF, data dwords, SOF, data dwords, and EOF instead of SOF, data dwords, EOF, SOF, data dwords, and EOF), then this state shall discard the Data Dword Received messages received before the subsequent SOF Received message.

This state shall discard the frame, send a Request Break message to the SL state machines (see 7.14) and shall terminate the state machine if:

a) this state receives more than 258 Data Dword Received messages after an SOF Received message and before an EOF Received message; or
b) this state receives fewer than 2 Data Dword Received messages after an SOF Received message and before an EOF Received message.

If this state receives an Invalid Dword Received message or an ERROR Received message after an SOF Received message and before an EOF Received message, then this state machine shall:

a) ignore the invalid dword or ERROR; or
b) discard the frame, send a Request Break message to the SL state machines (see 7.14) and shall terminate the state machine.

If the SMP request frame is received with a CRC error, this state shall discard the frame, send a Request Break message to the SL state machines (see 7.14) and shall terminate the state machine.

Otherwise, this state shall send a Frame Received confirmation to the port layer.

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SMP transmitter.

7.18.4.4.2.2 Transition SMP_TP1:Receive_Frame to SMP_TP2:Transmit_Frame

This transition shall occur after sending a Frame Received confirmation to the port layer.

7.18.4.4.3 SMP_TP2:Transmit_Frame state

If this state receives an SMP Transmit Break request, this state shall send a Request Break message to the SL state machines and terminate.

If this state receives a Tx Frame request, this state shall send a Transmit Frame message to the SMP transmitter, then wait for a Frame Transmitted message. After receiving a Frame Transmitted message, this state shall send a Frame Transmitted confirmation to the port layer, send a Request Close message to the SL state machines (see 7.14) and terminate.

After sending Transmit Frame message to the SMP transmitter, this state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SMP transmitter.

Port layer changes

8.2.3 PL_PM (port layer phy manager) state machine

8.2.3.1 PL_PM state machine overview

A PL_PM state machine:

a) receives messages from the PL_OC state machine;
b) sends requests to the link layer;
c) receives confirmations from the link layer;
d) sends confirmations to the transport layer;
e) sends messages to PL_OC state machine;
f) has an Arbitration Wait Time timer;
g) may have a Bus Inactivity Time Limit timer; and
h) may have Maximum Connect Time Limit timer.

This state machine consist of the following states:

a) PL_PM1:Idle (see 8.2.3.2) (initial state);
b) PL_PM2:Req_Wait (see 8.2.3.3);
c) PL_PM3:Connected (see 8.2.3.4); and
d) PL_PM4:Wait_For_Close (see 8.2.3.5).

After power on this state machine shall start in the PL_PM1:Idle state.

The PL_PM state machine shall maintain the timers listed in table 1.

Table 1 — PL_PM state machine timers

Timer	Initial value
Arbitration Wait Time timer	The arbitration wait time argument from a Retry Open message (see 8.2.2.3.1).
Bus Inactivity Time Limit timer	The value in the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1).
Maximum Connect Time Limit timer	Depending on the protocol used by the port: a) for SSP target ports, the value in the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1); or b) for SMP target ports, 2 ms.

Figure 4 shows part 1 of the PL_PM state machine.

