

To: T10 SAS Protocol Working Group

From: Brian Day

Subject: SAS 2.+ : 08-249 Link Layer Power Management

Revision History

Revision 0 - Initial draft

Related Documents

sas2r14 - Serial Attached SCSI - 2 Draft revision 14

08-015 - SAS: Add low power transceiver options (Gerry Houlder, Seagate Technology)

Overview

This proposal is a complement to 08-015. This proposal focusses on link layer changes to support lower power modes.

Proposed Changes

7.8.2 IDENTIFY address frame

Table 67 defines the IDENTIFY address frame format used for the identification sequence. The IDENTIFY address frame is sent by each logical phy after the phy reset sequence completes if the physical link is a SAS physical link. The IDENTIFY address frame sent by each logical phy in a physical phy shall be identical.

Table 67 — IDENTIFY address frame format

Byte\Bit	7	6	5	4	3	2	1	0	
0	Reserved	DEVICE TYPE			ADDRESS FRAME TYPE (0h)				
1	Reserved			REASON					
2	Reserved			SSP INITIATOR PORT	STP INITIATOR PORT	SMP INITIATOR PORT	Restricted (for OPEN address frame)		
3	Reserved			SSP TARGET PORT	STP TARGET PORT	SMP TARGET PORT	Restricted (for OPEN address frame)		
4	DEVICE NAME								
11	SAS ADDRESS								
12	SAS ADDRESS								
19	SAS ADDRESS								
20	PHY IDENTIFIER								
21	Reserved			SLUMBER CAPABLE	PARTIAL CAPABLE	INSIDE ZPSDS PERSISTENT	REQUESTED INSIDE ZPSDS	BREAK_REPLY CAPABLE	
22	Reserved								
27	Reserved								
28	(MSB)	CRC							
31								(LSB)	

The DEVICE TYPE field indicates the type of device containing the phy, and is defined in table 68.

Table 68 — DEVICE TYPE field

Code	Description
001b	End device
010b	Expander device
011b	Expander device compliant with a previous version of this standard
All others	Reserved

The ADDRESS FRAME TYPE field shall be set to 0h.

The REASON field indicates the reason for the link reset sequence and is defined in table 69.

Table 69 — REASON field

Code	Description
0h	Unknown reason
1h	Power on
2h	Hard reset (e.g., phy received a HARD_RESET primitive during the hard reset sequence)(see 4.4.2), or SMP PHY CONTROL function HARD RESET phy operation (see 10.4.3.28)
3h	SMP PHY CONTROL function LINK RESET phy operation, or TRANSMIT SATA PORT SELECTION SIGNAL phy operation (see 10.4.3.28)
4h	Loss of dword synchronization (see 6.9)
5h	After the multiplexing sequence completes, MUX (LOGICAL LINK 0) received in logical link 1 or MUX (LOGICAL LINK 1) received in logical link 0 (see 6.10).
6h	I_T nexus loss timer expired in the STP target port of an STP/SATA bridge when the phy was attached to a SATA device (see 4.5)
7h	Break Timeout Timer expired (see 7.12.8)
8h	Phy test function stopped (see 10.4.3.29)
9h	Expander device reduced functionality (see 4.6.8)
Ah - Fh	Reserved

An SSP INITIATOR PORT bit set to one indicates that an SSP initiator port is present. An SSP INITIATOR PORT bit set to zero indicates that an SSP initiator port is not present. Expander devices shall set the SSP INITIATOR PORT bit to zero.

An STP INITIATOR PORT bit set to one indicates that an STP initiator port is present. An STP INITIATOR PORT bit set to zero indicates that an STP initiator port is not present. Expander devices shall set the STP INITIATOR PORT bit to zero.

An SMP INITIATOR PORT bit set to one indicates that an SMP initiator port is present. An SMP INITIATOR PORT bit set to zero indicates that an SMP initiator port is not present. Expander devices may set the SMP INITIATOR PORT bit to one.

An SSP TARGET PORT bit set to one indicates that an SSP target port is present. An SSP TARGET PORT bit set to zero indicates that an SSP target port is not present. Expander devices shall set the SSP TARGET PORT bit to zero.

An STP TARGET PORT bit set to one indicates that an STP target port is present. An STP TARGET PORT bit set to zero indicates that an STP target port is not present. Expander devices shall set the STP TARGET PORT bit to zero.

An SMP TARGET PORT bit set to one indicates that an SMP target port is present. An SMP TARGET PORT bit set to zero indicates that an SMP target port is not present. Expander devices shall set the SMP TARGET PORT bit to one.

The DEVICE NAME field indicates the device name (see 4.2.4) of the SAS device or expander device transmitting the IDENTIFY address frame. A DEVICE NAME field set to 00000000 00000000h indicates the device name is not provided in this field.

NOTE 22 - In expander devices, the DEVICE NAME field, if not set to 00000000 00000000h, contains the same value as the SAS ADDRESS field.

For SAS ports, the SAS ADDRESS field indicates the port identifier (see 4.2.7) of the SAS port transmitting the IDENTIFY address frame. For expander ports, the SAS ADDRESS field indicates the device name (see 4.2.4) of the expander device transmitting the IDENTIFY address frame.

The PHY IDENTIFIER field indicates the phy identifier of the phy transmitting the IDENTIFY address frame.

The REQUESTED INSIDE ZPSDS bit indicates the value of the REQUESTED INSIDE ZPSDS bit in the zone phy information (see 4.9.3.1) at the time the IDENTIFY address frame is transmitted. If the phy transmitting the IDENTIFY address frame is contained in an end device, a non-zoning expander device, or a zoning expander device with zoning disabled, then the REQUESTED INSIDE ZPSDS bit shall be set to zero.

The INSIDE ZPSDS PERSISTENT bit indicates the value of the INSIDE ZPSDS PERSISTENT bit in the zone phy information (see 4.9.3.1) at the time the IDENTIFY address frame is transmitted. If the phy transmitting the IDENTIFY address frame is contained in an end device, a non-zoning expander device, or a zoning expander device with zoning disabled, then the INSIDE ZPSDS PERSISTENT bit shall be set to zero.

The BREAK_REPLY CAPABLE field indicates that the phy is capable of responding to received BREAK primitive sequences with a BREAK_REPLY primitive sequence (see 7.12.5).

[The PARTIAL CAPABLE and SLUMBER CAPABLE bits indicate that the phy is capable of the corresponding power management mode \(see xxx\).](#)

See 4.1.4 for additional requirements concerning the DEVICE TYPE field, SSP INITIATOR PORT bit, STP INITIATOR PORT bit, SMP INITIATOR PORT bit, SSP TARGET PORT bit, STP TARGET PORT bit, SMP TARGET PORT bit, and SAS ADDRESS field.

The CRC field is defined in 7.8.1.

7.14 SL (link layer for SAS logical phys) state machines

7.14.1 SL state machines overview

The SL (link layer for SAS logical phys) state machines controls connections, handling both connection requests (OPEN address frames), CLOSEs, and BREAKs. The SL state machines are as follows:

- a) SL_RA (receive OPEN address frame) state machine (see 7.14.3); and
- b) SL_CC (connection control) state machine (see 7.14.4).

All the SL state machines shall begin after receiving an Enable Disable SAS Link (Enable) message from the SL_IR state machines.

If a state machine consists of multiple states the initial state is as indicated in the state machine description.

[Editor's Note 0: Need to update figures with new states, and with Request Power Management request coming down, and Phy Power confirmation coming up.](#)

Figure 118 shows the SL state machines.

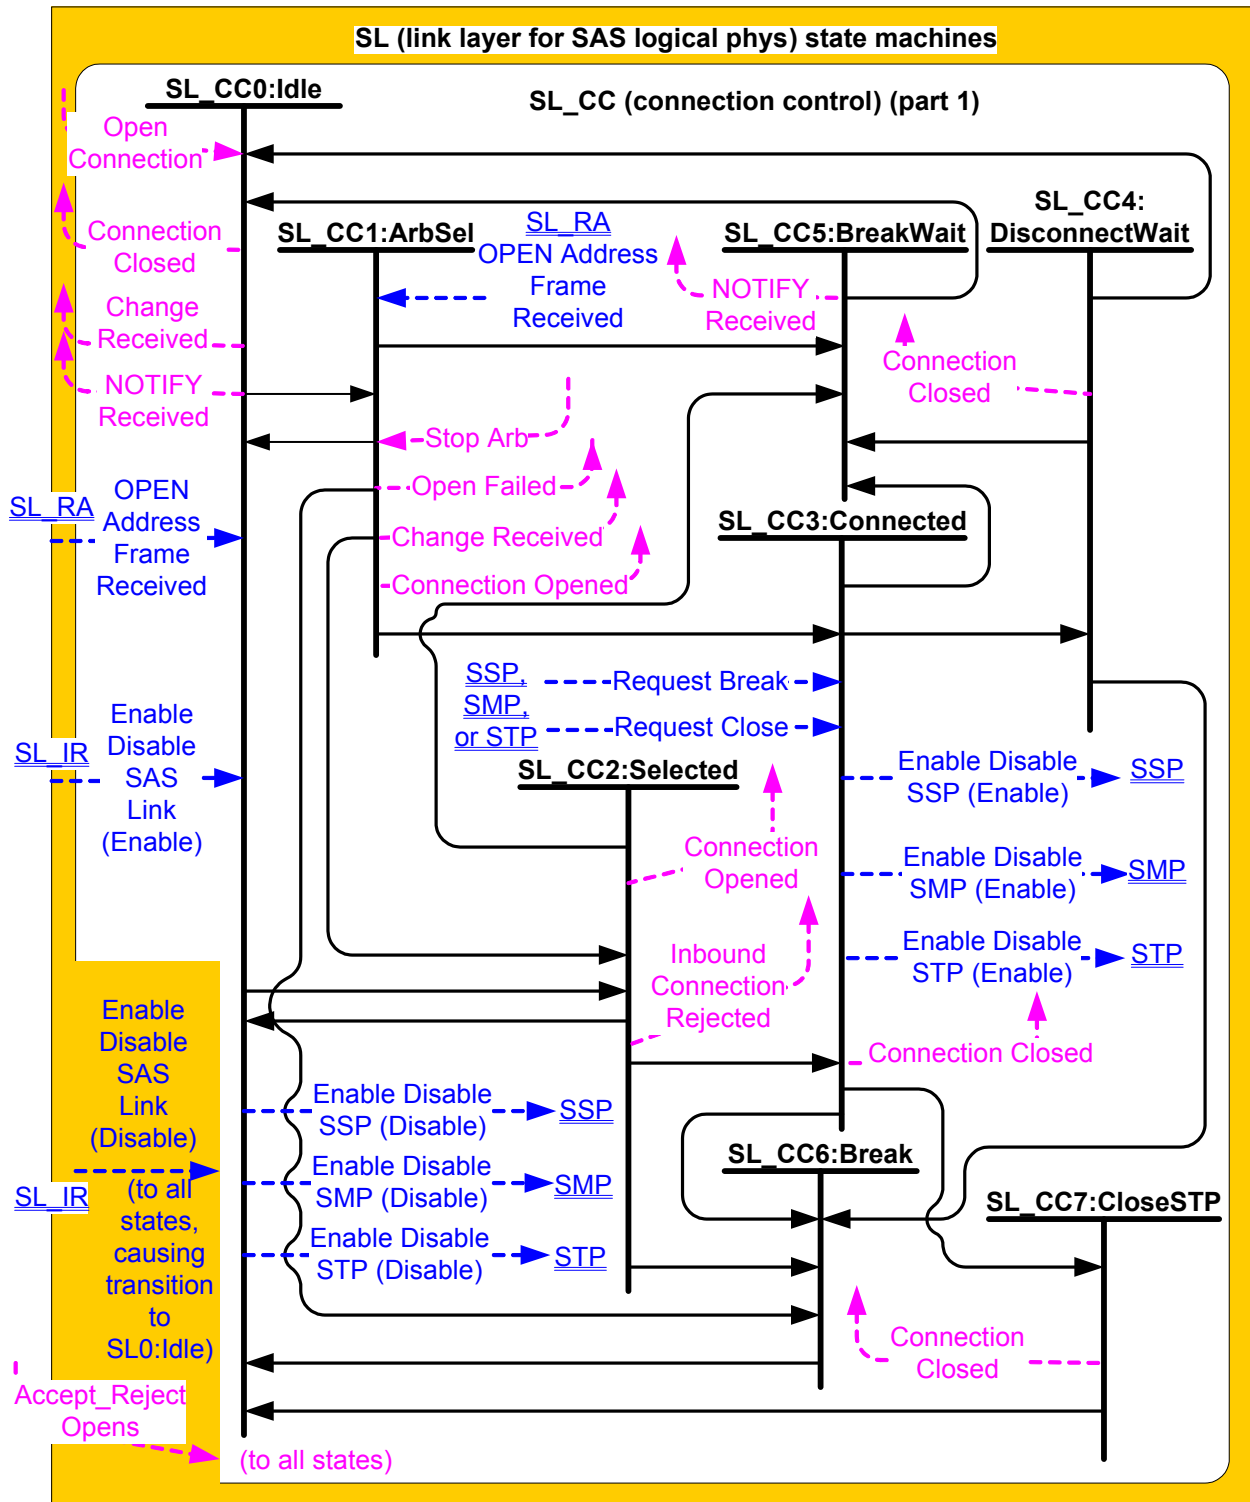


Figure 118 — SL (link layer for SAS logical phys) state machines (part 1)

Figure 119 shows the messages sent to the SL transmitter and received from the SL receiver.

