To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 3 November 2003
Subject: 03-346r2 SAS-1.1 Pathway recovery priority corrections

Revision history
Revision 0 (22 October 2003) First revision
Revision 1 (30 October 2003) Incorporated comments from Gil Romo (QLogic).
Revision 2 (3 November 2003) Incorporated comments from November SAS WG.

Related documents
sas1r01 - Serial Attached SCSI 1.1 revision 1

Overview
A few deadlock scenarios have been discovered and a few rules regarding pathway recovery have been proven to be unclear. They mostly relate to this paragraph in 7.12.4.4:

> The ECM shall instruct the arbitrating expander phy to reject the connection request by transmitting OPEN_REJECT (PATHWAY BLOCKED) when the Partial Pathway Timeout timer expires and the pathway recovery priority of the arbitrating expander phy (i.e., the expander phy requesting the connection) is less than the pathway recovery priority of all expander phys within the destination port with an arbitration status of WAITING_ON_PARTIAL.

and to the definitions of the Arbitrating confirmations (requirements for the ECM to follow) in the model section.

Problems include:

a) If a destination phy is reporting Phy Status (Blocked Partial Pathway) for a connection that came to it from the ECM via Transmit Open, the source of the fields used for that phy in pathway recovery priority comparisons is not defined (since there is no incoming OPEN address frame like there is for other requests). It should be the fields from the Transmit Open.
b) The pathway recovery comparison of the arbitrating phy vs. each of the destination phys needs to be “less than or equal to any” rather than “less than all.” With the source SAS address taken from the Transmit Open, the source SAS address can be the same, so a peer request might need to be retried. This request could be equal to a source SAS address in use at the destination port.
c) The connection rate field needs to be removed from the pathway recovery priority. If a wide expander port has a mix of link rates (some phys are 1.5 Gbps and others are 3 Gbps) and a deadlock is caused by a set of 3 Gbps connection requests, freeing a 1.5 Gbps phy may not resolve the deadlock. It might not occur until 6 Gbps connections are present, but it should be corrected now.
d) The Arbitrating confirmations from the ECM to the requesting phy consider these combinations of destination Phy Statuses:
   A) all are Partial Pathway, Blocked Partial Pathway, and/or Connected - return Arbitrating (Waiting on Connection)
   B) all are Partial Pathway - return Arbitrating (Waiting on Partial)
   C) all are Blocked Partial Pathway - return Arbitrating (Waiting on Blocked)

) but omit:

   A) all are Partial and/or Blocked - needs to return Arbitrating (Waiting on Partial).

e) Phy Status (Blocked Partial Pathway) needs to propagate through phys. If phy A is waiting on phy B which is waiting on phy C which is returning Phy Status (Blocked Partial Pathway), both phy A and phy B need to receive Phy Status (Blocked Partial Pathway) and run their PPTs. Phy Status (Partial Pathway) needs to be propagated the same way. For consistency, Phy Status (Connected) should be renamed to Phy Status (Connection), also triggered on receipt of Arbitrating (Waiting on Connection), and also be propagated. This becomes the first area in SAS that treats AIP (WAITING ON CONNECTION) differently from AIP (NORMAL) for more than basic forwarding.
f) The XL state machine needs to remember an incoming OPEN address frame that arrives at the same time as a Transmit Open request while it is in XL0:Idle. The Transmit Open takes precedence, but the incoming OPEN has to be remembered for eventual comparison with the outgoing OPEN (in

XL6:Open_Response_Wait). Other states like XL5:Forward_Open do this; XL0 just needs to match. This is just evident because of the bizarre edge-only messages used on the standard. In real hardware, reception of the OPEN is latched and attended to at the appropriate time.

SAS 1.0 impact: expanders following the SAS 1.0 rules might create deadlocks in complex topologies, especially with mixed physical link rates within wide ports. It should suffice to document the changes in SAS-1.1. All the problems relate to requirements placed on the ECM, which was not specified in detail like other parts of the expander model (e.g. the XL state machine).

Examples

Figure 1 shows a situation where the highest priority request is a blocked partial pathway. No changes are proposed; this just emphasizes that three independent comparisons (B to X, C to X, D to X) are run, not one comparison (B to C to D to X),

