

T10/04-029r2, SAS-1.1: Close Port Selector Detection Gap

To: T10 Technical Committee
From: Robert Sheffield (Robert.L.Sheffield@intel.com), Intel Corporation
Date: May 12, 2004
Subject: T10/04-029r2, SAS-1.1: Close Port Selector Detection Gap

Revision History

Revision 0 (December 29, 2003): first revision

Revision 1 (May 4, 2004): Use BROADCAST(CHANGE) interlock and a new LINK RATE NOT NEGOTIATED value for the NEGOTIATED PHYSICAL LINK RATE instead to establish positive reporting of port selector presence.

Revision 2 (May 12, 2004): Send BROADCAST(CHANGE) only when the ATTACHED SATA PORT SELECTOR bit changes state and clear the bit only on power-on or a transition from SP1 to SP0 when no COMWAKE is detected.

Related Documents

SAS1.1-r04 – Serial Attached SCSI-1.1 revision 04

Overview

The current text says that the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is cleared each time the SAS Phy state machine returns to the SP0:COMINIT state. This occurs in response to a reset event, or in response to a hot-plug timeout event (every 500 ms when no device is present). If a SATA port selector is present, during the short time between a transition to SP0:COMINIT and the subsequent detection of a SATA port selector (receipt of COMWAKE in response to the initial COMINIT), an SMP DISCOVER request may arrive, with the corresponding DISCOVER response indicating, incorrectly, that no SATA port selector is present.

There may be other information in the DISCOVER response that could be misinterpreted by an SMP initiator without knowing something about the state of the reported phy.

The suggested solution is to add another value to the NEGOTIATED PHYSICAL LINK RATE field in the DISCOVER response, 4h: Phy enabled, to indicate that the phy is enabled and has detected the presence of a SATA port selector, but has not yet achieved dword sync for any supported link rate. When this state exists, an SMP management application may recognize the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response indicates the presence of a Port Selector, and this information is valid whether the phy indicated is active or inactive. Also, to facilitate an effective interlock between the phy state and the SMP application, text is added to stipulate that a phy shall generate a BROADCAST(CHANGE) whenever a port selector appearance or removal is positively detected. Furthermore, text in the SP state machine OOB sequence is modified so that the ATTACHED SATA PORT SELECTOR bit is only set or cleared when the SP state machine makes a definitive determination about the presence or non-presence of a SATA port selector. The bit is cleared only upon a power-on or a hot-plug timeout transition from SP1 to SP0 and the bit was set upon entry to SP1 and no COMWAKE signal was detected while in SP1. Likewise the bit is only set [and BROADCAST(CHANGE) sent] when the bit was not previously set, but is set upon COMWAKE received in the SP0, SP1, SP3, or SP4 states.

Suggested Changes

Modify table 14 in subclause 4.6.6.5 BPP interface as follows:

Table 14 — Expander phy to BPP requests

Message	Description
Broadcast Event Notify (Phy Not Ready)	Request to transmit a BROADCAST (CHANGE) on all other ports because an expander phy's SP state machine transitioned from the SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready state to the SP0:OOB_COMINIT state (see 6.7) or because a virtual phy has been disabled (see 10.4.3.10). See 7.11.
Broadcast Event Notify (SATA Spinup Hold)	Request to transmit a BROADCAST (CHANGE) on all other ports because the SATA spinup hold state has been reached (see 6.7 and 6.9). See 7.11.
Broadcast Event Notify (Identification Sequence Complete)	Request to transmit a BROADCAST (CHANGE) on all other ports because an expander phy has completed the identification sequence (see 7.9) or because a virtual phy has been enabled (see 10.4.3.10). See 7.11.
Broadcast Event Notify (CHANGE Received)	Request to transmit a BROADCAST (CHANGE) on all other ports because a BROADCAST (CHANGE) was received.
Broadcast Event Notify (SATA Port Selector Change)	Request to transmit a BROADCAST (CHANGE) on all other ports because a SATA port selector appeared or was removed.
Broadcast Event Notify (RESERVED CHANGE Received)	Request to transmit a BROADCAST (RESERVED CHANGE) on all other ports because a BROADCAST (RESERVED CHANGE) was received.
Broadcast Event Notify (SES Received)	Request to transmit a BROADCAST (SES) on all other ports because a BROADCAST (SES) was received.
Broadcast Event Notify (RESERVED 1 Received)	Request to transmit a BROADCAST (RESERVED 1) on all other ports because a BROADCAST (RESERVED 1) was received.

