

T10/05-xxx revision 0

Date: March 04, 2005

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli); Mark Evans (Maxtor)

Subject: SAS 1.1 - Link and Port layer state machine fixes

1 Overview

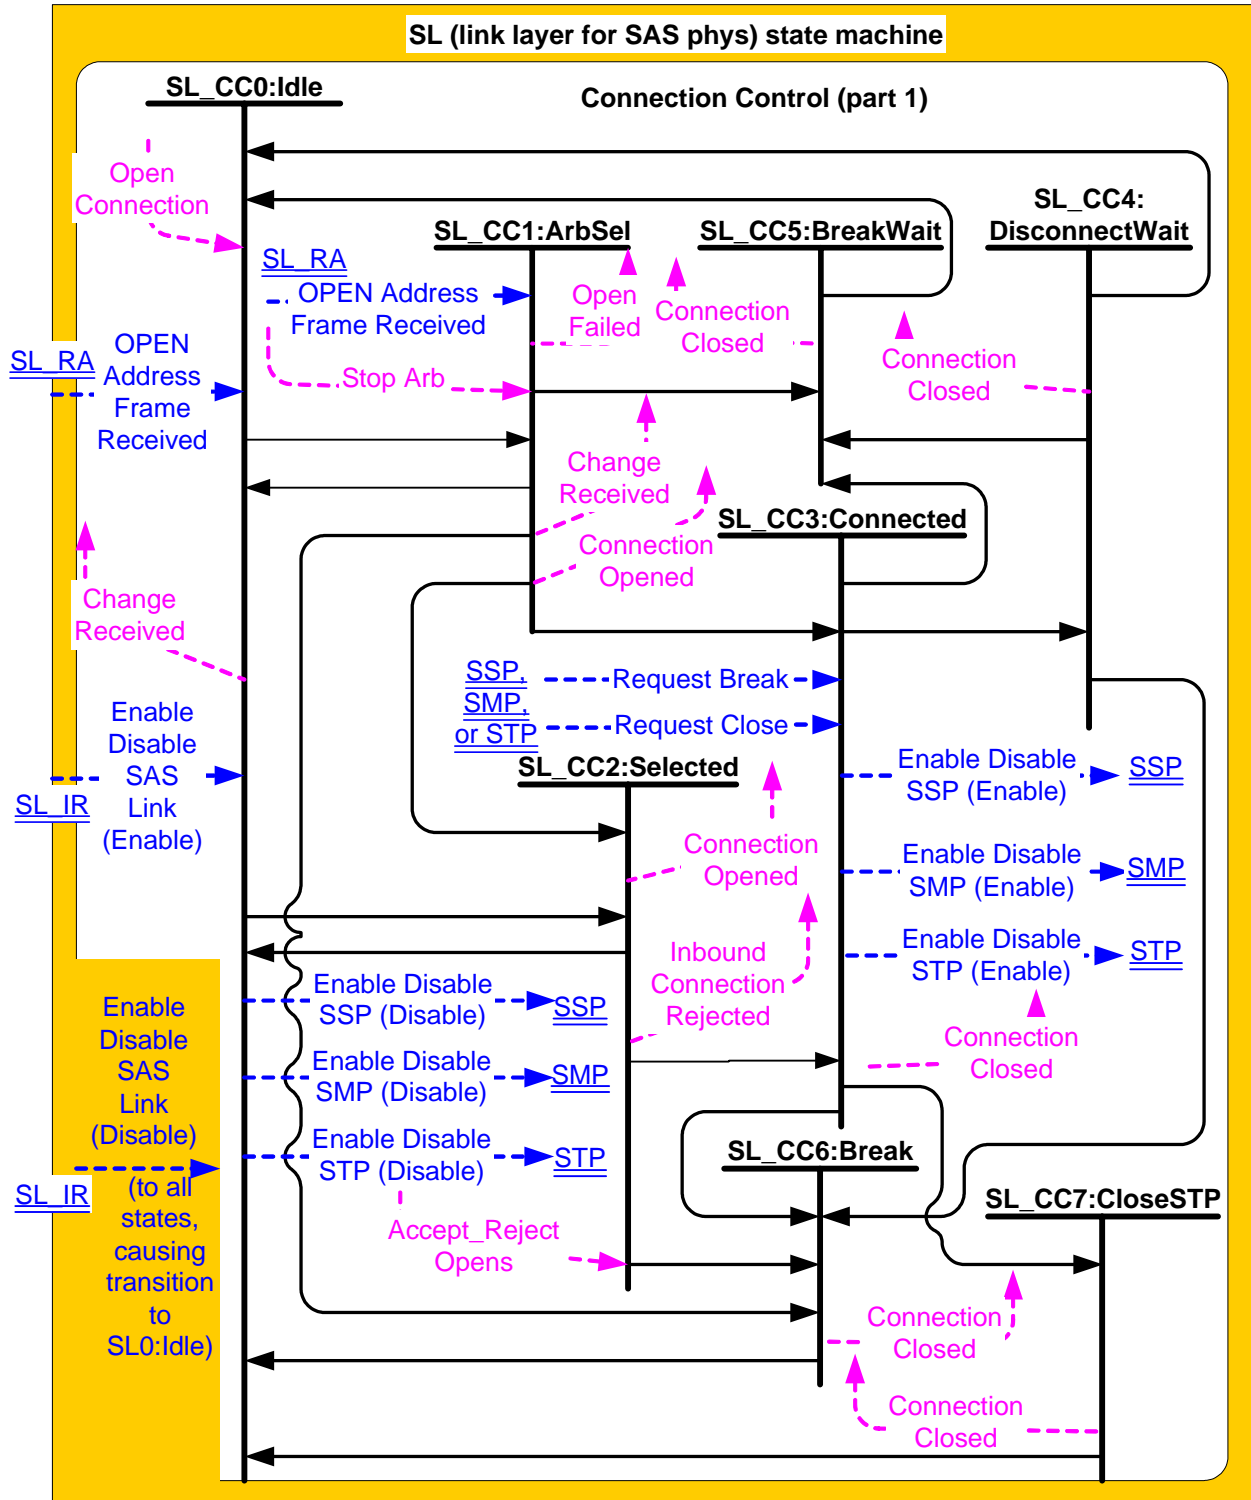
I received an email that pointed out that there are a few places in the SL state machine where it is possible to miss an Open Connection request resulting in a hang in the PL_PM2 state. While investigating this we found the following problems with the current SAS 1.1 draft.