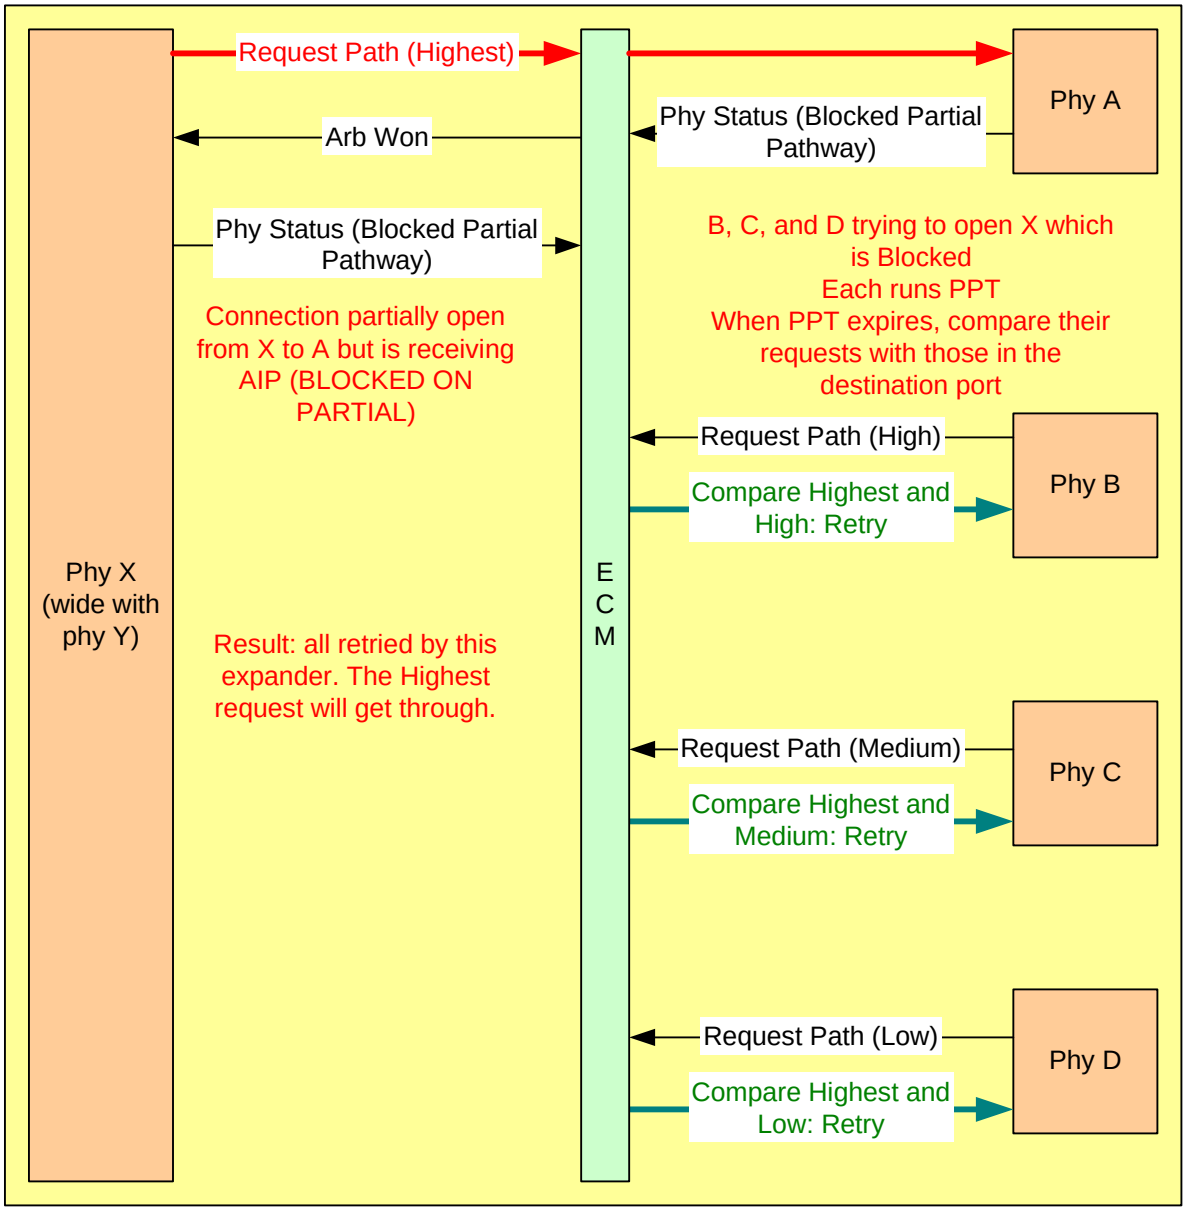

Figure 1 — Highest priority request is blocked

Figure 2 shows a situation where the lowest priority request is a blocked partial pathway. No changes are proposed; this just emphasizes that three independent comparisons (B to X, C to X, D to X) are run, not one comparison (B to C to D to X).