Modify subclause 6.7.3.2.1 to clear ATTACHED SATA PORT SELECTOR only on power up and to send the SATA Port Selector Change confirmation only if the ATTACHED SATA PORT SELECTOR bit was clear before receiving COMWAKE received:

6.7.3.2 SP0:OOB_COMINIT state

6.7.3.2.1 State description

This state is the initial state for this state machine.

Upon entry into this state, this state shall:

- a) send a Transmit COMINIT message to the SP transmitter;
- b) send a Stop DWS message to the SP_DWS state machine; and
- c) send a Phy Layer Not Ready confirmation to the link layer.

This state machine waits for receipt of a COMINIT Transmitted message and/or a COMINIT Detected message.

Any transition Entry to this state *due to a power on* shall set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to zero. If this state receives a COMWAKE Detected message and the phy supports attachment to a SATA device (i.e. the phy is attached to an STP/SATA bridge) and also supports attachment to a SATA port selector, *and the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero*, it shall set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to one *and send a SATA Port Selector Change confirmation to the link layer*.

Modify subclause 6.7.3.3.1 to send the SATA Port Selector Change confirmation only if the ATTACHED SATA PORT SELECTOR bit was clear before receiving COMWAKE received:

6.7.3.3 SP1:OOB_AwaitCOMX state

6.7.3.3.1 State description

Upon entry into this state, the Hot-Plug Timeout timer shall be initialized and started if this phy is:

- a) an expander phy; or
- b) an initiator phy or target phy implementing the Hot-Plug Timeout timer.

If this state receives COMWAKE Detected message and the phy supports attachment to a SATA device (i.e. the phy is attached to an STP/SATA bridge) and supports attachment to a SATA port selector, *and the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero*, it shall set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to one *and send a SATA Port Selector Change confirmation to the link layer*.

Modify subclause 6.7.3.3.2 to clear the ATTACHED SATA PORT SELECTOR bit and send the SATA Port Selector Change confirmation if the bit was not set on entry to SP1, and no COMWAKE was received while in SP1:

6.7.3.3.2 Transition SP1:OOB_AwaitCOMX to SP0:OOB_COMINIT

This transition shall occur if the Hot-Plug Timeout timer expires. *If the phy supports attachment to a SATA device (i.e. the phy is attached to an STP/SATA bridge) and supports attachment to a SATA port selector, and the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response was one upon entry to the SP1:OOB_AwaitCOMX state, and no COMWAKE detected message was received while in the SP1:OOB_AwaitCOMX state, this transition shall set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to zero and send a SATA Port Selector Change confirmation to the link layer.*

Modify subclause 6.7.3.5.1 to send the SATA Port Selector Change confirmation only if the ATTACHED SATA PORT SELECTOR bit was clear before receiving COMWAKE received:

6.7.3.5 SP3:OOB_AwaitCOMINIT_Sent state

6.7.3.5.1 State description

This state waits for a COMINIT Transmitted message.

If this state receives COMWAKE Detected message and the phy supports attachment to a SATA device (i.e. the phy is attached to an STP/SATA bridge) and supports attachment to a SATA port selector, and the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero, it shall set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to one and send a SATA Port Selector Change confirmation to the link layer.

Modify subclause 6.7.3.6.1 to send the SATA Port Selector Change confirmation only if the ATTACHED SATA PORT SELECTOR bit was clear before receiving COMWAKE received:

6.7.3.6 SP4:OOB_COMSAS state

6.7.3.6.1 State description

Upon entry into this state, this state shall send a Transmit COMSAS message to the SP transmitter.

This state waits for receipt of a COMSAS Transmitted message and/or a COMSAS Detected message.

If this state receives COMWAKE Detected message and the phy supports attachment to a SATA device (i.e. the phy is attached to an STP/SATA bridge) and supports attachment to a SATA port selector, and the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero, it shall set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to one and send a SATA Port Selector Change confirmation to the link layer.

Modify subclause 7.11 as follows:

7.11 SAS domain changes

After power on or receiving BROADCAST (CHANGE), an application client in each SAS initiator port should scan the SAS domain using the discover process (see 4.6.7.4) to search for SAS initiator devices, SAS target devices, and expander devices.