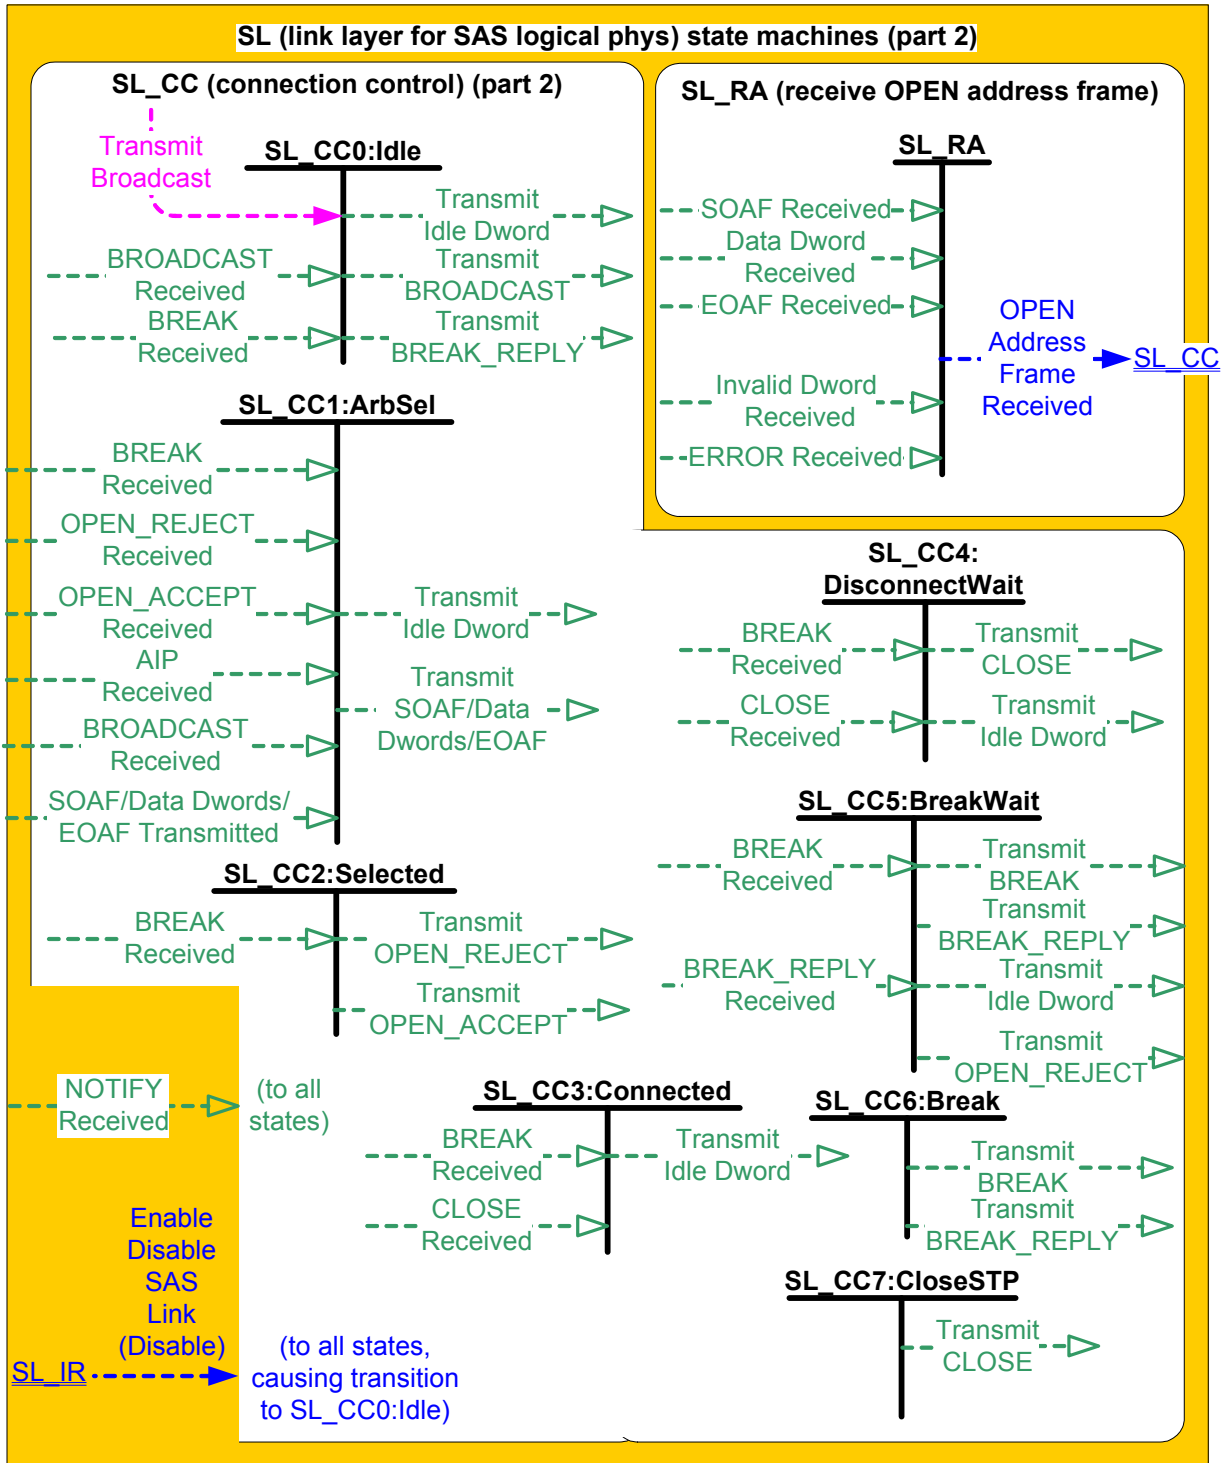


Figure 119 — SL (link layer for SAS logical phys) state machines (part 2)

7.14.2 SL transmitter and receiver

The SL transmitter receives the following messages from the SL state machines specifying primitive sequences, frames, and dwords to transmit:

- a) Transmit Idle Dword;
- b) Transmit SOAF/Data Dwords/EOAF;

- c) Transmit OPEN_ACCEPT;
- d) Transmit OPEN_REJECT with an argument indicating the specific type (e.g., Transmit OPEN_REJECT (Retry));
- e) Transmit BREAK;
- f) Transmit BREAK_REPLY;
- g) [Transmit PMREQ with an argument indicating the specific type \(e.g., Transmit PMREQ \(Partial\)\);](#)
- h) [Transmit PMACK;](#)
- i) [Transmit PMNAK;](#)
- j) Transmit BROADCAST; and
- k) Transmit CLOSE with an argument indicating the specific type (e.g., Transmit CLOSE (Normal)).

When the SL transmitter is requested to transmit a dword from any state within any of the SL state machines, it shall transmit that dword. If there are multiple requests to transmit, the following priority should be followed when selecting the dword to transmit:

- 1) BREAK_REPLY;
- 2) BREAK;
- 3) CLOSE;
- 4) OPEN_ACCEPT or OPEN_REJECT;
- 5) SOAF or data dword or EOAF; then
- 6) idle dword.

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

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

- a) SOAF/Data Dwords/EOAF Transmitted.

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

- a) SOAF Received;
- b) Data Dword Received;
- c) EOAF Received;
- d) BROADCAST Received with an argument indicating the specific type (e.g., BROADCAST Received (Change));
- e) BREAK Received;
- f) BREAK_REPLY Received;
- g) OPEN_ACCEPT Received;
- h) OPEN_REJECT Received with an argument indicating the specific type (e.g., OPEN_REJECT Received (No Destination));
- i) AIP Received;
- j) CLOSE Received with an argument indicating the specific type (e.g., CLOSE Received (Normal));
- k) NOTIFY Received (Power Loss Expected);
- l) [PMREQ Received with an argument indicated the specific type \(e.g., PMREQ Received \(Partial\)\);](#)
- m) ERROR Received; and
- n) Invalid Dword Received.

The SL receiver shall ignore all other dwords.

The SL transmitter relationship to other transmitters is defined in 4.3.2. The SL receiver relationship to other receivers is defined in 4.3.3.

7.14.3 SL_RA (receive OPEN address frame) state machine

The SL_RA state machine's function is to receive address frames and determine if the received address frame is an OPEN address frame and whether or not it was received successfully. This state machine consists of one state.

This state machine receives SOAFs, dwords of an OPEN address frames, and EOAFs.

This state machine shall ignore all messages except SOAF Received, Data Dword Received, and EOAF Received.

If this state machine receives a subsequent SOAF Received message after receiving an SOAF Received message but before receiving an EOAF Received message, then this state machine shall discard the Data Dword Received messages received before the subsequent SOAF Received message.

If this state machine receives more than eight Data Dword Received messages after an SOAF Received message and before an EOAF Received message, then this state machine shall discard the address frame.

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

- a) ignore the invalid dword or ERROR; or
- b) discard the address frame.

After receiving an EOAF Received message, this state machine shall check if the address frame is a valid OPEN address frame.

This state machine shall accept an address frame if:

- a) the ADDRESS FRAME TYPE field is set to 1h (i.e., OPEN);
- b) the number of data dwords between the SOAF and EOAF is 8; and
- c) the CRC field contains a good CRC.

Otherwise, this state machine shall discard the address frame. If the frame is not discarded then this state machine shall send a OPEN Address Frame Received message to the SL_CC0:Idle state and the SL_CC1:ArbSel state with an argument that contains all the data dwords received in the OPEN address frame.

7.14.4 SL_CC (connection control) state machine

7.14.4.1 SL_CC state machine overview

The state machine consists of the following states:

- a) SL_CC0:Idle (see 7.14.4.2)(initial state);
- b) SL_CC1:ArbSel (see 7.14.4.3);
- c) SL_CC2:Selected (see 7.14.4.4);
- d) SL_CC3:Connected (see 7.14.4.5);
- e) SL_CC4:DisconnectWait (see 7.14.4.6);
- f) SL_CC5:BreakWait (see 7.14.4.7);
- g) SL_CC6:Break (see 7.14.4.8); ~~and~~
- h) SL_CC7:CloseSTP (see 7.14.4.9);
- i) [SL_CC8:PM Request \(see xxx\)](#); and
- j) [SL_CC9:PM Quiet \(see xxx\)](#).

The state machine shall start in the SL_CC0:Idle state. The state machine shall transition to the SL_CC0:Idle state from any other state after receiving an Enable Disable SAS Link (Disable) message from the SL_IR state machines (see 7.9.4).

The SL_CC state machine receives the following messages from the SSP link layer state machine (see 7.16.8), the STP link layer state machine, and SMP link layer state machine (see 7.18.5):

- a) Request Break; and
- b) Request Close.

The SL_CC state machine sends the following messages to the SSP link layer state machine, the STP link layer state machine, and SMP link layer state machine:

- a) Enable Disable SSP (Enable);
- b) Enable Disable SSP (Disable);
- c) Enable Disable STP (Enable);
- d) Enable Disable STP (Disable);

- e) Enable Disable SMP (Enable); and
- f) Enable Disable SMP (Disable).

The SL_CC state machine receives the following messages from the SL_IR state machines (see 7.9.4):

- a) Enable Disable SAS Link (Enable); and
- b) Enable Disable SAS Link (Disable).

Any message received by a state that is not referred to in the description of that state or in this subclause shall be ignored.

If this state machine receives an Accept_Reject Opens (Accept SSP) request, then this state machine shall set the Reject SSP Opens state machine variable to NO. If this state machine receives an Accept_Reject Opens (Reject SSP) request, then this state machine shall set the Reject SSP Opens state machine variable to YES.