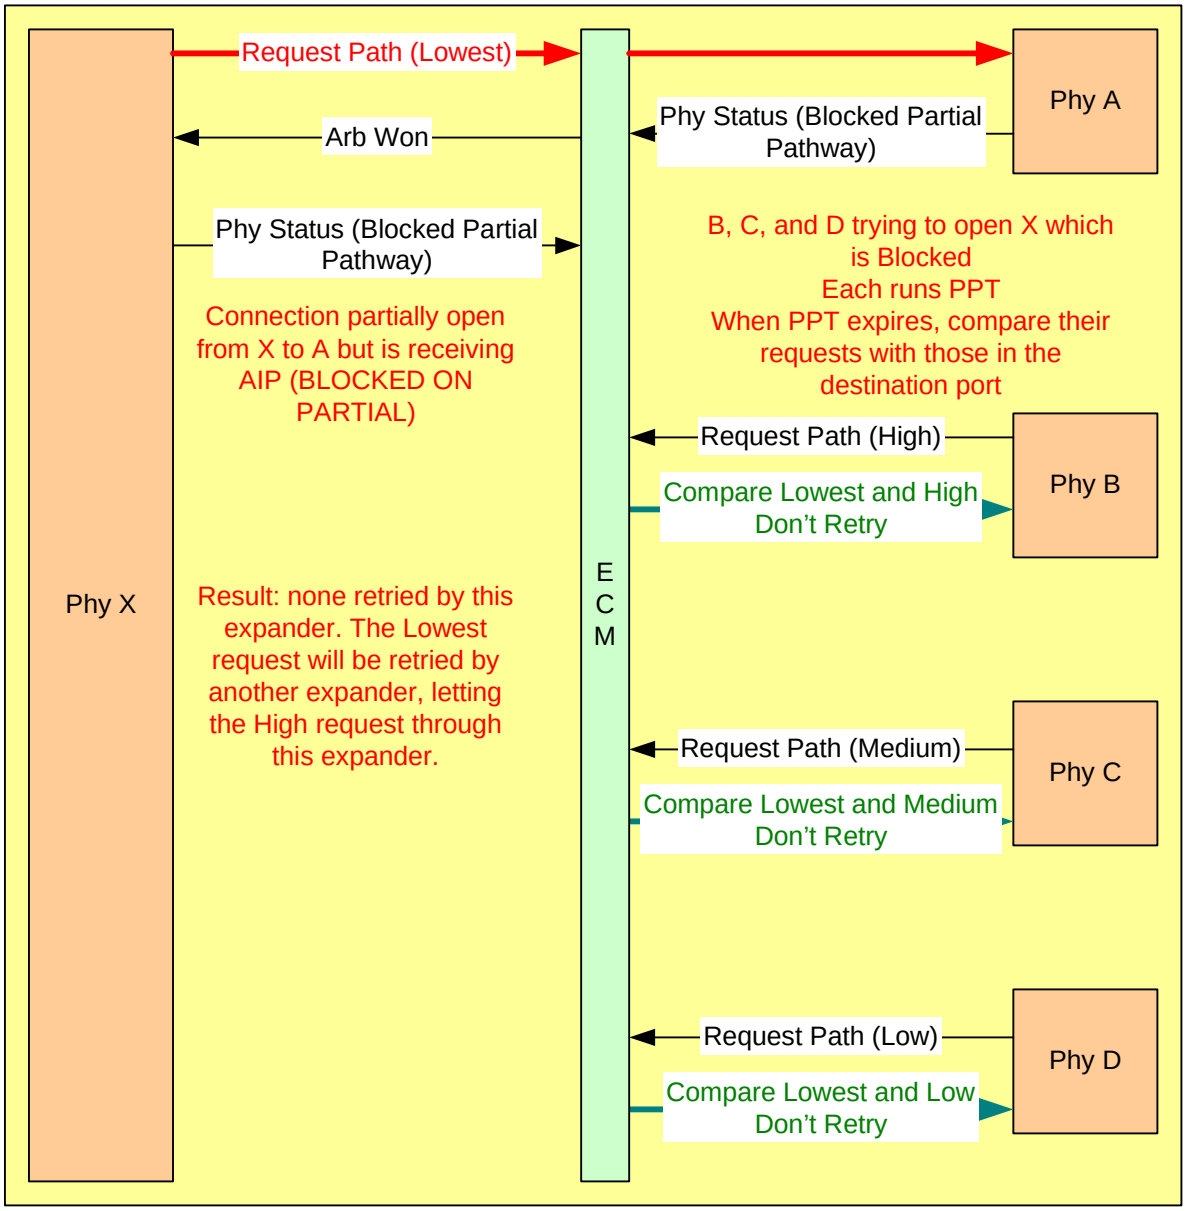

Figure 2 — Lowest priority request is blocked

Figure 3 shows the effect of comparing “all” vs any” when the destination is a wide port. The current rule is “all”; “any” is proposed instead to resolve deadlocks faster (at the expense of more retries).

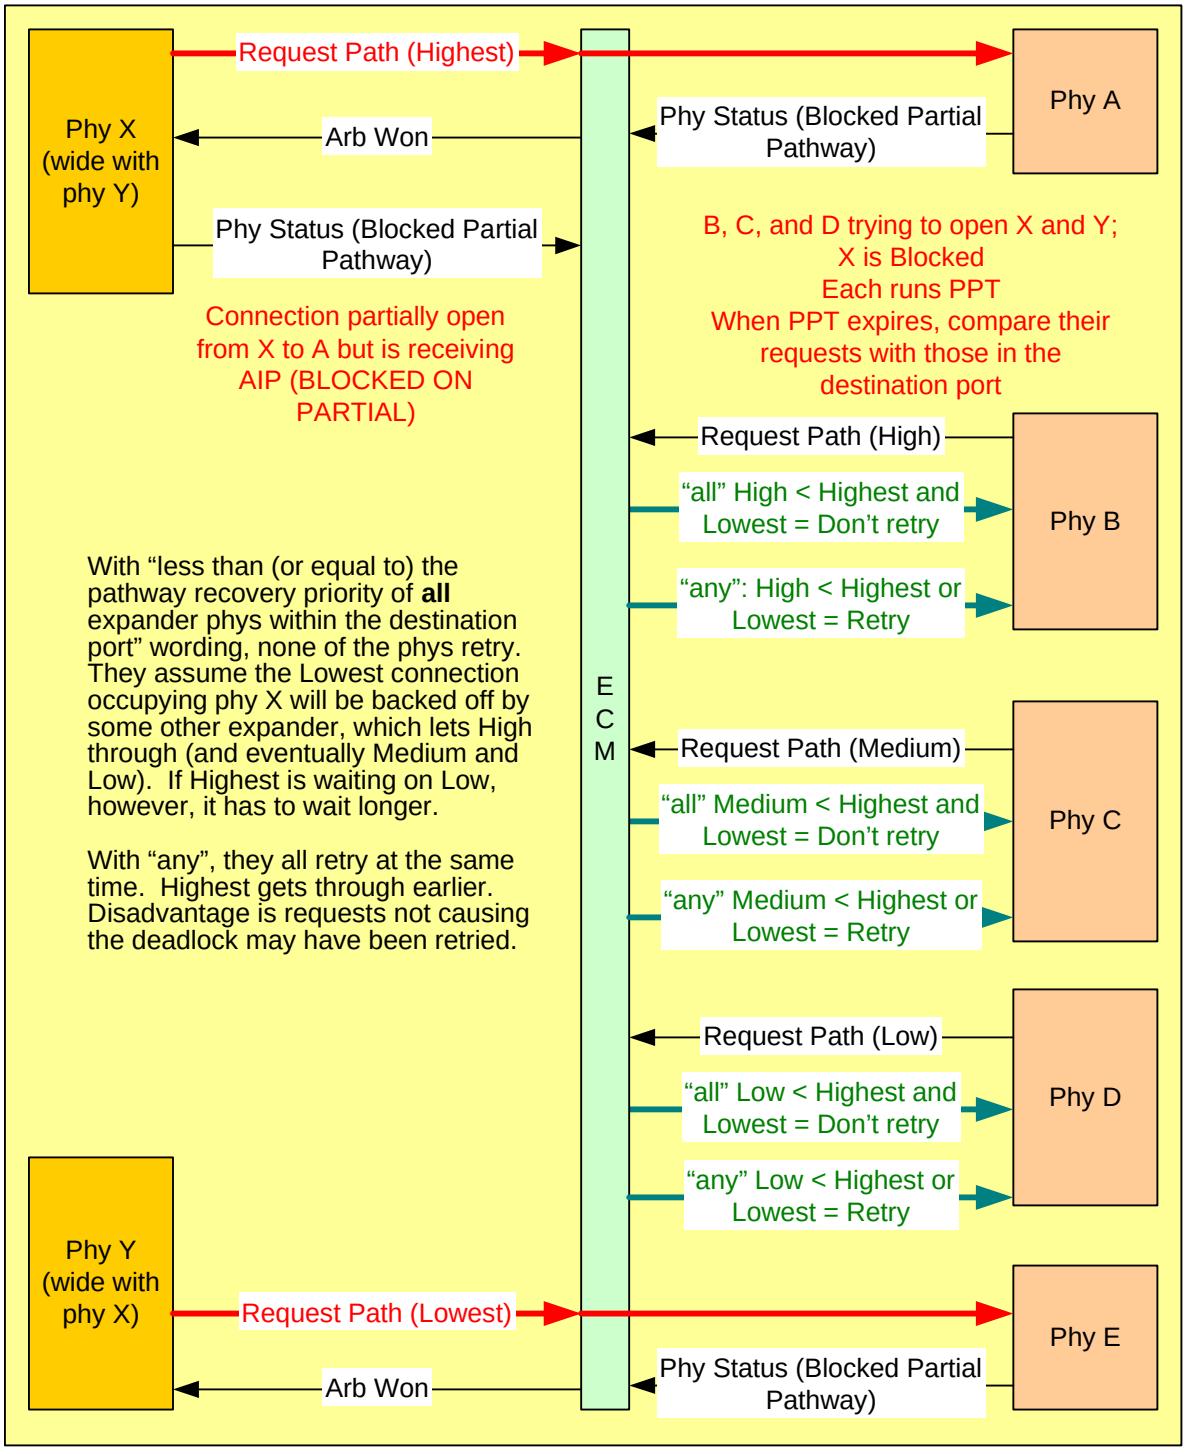

Figure 3 — All vs any

Figure 4 shows a case where the source addresses in pathway recovery priority comparisions are the same, so “less than or equal” is needed rather than “less than”.