- In the SL_CC5:BreakWait state there is a Connection Closed confirmation shown in the state machine figure but there is no description in the text as to when to send it or which Connection Closed it.
- In the SL_CC2:Selected state there is an Inbound Connection Reject confirmation that is not shown in the port layer state machines.
- There is a period of time when the PL_PM state could send an Open Connection request that would be ignored because the SL state machine is not in SL_CC0:Idle. The result of this would be that PL_PM would hang in the PL_PM2:Req_Wait state. This is most likely to occur as the result of sending or receiving a BREAK.

The solution to these problems is to add a new argument to the Connection Closed confirmation in the SL_CC6:Break and SL_CC5:BreakWait states that would indicate that a Transition to Idle occurred as a result of a BREAK primitive. This would cause the PL_PM states to transition to PL_PM1:Idle thereby removing the hang.

2 Changes to SAS 1.1 Link Layer

Change figure 126 to add in a Connection Closed confirmation exiting from the SL_CC6:Break state as shown below.



2.0.0.1 SL_CC5:BreakWait state

2.0.0.1.1 State description

This state closes the connection if one is established and releases all resources associated with the connection.

T10/05-xxx revision 0

This state shall:

- 1) send a Transmit BREAK message to the SL transmitter; and
- 2) initialize and start the Break Timeout timer.

2.0.0.1.2 Transition SL_CC5:BreakWait to SL_CC0:Idle

This transition shall occur after [sending a Connection Closed \(Break to Idle\) confirmation to the port layer if:](#)

- a) ~~receiving~~ a BREAK Received message [is received](#);
- b) ~~receiving~~ an OPEN_REJECT Received message [is received](#); or
- c) the Break Timeout timer expires.