If this state machine receives an Accept_Reject Opens (Accept SMP) request, then this state machine shall set the Reject SMP Opens state machine variable to NO. If this state machine receives an Accept_Reject Opens (Reject SMP) request, then this state machine shall set the Reject SMP Opens state machine variable to YES.

If this state machine receives an Accept_Reject Opens (Accept STP) request, then this state machine shall set the Reject STP Opens state machine variable to NO. If this state machine receives an Accept_Reject Opens (Reject STP) request, then this state machine shall set the Reject STP Opens state machine variable to YES.

Any detection of an internal error shall cause the SL_CC state machine to transition to the SL_CC5:BreakWait state.

The SL_CC state machine shall maintain the timers listed in table 70.

Table 70 — SL_CC timers

Timer	Initial value
Open Timeout timer	1 ms
Close Timeout timer	1 ms
Break Timeout timer	1 ms
Power Management Request Timeout timer	1 ms

The SL_CC state machine shall maintain the state machine variables listed in table 71.

Table 71 — SL_CC state machine variables

State machine variable	Description
Reject SSP Opens	Used to determine if the SCSI application layer is permitting SSP connection requests to be accepted on this phy.
Reject SMP Opens	Used to determine if the management application layer is permitting SMP connection requests to be accepted on this phy.
Reject STP Opens	Used to determine if the ATA application layer is permitting STP connection requests to be accepted on this phy.

7.14.4.2 SL_CC0:Idle state

7.14.4.2.1 State description

This state is the initial state and is the state that is used when there is no connection pending or established.

Upon entry into this state, this state shall send:

- a) an Enable Disable SSP (Disable) message to the SSP link layer state machines;
- b) an Enable Disable SMP (Disable) message to the SMP link layer state machines;
- c) an Enable Disable STP (Disable) message to the STP link layer state machines; and
- d) a Connection Closed (Transition to Idle) confirmation to the port layer.

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL transmitter (see 7.4).

If a BROADCAST Received (Change) message, BROADCAST Received (Reserved Change 0) message, or BROADCAST Received (Reserved Change 1) message is received, this state shall send a Change Received confirmation to the management layer.

If a Transmit Broadcast request is received with any argument, this state shall send a Transmit BROADCAST message with the same argument to the SL transmitter.

If a BREAK Received message is received and the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), then this state shall send a Transmit BREAK_REPLY message to the SL transmitter.

After this state receives an Enable Disable SAS Link (Enable) confirmation, this state shall:

- a) set the Reject SSP Opens state machine variable to a vendor-specific default value (i.e., YES or NO);
- b) set the Reject SMP Opens state machine variable to a vendor-specific default value (i.e., YES or NO); and
- c) set the Reject STP Opens state machine variable to a vendor-specific default value (i.e., YES or NO).

If this state receives a NOTIFY Received (Power Loss Expected) message and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this state shall send a Notify Received (Power Loss Expected) confirmation to the port layer.

If a PMREQ Received message is received and power management is enabled (see xxx), then this state shall send a Transmit PMACK message to the SL transmitter. If a PMREQ Received message is received and power management is not enabled, then this state shall send a Transmit PMNAK message to the SL transmitter.

7.14.4.2.2 Transition SL_CC0:Idle to SL_CC1:ArbSel

This transition shall occur after receiving both an Enable Disable SAS Link (Enable) confirmation and an Open Connection request. The Open Connection request includes these arguments:

- 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.

7.14.4.2.3 Transition SL_CC0:Idle to SL_CC2:Selected

This transition shall occur after receiving both an Enable Disable SAS Link (Enable) confirmation and an OPEN Address Frame Received message.

7.14.4.2.4 Transition SL_CC0:Idle to SL_CC8:PM_Request

If power management is enabled (see xxx), this transition shall occur after receiving both an Enable Disable SAS Link (Enable) confirmation and a Request Power Management request. The transition shall include an argument with the same type as the request (e.g. Request Power Management (Partial)).

7.14.4.2.5 Transition SL_CC0:Idle to SL_CC9:PM_Quiet

This transition shall occur after sending a Transmit PMACK message to the SL transmitter.

7.14.4.3 SL_CC1:ArbSel state

7.14.4.3.1 State description

This state is used to make a connection request.

Upon entry into this state, this state shall:

- 1) request an OPEN address frame be transmitted by sending a Transmit SOAF/Data Dwords/EOAF message to the SL transmitter with the dwords containing the OPEN address frame with its fields set to the arguments received with the Open Connection request;
- 2) initialize and start the Open Timeout timer; and
- 3) request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL transmitter.

This state shall ignore OPEN_REJECT Received and OPEN_ACCEPT Received messages from the time a Transmit SOAF/Data Dwords/EOAF message is sent to the SL transmitter until an SOAF/Data Dwords/EOAF Transmitted message is received from the SL transmitter.

If a BROADCAST Received (Change) message, BROADCAST Received (Reserved Change 0) message, or BROADCAST Received (Reserved Change 1) message is received this state shall send a Change Received confirmation to the management layer.

If a PMREQ Received message is received this state shall send a Transmit PMNAK to the SL transmitter.

Editor's Note 1: A different option is just have have the OPEN address frame act as a PMNAK as well, without transmitting a a PMNAK at all. But I prefer sending PMNAK, as this is more tolerant of bit errors in case the SOAF / etc is corrupted.

If an AIP Received message is received after requesting the OPEN address frame be transmitted, this state shall reinitialize and restart the Open Timeout timer. The state machine shall not enforce a limit on the number of AIPs received.

If this state receives an OPEN_REJECT Received message listed in table 72 after transmitting the OPEN address frame, this state shall send the corresponding Open Failed confirmation listed in table 72 to the port layer.

Table 72 — OPEN_REJECT Received message to Open Failed confirmation mapping

OPEN_REJECT Received message	Open Failed confirmation
OPEN_REJECT Received (Bad Destination)	Open Failed (Bad Destination)
OPEN_REJECT Received (Connection Rate Not Supported)	Open Failed (Connection Rate Not Supported)
OPEN_REJECT Received (Protocol Not Supported)	Open Failed (Protocol Not Supported)
OPEN_REJECT Received (Reserved Abandon 1)	Open Failed (Reserved Abandon 1)
OPEN_REJECT Received (Reserved Abandon 2)	Open Failed (Reserved Abandon 2)
OPEN_REJECT Received (Reserved Abandon 3)	Open Failed (Reserved Abandon 3)
OPEN_REJECT Received (STP Resources Busy)	Open Failed (STP Resources Busy)
OPEN_REJECT Received (Wrong Destination)	Open Failed (Wrong Destination)
OPEN_REJECT Received (Zone Violation)	Open Failed (Zone Violation)
OPEN_REJECT Received (No Destination)	Open Failed (No Destination)
OPEN_REJECT Received (Pathway Blocked)	Open Failed (Pathway Blocked)
OPEN_REJECT Received (Reserved Continue 0)	Open Failed (Reserved Continue 0)
OPEN_REJECT Received (Reserved Continue 1)	Open Failed (Reserved Continue 1)
OPEN_REJECT Received (Reserved Initialize 0)	Open Failed (Reserved Initialize 0)
OPEN_REJECT Received (Reserved Initialize 1)	Open Failed (Reserved Initialize 1)
OPEN_REJECT Received (Reserved Stop 0)	Open Failed (Reserved Stop 0)
OPEN_REJECT Received (Reserved Stop 1)	Open Failed (Reserved Stop 1)
OPEN_REJECT Received (Retry)	Open Failed (Retry)

7.14.4.3.2 Transition SL_CC1:ArbSel to SL_CC0:Idle

This transition shall occur after sending an Open Failed confirmation.

7.14.4.3.3 Transition SL_CC1:ArbSel to SL_CC2:Selected

This transition shall occur after receiving a SOAF/Data Dwords/EOAF Transmitted message if:

- a) one or more AIP Received messages have been received before an OPEN Address Frame Received message is received (i.e., the incoming OPEN address frame overrides the outgoing OPEN address frame); or
- b) no AIP Received messages have been received before an OPEN Address Frame Received message is received, and the arbitration fairness rules (see 7.12.3) indicate the received OPEN address frame overrides the outgoing OPEN address frame.

The arbitration fairness comparison shall compare:

- a) the value of the arbitration wait time argument in the Open Connection request for the outgoing OPEN address frame; and
- b) the value of the ARBITRATION WAIT TIME field received in the incoming OPEN address frame.

7.14.4.3.4 Transition SL_CC1:ArbSel to SL_CC3:Connected

This transition shall occur if this state receives an SOAF/Data Dwords/EOAF Transmitted message followed by an OPEN_ACCEPT Received message.

If the PROTOCOL field in the transmitted OPEN address frame was set to STP, then this state shall send a Connection Opened (STP, Source Opened) confirmation to the port layer before the transition. This transition shall include an Open STP Connection argument. At this point an STP connection has been opened between the source phy and the destination phy.

If the PROTOCOL field in the transmitted OPEN address frame was set to SSP, then this state shall send a Connection Opened (SSP, Source Opened) confirmation to the port layer before the transition. This transition shall include an Open SSP Connection argument. At this point an SSP connection has been opened between the source phy and the destination phy.

If the PROTOCOL field in the transmitted OPEN address frame was set to SMP, then this state shall send a Connection Opened (SMP, Source Opened) confirmation to the port layer before the transition. This transition shall include an Open SMP Connection argument. At this point an SMP connection has been opened between the source phy and the destination phy.

7.14.4.3.5 Transition SL_CC1:ArbSel to SL_CC5:BreakWait

This transition shall occur after receiving a SOAF/Data Dwords/EOAF Transmitted message if a BREAK Received message has not been received and after:

- a) a Stop Arb request is received and after sending an Open Failed (Port Layer Request) confirmation to the port layer;
- b) there is no response to the OPEN address frame before the Open Timeout timer expires and after sending an Open Failed (Open Timeout Occurred) confirmation to the port layer; or
- c) a NOTIFY Received (Power Loss Expected) message is received.

If a NOTIFY Received (Power Loss Expected) message was received and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this transition shall include a Power Loss Expected argument.

7.14.4.3.6 Transition SL_CC1:ArbSel to SL_CC6:Break

This transition shall occur after:

- a) receiving a SOAF/Data Dwords/EOAF Transmitted message;
- b) receiving a BREAK Received message; and
- c) sending an Open Failed (Break Received) confirmation to the port layer.

7.14.4.4 SL_CC2:Selected state

7.14.4.4.1 State description

This state completes the establishment of an SSP, SMP, or STP connection when an incoming connection request has won arbitration by sending a Transmit OPEN_ACCEPT message, or rejects opening a connection by sending a Transmit OPEN_REJECT message to the SL transmitter.

This state shall respond to an incoming OPEN address frame using the following rules:

- 1) If the OPEN address frame DESTINATION SAS ADDRESS field does not match the SAS address of this port, this state shall send a Transmit OPEN_REJECT (Wrong Destination) message to the SL transmitter (see 7.14.4.4.2);
- 2) If the OPEN address frame INITIATOR PORT bit, PROTOCOL field, FEATURES field, and/or INITIATOR CONNECTION TAG field are set to values that are not supported (e.g., a connection request from an SMP target port), this state shall send a Transmit OPEN_REJECT (Protocol Not Supported) message to the SL transmitter (see 7.14.4.4.2);
- 3) If the OPEN address frame CONNECTION RATE field is set to a connection rate that is not supported, this state shall send a Transmit OPEN_REJECT (Connection Rate Not Supported) message to the SL transmitter (see 7.14.4.4.2);
- 4) If the OPEN address frame PROTOCOL field is set to STP, the STP target port supports affiliations, and the source SAS address is not that of an STP initiator port with an affiliation established or the source SAS address is not that of an STP initiator port with ATA task file register set resources (see 7.17.4),

- this state shall send a Transmit OPEN_REJECT (STP Resources Busy) message to the SL transmitter (see 7.14.4.4.2);
- 5) If the OPEN address frame PROTOCOL field is set to SSP and the Reject SSP Opens state machine variable is set to YES, this state shall send a Transmit OPEN_REJECT (Retry) message to the SL transmitter (see 7.14.4.4.2);
 - 6) If the OPEN address frame PROTOCOL field is set to SMP and the Reject SMP Opens state machine variable is set to YES, this state shall send a Transmit OPEN_REJECT (Retry) message to the SL transmitter (see 7.14.4.4.2);
 - 7) If the OPEN address frame PROTOCOL field is set to STP and the Reject STP Opens state machine variable is set to YES, this state shall send a Transmit OPEN_REJECT (Retry) message to the SL transmitter (see 7.14.4.4.2);
 - 8) If the OPEN address frame PROTOCOL field is set to SSP and the Reject SSP Opens state machine variable is set to NO, this state shall send a Transmit OPEN_ACCEPT message to the SL transmitter and send a Connection Opened (SSP, Destination Opened) confirmation to the port layer (see 7.14.4.4.3);
 - 9) If the OPEN address frame PROTOCOL field is set to SMP and the Reject SMP Opens state machine variable is set to NO, this state shall send a Transmit OPEN_ACCEPT message to the SL transmitter and send a Connection Opened (SMP, Destination Opened) confirmation to the port layer (see 7.14.4.4.3); or
 - 10) If the OPEN address frame PROTOCOL field is set to STP and the Reject STP Opens state machine variable is set to NO, this state shall send a Transmit OPEN_ACCEPT message to the SL transmitter and send a Connection Opened (STP, Destination Opened) confirmation to the port layer (see 7.14.4.4.3).