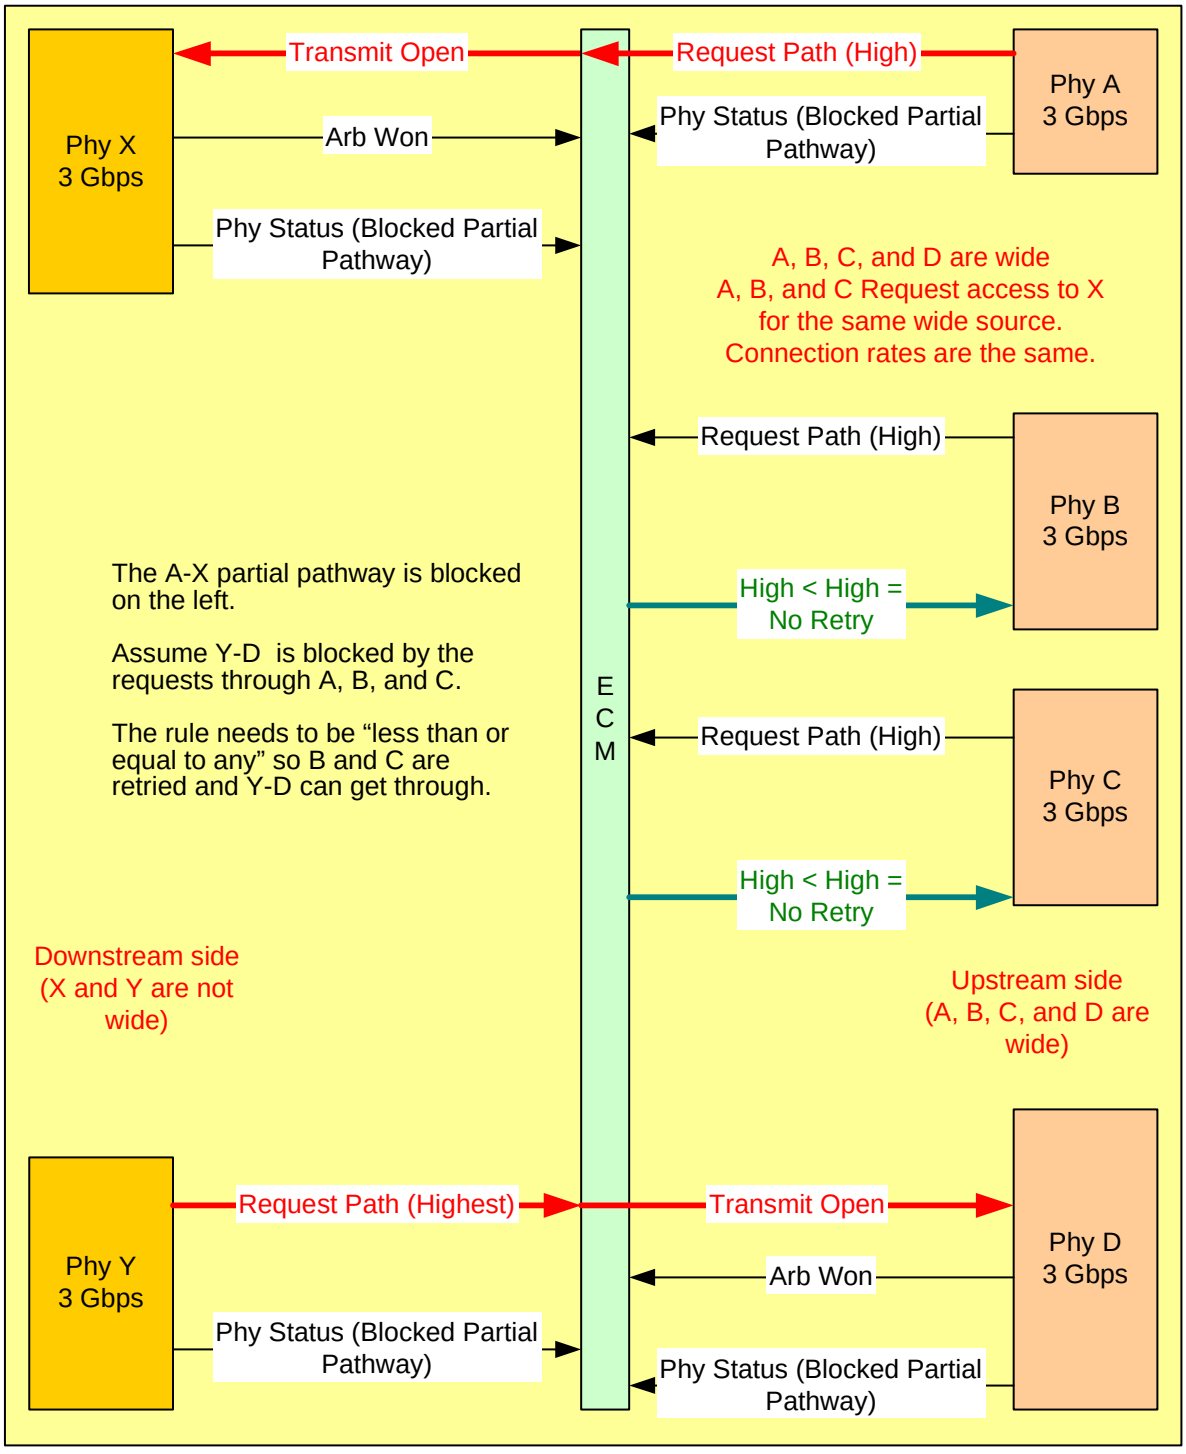

Figure 4 — Less than or equal

Figure 5 shows a case where the connection rate needs to be ignored or the connection retried might not break a deadlock. Notice this example requires 6 Gbps which is not yet defined in SAS.

