

To: T10 Technical Committee
From: Robert Sheffield (Robert.L.Sheffield@intel.com), Intel Corporation
Date: November 10, 2003
Subject: T10/03-273r2 SAS-1.1 Support for SATA Port Selector

Revision History

Revision 0 (August 8, 2003) first revision

Revision 1 (October 20, 2003): Incorporate feedback from September SAS Protocol WG

- Change reference to the SATA-II Port Selector document instead of T10 posting.
- Move footnote-3 under paragraph 3.1.x about STP initiator cooperation to section 10.3 ATA Application Layer
- Change SP3:OOB_COMSAS to SP4:OOB_COMSAS in paragraph 3.1.x
- Change the ATTACHED PORT SELECTOR bit from bit-4 byte-15 in the DISCOVER response to bit-7 (to separate it from the bits that have counterparts in byte-14).
- Use digits instead of words for bit values in Table x - Decode of ATTACHED SATA (PORT SELECTOR/DEVICE) Bits
- Add formal definition of the SATA port selection signal instead of referring to a paragraph in the SATA-II specification.
- Modify error reporting to provide consistency with affiliations.
- Use appropriate table footnote format (i.e. – a, b, c... immediately after table).
- Remove CLEAR ATTACHED SATA PORT SELECTOR phy operation.
- Include updated SP state machine reflecting transitions involved with SATA Port Selector interoperability.
- Miscellaneous editorial changes.

Revision 2 (November 10, 2003): Modify as directed by November SAS Protocol WG for approval.

Related Documents

SAS1-r01 – Working Draft Serial Attached SCSI –1.1 (SAS-1.1) revision 1.

Serial ATA II: Port Selector Revision 1.0, 28-July-2003 (a.k.a. the SATA-II Port Selector Specification)

Problem Description

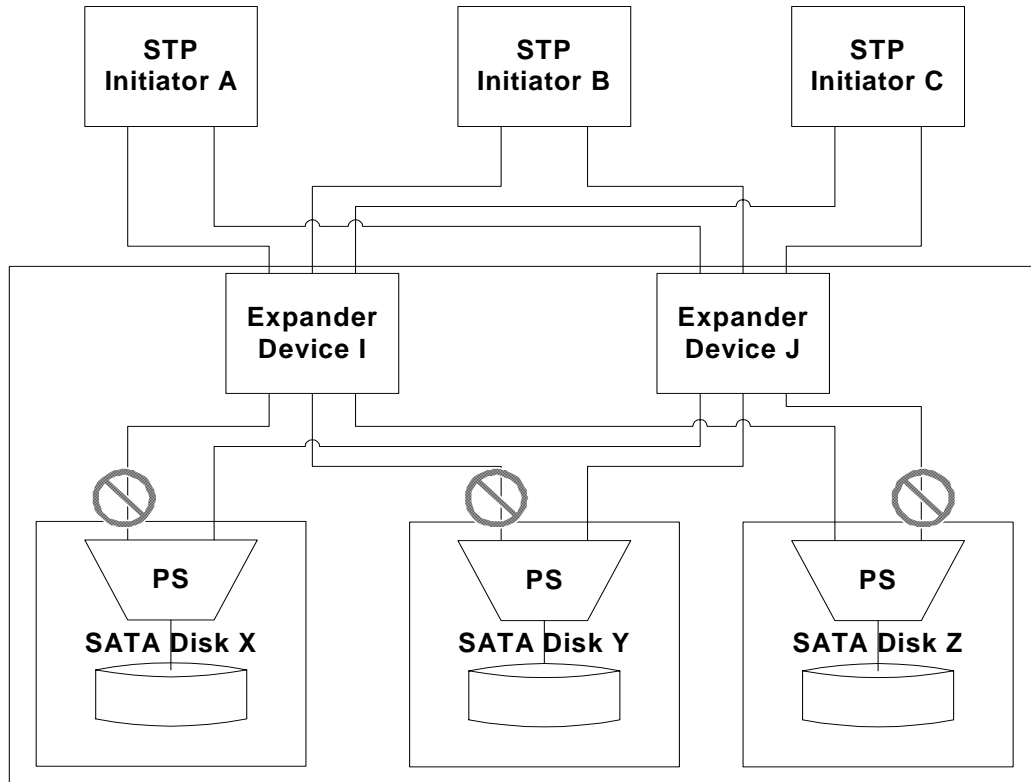
SAS is defined so as to support Serial-ATA and Serial-ATA-II extensions to SATA. This is to provide the infrastructure to allow storage systems suppliers to make the choice of which devices to supply with storage subsystems (SATA or SAS-SSP) as a procurement specification rather than as a design constraint. Many such storage subsystems are designed to provide alternate path capability to make systems tolerant of expander device failure, or failure of other components along the path between initiators and the storage devices. This is possible in SAS because SSP protocol and SCSI define the framework to support dual-port SSP target devices.

