

**To: T10 SAS Protocol Working Group**

**From: Brian Day (LSI)**

**Subject: SAS 2: Only issue BROADCAST(CHANGE) for initial Spinup Hold**

### **Revision History**

Revision 0 - Initial draft (September 12, 2007)

Revision 1 - Incorporates comments from working group meeting (September 19, 2007)

### **Related Documents**

sas2r11 - Serial Attached SCSI - 2 Draft revision 11

06-478r2 - SAS 2: Changes for the NEGOTIATED PHYSICAL LINK RATE revision 2

### **Overview**

There is currently a rule that expanders shall issue a BROADCAST(CHANGE) each time the phy SP state machine reaches the SP26: SATA\_SpinupHold state. There is also a statement in section 4.7.1 about the discovery process that says:

“The discover process may be aborted prior to completion and restarted if there is an indication that it may be based on incorrect information (e.g., reception of a Broadcast (Change) or a change in the EXPANDER CHANGE COUNT field returned in an SMP response).”

However, SATA drives that implement Asynchronous Signal Recovery (ASR), will generate COMINIT signals very frequently (usually around 10ms), and cause the SP state machine to cycle back through the SP0 to SP26 states. As currently defined, this causes a large number of BROADCAST(CHANGE) primitives to be transmitted by the expander. If the discover process is aborted each time these BROADCASTs are received, the discover process could potentially not complete.

This proposal modifies the rule of when expanders generate BROADCAST(CHANGE) so that it occurs only for the “initial” entry into SP26: SATA\_SpinupHold state.

### **Proposed Changes**

## 6.8.7 SATA spinup hold state SP26:SATA\_SpinupHold

### 6.8.7.1 State description

Figure 1 shows the SP26:SATA\_SpinupHold state.

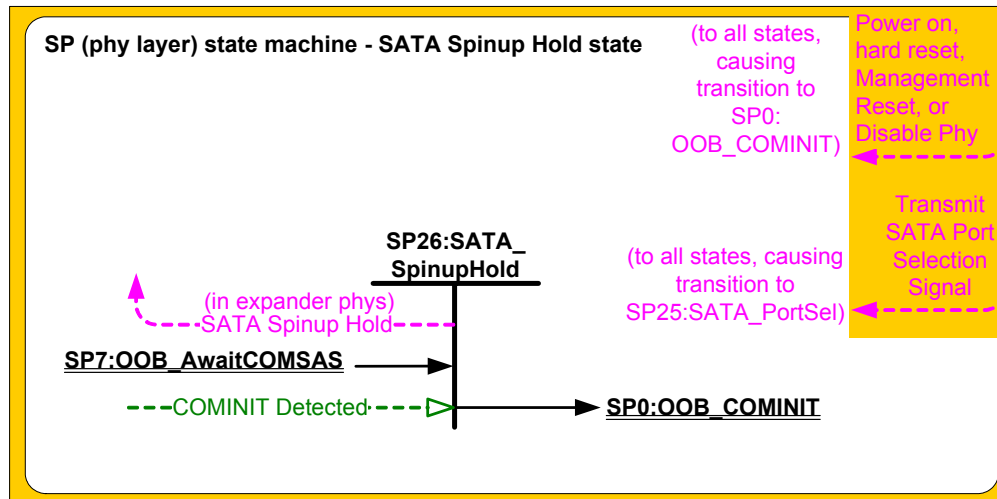


Figure 1 — SP (phy layer) state machine - SATA spinup hold state

~~Upon entry into this state, this state shall set the ResetStatus state machine variable to SPINUP\_HOLD.~~

~~If this state machine is in an expander phy, this state shall send a SATA Spinup Hold confirmation to the link layer.~~

Upon entry into this state, this state shall:

- a) if the ResetStatus state machine variable is set to SPINUP\_HOLD, not change the ResetStatus state machine variable; or
- b) if the ResetStatus state machine variable is not set to SPINUP\_HOLD:
  - A) set the ResetStatus state machine variable to SPINUP\_HOLD; and
  - B) if this state machine is in an expander phy, send a SATA Spinup Hold confirmation to the link layer.