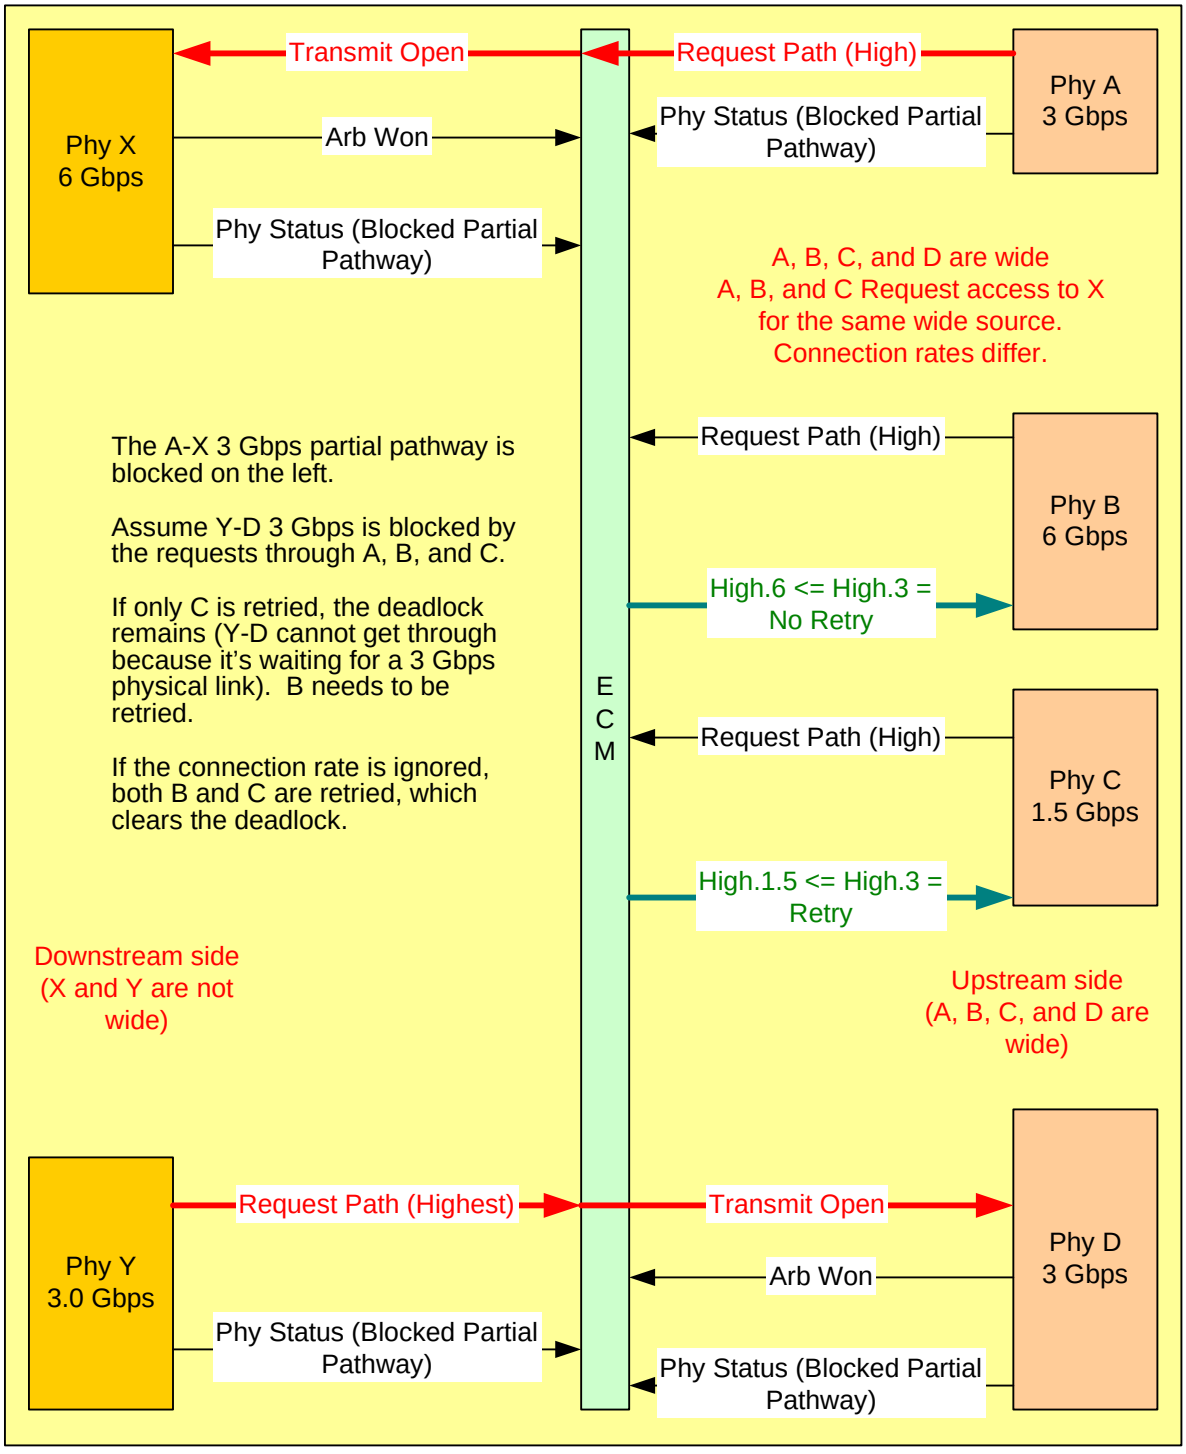

Figure 5 — Connection rate example

Figure 6 shows a case where Phy Status (Partial Pathway) and Phy Status (Blocked Partial Pathway) need to be forwarded from phy to phy.

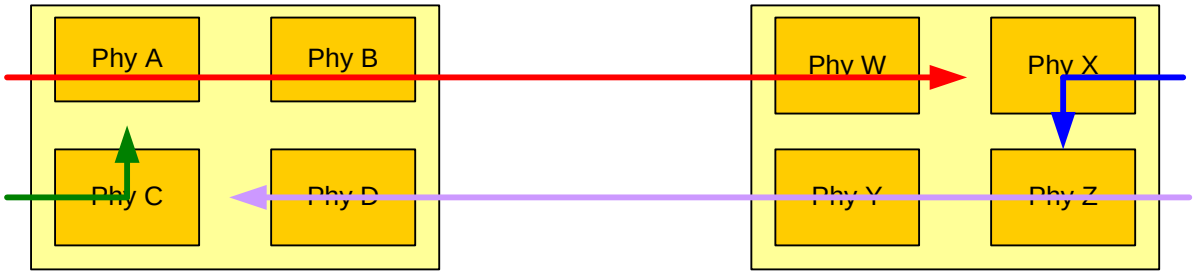

Left-side expander close-up (right side is just a mirror image of this):