If this state sends a Transmit OPEN_REJECT message to the SL transmitter, it shall also send an Inbound Connection Rejected confirmation to the port layer.

NOTE 23 - Possible livelock scenarios occur if the BREAK_REPLY method of responding to BREAK primitive sequences is disabled and a SAS logical phy transmits BREAK to abort a connection request (e.g., if its Open Timeout timer expires). SAS logical phys should respond to OPEN Address frames faster than 1 ms to reduce susceptibility to this problem.

[If a PMREQ Received message is received this state shall send a Transmit PMNAK to the SL transmitter.](#)

7.14.4.4.2 Transition SL_CC2:Selected to SL_CC0:Idle

This transition shall occur after this state sends a Transmit OPEN_REJECT message to the SL transmitter.

7.14.4.4.3 Transition SL_CC2:Selected to SL_CC3:Connected

This transition shall occur after sending a Connection Opened confirmation.

This transition shall include:

- a) an Open SSP Connection, Open STP Connection, or Open SMP Connection argument based on the requested protocol; and
- b) the received OPEN address frame.

7.14.4.4.4 Transition SL_CC2:Selected to SL_CC5:BreakWait

If the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this transition shall occur after receiving a NOTIFY Received (Power Loss Expected) message and shall include a Power Loss Expected argument.

7.14.4.4.5 Transition SL_CC2:Selected to SL_CC6:Break

This transition shall occur after a BREAK Received message is received.

7.14.4.5 SL_CC3:Connected state

7.14.4.5.1 State description

This state enables the SSP, STP, or SMP link layer state machine to transmit dwords during a connection. See 7.13 for details on rate matching during the connection.

If this state is entered from SL_CC1:ArbSel state or the SL_CC2:Selected state with an argument of Open SMP Connection then this state shall send an Enable Disable SMP (Enable) message to the SMP link layer state machines (see 7.18.5).

If this state is entered from SL_CC1:ArbSel state or the SL_CC2:Selected state with an argument of Open SSP Connection then this state shall send an Enable Disable SSP (Enable) message to the SSP link layer state machines (see 7.16.8).

If this state is entered from SL_CC1:ArbSel state or the SL_CC2:Selected state with an argument of Open STP Connection then this state shall send an Enable Disable STP (Enable) message to the STP link layer state machines (see 7.17.8).

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL transmitter until the SSP, SMP, or STP link layer state machine starts transmitting.

A CLOSE Received message may be received at any time while in this state, but shall be ignored during SSP and SMP connections. If a CLOSE Received (Clear Affiliation) is received during an STP connection, this state shall clear any affiliation (see 7.17.4).

[Editor's Note 2: PMREQs are ignored in Connected state.](#)

7.14.4.5.2 Transition SL_CC3:Connected to SL_CC4:DisconnectWait

This transition shall occur if a Request Close message is received.

7.14.4.5.3 Transition SL_CC3:Connected to SL_CC5:BreakWait

This transition shall occur after sending a Connection Closed (Break Requested) confirmation to the port layer if:

- a) a Request Break message is received and a BREAK Received message has not been received; or
- b) a NOTIFY Received (Power Loss Expected) message is received.

If a NOTIFY Received (Power Loss Expected) message was received and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this state shall include a Power Loss Expected argument.

7.14.4.5.4 Transition SL_CC3:Connected to SL_CC6:Break

This transition shall occur if a BREAK Received message is received and after sending a Connection Closed (Break Received) confirmation to the port layer.

7.14.4.5.5 Transition SL_CC3:Connected to SL_CC7:CloseSTP

This transition shall occur if a CLOSE Received message is received during an STP connection.

7.14.4.6 SL_CC4:DisconnectWait state

7.14.4.6.1 State description

This state closes the connection and releases all resources associated with the connection.

Upon entry into this state, this state shall:

- 1) send a Transmit CLOSE (Normal) message or Transmit CLOSE (Clear Affiliation) message to the SL transmitter (see 7.17.6); and
- 2) initialize and start the Close Timeout timer.

A CLOSE Received message may be received at any time while in this state. If a CLOSE Received (Clear Affiliation) is received during an STP connection, this state shall clear any affiliation (see 7.17.4).

NOTE 24 - Possible livelock scenarios occur if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled and a SAS logical phy transmits BREAK to break a connection (e.g., if its Close Timeout timer expires). SAS logical phys should respond to CLOSE faster than 1 ms to reduce susceptibility to this problem.

[Editor's Note 3: PMREQs are ignored in DisconnectWait state.](#)

7.14.4.6.2 Transition SL_CC4:DisconnectWait to SL_CC0:Idle

This transition shall occur after:

- a) sending a Transmit CLOSE message to the SL transmitter;
- b) receiving a CLOSE Received message; and
- c) sending a Connection Closed (Normal) confirmation to the port layer.

7.14.4.6.3 Transition SL_CC4:DisconnectWait to SL_CC5:BreakWait

This transition shall occur if a NOTIFY Received (Power Loss Expected) message is received or if:

- a) a BREAK Received message has not been received;
- b) no CLOSE Received message is received in response to a Transmit CLOSE message before the Close Timeout timer expires; and
- c) after sending a Connection Closed (Close Timeout) confirmation to the port layer.

If a NOTIFY Received (Power Loss Expected) message was received and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this state shall include a Power Loss Expected argument.

7.14.4.6.4 Transition SL_CC4:DisconnectWait to SL_CC6:Break

This transition shall occur after receiving a BREAK Received message and after sending a Connection Closed (Break Received) confirmation to the port layer.

7.14.4.7 SL_CC5:BreakWait state

7.14.4.7.1 State description

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

Upon entry into this state, this state shall:

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

If this state is entered with a Power Loss Expected message or this state receives a NOTIFY Received (Power Loss Expected) message and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this state shall send a Notify Received (Power Loss Expected) confirmation to the port layer.

If a BREAK Received message is received and the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), then this state shall send a Transmit BREAK_REPLY message to the SL transmitter.

NOTE 25 - Some SAS logical phys compliant with previous versions of this standard send a Transmit OPEN_REJECT (Retry) message to the SL transmitter in response to each OPEN Address Frame Received message received while in this state.

[Editor's Note 4: PMREQs are ignored in BreakWait state.](#)

7.14.4.7.2 Transition SL_CC5:BreakWait to SL_CC0:Idle

This transition shall occur after:

- a) receiving a BREAK_REPLY Received message if the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5);
- b) receiving a BREAK Received message if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled (see 7.12.5); or
- c) the Break Timeout timer expires.

7.14.4.8 SL_CC6:Break state

7.14.4.8.1 State description

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

Upon entry into this state:

- a) if the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), this state shall send a Transmit BREAK_REPLY message to the SL transmitter (see 7.14.4.8.2); and
- b) if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled (see 7.12.5), this state shall send a Transmit BREAK message to the SL transmitter.

If this state receives a NOTIFY Received (Power Loss Expected) message and the SAS port that contains this state machine supports NOTIFY (Power Loss Expected) (e.g., the SAS port is a SSP target port), then this state shall send a Notify Received (Power Loss Expected) confirmation to the port layer.

[Editor's Note 5: PMREQs are ignored in Break state.](#)

7.14.4.8.2 Transition SL_CC6:Break to SL_CC0:Idle

This transition shall occur after sending a Transmit BREAK message or a Transmit BREAK_REPLY message to the SL transmitter.

7.14.4.9 SL_CC7:CloseSTP state

7.14.4.9.1 State description

This state closes an STP connection and releases all resources associated with the connection.

Upon entry into this state, this state shall:

- 1) send a Transmit CLOSE (Normal) message or Transmit CLOSE (Clear Affiliation) message to the SL transmitter (see 7.17.6); and
- 2) send a Connection Closed (Normal) confirmation to the port layer (see 7.14.4.9.2).

NOTE 26 - Possible livelock scenarios occur if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled and a SAS logical phy transmits BREAK to break a connection (e.g., if its Close Timeout timer expires). SAS logical phys should respond to CLOSE faster than 1 ms to reduce susceptibility to this problem.

Editor's Note 6: PMREQs are ignored in CloseSTP state.

7.14.4.9.2 Transition SL_CC7:CloseSTP to SL_CC0:Idle

This transition shall occur after sending a Connection Closed (Normal) confirmation to the port layer.

7.14.4.10 SL_CC8:PM_Request state

7.14.4.10.1 State description

This state generates the PMREQ, and waits for an acknowledgement.

Upon entry into this state, this state shall:

- a) send a Transmit PMREQ message to the SL transmitter; and
- b) initialize and start the Power Management Request Timeout timer.

If this state receives a PMREQ Received message, and this phy has lower power management priority (see xxx), then this state shall send a Transmit PMACK message to the SL transmitter.

Editor's Note 7: New paragraph to be added somewhere yet. The concept of power management priority above is to determine who is responsible to PMACK/PMNAK a power management request when the PMREQs cross on the wire. One initial thought is just the link with higher SAS address wins. Also consider when the PMREQs are not for same power state... should partial always win?

If this state receives an Open Connection request, this state shall send an Open Failed (Power Management Pending) confirmation to the port layer.

Editor's Note 8: The intention is that this Open Failed confirmation would then get turned into a Retry Open in the port layer so that the request can be retried. The point here is that we don't want a link to send an OPEN frame after it has sent a PMREQ until after we leave this state (while the PMREQ hasn't been acknowledged yet). This port layer change isn't included in this proposal yet.

7.14.4.10.2 Transition SL_CC8:PM_Request to SL_CC9:PM_Quiet

This transition shall occur after:

- a) sending a Transmit PMACK message to the SL transmitter; or
- b) receiving a PMACK Received message.

7.14.4.10.3 Transition SL_CC8:PM_Request to SL_CC0:Idle

This transition shall occur after:

- a) a PMNAK Received message is received; or
- b) the Power Management Request Timeout timer expires.

7.14.4.10.4 Transition SL_CC8:PM_Request to SL_CC2:Selected

This transition shall occur after receiving an OPEN Address Frame Received message.

7.14.4.11 SL_CC9:PM_Quiet state

7.14.4.11.1 State description

This state puts the phy into a low power state.

Upon entry into this state, this state shall send a Power Management (Enter Partial) or Power Management (Enter Slumber) to the phy layer.

If this state receives an Open Connection request, this state shall send a Power Management (Exit) request to the phy layer.

7.14.4.11.2 Transition SL_CC9:PM_Quiet to SL_CC0:Idle

This transition shall occur after a Phy Power (Wakeup) confirmation is received and an Open Connection request has not been received.

Editor's Note 9: This confirmation isn't listed in 08-015. This is an indication from the phy that it has gone back to the Phy Ready state.

7.14.4.11.3 Transition SL_CC9:PM_Quiet to SL_CC1:ArbSel

This transition shall occur after receiving both a Phy Power (Wakeup) confirmation and an Open Connection request. The Open Connection request includes these arguments:

- 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.

7.15 XL (link layer for expander logical phys) state machine

7.15.1 XL state machine overview

The XL state machine controls the flow of dwords on the logical link and establishes and maintains connections with another XL state machine as facilitated by the expander function (e.g., the ECM and ECR).

This state machine consists of the following states:

- a) XL0:Idle (see 7.15.3)(initial state);
- b) XL1:Request_Path (see 7.15.4);
- c) XL2:Request_Open (see 7.15.5);
- d) XL3:Open_Confirm_Wait (see 7.15.6);
- e) XL4:Open_Reject (see 7.15.7);
- f) XL5:Forward_Open (see 7.15.8);
- g) XL6:Open_Response_Wait (see 7.15.9);
- h) XL7:Connected (see 7.15.10);
- i) XL8:Close_Wait (see 7.15.11);
- j) XL9:Break (see 7.15.12); **and**
- k) XL10:Break_Wait (see 7.15.13);
- l) XL11:PM_Request (see xxx); and
- m) XL12:PM_Quiet (see xxx).