Likewise, SATA-II defines the framework for construction of a SATA domain that provides fail-over path capability through use of the SATA port selector. It is a natural expectation that if SAS, in general, provides the framework to support path fail-over capability, and SATA provides the framework to support path fail-over, then the same capability should be available for suitably equipped SATA devices in a SAS domain. This proposal addresses only the applicability of the SATA port selector in the attachment of SATA devices in a SAS domain. Nothing in this proposal identifies anything related to the operation of a SATA port selector with SAS target devices.

The SATA port selector is defined in SATA-II to provide a means for a SATA host to access a SATA device via an alternate path in case a primary path fails. The port selector specification provides both in-band (OOB signaling-based) and side-band methods for engaging the alternate path. This proposal identifies modifications to the SAS OOB sequence to provide detection of a

SATA port selector attached behind a STP/SATA bridge in an expander device, and defines extensions to SMP protocol to allow an STP initiator to discover and control the port selector.

Figure xx shows an example of an applicable configuration with SATA port selectors interposed between a set of SATA devices and a redundant pair of expander devices under the control of STP initiators in a SAS domain.



Each port-selector is always engaged along either one of two possible paths to an attached SATA disk. In this configuration, STP initiators access SATA devices X and Y via expander device J, and STP initiators access SATA device Z via expander device I. The SATA port selector provides distinct OOB signaling to allow the SAS expander to detect its presence on either the active SATA port selector phy or the inactive SATA port selector phy. The SATA port selector also responds to distinct OOB signaling from the inactive phy to engage that path as the active phy, causing the previously active phy to become inactive. All of the distinct OOB signaling for the SATA port selector is generated through varying timing and sequencing of currently defined COMINIT and COMWAKE signals.

Relevant Attributes of the SATA Port Selector

The latest version of the SATA port selector described in SATA-II defines elements that assist in the discovery and control of the SATA port selector in a SATA domain. These same attributes also enable effective management of the SATA port selector in a SAS domain. The following SATA port selector attributes facilitate operation in a SAS domain:

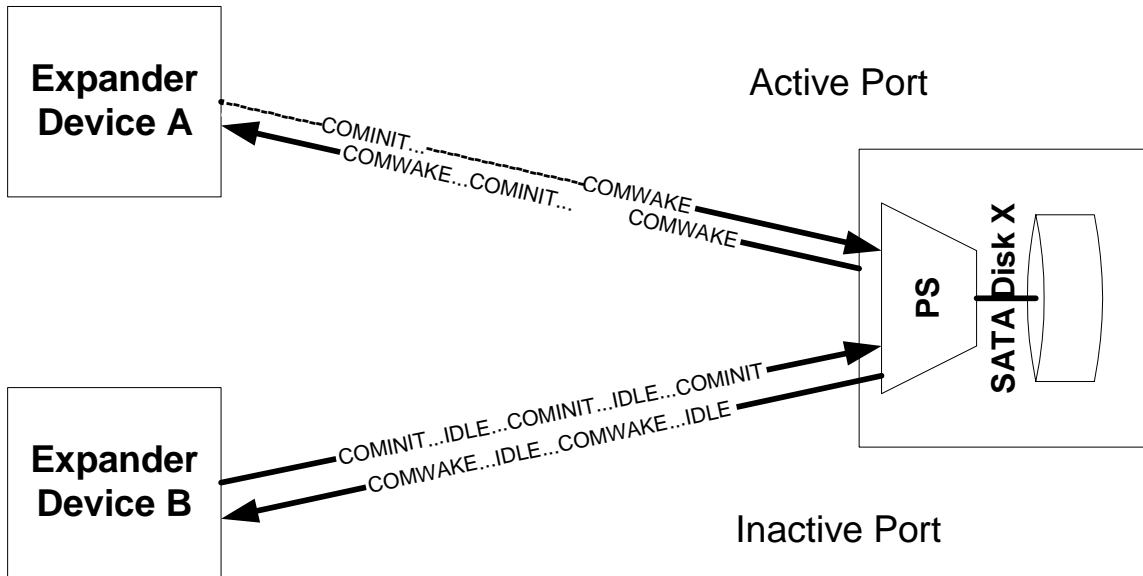
- SATA port selector responds to COMINIT¹ from the host with COMWAKE

¹ The SATA specification uses the term COMRESET to indicate when this signal originates from the Host, but since it's identical to COMINIT and the SAS convention is to use COMINIT both for the signal originated by the SAS initiator or expander device and the signal originated by the SATA device, the term COMINIT is used throughout this proposal, even where SATA would use the term COMRESET.

- On inactive phy: Always respond with COMWAKE
- On active phy: Always respond with COMWAKE followed by COMINIT
- After power-on, the first phy to initialize becomes the active phy.

Summary of SAS Modifications to Accommodate SATA Port Selector

The figure below shows the OOB signals exchanged between the SATA port selector and two expander devices on the active and inactive phys of the port Selector. The initial COMWAKE transmitted by the SATA port selector on both the active and inactive phys that provide an indication to the expander device that the SATA port selector is present.



Expander devices need some simple extensions to deal effectively with the SATA port selector:

- Modify expander SP state machine (SP0:OOB_COMINIT and other SP states) so that:
 - It recognizes COMWAKE and ACTIVE/INACTIVE state of SATA port selector, and
 - The SAS/SATA bridge logs SATA port selector state in DISCOVER response;
- Modify the SMP DISCOVER response to
 - Add a new ATTACHED SATA PORT SELECTOR bit;
- Modify the SMP Phy Control commands to:
 - Add the TRANSMIT SATA PORT SELECTION SIGNAL subcommand to cause the expander device phy to send the fail-over OOB sequence to the attached SATA port selector (a variant of HARD RESET);

Changes to apply to SAS-1.1r0

Add the following definition to sub clause 3.1Definitions:

3.1.x SATA port selector: A device defined by *Serial ATA II: Port Selector*, Revision 1.0, 28-July-2003, that connects two SATA hosts to a common SATA device, and provides the means for either SATA host to access the device at any given time, but not both simultaneously. This standard defines protocol elements needed to use the SATA port selector in a SAS domain.

Add the following text and figure after the last paragraph in sub clause 6.5 Out of band (OOB) signals:




Figure 59 – SATA port selection signal

The SATA PORT SELECTION SIGNAL shown in figure 59 causes the attached SATA port selector to select the attached phy as the active phy (see the SATA-II Port Selector Specification). The generated SATA PORT SELECTION SIGNAL shall be composed of 5 COMINIT signals, each starting a specified time interval, T1 or T2 as shown in Figure 59, after the start of the previous COMINIT signal. The values of T1 and T2 shall be as shown in Table 1.

Table 1 - Intervals between COMINIT in the SATA PORT SELECTION signal

Interval	Value (OOBI)
T1	3.0x10 ⁶
T2	12x10 ⁶

Add the following text to the bottom of the **SP0:OOB_COMINIT** sub clause – 6.7.3.2.1 STATE description

Any transition to this state shall cause the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to be set 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 supports attachment to a SATA port selector, it shall cause the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to be set to one.

