To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@intel.com)
Date: 17 April 2007
Subject: 07-189r0: SAS-2: Allow omission of AIP for internal destination ports

Revision history

Revision 0 (17 April 2007) First revision

Related documents

SAS-2-r09 - Serial Attached SCSI - 2 (SAS-2) revision 09

Overview

The last paragraph in 7.12.4.3 Arbitration status states that, upon receiving an OPEN address frame:

> Expander devices shall transmit an AIP (e.g., an AIP (NORMAL)) within 128 dwords of receiving an OPEN address frame.

This is a signal from the expander device receiving the OPEN address frame to the expander device that transmitted the OPEN address frame that the possible condition of OPEN address frames crossing on the wire no longer exists, and that responsibility for further arbitration (if any) is held by the expander device that received the OPEN address frame (i.e., the expander device that transmitted the OPEN address frame can release arbitration-related resources like the AWT).

A SAS end device that receives an OPEN address frame does return an AIP because it will return either an OPEN_ACCEPT or OPEN_REJECT, and does not need to perform further arbitration to secure a pathway to the destination port.

There are cases where the destination of an OPEN address frame is within an expander device (e.g., an integrated SES controller, or an SMP target port). When an OPEN address frame is received with the destination address specifying an integrated port, there is no need to transmit an AIP after the OPEN address frame is received.

This proposal suggests SAS-2 changes to allow an expander device to omit transmitting an AIP if the destination SAS port of the OPEN address frame is within the expander device and there is no subsequent arbitration for internal pathways required to establish a connection to the destination SAS port.

Suggested changes

Modify subclause 7.2.5.1 as follows:

7.2.5.1 AIP (Arbitration in progress)

AIP is sent by an expander device after a connection request while arbitrating for a pathway to specify that the connection request is being processed and specify the status of the connection request.

The versions of AIP representing different statuses are defined in table 100.

Table 100 — AIP primitives

Primitive	Description
AIP (NORMAL)	Expander device has accepted the connection request. This may be sent zero, one, or multiple times (see 7.12.4.3).
AIP (RESERVED 0)	Reserved. Processed the same as AIP (NORMAL).
AIP (RESERVED 1)	
AIP (RESERVED 2)	
AIP (WAITING ON CONNECTION)	Expander device has determined the routing for the connection request, but either the destination phys are all being used for connections or there are insufficient routing resources to complete the connection request. This may be sent multiple times (see 7.12.4.3).
AIP (WAITING ON DEVICE)	Expander device has determined the routing for the connection request and forwarded it to the output physical link. This is sent one time (see 7.12.4.3).
AIP (WAITING ON PARTIAL)	Expander device has determined the routing for the connection request, but the destination phys are all busy with other partial pathways. This may be sent multiple times (see 7.12.4.3).
AIP (RESERVED WAITING ON PARTIAL)	Reserved. Processed the same as AIP (WAITING ON PARTIAL).

See 7.12 for details on connections.

Modify subclause 7.12.4.3as follows:

7.12.4.3 Arbitration status

Arbitration status shall be conveyed between expander devices and by expander devices to SAS endpoints using AIP (see 7.2.5.1). This status is used to monitor the progress of connection attempts and to facilitate pathway recovery as part of deadlock recovery.

The arbitration status of an expander phy is set to the last type of AIP received.

Before an expander device transmits AIP, it may have transmitted an OPEN address frame on the same physical link. Arbitration fairness dictates which OPEN address frame wins (see 7.12.3).

After an expander device transmits an AIP, it shall not transmit an OPEN address frame unless it has higher arbitration priority than the incoming connection request.

After transmitting an AIP primitive sequence, and expander device shall transmit at least one other dword (e.g., an idle dword) before transmitting another AIP primitive sequence.

Expander devices shall transmit at least one AIP every 128 dwords while transmitting AIP (NORMAL), AIP (WAITING ON PARTIAL), or AIP (WAITING ON CONNECTION).

NOTE 1 - Expander devices compliant with previous versions of this standard were not required to transmit three consecutive AIP primitives, as AIP was defined as a single primitive sequence (see 7.2.4.2) rather than as an extended primitive sequence (see 7.2.4.5).

When arbitrating for access to an outgoing expander port, expander ~~Expander~~ devices shall transmit an AIP (e.g., an AIP (NORMAL)) within 128 dwords of receiving an OPEN address frame. If the destination SAS address is within the expander device (e.g., an SMP target port) and there is a pathway to the destination SAS port that is not blocked by an active connection, the expander device may return OPEN_ACCEPT or OPEN_REJECT without transmitting an AIP.

Modify subclause 7.15.9 as follows:

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. During an STP connection, the expander device may expand or contract a repeated or continued primitive sequence (see 7.2.4).

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 if the OPEN address frame was transmitted to an external expander port;
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.

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.