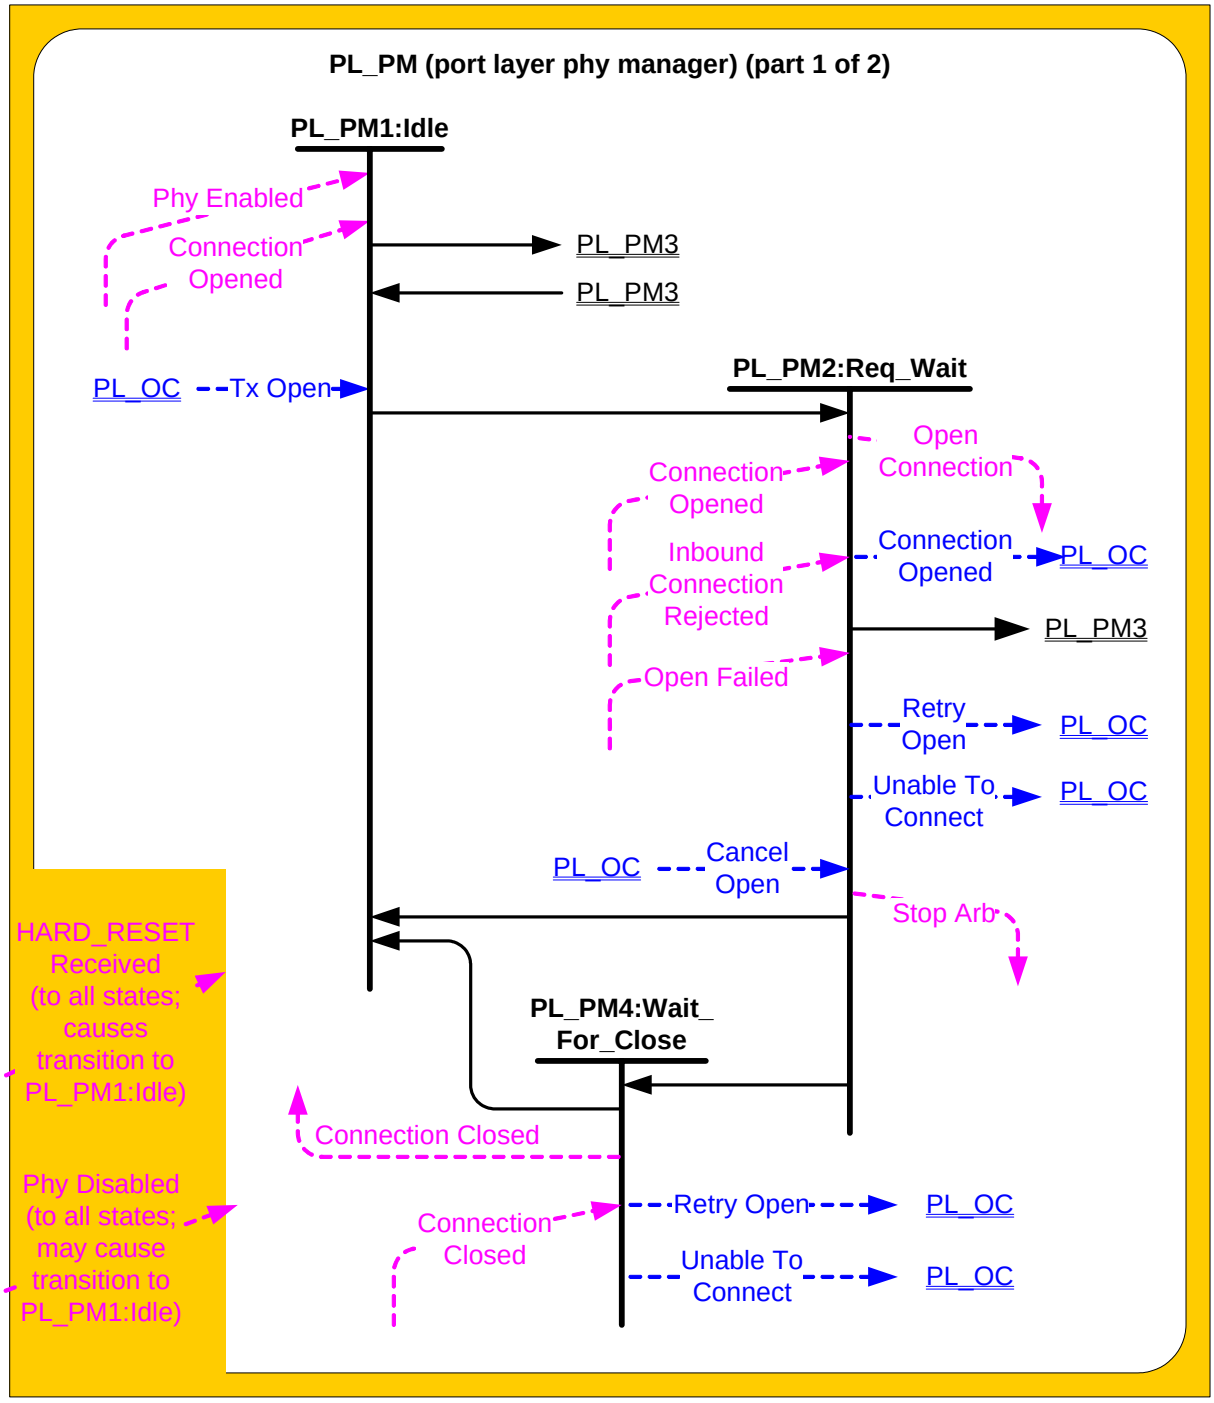

Figure 4 — PL_PM (port layer phy manager) state machine (part 1)

Figure 5 shows part 2 of the PL_PM state machine.