The XL state machine shall start in the XL0:Idle state. The XL state machine shall transition to the XL0:Idle state from any other state after receiving an Enable Disable SAS Link (Disable) message from the SL_IR state machines (see 7.9.4).

The XL state machine receives the following messages from the SL_IR state machines:

- a) Enable Disable SAS Link (Enable); and

- b) Enable Disable SAS Link (Disable).

Any message received by a state that is not referred to in the description of that state shall be ignored.

The XL state machine shall maintain the timers listed in table 73.

Table 73 — XL timers

Timer	Initial value
Partial Pathway Timeout timer	Partial pathway timeout value (see 7.12.4.4)
Break Timeout timer	1 ms
Power Management Request Timeout timer	1 ms

Editor's Note 10: Need to update figures with new states, and with Request Power Management request coming down, and Phy Power confirmation coming up.

Figure 120 shows several states in the XL state machine.

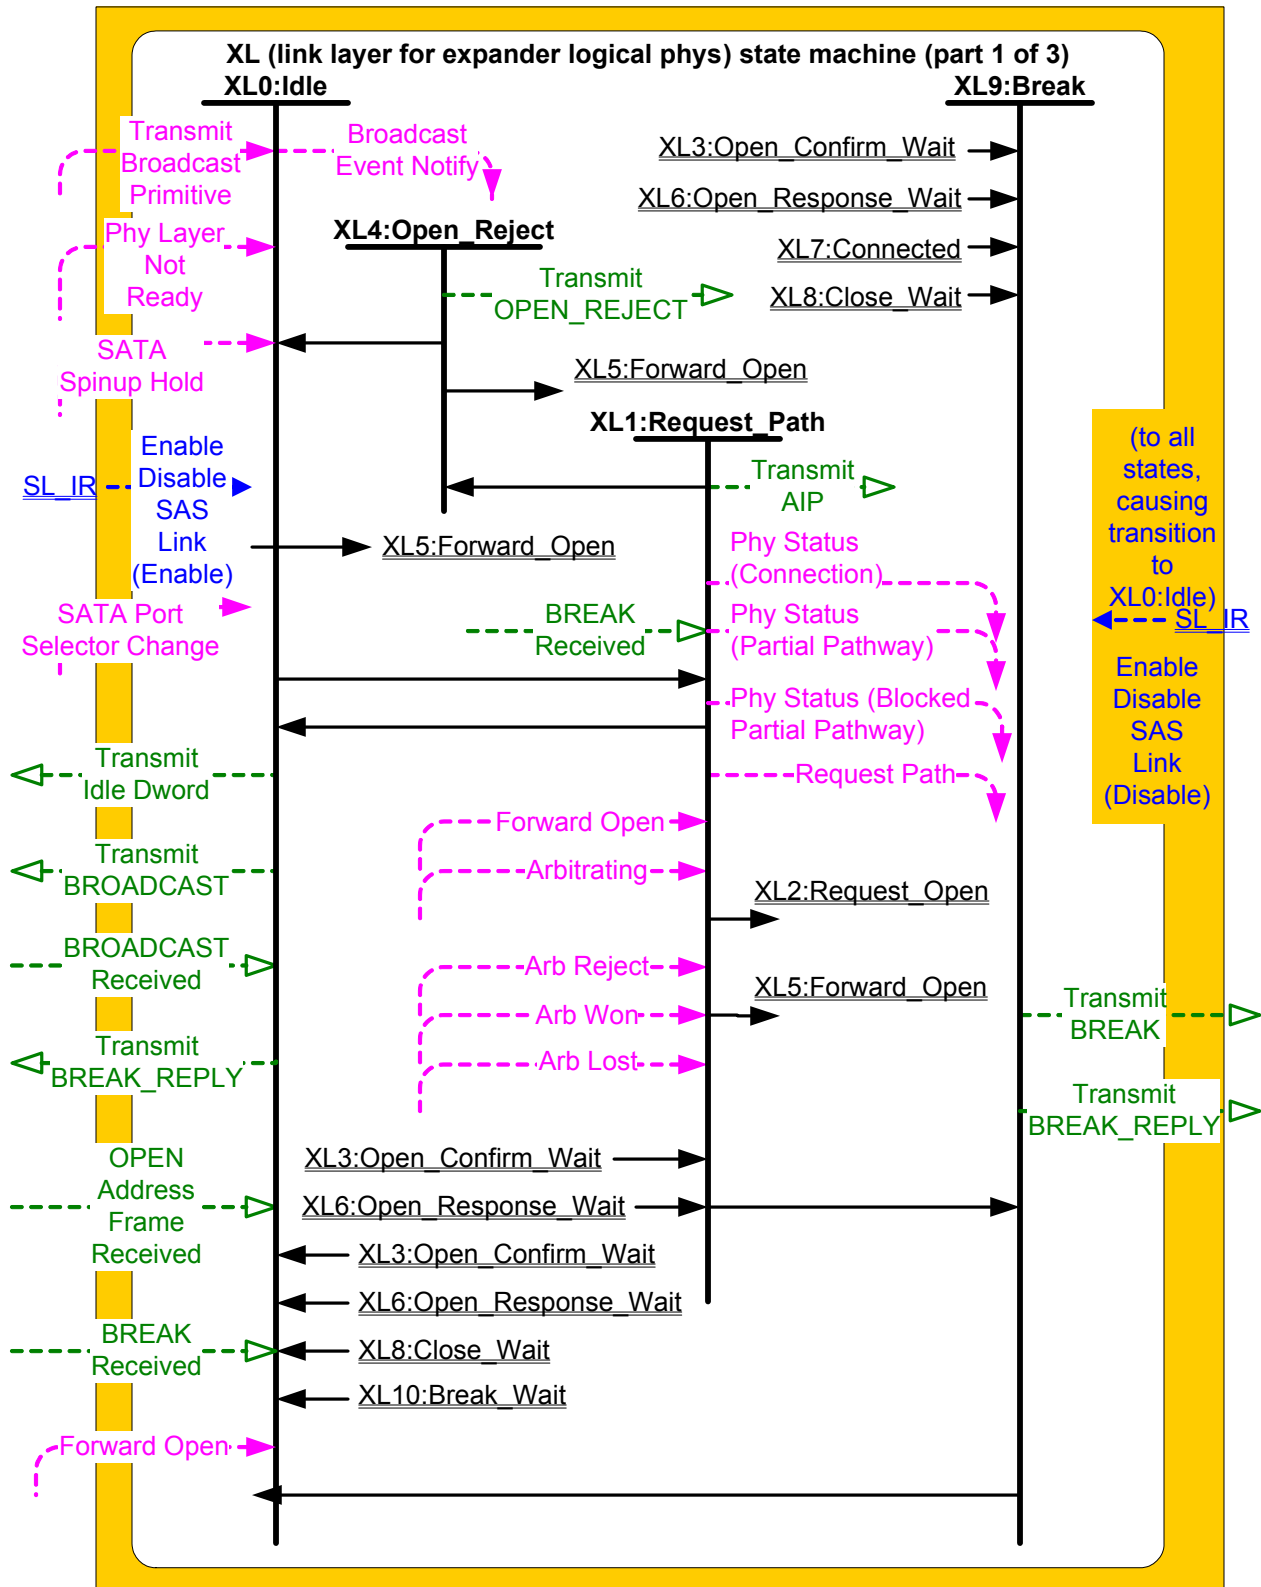


Figure 120 — XL (link layer for expander logical phys) state machine (part 1)

Figure 121 shows additional states in the XL state machine.

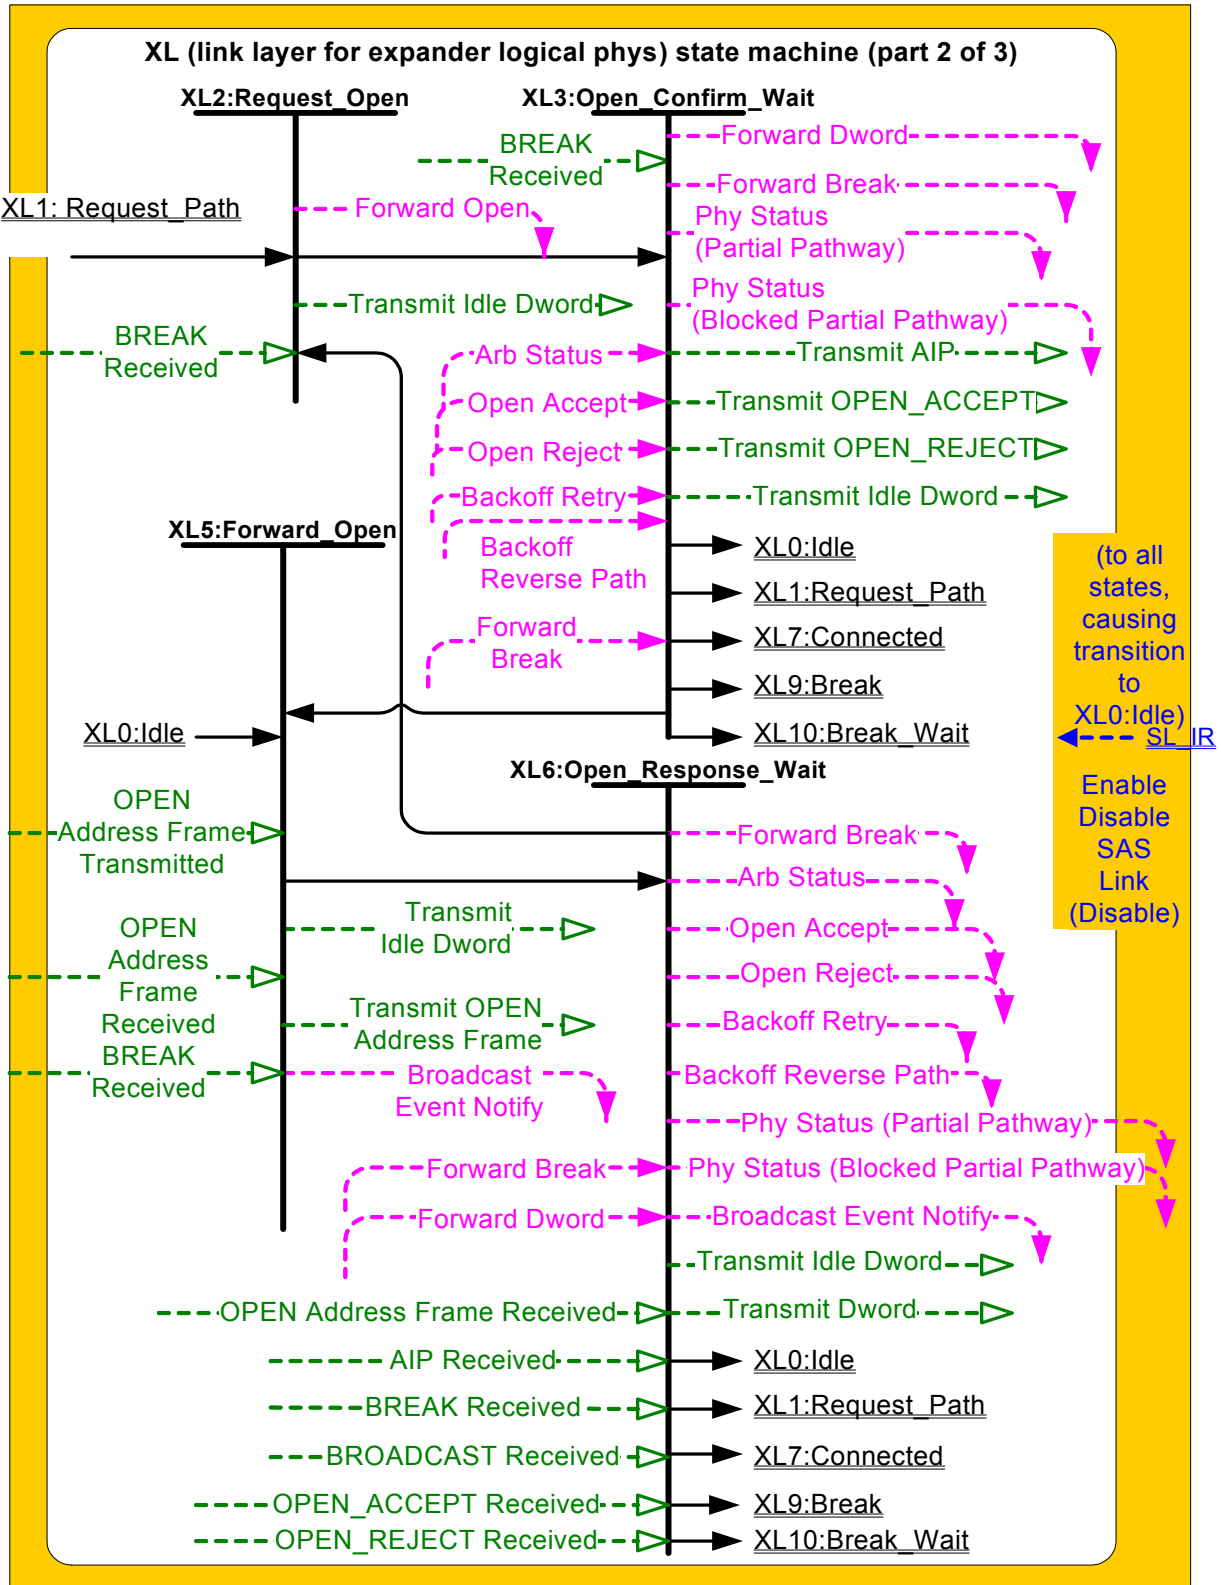


Figure 121 — XL (link layer for expander logical phys) state machine (part 2)

Figure 122 shows additional states in the XL state machine.