### 6.8.7.2 Transition SP26:SATA\_SpinupHold to SP0:OOB\_COMINIT

This transition shall occur if this state receives a COMINIT Detected message.

## 7.11 SAS domain changes (Broadcast (Change) usage)

After power on or receiving Broadcast (Change) via an SMP initiator port, the management application client should scan the SAS domain using the discover process (see 4.7) to search for SAS initiator devices, SAS target devices, and expander devices.

The expander device shall originate Broadcast (Change) from at least one phy in each expander port other than the expander port that is the cause for originating Broadcast (Change).

Expander devices shall originate Broadcast (Change) for the following expander phy-related 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 6.8);

NOTE 1 - This occurs when the expander phy is reset or disabled with the SMP PHY CONTROL function DISABLE, LINK RESET, HARD RESET, or TRANSMIT SATA PORT SELECTION SIGNAL phy operations (see 10.4.3.28) as well as when dword synchronization is unexpectedly lost;

- b) after a virtual phy has been disabled with the SMP PHY CONTROL function DISABLE phy operation or started processing a reset requested by the LINK RESET or HARD RESET phy operations (see 10.4.3.28);
- c) after an expander phy's SP state machine ~~reaches the SATA spinup hold state (see 6.8.7 and 6.11)~~ reaches the SP26:SATA\_SpinupHold state and sends a SATA Spinup Hold confirmation as defined in 6.8.7 and 6.11;
- d) after an expander phy's SP state machine receives a COMWAKE Detected message in states SP0:OOB\_COMINIT, SP1:OOB\_AwaitCOMX, SP3:OOB\_AwaitCOMINIT\_Sent, or SP4:OOB\_COMSAS if the ATTACHED SATA PORT SELECTOR bit is set to zero in the SMP DISCOVER response (see 10.4.3.9) prior to receiving the COMWAKE detected message (see 6.8.3 and table 341 in 10.4.3.9);
- e) after an expander phy's SP state machine transitions from the SP1:OOB\_AwaitCOMX state to the SP0:OOB\_COMINIT state if the ATTACHED SATA PORT SELECTOR bit was set to one in the SMP DISCOVER response 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 6.8.3.3.2);
- f) after an expander phy completes the link reset sequence (see 7.9);
- g) after a virtual phy has been enabled or completed processing a reset requested by the SMP PHY CONTROL function LINK RESET or HARD RESET phy operations (see 10.4.3.28);
- h) after an STP/SATA bridge receives an initial Register - Device to host FIS (see 9.3.1); and

In zoning expander devices with zoning enabled, forwarding Broadcasts is subject to restrictions defined in 4.9.5.

In zoning expander devices with zoning enabled, a Broadcast (Change) for an expander phy-related reason shall be originated from the source zone group of the expander phy causing the Broadcast (Change).

Expander devices shall originate Broadcast (Change) for the following expander device-related reasons:

- a) after a self-configuring expander device has changed its CONFIGURING bit from one to zero in the SMP REPORT GENERAL response (see 10.4.3.3) as described in 4.7.2. In zoning expander devices with zoning enabled, the source zone group shall be 1;
- b) after a locked expander device is unlocked (see 4.9.6.5 and 10.4.3.23), with the source zone group as specified in 4.9.6.5 and 10.4.3.23.

Expander devices shall forward Broadcast (Change) for the following reasons:

- a) after an expander phy receives Broadcast (Change).

For a virtual phy, if there is any time after a reset is originated during which connection requests to the attached SAS address result in connection responses of OPEN\_REJECT (NO DESTINATION), the expander device shall originate the Broadcast (Change) twice, once at the start of the reset (i.e., when the SAS address becomes unavailable) and once at its completion (i.e., when the SAS address becomes available). If there is no such time window, the expander device shall originate the Broadcast (Change) once.

SAS initiator ports may originate Broadcast (Change) 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 receives Broadcast (Change) shall follow the SAS initiator device rules (see 7.9.2) to discover and configure the topology.

An expander device that receives Broadcast (Change) shall follow the expander device rules (see 7.9.3) to discover and configure the topology.

See 10.4.3.3 for details on counting Broadcast (Change) origination in an expander device.