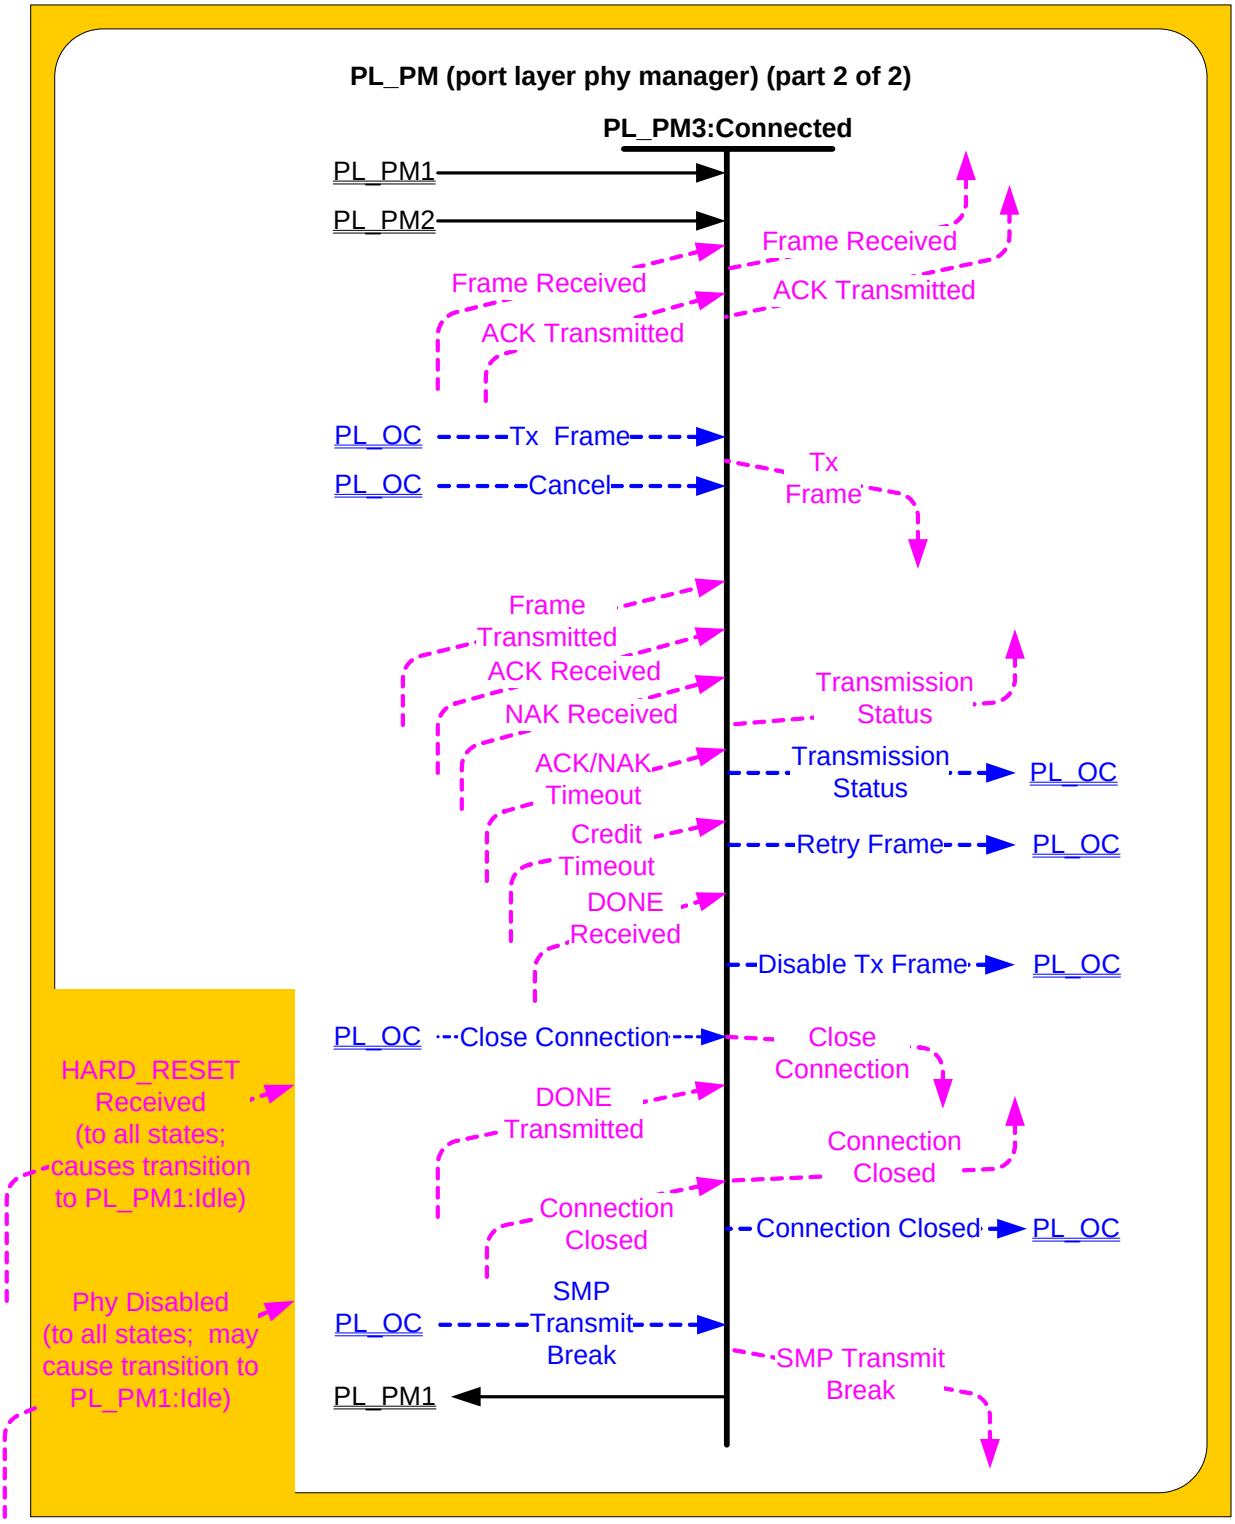

Figure 5 — PL_PM (port layer phy manager) state machine (part 2)

8.2.3.2 PL_PM1:Idle state

8.2.3.2.1 PL_PM1:Idle state description

This is the initial state of the PL_PM state machine.