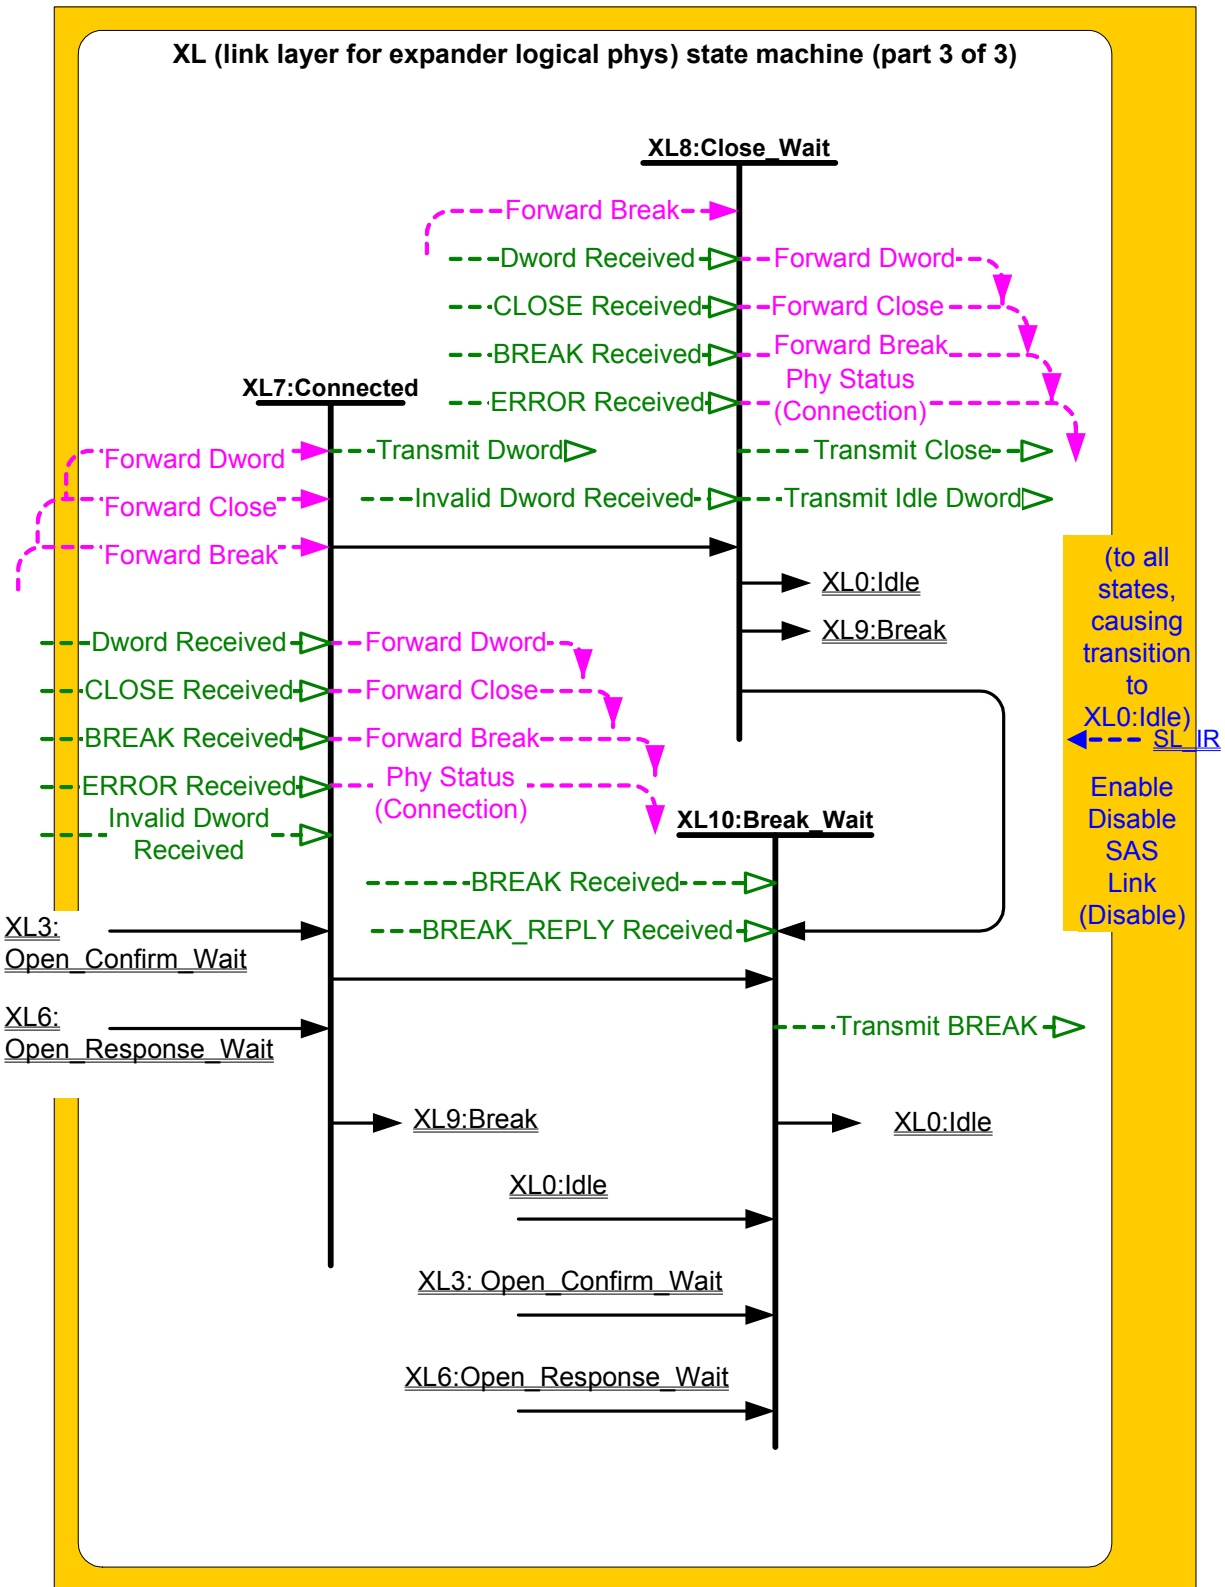


Figure 122 — XL (link layer for expander logical phys) state machine (part 3)

7.15.2 XL transmitter and receiver

The XL transmitter receives the following messages from the XL state machine specifying primitive sequences, frames, and dwords to transmit:

- a) Transmit Idle Dword;
- b) Transmit AIP with an argument indicating the specific type (e.g., Transmit AIP (Normal));
- c) Transmit BREAK;
- d) Transmit BREAK_REPLY;
- e) Transmit BROADCAST with an argument indicating the specific type (e.g., Transmit BROADCAST (Change));
- f) Transmit CLOSE with an argument indicating the specific type (e.g., Transmit CLOSE (Normal));
- g) Transmit OPEN_ACCEPT;
- h) Transmit OPEN_REJECT, with an argument indicating the specific type (e.g., Transmit OPEN_REJECT (No Destination));
- i) [Transmit PMREQ with an argument indicating the specific type \(e.g., Transmit PMREQ \(Partial\)\);](#)
- j) [Transmit PMACK;](#)
- k) [Transmit PMNAK;](#)
- l) Transmit OPEN Address Frame; and
- m) Transmit Dword.

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

- a) OPEN Address Frame Transmitted.

The XL transmitter shall ensure physical link rate tolerance management requirements are met (see 7.3) while originating dwords.

The XL transmitter shall ensure physical link rate tolerance management requirements are met while forwarding dwords (i.e., during a connection) by inserting or deleting as many deletable primitives as required to match the transmit and receive connection rates (see 7.3.2).

The XL transmitter shall ensure physical link rate tolerance management requirements are met (see 7.3) during and after switching from forwarding dwords to originating dwords, including, for example:

- a) when transmitting BREAK;
- b) when transmitting BREAK_REPLY;
- c) when transmitting CLOSE;
- d) when transmitting an idle dword after closing a connection (i.e., after receiving BREAK, BREAK_REPLY, or CLOSE);
- e) while transmitting a SATA frame to a SAS logical link during an STP connection, when transmitting the first SATA_HOLD in response to detection of SATA_HOLD; and
- f) while receiving dwords of a SATA frame from a SAS logical link during an STP connection, when transmitting SATA_HOLD.

NOTE 27 - The XL transmitter may always insert a deletable primitive before transmitting a BREAK, BREAK_REPLY, CLOSE, or SATA_HOLD to meet physical link rate tolerance management requirements.

The XL transmitter shall insert a deletable primitive before switching from originating dwords to forwarding dwords, including, for example:

- a) when transmitting OPEN_ACCEPT;
- b) when transmitting the last idle dword before a connection is established (i.e., after receiving OPEN_ACCEPT);
- c) while transmitting a SATA frame to a SAS logical link during an STP connection, when transmitting the last dword from the STP flow control buffer in response to release of SATA_HOLD;
- d) while transmitting a SATA frame to a SAS logical link during an STP connection, when transmitting the last SATA_HOLD in response to release of SATA_HOLD (e.g., if the STP flow control buffer is empty); and
- e) while receiving dwords of a SATA frame from a SAS logical link during an STP connection, when transmitting the last SATA_HOLD.

NOTE 28 - This ensures that physical link rate tolerance management requirements are met, even if the forwarded dword stream does not include a deletable primitive until the last possible dword.

The XL transmitter shall ensure rate matching requirements are met during a connection (see 7.13).

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

The XL receiver sends the following messages to the XL state machine indicating primitive sequences, frames, and dwords received from the SP_DWS receiver (see 6.9.2):

- a) AIP Received with an argument indicating the specific type (e.g., AIP Received (Normal));
- b) BREAK Received;
- c) BREAK_REPLY Received;
- d) BROADCAST Received;
- e) CLOSE Received;
- f) OPEN_ACCEPT Received;
- g) OPEN_REJECT Received;
- h) OPEN Address Frame Received;
- i) [PMREQ Received with an argument indicated the specific type \(e.g., PMREQ Received \(Partial\)\)](#);
- j) Dword Received with an argument indicating the data dword or primitive received; and
- k) Invalid Dword Received.

The XL receiver shall ignore all other dwords.

While receiving an address frame, if the XL receiver receives an invalid dword or ERROR, then the XL receiver shall:

- a) ignore the invalid dword or ERROR; or
- b) discard the address frame.

The XL transmitter relationship to other transmitters is defined in 4.3.2. The XL receiver relationship to other receivers is defined in 4.3.3.

7.15.3 XL0:Idle state

7.15.3.1 State description

This state is the initial state and is the state that is used when there is no connection pending or established.

If a Phy Layer Not Ready confirmation is received, this state shall send a Broadcast Event Notify (Phy Not Ready) request to the BPP.

If a SATA Spinup Hold confirmation is received, this state shall send a Broadcast Event Notify (SATA Spinup Hold) request to the BPP.

If an Enable Disable SAS Link (Enable) message is received, this state shall send a Broadcast Event Notify (Identification Sequence Complete) request to the BPP.

If a SATA Port Selector Change confirmation is received, this state shall send a Broadcast Event Notify (SATA Port Selector Change) request to the BPP.

If a BROADCAST Received message is received, this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., Change Received).

If a Transmit Broadcast indication is received, this state shall send a Transmit BROADCAST message to the XL transmitter with an argument specifying the specific type from the Transmit Broadcast indication. Otherwise, this state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

If a BREAK Received message is received and the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), then this state shall send a Transmit BREAK_REPLY message to the XL transmitter.