2.0.0.2 SL_CC6:Break state

2.0.0.2.1 State description

This state closes any connection and releases all resources associated with this connection.

This state shall send a Transmit BREAK message to the SL transmitter.

2.0.0.2.2 Transition SL_CC6:Break to SL_CC0:Idle

This transition shall occur after sending a Transmit BREAK message to the SL transmitter [and sending a Connection Closed \(Break to Idle\) confirmation to the port layer.](#)

3 Changes to SAS 1.1 Port Layer

Add to figure 147 an Inbound Connection Rejection confirmation going to the PL_PM2:Req_Wait state. As shown in the figure below.

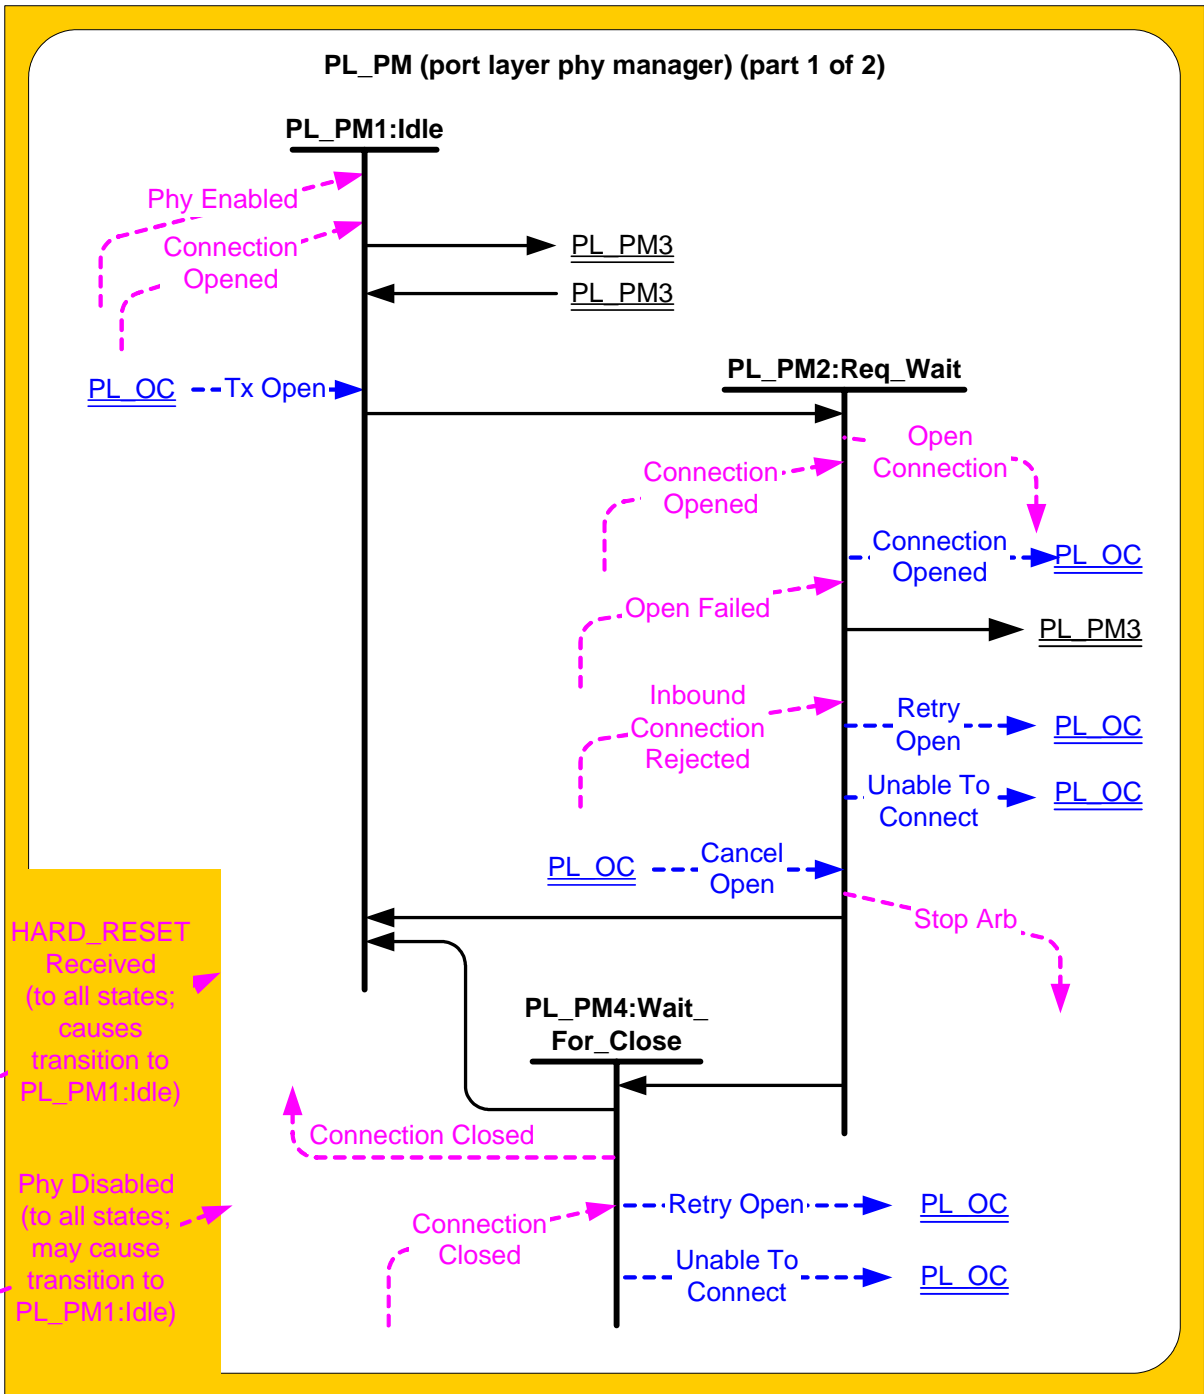


Figure 1 — PL_PM (port layer phy manager) state machine (part 1)

Figure 2 shows part 2 of the PL_PM state machine.