8.2.3.2.2 Transition PL_PM1:Idle to PL_PM2:Req_Wait

This transition shall occur after:

a) a Phy Enabled confirmation is received; and
b) a Tx Open message is received.

8.2.3.2.3 Transition PL_PM1:Idle to PL_PM3:Connected

This transition shall occur after a Connection Opened confirmation is received.

8.2.3.3 PL_PM2:Req_Wait state

8.2.3.3.1 PL_PM2:Req_Wait state overview

This state sends an Open Connection request to the link layer and waits for a confirmation. This state sends and receives connection management messages to and from the PL_OC state machine.

If this state receives a HARD_RESET Received confirmation, then this state shall terminate all operations.

8.2.3.3.2 PL_PM2:Req_Wait establishing a connection

Upon entry into this state, this state shall:

a) create an Arbitration Wait Time timer;
b) initialize the Arbitration Wait Time timer to the arbitration wait time argument received with the Tx Open message;
c) start the Arbitration Wait Time timer; and
d) send an Open Connection request to the link layer.

The Open Connection request shall contain the following arguments from the Tx Open message to be used in an OPEN address frame:

a) initiator port bit;
b) protocol;
c) connection rate;
d) initiator connection tag;
e) destination SAS address;
f) source SAS address;
g) pathway blocked count; and
h) arbitration wait time.

8.2.3.3.3 PL_PM2:Req_Wait connection established

If this state receives a Connection Opened confirmation, then this state shall send a Connection Opened message to the PL_OC state machine.

If this state receives a Connection Opened confirmation and the confirmation was not in response to an Open Connection request from this state (i.e., the connection was established in response to an OPEN address frame from another device), then this state shall discard any Open Connection request and send a Retry Open message to the PL_OC state machine. If the Connection Opened confirmation was from the destination of the Open Connection request, then this state shall send a Retry Open (Opened By Destination) message. If the Connection Opened confirmation was from a destination other than the destination of the Open Connection request, then this state shall send a Retry Open (Opened By Other) message.

A Retry Open (Opened By Destination) or Retry Open (Opened By Other) message shall contain the following arguments:

a) initiator port bit set to the value received with the Tx Open message;
b) protocol set to the value received with the Tx Open message;
c) connection rate set to the value received with the Tx Open message;
d) initiator connection tag set to the value received with the Tx Open message;
e) destination SAS address set to the value received with the Tx Open message;

f) source SAS address set to the value received with the Tx Open message;
g) pathway blocked count set to the value received with the Tx Open message; and
h) arbitration wait time set to the value of the Arbitration Wait Time timer.

8.2.3.3.4 PL_PM2:Req_Wait unable to establish a connection

If this state receives one of the Open Failed confirmations listed in table 2, then this state shall send either a Retry Open message or an Unable To Connect message to the PL_OC state machine.

Table 2 defines the message to be sent for each Open Failed confirmation.

Table 2 — Messages from Open Failed confirmations

Confirmation received	Message to be sent to PL_OC
Open Failed (Pathway Blocked)	Retry Open (Pathway Blocked)
Open Failed (Retry)	Retry Open (Retry)
Open Failed (No Destination)	Retry Open (No Destination)
Open Failed (Bad Destination)	Unable To Connect (Bad Destination)
Open Failed (Connection Rate Not Supported)	Unable To Connect (Connection Rate Not Supported)
Open Failed (Protocol Not Supported)	Unable To Connect (Protocol Not Supported)
Open Failed (STP Resources Busy)	Unable To Connect (STP Resources Busy)
Open Failed (Wrong Destination)	Unable To Connect (Wrong Destination)

If this state receives an Inbound Connection Rejected confirmation after sending an Open Connection request, then this state shall discard the Open Connection request and send a Retry Open (Collided) message to the PL_OC state machine.

A Retry Open message shall include the following arguments:

a) initiator port bit set to the value received with the Tx Open message;
b) protocol set to the value received with the Tx Open message;
c) connection rate set to the value received with the Tx Open message;
d) initiator connection tag set to the value received with the Tx Open message;
e) destination SAS address set to the value received with the Tx Open message;
f) source SAS address set to the value received with the Tx Open message;
g) pathway blocked count argument set to the value received with the Tx Open message; and
h) arbitration wait time set to the value of the Arbitration Wait Time timer.

An Unable To Connect message shall include the following arguments:

a) initiator connection tag set to the value received with the Tx Open message;
b) destination SAS address set to the value received with the Tx Open message; and
c) source SAS address set to the value received with the Tx Open message.