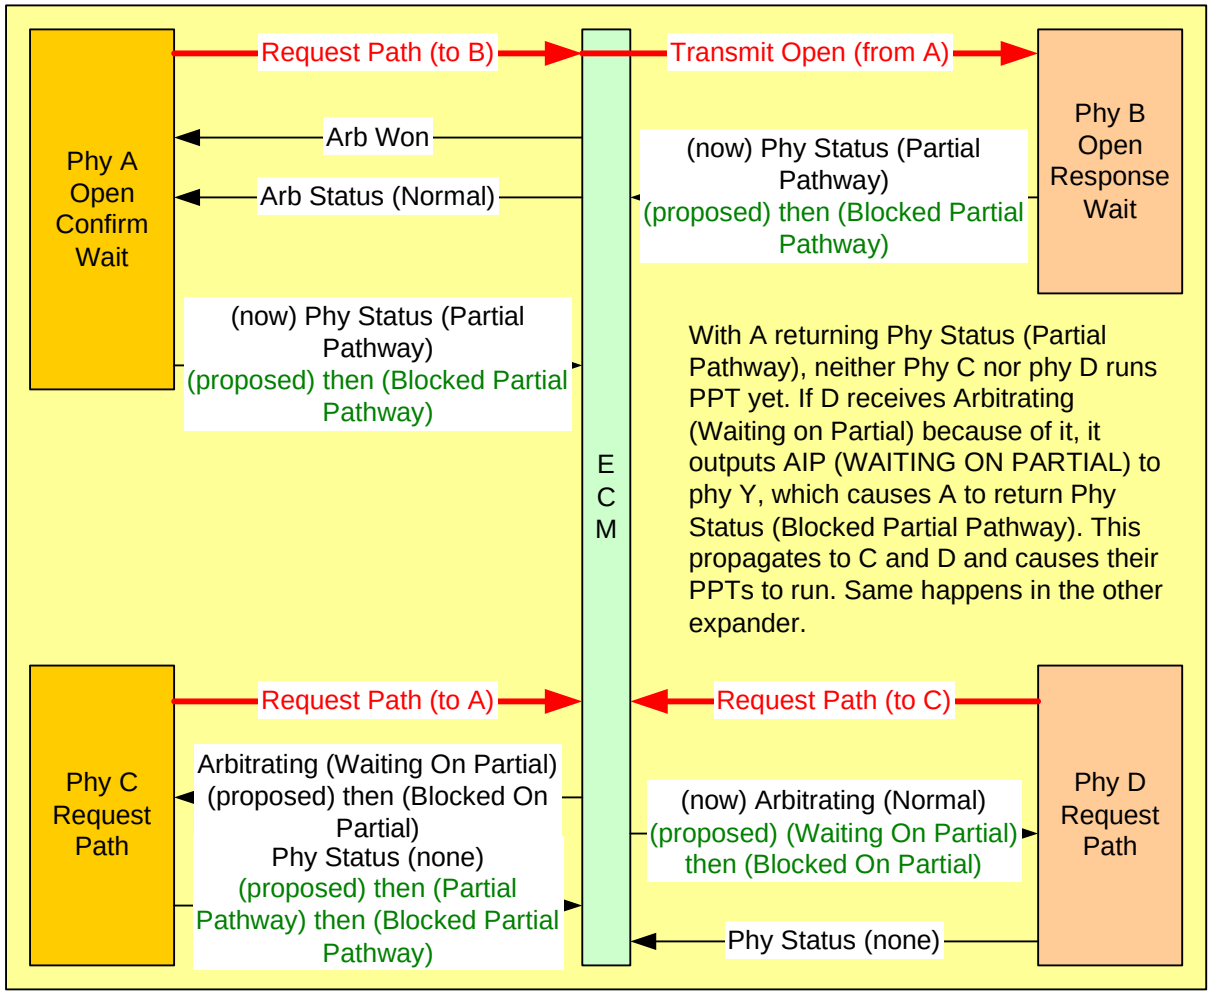

Figure 6 — Forward Phy Status from phy to phy

Suggested changes

4.6.6.3 ECM interface

Table 1 describes the requests from an expander phy to the ECM.

Table 1 — Expander phy to ECM requests

Message	Description
Request Path (arguments)	Request for a connection.

Table 2 describes the responses from an expander phy to the ECM.

Table 2 — Expander phy to ECM responses

Message	Description
Phy Status (Partial Pathway)	Response meaning that an expander phy: a) ~~contains a~~is being used for an unblocked partial pathway (i.e., the expander phy is in the XL3:Open_Confirm_Wait state or XL6:Open_Response_Wait state and has not transmitted or received an AIP (WAITING ON PARTIAL)); or b) has sent a Request Path request to the ECM and is receiving Arbitrating (Waiting On Partial) from the ECM.
Phy Status (Blocked Partial Pathway)	Response meaning that an expander phy: a) ~~contains a~~ is being used for a blocked partial pathway ~~and the most recent AIP received or indicated is AIP (WAITING ON PARTIAL).~~(i.e., the expander phy is in the XL3:Open_Confirm_Wait state or XL6:Open_Response_Wait state and has transmitted or received an AIP (WAITING ON PARTIAL)); or b) has sent a Request Path request to the ECM and is receiving Arbitrating (Blocked On Partial) from the ECM.
Phy Status (Connection~~ed~~)	Response meaning that an expander phy: a) is being used for a connection (i.e., the expander phy is in the XL7:Connected or XL8:Close_Wait state and is receiving Arbitrating (Waiting On Connection) from the ECM); b) has sent a Request Path request to the ECM and is receiving Arbitrating (Waiting On Connection) from the ECM.

Table 3 describes the confirmations from the ECM to an expander phy. These confirmations are sent in confirmation of a Request Path request.