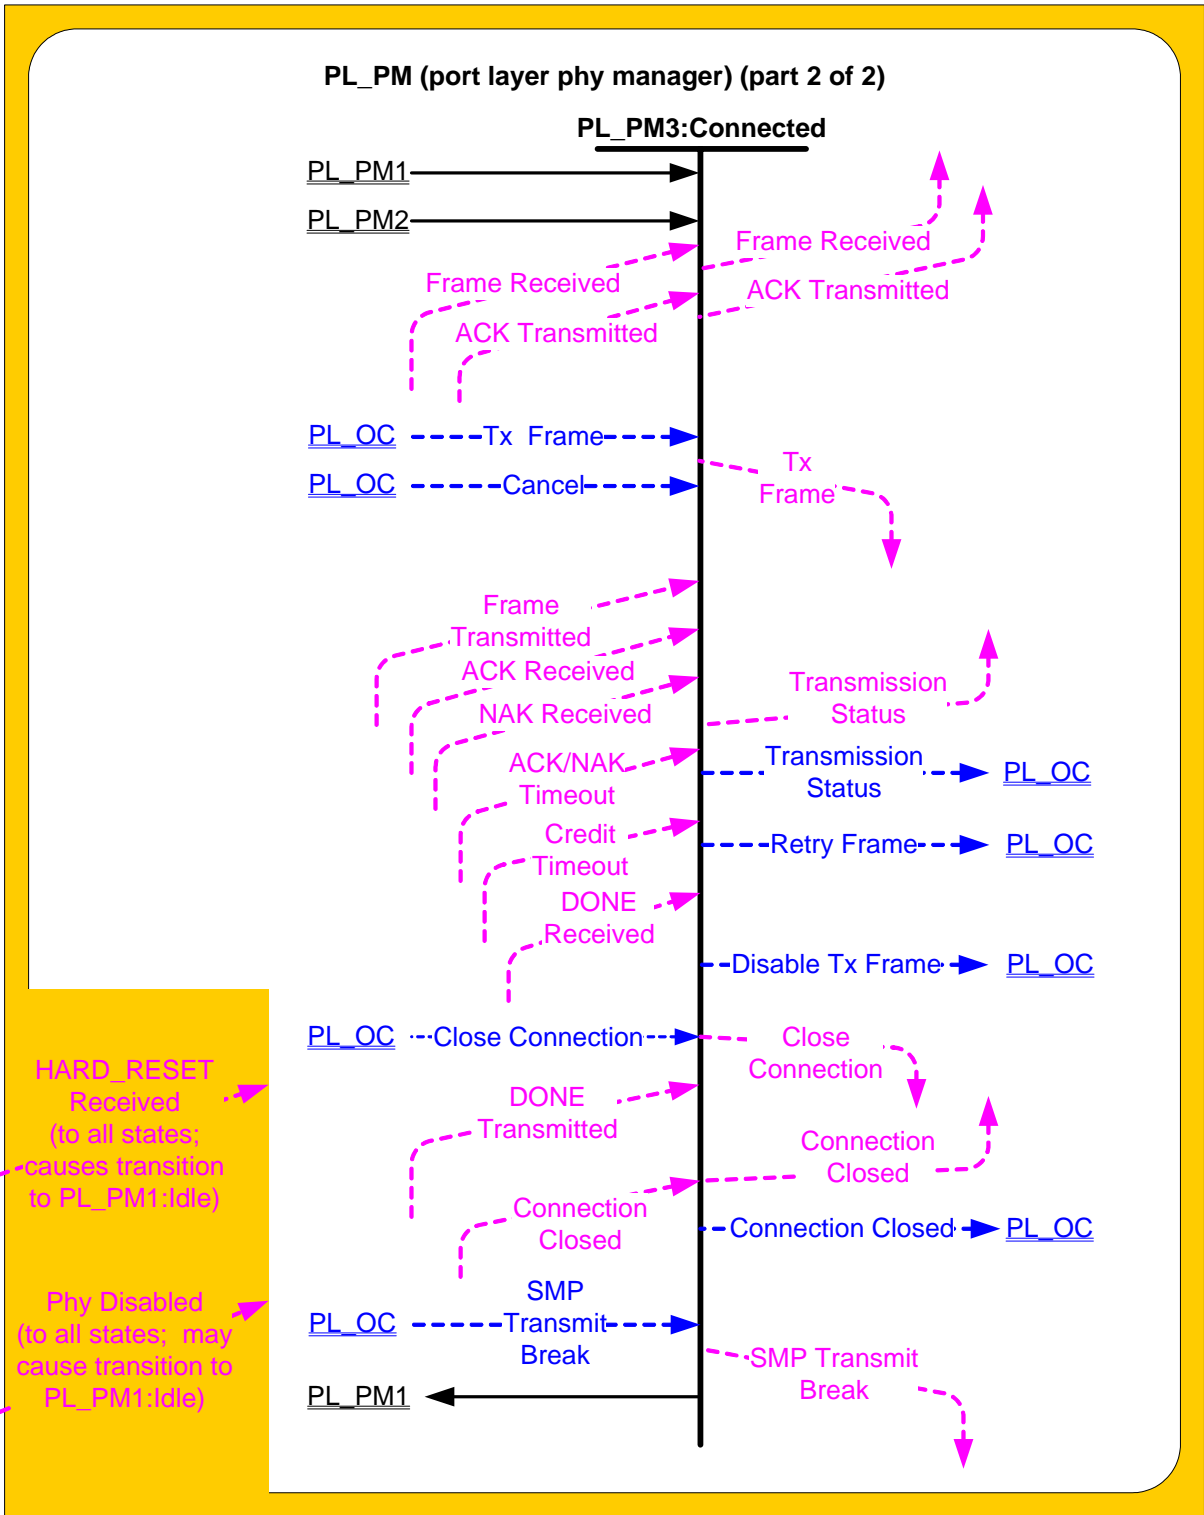


Figure 2 — PL_PM (port layer phy manager) state machine (part 2)

3.0.0.1 PL_PM1:Idle state

3.0.0.1.1 PL_PM1:Idle state description

This is the initial state of the PL_PM state machine.