8.2.3.3.5 PL_PM2:Req_Wait connection management

If this state receives a Cancel Open message and a Connection Opened confirmation has not been received, then this state shall send a Stop Arb request to the link layer.

8.2.3.3.6 Transition PL_PM2:Req_Wait to PL_PM1:Idle

This transition shall occur after:

a) a Retry Open message is sent to the PL_OC state machine;
b) an Unable To Connect message is sent to the PL_OC state machine;

c) all operations have been terminated after a HARD_RESET Received confirmation is received; or
d) a Phy Disabled confirmation is received.

8.2.3.3.7 Transition PL_PM2:Req_Wait to PL_PM3:Connected

This transition shall occur after a Connection Opened confirmation is received.

8.2.3.3.8 Transition PL_PM2:Req_Wait to PL_PM4:Wait_For_Close

This transition shall occur after one of the following confirmations is received:

a) an Open Failed (Open Timeout Occurred);
b) an Open Failed (Break Received); or
c) an Open Failed (Port Layer Request).

8.2.3.4 PL_PM3:Connected state

8.2.3.4.1 PL_PM3:Connected state description

If:

a) the protocol for the connection is SSP, ~~and this state is in~~the port is an SSP target port, and the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) is not set to zero; or
b) the protocol for the connection is SMP and the port is an SMP target port,

then, upon entry into this state, this state shall:

a) create a Maximum Connect Time Limit timer;
b) initialize the Maximum Connect Time Limit timer as specified in table xx (see 8.2.3.1); and
c) start the Maximum Connect Time Limit timer.

~~If the protocol for the connection is SSP, and this state is in an SSP target port, and the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) is set to zero, then this state shall not create a Maximum Connect Time Limit timer (i.e., there is no maximum connect time limit).~~

Editor's Note 1: the "shall not" rule is unnecessary

Other SAS ports may implement a Maximum Connect Time Limit timer in a vendor-specific manner.

If the protocol for the connection is SSP, ~~and~~ this state is in an SSP target port, and the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) is set to a non-zero value, then, upon entry into this state, this state shall:

a) create a Bus Inactivity Time Limit timer;
b) initialize the Bus Inactivity Time Limit timer as specified in table xx (see 8.2.3.1); and
c) ~~not~~ start the Bus Inactivity Time Limit timer.

Editor's Note 2: let the Bus Inactivity Time Limit timer run starting at the beginning of a connection rather than wait for the first ACK/NAK to start it.

~~If the protocol for the connection is SSP, and this state is in an SSP target port, and the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) is set to zero, then this state shall not create a Bus Inactivity Time Limit timer (i.e., there is no maximum bus inactivity time limit).~~

Editor's Note 3: the "shall not" rule is unnecessary

Other SAS ports may implement a Bus Inactivity Time Limit timer in a vendor-specific manner.

If a Bus Inactivity Time Limit timer has been created and this state receives a Tx Frame message, then this state shall:

a) stop the Bus Inactivity Time Limit timer, if it is running; and
b) initialize the Bus Inactivity Time Limit timer as specified in table xx (see 8.2.3.1).

If this state receives a Tx Frame message, this state shall send a Tx Frame request to the link layer. The following arguments from the Tx Frame message shall be included with the Tx Frame request:

a) the frame to be transmitted; and
b) if this state is in an SSP port, Balance Required or Balance Not Required.

For STP connections, this state connects the STP transport layer to the STP link layer.

If a Bus Inactivity Time Limit timer ~~has been created and this state receives an ACK Received or NAK Received confirmation, then this state shall start the Bus Inactivity Time Limit timer. If the Bus Inactivity Time Limit timer~~ expires ~~before this state receives a Tx Frame message,~~ then ~~this state shall send a Close Connection request to the link layer~~.

a) if the connection is SSP and there is no Tx Frame request outstanding (i.e., this state is not waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer; and
b) if the connection is SSP and there is a Tx Frame request outstanding (i.e., this state is waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer after receiving an ACK Received or NAK Received confirmation.

Editor's Note 4: let the Bus Inactivity Time Limit timer run starting at the beginning of a connection rather than wait for the first ACK/NAK to start it. Split into a) b) format to facilitate also using this timer for STP (see 05-306) and to parallel the changes in the next paragraph.

If a Maximum Connect Time Limit timer ~~has been created and this state receives an ACK Received or NAK Received confirmation, then this state shall check the Maximum Connect Time Limit timer. If the Maximum Connect Time Limit timer has expired~~ expires:

a) if the connection is SSP and there is no Tx Frame request outstanding (i.e., this state is not waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer; ~~and~~
b) if the connection is SSP and there is a Tx Frame request outstanding (i.e., this state is waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer after receiving an ACK Received or NAK Received confirmation; and
c) if the connection is SMP, then this state shall send an SMP Transmit Break request to the link layer.