The expander device shall transmit BROADCAST (CHANGE) from at least one phy in each expander port other than the expander port that is the cause for transmitting BROADCAST (CHANGE).

Expander devices shall transmit BROADCAST (CHANGE) for the following reasons:

- a) after an expander phy's SP state machine transitions from the SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready state to the SP0:OOB_COMINIT state (see 6.7);
NOTE 23 - this occurs when the expander phy is reset or disabled with the SMP PHY CONTROL function (see 10.4.3.10) as well as when dword synchronization is unexpectedly lost;
- b) after a virtual phy has been disabled (see 10.4.3.10);
- c) after the SATA spinup hold state has been reached (see 6.9);
- d) after an expander phy's SP state machine receives a COMWAKE detected message in states SP0, SP1, SP3, or SP4 if the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero prior to receiving the COMWAKE detected message (see 6.7.3 and 10.4.3.5, table 157);
- e) after an expander phy's SP state machine transitions from the SP1:OOB_AwaitCOMX state to the SP0:OOB_COMINIT state if the value of the SATA PORT SELECTOR bit in the DISCOVER response was one upon entry to SP1:OOB_AwaitCOMX, and if no COMWAKE detected message was received while in SP1:OOB_AwaitCOMX before the transition to SP0:OOB_COMINIT (see 6.7.3.3.2);

T10/04-029r2, SAS-1.1: Close Port Selector Detection Gap

- d)f) after the link reset sequence completes (see 7.9);
- e)g) after a virtual phy has been enabled (see 10.4.3.10);
- f)h) after a self-configuring expander device has completed configuration and has changed its CONFIGURING bit from one to zero in the SMP REPORT GENERAL function (see 10.4.3.3); and
- g)i) after the expander device receives BROADCAST (CHANGE).

BROADCAST (CHANGE) may be sent by SAS initiator ports to force other SAS initiator ports and expander ports to re-run the discover process, but should not be sent by SAS target ports.

A SAS initiator port that detects BROADCAST (CHANGE) shall follow the SAS initiator device rules (see 7.9.2) to discover the topology.

A fanout expander device that detects BROADCAST (CHANGE) shall follow the fanout device rules (see 7.9.3) to discover the topology.

An edge expander device that detects BROADCAST (CHANGE) shall follow the edge device rules (see 7.9.4). See 10.4.3.3 for details on counting BROADCAST (CHANGE) generation in an expander device.

Modify table 155 as follows:

Table 155 — Negotiated physical link rate

Code	Description
0h	Phy is enabled; unknown physical link rate. ^a
1h	Phy is disabled.
2h	Phy is enabled; the phy obtained dword synchronization for at least one physical link rate during speed negotiation (either SAS or SATA), but the speed negotiation sequence failed (i.e., the last speed negotiation window, using a physical link rate expected to succeed, failed). These failures may be logged in the SMP REPORT PHY ERROR LOG function (see 10.4.3.6) and/or the Protocol-Specific log page (see 10.2.7.1).
3h	Phy is enabled; detected a SATA device and entered the SATA spinup hold state. The LINK RESET and HARD RESET operations in the SMP PHY CONTROL function (see 10.4.3.10) may be used to release the phy. This field shall be updated to this value after the SATA OOB sequence completes if SATA spinup hold is supported.
4h	Phy is enabled; detected a SATA Port Selector device attached to the phy. The link rate has not been negotiated since the last time the phy's SP state machine entered the SP0:OOB_COMINIT state. The SATA spinup hold state has not been entered since the last time the phy's SP state machine entered the SP0:OOB_COMINIT state. The value in this field may change to 2h, 3h, 8h, or 9h. Presence of a SATA port selector shall be indicated by the ATTACHED SATA PORT SELECTOR bit in the SMP DISCOVER response.
8h	Phy is enabled; This field shall be updated to this value after the speed negotiation sequence completes indicating a negotiated physical link rate of 1,5 Gbps.
9h	Phy is enabled; This field shall be updated to this value after the speed negotiation sequence completes indicating a negotiated physical link rate of 3,0 Gbps.
All others	Reserved.
^a This code may be used by an application client in its local data structures to indicate an unknown negotiated physical link rate (e.g., before the discover process has queried the phy).	