Table 3 — ECM to expander phy confirmations

Message	Description
Arbitrating (Normal)	Confirmation that the ECM has received the Request Path request.
Arbitrating (Waiting On Partial)	Confirmation that the ECM has determined that: a) ~~at least one destination phy matches)~~ there is a destination port capable of routing to the requested destination SAS address; ~~and~~ b) at least one phy within the destination port supports the requested connection rate; c) ~~all~~each of the phys within the destination port ~~are~~ is returning a Phy Status (Partial Pathway) or Phy Status (Blocked Partial Pathway) response.; and d) at least one of the phys within the destination port is returning a Phy Status (Partial Pathway) response.
Arbitrating (Blocked On Partial)	Confirmation that the ECM has determined that: a) ~~at least one destination phy matches~~ there is a destination port capable of routing to the requested destination SAS address; ~~and~~ b) at least one phy within the destination port supports the requested connection rate; c) ~~all~~each of the phys within the destination port ~~are~~ is returning a Phy Status (Blocked Partial Pathway) response.
Arbitrating (Waiting On Connection)	Confirmation that the ECM has determined that: a) ~~at least destination phy matches~~ there is a destination port capable of routing to the requested destination SAS address; b) at least one phy within the destination port supports the requested connection rate; c) ~~no phys within the destination port are available~~each of the phys within the destination port is returning a Phy Status (Partial Pathway), Phy Status (Blocked Partial Pathway), or Phy Status (Connection) response; and d) at least one of the phys within the destination port ~~are~~ is returning a Phy Status (Connection~~ed~~) response.
Arb Won	Confirmation that an expander phy has won path arbitration.
Arb Lost	Confirmation that an expander phy has lost path arbitration.
Arb Reject (No Destination)	Confirmation that the ECM did not find an operational expander phy ~~configured to match~~capable of routing to the requested destination SAS address.
Arb Reject (Bad Destination)	Confirmation that the ECM has determined that the requested destination SAS address maps back to the requesting port.
Arb Reject (Bad Connection Rate)	Confirmation that the ECM has determined that ~~at least one destination phy matches~~there is a destination port capable of routing to the requested destination SAS address but no phys within the destination port are configured to support the requested connection rate.
Arb Reject (Pathway Blocked)	Confirmation that the ECM has determined that the requesting expander phy shall back off according to SAS pathway recovery rules.

4.6.6.4 ECR interface

Table 4 describes the requests from an expander phy to the ECR and the corresponding indications from the ECR to another expander phy.

Table 4 — Expander phy to ECR to expander phy requests and indications

Message	Description
Transmit Open (arguments)	Request/indication to transmit an OPEN address frame.
Transmit Close	Request/indication to transmit a CLOSE.
Transmit Break	Request/indication to transmit a BREAK.
Transmit Dword	Request/indication to transmit a dword.

Table 5 describes the responses from an expander phy to the ECR and the corresponding confirmations from the ECR to another expander phy. These responses are sent in response to a Transmit Open indication.

Table 5 — Expander phy to ECR to expander phy responses and confirmations

Message	Description
Arb Status (Normal)	Confirmation/response that AIP (NORMAL) has been received.
Arb Status (Waiting On Partial)	Confirmation/response that AIP (WAITING ON PARTIAL) has been received.
Arb Status (Waiting On Connection)	Confirmation/response that AIP (WAITING ON CONNECTION) has been received.
Arb Status (Waiting On Device)	Confirmation/response that AIP (WAITING ON DEVICE) has been received.
Open Accept	Confirmation/response that OPEN_ACCEPT has been received.
Open Reject	Confirmation/response that OPEN_REJECT has been received.
Backoff Retry	Confirmation/response that: a) a higher priority OPEN address frame has been received (see 7.12.3); and b) the source SAS address and connection rate of the received OPEN address frame are not equal to the destination SAS address and connection rate of the transmitted OPEN address frame.
Backoff Reverse Path	Confirmation/response that: a) a higher priority OPEN address frame has been received (see 7.12.3); and b) the source SAS address and connection rate of the received OPEN address frame are equal to the destination SAS address and connection rate of the transmitted OPEN address frame.

7.12.4.4 Pathway recovery

Pathway recovery provides a means to abort connection requests in order to prevent deadlock using pathway recovery priority comparisons. Pathway recovery priority compares the OPEN address frames of the blocked connection requests as described in table 6.

Table 6 — Pathway recovery priority

Bits ~~75-68~~ 71-64 (~~75~~ 71 is MSB)	Bits ~~67-5~~ 63-0 (0 is LSB)	~~Bits 3-0 (0 is LSB)~~
PATHWAY BLOCKED COUNT field value	SOURCE SAS ADDRESS field value	~~CONNECTION RATE field value~~

When the Partial Pathway Timeout timer for an arbitrating expander phy expires (i.e., reaches a value of zero), the ECM shall determine whether to continue the connection request or to abort the connection request.

The ECM shall instruct the arbitrating expander phy to reject the connection request by transmitting OPEN_REJECT (PATHWAY BLOCKED) when: [put into an a/b list]

a) the Partial Pathway Timeout timer expires; and
b) the pathway recovery priority of the arbitrating expander phy (i.e., the expander phy requesting the connection) is less than or equal to the pathway recovery priority of ~~all~~ any of the expander phys within the destination port ~~with an arbitration status of WAITING_ON_PARTIAL~~that are sending Phy Status (Blocked Partial Pathway) responses to the ECM.

The pathway blocked count and source SAS address values used to form the pathway recovery priority of a destination phy are those of the Request Path request if the phy sent a Request Path request to the ECM or those of the Transmit Open indication if the phy received a Transmit Open indication from the ECR.

[7.15.1 In the XL state machine figures, change Phy Status (Connected) to Phy Status (Connection)]

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 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.

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 Transmit Open indication is not being received; and
c) an OPEN Address Frame Received message is received.

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 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 Transmit Open indication is received.

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

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

[Editor's note: ...even through the Transmit Open is being serviced (it has to be compared later). Both transitions was missing the bizarre standardese way to remember that an incoming OPEN address frame arrived.]

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 a Transmit AIP (Normal) ~~parameter~~message to the XL transmitter.

If an Arbitrating (Waiting On Partial) or Arbitrating (Blocked On Partial) confirmation is received, this state shall repeatedly send a Transmit AIP (Waiting On Partial) ~~parameter~~message to the XL transmitter.

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

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

If an Arbitrating (Waiting On Connection) confirmation is received, this state shall repeatedly send a Transmit AIP (Waiting On Connection) ~~parameter~~message to the XL transmitter.

If an Arbitrating (Waiting On Connection) confirmation is received, this state shall repeatedly send a Phy Status (Connection) message 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; otherwise, 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) destination SAS address;
b) source SAS address;
c) protocol;
d) connection rate;
e) arbitration wait time;
f) initiator port bit;
g) initiator connection tag;
h) pathway blocked count;
i) partial pathway timeout status; and
j) 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, timeout status is conveyed to the expander connection manager via the partial pathway timeout status argument in the Request Path request.

7.15.4.2 Transition XL1:Request_Path to XL2:Request_Open

This transition shall occur after receiving an Arb Won confirmation.