Editor's Note 5: As written, the state machine would never close an SSP connection if the phy never chose to transmit a frame. It should do so - the other side might be broken and leave the connection hanging. It is important that the phy wait for an ACK or NAK after it has a frame outstanding, however, to allow it to choose DONE (ACK/NAK TIMEOUT) if appropriate. Split into a)b) format to facilitate using this timer for SMP and STP (in 05-306).

If this state receives a Tx Frame message after sending a Close Connection request but before receiving a Connection Closed confirmation, then this state shall send a Retry Frame message to the PL_OC state machine.

If this state receives a Frame Received confirmation, then this state shall send a Frame Received confirmation to the transport layer. The confirmation shall include the arguments received with the confirmation (e.g., the frame).

If this state receives an ACK Transmitted confirmation, then this state shall send an ACK Transmitted confirmation to the transport layer including the tag of the frame that was ACKed.

If this state receives a Frame Transmitted confirmation, then this state shall send a Transmission Status (Frame Transmitted) confirmation to the transport layer.

If this state receives an ACK Received confirmation, then this state shall send a Transmission Status (ACK Received) confirmation to the transport layer.

If this state receives a NAK Received confirmation, then this state shall send a Transmission Status (NAK Received) confirmation to the transport layer.

If this state receives an ACK/NAK Timeout confirmation, then this state shall send a Transmission Status (ACK/NAK Timeout) confirmation to the transport layer.

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

a) discard all Tx Frame requests for the specified tag;
b) send a Transmission Status (Cancel Acknowledge) confirmation to the transport layer including the destination SAS address and the tag as arguments; and
c) discard any subsequent confirmations for previous Tx Frame requests sent for the tag.

If this state receives a Close Connection message from the PL_OC state machine, then this state shall send a Close Connection request to the link layer.

If this state receives one of the following:

a) a Connection Closed (Normal) confirmation;
b) a Connection Closed (Close Timeout) confirmation;
c) a Connection Closed (Break Requested) confirmation;
d) a Connection Closed (Break Received) confirmation; or
e) a Connection Closed (Transition to Idle) confirmation,

then this state shall send a Connection Closed message to the PL_OC state machine including the argument received with the confirmation.

If this state receives a Connection Closed (Transition to Idle) confirmation after receiving:

a) a Connection Closed (Break Received) confirmation; or
b) a Connection Closed (Break Requested) confirmation,

then this state shall send a Transmission Status (Break Received) confirmation to the transport layer.

If this state receives a Connection Closed (Normal) confirmation, a Connection Closed (Transition to Idle) confirmation, or a Phy Disabled confirmation after sending a Transmission Status (Frame Transmitted) confirmation, but before this state receives an ACK Received or NAK Received confirmation, then this state shall send a Transmission Status (Connection Lost Without ACK/NAK) confirmation to the transport layer.

If this state receives a Connection Closed (Normal) confirmation, a Connection Closed (Transition to Idle) confirmation, or a Phy Disabled confirmation after sending a Tx Frame request but before receiving a Frame Transmitted confirmation, then this state shall send a Retry Frame message to the PL_OC state machine.

If this state receives a Connection Closed confirmation during an SMP connection, this state shall send a Connection Closed confirmation to the transport layer.

If this state receives a Credit Timeout confirmation, then this state shall send a Retry Frame message to the PL_OC state machine.

A Retry Frame message shall include the following arguments from the Tx Frame message:

a) initiator port bit;
b) protocol;
c) connection rate;
d) initiator connection tag;
e) destination SAS address;
f) source SAS address; and

g) frame.

After this state receives a DONE Received (Normal) or DONE Received (Credit Blocked) confirmation, if it does not receive a Tx Frame message within 1 ms, then this state shall send a Disable Tx Frames message to the PL_OC state machine.

If this state receives a DONE Received (ACK/NAK Timeout) or DONE Transmitted confirmation, then this state shall send a Disable Tx Frames message to the PL_OC state machine.

If this state receives an SMP Transmit Break message, then this state shall send an SMP Transmit Break request to the link layer.

Editor's Note 6: SMP Transmit Break is in the figure but not in the text

If this state receives a HARD_RESET Received confirmation, then this state machine shall terminate all operations.

8.2.3.4.2 Transition PL_PM3:Connected to PL_PM1:Idle

This transition shall occur after:

a) a Connection Closed (Transition to Idle) message is sent to the PL_OC state machine; or
b) all operations are terminated after a HARD_RESET Received confirmation is received.

8.2.3.5 PL_PM4:Wait_For_Close state

8.2.3.5.1 PL_PM4:Wait_For_Close state description

After receiving a Connection Closed (Transition to Idle) confirmation, if this state was entered as the result of the PL_PM2:Req_Wait state receiving an Open Failed (Open Timeout Occurred) confirmation, then this state shall send a Retry Open (Open Timeout Occurred) message to the PL_OC state machine. The Retry Open message shall include the following arguments:

a) initiator port bit set to the value received with the Tx Open message;
b) protocol set to the value received with the Tx Open message;
c) connection rate set to the value received with the Tx Open message;
d) initiator connection tag set to the value received with the Tx Open message;
e) destination SAS address set to the value received with the Tx Open message;
f) source SAS address set to the value received with the Tx Open message;
g) pathway blocked count argument set to the value received with the Tx Open message; and
h) arbitration wait time set to the value of the Arbitration Wait Time timer.