Modify Table 157 as follows:

Table 157 — ATTACHED SATA PORT SELECTOR and ATTACHED SATA DEVICE bits

ATTACHED SATA PORT SELECTOR bit value ^{a b}	ATTACHED SATA DEVICE bit value ^{c d}	Description
0	0	Neither a SATA port selector nor a SATA device is attached and ready on the selected phy.
0	1	The attached phy is a SATA device phy. No SATA port selector is present (i.e., the SP state machine did not detect COMWAKE in response to the initial COMINIT, but sequenced through the normal (non-SATA port selector) SATA device OOB sequence).
1	0	The attached phy is a SATA port selector host phy, and either: a) the attached phy is the inactive host phy, or b) the attached phy is the active host phy and a SATA device is either not present or not ready behind the SATA port selector (i.e., the SP state machine detected COMWAKE while waiting for COMINIT).
1	1	The attached phy is a SATA port selector's active host phy and a SATA device is present behind the SATA port selector (i.e., the SP state machine detected COMWAKE while waiting for COMINIT, timed out waiting for COMSAS, and exchanged COMWAKE with an attached SATA device).

^a This bit is invalid if the value in NEGOTIATED PHYSICAL LINK RATE is 0h or 1h.
^b Whenever this bit changes, the phy shall generate a BROADCAST(CHANGE) notification.
^{ac} For the purposes of this bit, the SATA port selector is not considered a SATA device.
^{bd} This bit shall be updated after the SATA OOB sequence completes and before the SATA speed negotiation sequence begins (i.e., at SATA spin-up hold time (see 6.9)).

Modify figure 90 in subclause 7.15.1 as follows:

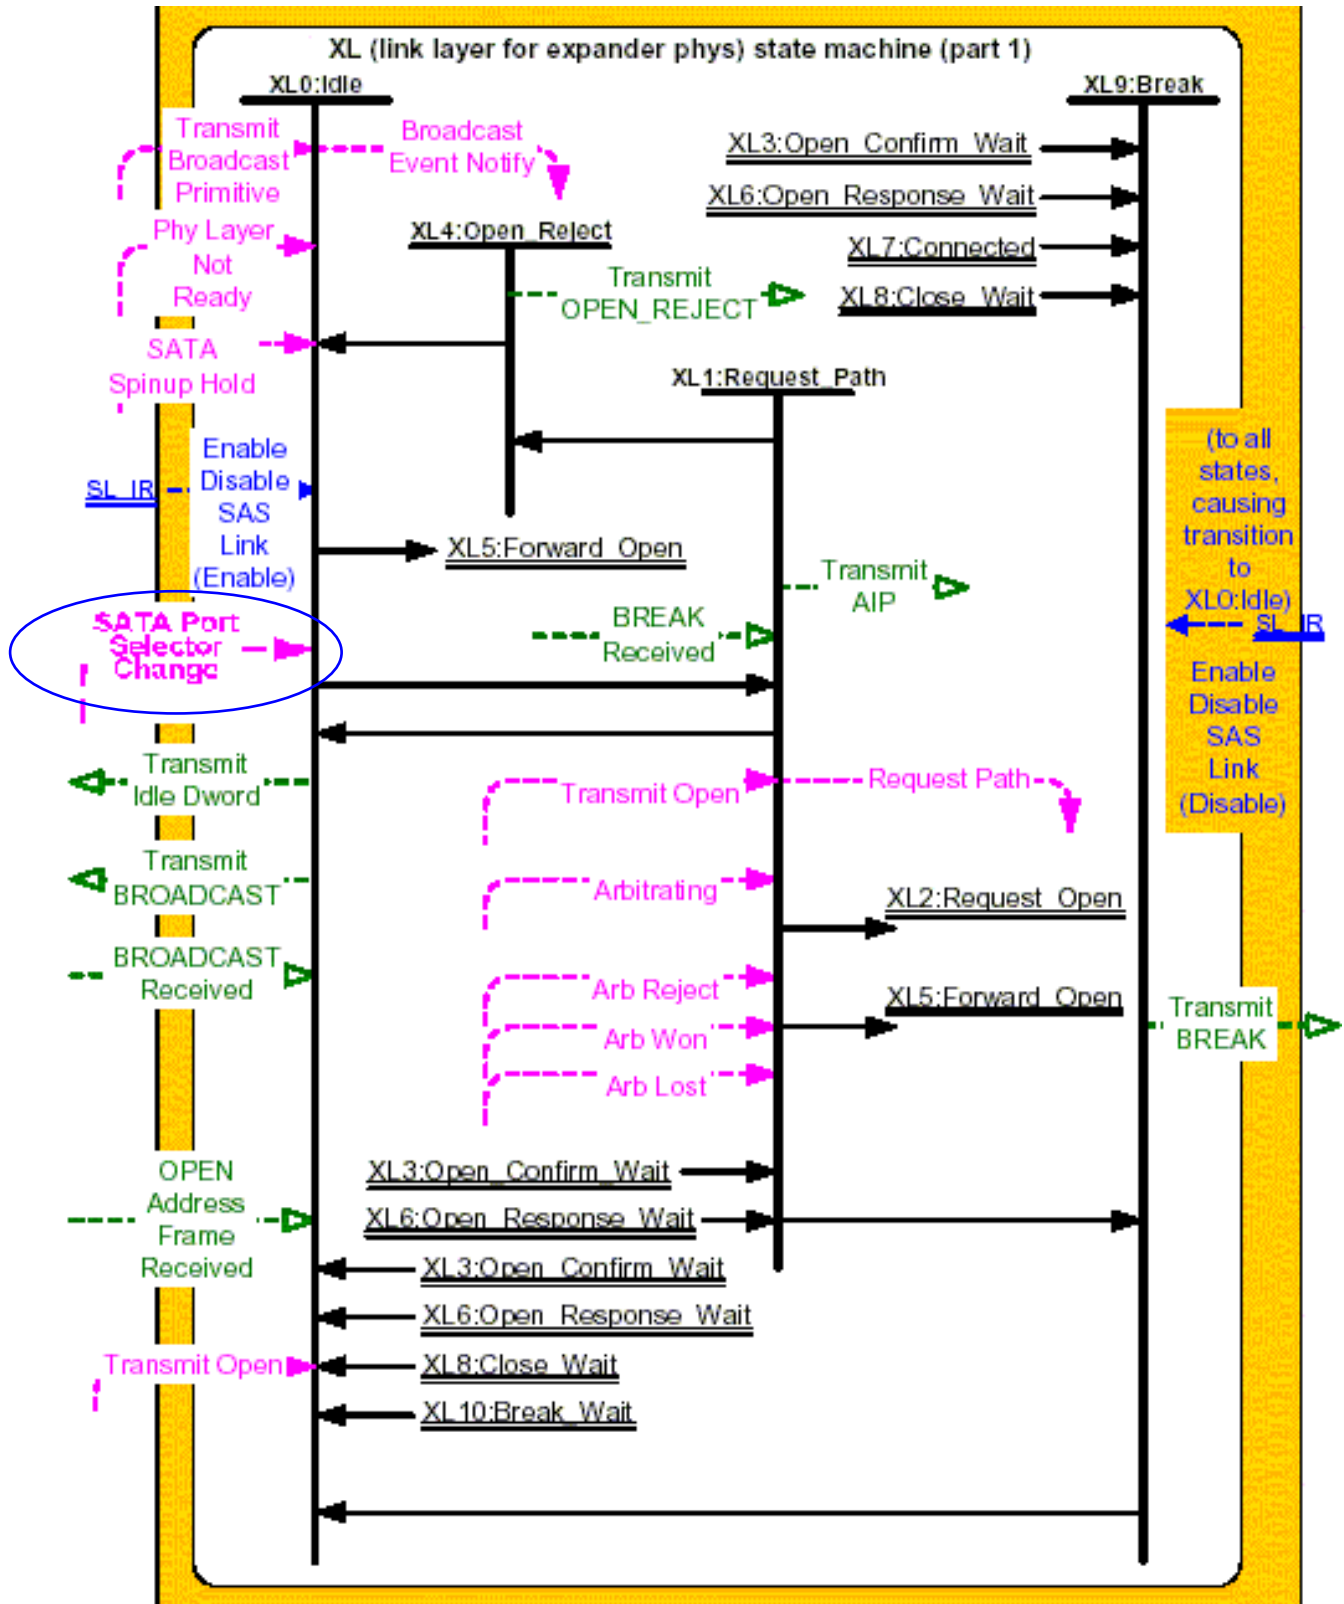


Figure 90 — XL (link layer for expander phys) state machine (part 1)

Modify subclause 7.15.3.1 as follows:

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.