If a PMREQ Received message is received and power management is enabled (see xxx), then this state shall send a Transmit PMACK message to the SL transmitter. If a PMREQ Received message is received and power management is not enabled, then this state shall send a Transmit PMNAK message to the SL transmitter.

7.15.3.2 Transition XL0:Idle to XL1:Request_Path

This transition shall occur if:

- a) an Enable Disable SAS Link (Enable) message has been received;
- b) a Forward Open indication is not being received; and
- c) an OPEN Address Frame Received message is received.

This state shall include an OPEN Address Frame Received argument with the transition.

7.15.3.3 Transition XL0:Idle to XL5:Forward_Open

This transition shall occur if:

- a) an Enable Disable SAS Link (Enable) message has been received; and
- b) a Forward Open indication is received.

This transition shall include a set of arguments containing the arguments received in the Forward Open indication.

If an OPEN Address Frame Received message is received, this state shall include an OPEN Address Frame Received argument with the transition.

7.15.3.3.1 Transition XL0:Idle to XL11:PM_Request

If power management is enabled (see xxx), this transition shall occur after receiving both an Enable Disable SAS Link (Enable) confirmation and a Request Power Management request. The transition shall include an argument with the same type as the request (e.g. Request Power Management (Partial)).

7.15.3.3.2 Transition XL0:Idle to XL12:PM_Quiet

This transition shall occur after sending a Transmit PMACK message to the SL transmitter.

7.15.4 XL1:Request_Path state

7.15.4.1 State description

This state is used to arbitrate for connection resources and to specify the destination of the connection.

If an Arbitrating (Normal) confirmation is received, this state shall repeatedly send Transmit AIP (Normal) and Transmit Idle Dword messages to the XL transmitter in accordance with AIP transmission rules (see 7.12.4.3).

If an Arbitrating (Waiting On Partial) or Arbitrating (Blocked On Partial) confirmation is received, this state shall repeatedly send Transmit AIP (Waiting On Partial) and Transmit Idle Dword messages to the XL transmitter in accordance with AIP transmission rules (see 7.12.4.3).

If an Arbitrating (Waiting On Partial) confirmation is received, this state shall repeatedly send a Phy Status (Partial Pathway) response to the ECM.

If an Arbitrating (Blocked On Partial) confirmation is received, this state shall repeatedly send a Phy Status (Blocked Partial Pathway) response to the ECM.

If an Arbitrating (Waiting On Connection) confirmation is received, this state shall repeatedly send Transmit AIP (Waiting On Connection) and Transmit Idle Dword messages to the XL transmitter in accordance with AIP transmission rules (see 7.12.4.3).

If an Arbitrating (Waiting On Connection) confirmation is received, this state shall repeatedly send a Phy Status (Connection) response to the ECM.

If this state is entered from the XL6:Open_Response_Wait state, the Retry Priority Status argument shall be set to IGNORE AWT. If this state is entered from any other state, the Retry Priority Status argument shall be set to NORMAL.

Upon entry into this state, this state shall send a Request Path request to the ECM with the following arguments:

- 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;
- h) arbitration wait time; and
- i) retry priority status.

This state maintains the Partial Pathway Timeout timer.

If the Partial Pathway Timeout timer is not already running, the Partial Pathway Timeout timer shall be initialized and started when an Arbitrating (Blocked On Partial) confirmation is received.

If the Partial Pathway Timeout timer is already running, the Partial Pathway Timeout timer shall continue to run if an Arbitrating (Blocked On Partial) confirmation is received.

The Partial Pathway Timeout timer shall be stopped when one of the following confirmations is received:

- a) Arbitrating (Waiting On Partial); or
- b) Arbitrating (Waiting On Connection).

If the Partial Pathway Timeout timer expires, this state shall send a Partial Pathway Timeout Timer Expired request to the ECM.

[Editor's Note 11: PMREQs are ignored in Request_Path state.](#)

7.15.4.2 Transition XL1:Request_Path to XL0:Idle

This transition shall occur if:

- a) a BREAK Received message has not been received; and
- b) an Arb Lost confirmation is received.

7.15.4.3 Transition XL1:Request_Path to XL2:Request_Open

This transition shall occur if:

- a) a BREAK Received message has not been received; and
- b) an Arb Won confirmation is received.

7.15.4.4 Transition XL1:Request_Path to XL4:Open_Reject

This transition shall occur if:

- a) a BREAK Received message has not been received; and
- b) an Arb Reject confirmation is received.

This transition shall include an Arb Reject argument corresponding to the Arb Reject confirmation.

7.15.4.5 Transition XL1:Request_Path to XL5:Forward_Open

This transition shall occur if a Forward Open indication is received and none of the following confirmations have been received:

- a) Arbitrating (Normal);
- b) Arbitrating (Waiting On Partial);
- c) Arbitrating (Blocked On Partial);
- d) Arbitrating (Waiting On Connection);
- e) Arb Won;
- f) Arb Lost;
- g) Arb Reject (No Destination);
- h) Arb Reject (Bad Destination);
- i) Arb Reject (Connection Rate Not Supported);
- j) Arb Reject (Zone Violation);
- k) Arb Reject (Pathway Blocked); or
- l) Arb Reject (Retry).

This transition shall include:

- a) an OPEN Address Frame Received argument containing the arguments received in the Forward Open indication; and
- b) a BREAK Received argument if a BREAK Received message was received.

7.15.4.6 Transition XL1:Request_Path to XL9:Break

This transition shall occur after receiving a BREAK Received message if a Forward Open indication has not been received.

7.15.5 XL2:Request_Open state

7.15.5.1 State description

This state is used to forward an OPEN address frame through the ECR to a destination phy.

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

Upon entry into this state, this state shall send a Forward Open request to the ECR, received by the destination phy as a Forward Open indication (see 7.15.5.2). The arguments to the Forward Open request are:

- a) initiator port bit;
- b) protocol;
- c) features;
- d) connection rate;
- e) initiator connection tag;
- f) destination SAS address;
- g) source SAS address;
- h) source zone group;
- i) pathway blocked count;
- j) arbitration wait time; and
- k) more compatible features.

[Editor's Note 12: PMREQs are ignored in Request_Open state.](#)

7.15.5.2 Transition XL2:Request_Open to XL3:Open_Confirm_Wait

This transition shall occur after sending a Forward Open request to the ECR.

If a BREAK Received message is received, this state shall include a BREAK Received argument with the transition.

7.15.6 XL3:Open_Confirm_Wait state

7.15.6.1 State description

This state waits for confirmation for an OPEN address frame sent on a destination phy.

This state shall send the following messages to the XL transmitter:

- a) Transmit AIP (Normal) when an Arb Status (Normal) confirmation is received;
- b) Transmit AIP (Waiting On Partial) when an Arb Status (Waiting On Partial) confirmation is received;
- c) Transmit AIP (Waiting On Connection) when an Arb Status (Waiting On Connection) confirmation is received;
- d) Transmit AIP (Waiting On Device) when an Arb Status (Waiting On Device) confirmation is received;
- e) Transmit OPEN_ACCEPT when an Open Accept confirmation is received (see 7.15.6.5);
- f) Transmit OPEN_REJECT when an Open Reject confirmation is received with the argument from the Open Reject confirmation, after releasing path resources (see 7.15.6.2); or
- g) request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages when none of the previous conditions are present.

If a Backoff Retry confirmation is received, this state shall release path resources.

If a BREAK Received message is received or a BREAK Received argument is included in the transition into this state, this state shall send a Forward Break request to the ECR (see 7.15.6.6).

This state shall repeatedly send a Phy Status (Partial Pathway) response to the ECM until an Arb Status (Waiting On Partial) confirmation is received. After an Arb Status (Waiting on Partial) confirmation is received, this state shall repeatedly send a Phy Status (Blocked Partial Pathway) response to the ECM.

If a Dword Received message is received containing a valid dword except a BREAK primitive, this state shall send a Forward Dword request to the ECR containing that dword.

If:

- a) an Invalid Dword Received message is received; and
- b) the expander logical phy is forwarding to an expander logical phy attached to a SAS logical link,

then the expander logical phy shall:

- a) send an ERROR primitive with the Forward Dword request instead of the invalid dword; or
- b) delete the invalid dword.

If:

- a) an ERROR primitive is received with the Dword Received message or an Invalid Dword Received message is received; and
- b) the expander logical phy is forwarding to an expander phy attached to a SATA physical link,

then the expander logical phy shall:

- a) send a SATA_ERROR with the Forward Dword request instead of the invalid dword or ERROR primitive; or
- b) delete the ERROR primitive or invalid dword.

[Editor's Note 13: PMREQs are ignored in Open_Confirm_Wait state.](#)

7.15.6.2 Transition XL3:Open_Confirm_Wait to XL0:Idle

This transition shall occur after sending a Transmit OPEN_REJECT message to the XL transmitter if:

- a) a BREAK Received message has not been received; and

- b) a BREAK Received argument was not included in the transition into this state.

7.15.6.3 Transition XL3:Open_Confirm_Wait to XL1:Request_Path

This transition shall occur after receiving a Backoff Retry confirmation, after releasing path resources if:

- a) a BREAK Received message has not been received; and
- b) a BREAK Received argument was not included in the transition into this state.

7.15.6.4 Transition XL3:Open_Confirm_Wait to XL5:Forward_Open

This transition shall occur after receiving a Backoff Reverse Path confirmation if:

- a) a BREAK Received message has not been received; and
- b) a BREAK Received argument was not included in the transition into this state.

The transition shall include the Backoff Reverse Path arguments (i.e., the OPEN address frame).

7.15.6.5 Transition XL3:Open_Confirm_Wait to XL7:Connected

This transition shall occur after sending a Transmit OPEN_ACCEPT message to the XL transmitter if:

- a) a BREAK Received message has not been received; and
- b) a BREAK Received argument was not included in the transition into this state.

7.15.6.6 Transition XL3:Open_Confirm_Wait to XL9:Break

This transition shall occur after sending a Forward Break request to the ECR.

7.15.6.7 Transition XL3:Open_Confirm_Wait to XL10:Break_Wait

This transition shall occur after receiving a Forward Break indication if:

- a) a BREAK Received message has not been received; and
- b) a BREAK Received argument was not included in the transition into this state.

7.15.7 XL4:Open_Reject state

7.15.7.1 State description

This state is used to reject a connection request.

This state shall send one of the following messages to the XL transmitter (see 7.15.7.2):

- a) a Transmit OPEN_REJECT (No Destination) message when an Arb Reject (No Destination) argument is received with the transition into this state;
- b) a Transmit OPEN_REJECT (Bad Destination) message when an Arb Reject (Bad Destination) argument is received with the transition into this state;
- c) a Transmit OPEN_REJECT (Connection Rate Not Supported) message when an Arb Reject (Connection Rate Not Supported) argument is received with the transition into this state;
- d) a Transmit OPEN_REJECT (Zone Violation) message when an Arb Reject (Zone Violation) argument is received with the transition into this state;
- e) a Transmit OPEN_REJECT (Pathway Blocked) message when an Arb Reject (Pathway Blocked) argument is received with the transition into this state; or
- f) a Transmit OPEN_REJECT (Retry) message when an Arb Reject (Retry) argument is received with the transition into this state.

[Editor's Note 14: PMREQs are ignored in Open_Reject state.](#)

7.15.7.2 Transition XL4:Open_Reject to XL0:Idle

This transition shall occur after sending a Transmit OPEN_REJECT message to the XL transmitter.

7.15.7.3 Transition XL4:Open_Reject to XL5:Forward_Open

This transition shall occur if a Forward Open indication is received. This transition shall include an OPEN Address Frame Received argument containing the arguments received in the Forward Open indication.

7.15.8 XL5:Forward_Open state

7.15.8.1 State description

This state is used to transmit an OPEN address frame passed with the transition into this state.

If a BROADCAST Received message is received, this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., Change Received).

Upon entry into this state, this state shall send a Transmit OPEN Address Frame message to the XL transmitter with the fields set to the values specified with the transition into this state.

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

If a PMREQ Received message is received this state shall send a Transmit PMNAK to the SL transmitter.

7.15.8.2 Transition XL5:Forward_Open to XL6:Open_Response_Wait

This transition shall occur after receiving an OPEN Address Frame Transmitted message.

If an OPEN Address Frame Received message or argument is received, this state shall include an OPEN Address Frame Received argument with the transition.

If a BREAK Received message or argument is received, this state shall include a BREAK Received argument with the transition.