If this state receives a Connection Closed confirmation and the connection request was for an SMP connection, this state shall send a Connection Closed confirmation to the transport layer.

After receiving a Connection Closed (Transition to Idle) confirmation, if this state was entered after the PL_PM2:Req_Wait state received an Open Failed (Port Layer Request) confirmation (i.e., as the result of the PL_PM2:Req_Wait state sending a Stop Arb request), then this state shall send an Unable to Connect (Port Layer Request) message to the PL_OC state machine.

After receiving a Connection Closed (Transition to Idle) confirmation, if this state was entered as the result of the PL_PM2:Req_Wait state receiving an Open Failed (Break Received) confirmation, then this state shall send an Unable to Connect (Break Received) message to the PL_OC state machine.

The Unable To Connect message shall include the following arguments:

a) initiator connection tag set to the value received with the Tx Open message;
b) destination SAS address set to the value received with the Tx Open message; and
c) source SAS address set to the value received with the Tx Open message.

If this state receives a HARD_RESET Received confirmation, then this state shall terminate all operations.

8.2.3.5.2 Transition PL_PM4:Wait_For_Close to PL_PM1:Idle

This transition shall occur after:

a) a Retry Open or Unable To Connect message is sent to the PL_OC state machine; or
b) all operations are terminated after a HARD_RESET Received confirmation is received.

Transport layer changes

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 start in the MT_IP1:Idle state.

~~The MT_IP state machine shall maintain the timers listed in table 3.~~

~~**Table 3 — MT_IP timers**~~

~~Timer~~	~~Initial value~~
~~SMP Frame Receive Timeout timer~~	~~Vendor specific~~

Editor's Note 7: the port layer Maximum Connect Time Limit timer replaces the SMP Frame Receive Timeout timer. The transport layer should not be running timers.

Figure 6 describes the MT_IP state machine.