Add the following text to the end of the descriptions of each of the following states:

- **SP1:OOB_AwaitCOMX** sub clause – 6.7.3.3.1 STATE description
- **SP3:OOB_AwaitCOMINITSent** sub clause – 6.7.3.5.1 STATE description
- **SP4:OOB_COMSAS** sub clause – 6.7.3.6.1 STATE description

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, it shall cause the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to be set to one.

Add sub clause 6.7.6 as follows:

6.7.6 SATA_PortSelectionSignalPending state

Figure 70 shows the SATA_PortSelectionSignalPending state. This state controls transmission of the SATA Port Selection signal when a specified phy processes an SMP Phy Control (TRANSMIT SATA PORT SELECTION SIGNAL) request.

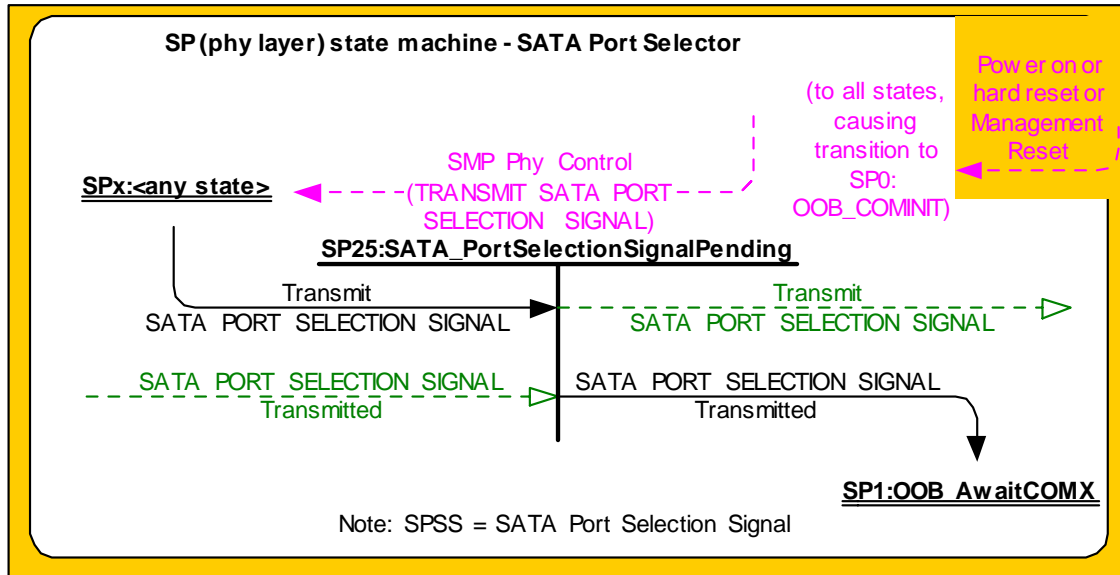


Figure 70 - SP (phy layer) state machine – SATA_PortSelectionSignalPending state

6.7.6.1 Transition SPx:<any state> to SP25:SATA_PortSelectionSignalPending

If the phy supports attachment of a SATA device and attachment of a SATA Port Selector, a transition shall occur from any SP state to this state upon receipt of an SMP Phy Control function for the phy specifying a PHY OPERATION of TRANSMIT SATA PORT SELECTION SIGNAL. The phy shall transmit the SATA PORT SELECTION SIGNAL (see table 163 – phy operation and sub clause 6.5 – out of band signals). This transition shall clear the ATTACHED SATA PORT SELECTOR bit in the DISCOVER Response.

6.7.6.2 Transition SP25:SATA_PortSelectionSignalPending to SP1:OOB_AwaitCOMX

This transition shall occur when the phy completes transmission of the SATA PORT SELECTION signal (SATA Port Selection Signal Transmitted).