3.0.0.1.2 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.

3.0.0.1.3 Transition PL_PM1:Idle to PL_PM3:Connected

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

3.0.0.2 PL_PM2:Req_Wait state

3.0.0.2.1 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.

3.0.0.2.2 PL_PM2:Req_Wait establishing a connection

After this state receives a Tx Open message, 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.

3.0.0.2.3 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.

~~If this state receives a Incoming 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 (Opened By Destination); or Retry Open (Opened By Other); ~~or Retry Open (Collided)~~ 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.

3.0.0.2.4 PL_PM2:Req_Wait unable to establish a connection

If this state receives ~~an~~ one of the Open Failed confirmations listed in table 1, then this state shall send either a Retry Open message or an Unable To Connect message to the PL_OC state machine.

Table 1 defines the message to be sent for each Open Failed confirmation.

Table 1 — 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 (Break Received)	Unable To Connect (Break Received)
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.

3.0.0.2.5 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.

3.0.0.2.6 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.

3.0.0.2.7 Transition PL_PM2:Req_Wait to PL_PM3:Connected

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

3.0.0.2.8 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).

3.0.0.3 PL_PM3:Connected state**3.0.0.3.1 PL_PM3:Connected state description**

If the protocol for the connection is SSP, and this state is in an SSP target port, and the value in the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.6.1) is not zero, 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; 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 value in the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.6.1) is zero, then this state shall not create a Maximum Connect Time Limit timer (i.e., there is no maximum connect time limit).

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 value in the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.6.1) is not zero, 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; and
- c) not start the Bus Inactivity Time Limit timer.

If the protocol for the connection is SSP, and this state is in an SSP target port, and the value in the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.6.1) is zero, then this state shall not create a Bus Inactivity Time Limit timer (i.e., there is no maximum bus inactivity time limit).

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.

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.

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, then this state shall send a Close Connection request to the link layer.

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.

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.

This state shall wait to receive a Connection Closed (Break to Idle) confirmation after receiving one of the following:

- a) a Connection Closed (Close Timeout) confirmation;
- b) a Connection Closed (Break Requested) confirmation; or
- c) a Connection Closed (Break Received) confirmation.

If this state receives a Connection Closed (Break to Idle) confirmation after receiving a Connection Closed (Close Timeout) confirmation or a Connection Closed (Break Requested) confirmation, then this state shall send a Connection Closed message to the PL_OC state machine.

If this state receives a Connection Closed (Break to Idle) confirmation after receiving a Connection Closed (Break Received) confirmation or 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 (Break 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 (Break to Idle) confirmation, or a Phy Disabled confirmation after sending a Transmit 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 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.

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

If this state receives a Connection Closed ([Normal](#)) confirmation, then this state shall send a Connection Closed message to the PL_OC state machine.

[If this state receives a Connection Closed \(Break to Idle\) confirmation, then this state shall send a Connection Closed 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 HARD_RESET Received confirmation, then this state machine shall terminate all operations.

3.0.0.3.2 Transition PL_PM3:Connected to PL_PM1:Idle

This transition shall occur after:

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

3.0.0.4 PL_PM4:Wait_For_Close state

3.0.0.4.1 PL_PM4:Wait_For_Close state description

After this state receives a Connection Closed ([Break to Idle](#)) confirmation, and 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 during an SMP connection, this state shall send a Connection Closed confirmation to the transport layer.

After this state receives a Connection Closed ([Break to Idle](#)) confirmation, and 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 this state receives a Connection Closed (Break to Idle) confirmation, and 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 a 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.

3.0.0.4.2 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.