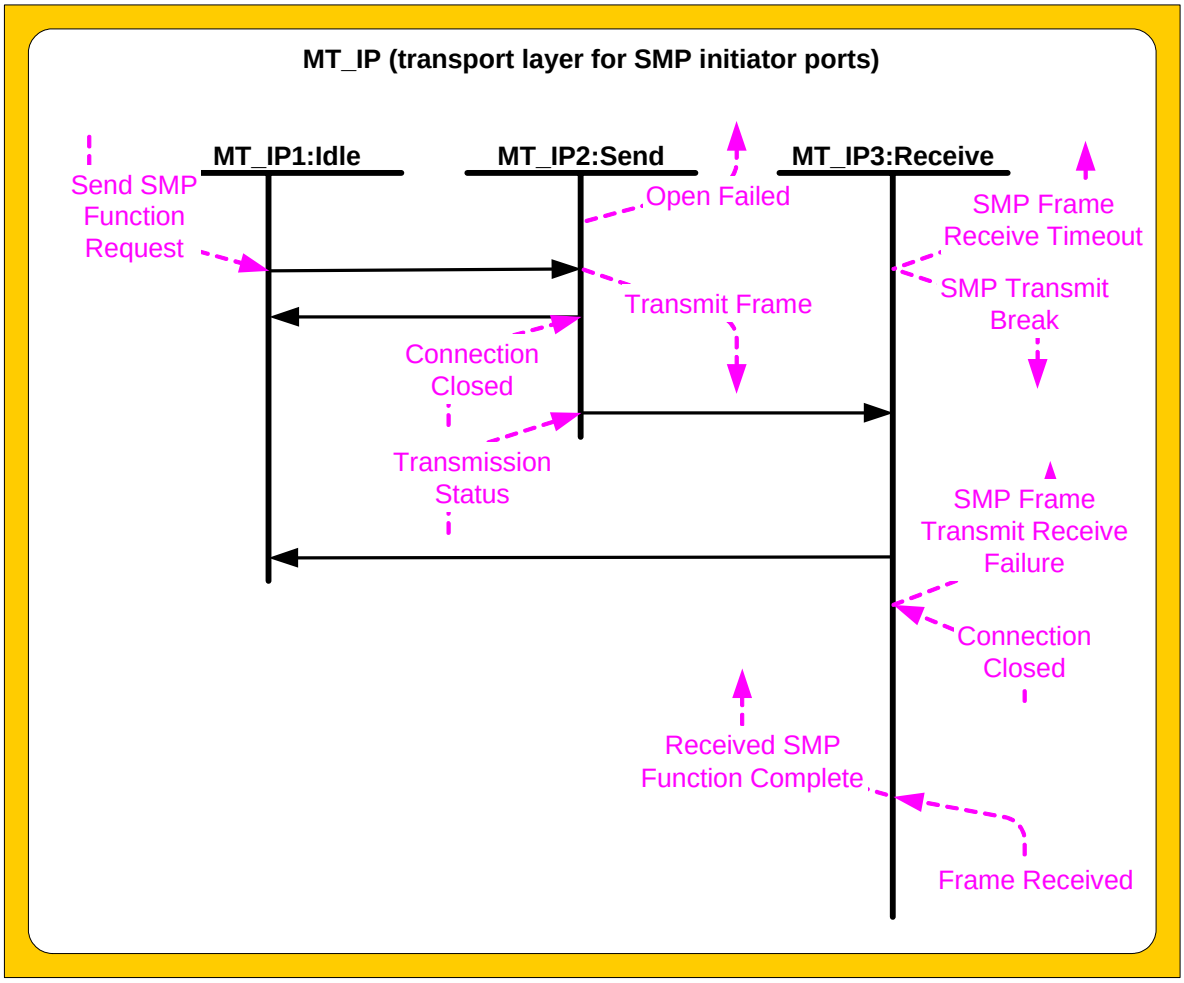

Figure 6 — MT_IP (transport layer for SMP initiator ports) state machine

Editor's Note 8: Delete the SMP Transmit Break request from MT_IP3:Receive state

9.4.5.2.2 MT_IP1:Idle state

9.4.5.2.2.1 State description

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

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 in the transition into this state:

a) request bytes;

and sends a Transmit Frame request to the port layer with the following arguments:

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

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 a Frame Received confirmation is received and the SMP frame type is not equal to 41h, this state shall send a SMP Frame Transmit Receive Failure confirmation to the management application 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).

This state machine shall start in the MT_TP1:Idle state.

Figure 7 describes the MT_TP state machine.

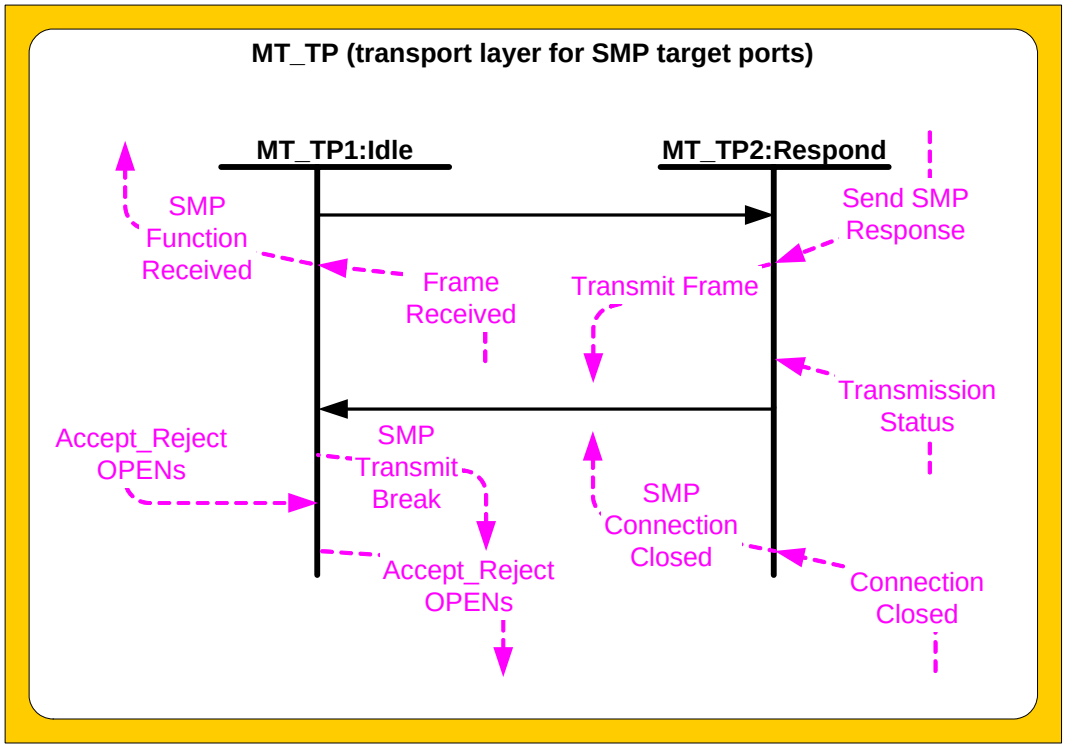

Figure 7 — MT_TP (transport layer for SMP target ports) state machine

The MT_TP state machine shall comply with the time limits listed in table 3.

Table 4 — MT_TP time limits

Time limit	Value	Description
SMP Response time limit	1 900 µs	Maximum time from receiving an SMP_REQUEST frame to transmitting an SMP_RESPONSE frame.

9.4.5.3.2 MT_TP1:Idle state

9.4.5.3.2.1 State description

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

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 within the SMP Response time limit (see table 4 in 9.4.5.3.1).

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.