7.15.9 XL6:Open_Response_Wait state

7.15.9.1 State description

This state waits for a response to a transmitted OPEN address frame and determines the appropriate action to take based on the response.

This state shall either:

- a) request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter, honoring ALIGN insertion rules for rate matching and physical link rate tolerance management; or
- b) send Transmit Dword messages to the XL transmitter to transmit all dwords received with Forward Dword indications.

If a BROADCAST Received message is received before an AIP Received message is received this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., Broadcasts).

This state shall send the following responses to the ECR, which are received by the source phy as confirmations:

- a) an Open Accept response when an OPEN_ACCEPT Received message is received (see 7.15.9.5);
- b) an Open Reject response when an OPEN_REJECT Received message is received, after releasing any path resources (see 7.15.9.2);
- c) a Backoff Retry response, after releasing path resources (see 7.15.9.3), when:
 - A) an AIP Received message has not been received;
 - B) an OPEN Address Frame Received message is received or an OPEN Address Frame Received argument is included in the transition into this state containing a higher priority OPEN address frame according to the arbitration fairness comparison (see 7.12.3); and
 - C) the destination SAS address and connection rate of the received OPEN address frame are not equal to the source SAS address and connection rate of the transmitted OPEN address frame;

- d) a Backoff Retry response, after releasing path resources (see 7.15.9.3), when:
 - A) an AIP Received message has been received;
 - B) an OPEN Address Frame Received message is received or an OPEN Address Frame Received argument is included in the transition into this state; and
 - C) the destination SAS address and connection rate of the received OPEN address frame are not equal to the source SAS address and connection rate of the transmitted OPEN address frame;
 - e) a Backoff Reverse Path response (see 7.15.9.4) when:
 - A) an AIP Received message has not been received,
 - B) an OPEN Address Frame Received message is received or an OPEN Address Frame Received argument is included in the transition into this state containing a higher priority OPEN address frame according to the arbitration fairness comparison (see 7.12.3); and
 - C) the destination SAS address and connection rate of the received OPEN address frame are equal to the source SAS address and connection rate of the transmitted OPEN address frame;
- and
- f) a Backoff Reverse Path response (see 7.15.9.4) when:
 - A) an AIP Received message has been received;
 - B) an OPEN Address Frame Received message is received or an OPEN Address Frame Received argument is included in the transition into this state; and
 - C) the destination SAS address and connection rate of the received OPEN address frame are equal to the source SAS address and connection rate of the transmitted OPEN address frame.

A Backoff Reverse Path response shall include the contents of the OPEN Address Frame Received message or argument.

This state shall send the following responses to the ECR, which are received by the source phy as confirmations:

- a) an Arb Status (Waiting On Device) response upon entry into this state;
- b) an Arb Status (Normal) response when an AIP Received (Normal) message is received;
- c) an Arb Status (Waiting On Partial) response when an AIP Received (Waiting On Partial) message is received;
- d) an Arb Status (Waiting On Connection) response when an AIP Received (Waiting On Connection) message is received; and
- e) an Arb Status (Waiting On Device) response when an AIP Received (Waiting On Device) message is received.

If a BREAK Received message is received or a BREAK Received argument is included in the transition into this state, this state shall send a Forward Break request to the ECR (see 7.15.9.6).

This state shall repeatedly send a Phy Status (Partial Pathway) response to the ECM until an AIP Received (Waiting On Partial) message is received. After an AIP Received (Waiting On Partial) message is received, this state shall repeatedly send a Phy Status (Blocked Partial Pathway) response to the ECM.

[If a PMREQ Received message is received this state shall send a Transmit PMNAK to the SL transmitter.](#)

7.15.9.2 Transition XL6:Open_Response_Wait to XL0:Idle

This transition shall occur after sending an Open Reject response to the ECR.

7.15.9.3 Transition XL6:Open_Response_Wait to XL1:Request_Path

This transition shall occur after sending a Backoff Retry response to the ECR.

7.15.9.4 Transition XL6:Open_Response_Wait to XL2:Request_Open

This transition shall occur after sending a Backoff Reverse Path response to the ECR.

7.15.9.5 Transition XL6:Open_Response_Wait to XL7:Connected

This transition shall occur after sending an Open Accept response to the ECR.

7.15.9.6 Transition XL6:Open_Response_Wait to XL9:Break

This transition shall occur after sending a Forward Break response to the ECR.

7.15.9.7 Transition XL6:Open_Response_Wait to XL10:Break_Wait

This transition shall occur after receiving a Forward Break indication if:

- a) a BREAK Received message has not been received; and
- b) a BREAK Received argument was not included in the transition into this state.

7.15.10 XL7:Connected state**7.15.10.1 State description**

This state provides a full-duplex circuit between two phys within an expander device.

This state shall send Transmit Dword messages to the XL transmitter to transmit all dwords received with Forward Dword indications. During an STP connection, the expander device may expand or contract a repeated or continued primitive sequence (see 7.2.4).

If this state has not sent a Forward Close request to the ECR, this state shall send Forward Dword requests to the ECR containing each valid dword except BREAK and CLOSE primitives received with Dword Received messages. During an STP connection, the expander device may expand or contract a repeated or continued primitive sequence (see 7.2.4).

If:

- a) an Invalid Dword Received message is received; and
- b) the expander phy is forwarding to an expander logical phy attached to a SAS logical link,

the expander logical phy shall:

- a) send an ERROR primitive with the Forward Dword request instead of the invalid dword; or
- b) delete the invalid dword.

If:

- a) an ERROR primitive is received with the Dword Received message or an Invalid Dword Received message is received; and
- b) the expander phy is forwarding to an expander logical phy attached to a SATA phy,

the expander logical phy shall:

- a) send a SATA_ERROR with the Forward Dword request instead of the invalid dword or ERROR primitive; or
- b) delete the ERROR primitive or invalid dword.

If a CLOSE Received message is received, this state shall send a Forward Close request to the ECR with the argument from the CLOSE Received message.

If a BREAK Received message is received, this state shall send a Forward Break request to the ECR (see 7.15.10.3).

This state shall repeatedly send a Phy Status (Connection) response to the ECM.

[Editor's Note 15: PMREQs are ignored in Connected state.](#)

7.15.10.2 Transition XL7:Connected to XL8:Close_Wait

This transition shall occur after receiving a Forward Close indication if a BREAK Received message has not been received.

7.15.10.3 Transition XL7:Connected to XL9:Break

This transition shall occur after sending a Forward Break request to the ECR.

7.15.10.4 Transition XL7:Connected to XL10:Break_Wait

This transition shall occur after receiving a Forward Break indication if a BREAK Received message has not been received.

7.15.11 XL8:Close_Wait state**7.15.11.1 State description**

This state closes a connection and releases path resources.

Upon entry into this state, this state shall send a Transmit CLOSE message to the XL transmitter with the argument from the Forward Close indication, then shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

NOTE 29 - Possible livelock scenarios occur if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled and a phy transmits BREAK to break a connection (e.g., if its Close Timeout timer expires). Phys should respond to CLOSE faster than 1 ms to reduce susceptibility to this problem.

If a Dword Received message is received containing a valid dword except a BREAK or CLOSE primitive, this state shall send a Forward Dword request to the ECR containing that dword. During an STP connection, the expander device may expand or contract a repeated or continued primitive sequence (see 7.2.4).

If:

- a) an Invalid Dword Received message is received; and
- b) the expander logical phy is forwarding to an expander logical phy attached to a SAS logical link,

the expander logical phy shall:

- a) send an ERROR primitive with the Forward Dword request instead of the invalid dword; or
- b) delete the invalid dword.

If:

- a) an ERROR primitive is received with the Dword Received message or an Invalid Dword Received message is received; and
- b) the expander logical phy is forwarding to an expander phy attached to a SATA physical link,

the expander logical phy shall:

- a) send a SATA_ERROR with the Forward Dword request instead of the invalid dword or ERROR primitive; or
- b) delete the ERROR primitive or invalid dword.

If a CLOSE Received message is received, this state shall release path resources and send a Forward Close request to the ECR with the argument from the CLOSE Received message (see 7.15.11.2).

If a BREAK Received message is received, this state shall send a Forward Break request to the ECR (see 7.15.11.3).

This state shall repeatedly send a Phy Status (Connection) response to the ECM.

Editor's Note 16: PMREQs are ignored in Close_Wait state.

7.15.11.2 Transition XL8:Close_Wait to XL0:Idle

This transition shall occur after sending a Forward Close request to the ECR.

7.15.11.3 Transition XL8:Close_Wait to XL9:Break

This transition shall occur after sending a Forward Break request to the ECR.

7.15.11.4 Transition XL8:Close_Wait to XL10:Break_Wait

This transition shall occur after receiving a Forward Break indication if a BREAK Received message has not been received.

7.15.12 XL9:Break state**7.15.12.1 State description**

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

This state shall send a Transmit BREAK message to the XL transmitter (see 7.15.12.2).

This state shall:

- a) send a Transmit BREAK_REPLY message to the XL transmitter if the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5); and
- b) send a Transmit BREAK message to the XL transmitter if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled (see 7.12.5).

[Editor's Note 17: PMREQs are ignored in Break state.](#)

7.15.12.2 Transition XL9:Break to XL0:Idle

This transition shall occur after sending a Transmit BREAK message or a Transmit BREAK_REPLY message to the XL transmitter.

7.15.13 XL10:Break_Wait state**7.15.13.1 State description**

This state closes the connection if there is one and releases path resources associated with the connection.

Upon entry into this state, this state shall:

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

If a BREAK Received message is received and the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), then this state shall send a Transmit BREAK_REPLY message to the XL transmitter.

[Editor's Note 18: PMREQs are ignored in Break_Wait state.](#)

7.15.13.2 Transition XL10:Break_Wait to XL0:Idle

This transition shall occur after:

- a) a BREAK_REPLY Received message is received if the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5);

- b) a BREAK Received message is received if the BREAK_REPLY method of responding to received BREAK primitive sequences is disabled (see 7.12.5); or
- c) the Break Timeout timer expires.

7.15.13.3 XL11:PM_Request state

7.15.13.3.1 State description

This state generates the PMREQ, and waits for an acknowledgement.

Upon entry into this state, this state shall:

- a) send a Transmit PMREQ message to the SL transmitter; and
- b) initialize and start the Power Management Request Timeout timer.

If this state receives a PMREQ Received message, and this phy has lower power management priority (see xxx), then this state shall send a Transmit PMACK message to the SL transmitter.

If this state receives an Forward Open request, this state shall send a Backoff Retry confirmation to the ECR.

Editor's Note 19: During the time waiting for a response to the request, send back a Backoff Retry to the ECR. This will result in continued AIPs back to the OPENing device.

7.15.13.3.2 Transition XL11:PM_Request to XL12:PM_Quiet

This transition shall occur after:

- a) sending a Transmit PMACK message to the SL transmitter; or
- b) receiving a PMACK Received message.

7.15.13.3.3 Transition XL11:PM_Request to XL12:Idle

This transition shall occur after:

- a) a PMNAK Received message is received; or
- b) the Power Management Request Timeout timer expires.

7.15.13.3.4 Transition XL11:PM_Request to XL5:Forward

This transition shall occur after receiving an OPEN Address Frame Received message.

7.15.13.4 XL12:PM_Quiet state

7.15.13.4.1 State description

This state puts the phy into a low power state.

Upon entry into this state, this state shall send a Power Management (Enter Partial) or Power Management (Enter Slumber) to the phy layer.

If this state receives an Forward Open request, this state shall:

- a) send a Power Management (Exit) request to the phy layer; and
 - A) if the phy is in a slumber power state, send an Open Reject Retry confirmation to the ECR; or
 - B) if the phy is in a partial power state, send a Backoff Retry confirmation to the ECR.

Editor's Note 20: When in slumber, get the connection request rejected. When in partial, do the backoff retry to continue the AIPs. Both of these will result in a "new" Forward Open to happen at later time from the ECR.

7.15.13.4.2 Transition XL12:PM_Quiet to XL0:Idle

This transition shall occur after a Phy Power (Wakeup) confirmation is received.

Open Issues

This proposal is written from single link perspective, but there can be two links contained in a single phy when multiplexing is enabled. My thought is that the phy layer doesn't actually go into the partial / slumber states until both link layers have gone into the PM_Quiet state. There needs to be some (probably minor) changes to 08-015 to accommodate that.

Also with multiplexing, possible that each link could request different power states. Probably need to have some form of resolution criteria. Having partial take priority is the "safest" since recovery time is within the slumber requirements.

Currently, proposal does not have expander doing any "filtering" of the forwarded dwords. This would need to be added yet so that PMREQs/etc are not forwarded through expanders.