7.15.4.3 Transition XL1:Request_Path to XL4:Open_Reject

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

7.15.4.4 Transition XL1:Request_Path to XL0:Idle

This transition shall occur after receiving an Arb Lost confirmation.

7.15.4.5 Transition XL1:Request_Path to XL9:Break

This transition shall occur receiving a BREAK Received message.

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 Transmit Open request to the ECR, received by the destination phy as a Transmit Open indication. The arguments to the Transmit Open request are:

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

7.15.5.2 Transition XL2:Request_Open to XL3:Open_Confirm_Wait

This transition shall occur after sending a Transmit Open request.

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 to 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;
f) Transmit OPEN_REJECT when an Open Reject confirmation is received with the argument from the Open Reject confirmation, after releasing path resources; 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, this state shall send a Transmit Break request to the ECR.

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

7.15.6.2 Transition XL3:Open_Confirm_Wait to XL0:Idle

This transition shall occur after sending a Transmit OPEN_REJECT message.

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.

7.15.6.4 Transition XL3:Open_Confirm_Wait to XL5:Forward_Open

This transition shall occur after receiving a Backoff Reverse Path confirmation.

7.15.6.5 Transition XL3:Open_Confirm_Wait to XL7:Connected

This transition shall occur after sending a Transmit OPEN_ACCEPT message.

7.15.6.6 Transition XL3:Open_Confirm_Wait to XL9:Break

This transition shall occur after sending a Transmit Break request.

7.15.6.7 Transition XL3:Open_Confirm_Wait to XL10:Break_Wait

This transition shall occur after receiving a Transmit Break indication.

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:

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 (Bad Connection Rate) argument is received with the transition into this state; or
d) a Transmit OPEN_REJECT (Pathway Blocked) message when an Arb Reject (Pathway Blocked) argument is received with the transition into this state.

7.15.7.2 Transition XL4:Open_Reject to XL0:Idle

This transition shall occur after OPEN_REJECT has been transmitted.

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.

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

If a BREAK Received message 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 request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

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., CHANGE Received).

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

a) an Open Accept response when an OPEN_ACCEPT Received message is received;
b) an Open Reject response when an OPEN_REJECT Received message is received, after releasing any path resources;
c) a Backoff Retry response when an AIP Received message has not been received, 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 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, after releasing path resources;
d) a Backoff Retry response when an AIP Received message has been received and an OPEN Address Frame Received message is received or an OPEN Address Frame Received argument is included in the transition into this state, and 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, after releasing path resources;
e) a Backoff Reverse Path response an AIP Received message has not been received, 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 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 when an AIP Received message has been received, an OPEN Address Frame Received message is received or an OPEN Address Frame Received argument is included in the transition into this state, and 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.

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

a) an Arb Status (Waiting On Device) response when an AIP Received message has not been received;
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 Transmit Break request to the ECR.

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

7.15.9.2 Transition XL6:Open_Response_Wait to XL0:Idle

This transition shall occur after sending an Open Reject response.

7.15.9.3 Transition XL6:Open_Response_Wait to XL1:Request_Path

This transition shall occur after sending a Backoff Retry response, after releasing path resources.

7.15.9.4 Transition XL6:Open_Response_Wait to XL2:Request_Open

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

7.15.9.5 Transition XL6:Open_Response_Wait to XL7:Connected

This transition shall occur after sending an Open Accept response.

7.15.9.6 Transition XL6:Open_Response_Wait to XL9:Break

This transition shall occur after sending a Transmit Break response.

7.15.9.7 Transition XL6:Open_Response_Wait to XL10:Break_Wait

This transition shall occur after receiving a Transmit Break indication.

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 Transmit Dword indications.

This state shall send Transmit Dword requests to the ECR containing each valid dword except BREAK and CLOSE primitives received with Dword Received messages.

If an invalid dword is received with the Dword Received message and the expander phy is forwarding to an expander phy attached to a SAS physical link, the expander phy shall send an ERROR primitive with the Transmit Dword request instead of the invalid dword.

If an invalid dword or an ERROR primitive is received with Dword Received message and the expander phy is forwarding to an expander phy attached to a SATA physical link, the expander phy shall send a SATA_ERROR primitive with the Transmit Dword request instead of the invalid dword or ERROR primitive.

If a CLOSE Received message is received, this state shall send a Transmit 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 Transmit Break request to the ECR.

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

7.15.10.2 Transition XL7:Connected to XL8:Close_Wait

This transition shall occur after receiving a Transmit Close indication.

7.15.10.3 Transition XL7:Connected to XL9:Break

This transition shall occur after sending a Transmit Break request.

7.15.10.4 Transition XL7:Connected to XL10:Break_Wait

This transition shall occur after receiving a Transmit Break indication.

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 Transmit Close indication, then shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

If a Dword Received message is received containing a valid dword except a BREAK or CLOSE primitive, this state shall send Transmit Dword requests to the ECR containing that dword.

If an invalid dword is received with the Dword Received message and the expander phy is forwarding to an expander phy attached to a SAS physical link, the expander phy shall send an ERROR primitive with the Transmit Dword request instead of the invalid dword.

If an invalid dword or an ERROR primitive is received with Dword Received message and the expander phy is forwarding to an expander phy attached to a SATA physical link, the expander phy shall send a SATA_ERROR primitive with the Transmit Dword request instead of the invalid dword or ERROR primitive.

If a CLOSE Received message is received, this state shall release path resources and send a Transmit 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 Transmit Break request to the ECR.

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

7.15.11.2 Transition XL8:Close_Wait to XL0:Idle

This transition shall occur after sending a Transmit Close request.

7.15.11.3 Transition XL8:Close_Wait to XL9:Break

This transition shall occur after sending a Transmit Break request.

7.15.11.4 Transition XL8:Close_Wait to XL10:Break_Wait

This transition shall occur after a Transmit Break indication is received.

7.15.12 XL9:Break state

7.15.12.1 State description

This state closes any connection and releases path resources.

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

7.15.12.2 Transition XL9:Break to XL0:Idle

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

7.15.13 XL10:Break_Wait state

7.15.13.1 State description

This state closes any connection and releases path resources.

This state shall send a Transmit BREAK message to the XL transmitter. After transmitting the BREAK this state shall initialize and start the Break Timeout timer.

7.15.13.2 Transition XL10:Break_Wait to XL0:Idle

This transition shall occur after a BREAK Received message is received or after the Break Timeout timer expires.