In sub clause 10.4.3.2 – SMP function response frame format, modify Table 142 – Function Results in the cell corresponding to code 12h and SMP function(s) as follows:

Change:

REPORT PHY SATA

to:

REPORT PHY SATA, PHY CONTROL (CLEAR AFFILIATION) and PHY CONTROL (TRANSMIT SATA PORT SELECTION SIGNAL)

Table 147 – Discover Response in section 10.4.3.5 DISCOVER function, add bit-7 in byte-15: ATTACHED SATA PORT SELECTOR.

Add the following table in section 10.4.3.5 DISCOVER function to describe the ATTACHED SATA PORT SELECTOR and ATTACHED SATA DEVICE bits.

Table x describes the ATTACHED SATA PORT SELECTOR bit and the ATTACHED SATA DEVICE bit.

TABLE x - Decode of ATTACHED SATA (PORT SELECTOR/DEVICE) Bits

ATTACHED SATA		Description
PORT SELECTOR	DEVICE ^{a,b}	
0	0	Neither a SATA port selector nor a SATA device is attached and ready on the selected phy.
0	1	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. No SATA port selector is present and a SATA device is attached to the selected phy.
1	0	The SP state machine detected COMWAKE while waiting for COMINIT. A SATA port selector is present and either the port is the inactive path, or it is the active path and the attached SATA device is either not present or not ready.
1	1	The SP state machine detected COMWAKE while waiting for COMINIT, timed out waiting for COMSAS, and exchanged COMWAKE with an attached SATA device. The SATA device is connected through the active port of the port selector.
^a For the purposes of this bit, the SATA port selector is not considered a SATA device.		
^b 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 sub-clause 6.9)).		

Delete the 7th paragraph on page 342.

~~An ATTACHED SATA DEVICE bit set to one indicates a SATA device port is attached. An ATTACHED SATA DEVICE bit set to zero indicates a SATA device port is not attached. This bit shall be updated after the SATA OOB sequence completes and before the SATA speed negotiation sequence begins (i.e., at SATA spinup hold time (see 6.9)).~~

Add a new PHY CONTROL PHY OPERATION: TRANSMIT SATA PORT SELECTION SIGNAL

Table 162 – Phy operation

Code	Operation	Description
07h	TRANSMIT SATA PORT SELECTION SIGNAL	If the expander phy incorporates an STP/SATA bridge and supports attachment of a SATA port selector, this command shall cause the phy to transmit the SATA PORT SELECTION SIGNAL (see sub clause 6.5 – Out of band signals) that causes the SATA port selector to engage the attached phy as the active phy (and consequently render the alternate phy inactive). If the phy has an active affiliation it shall be cleared. If the expander phy does not implement an STP/SATA bridge or does not support attachment of a SATA port selector, then the expander shall respond with an SMP FUNCTION RESULT of PHY DOES NOT SUPPORT SATA. If this command is issued to a phy that does support attachment of a SATA port selector but is attached to a SAS device, it shall return an SMP FUNCTION RESULT of SMP FUNCTION FAILED.

Add the following text to the end of sub-clause 10.4.3.10 - PHY CONTROL Function.

If the indicated expander phy has an STP/SATA bridge and supports attachment of a SATA port selector, a PHY OPERATION of TRANSMIT SATA PORT SELECTION SIGNAL shall cause the identified phy to transmit the SATA PORT SELECTION SIGNAL. If an affiliation is active for the phy it shall be cleared.

During transmission of the SATA PORT SELECTION SIGNAL, the SP (phy layer) state machine shall suppress the COMINIT Detected and COMWAKE Detected messages. Following transmission of the SATA port selection signal, the SP state machine shall transition to the SP1:OOB_AwaitCOMX state, set the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to zero, and cease suppression of the COMINIT Detected and COMWAKE Detected messages.

Add the following sub-clause to Informative Annex F.4:

F.4 Considerations for STP Initiator Ports and SATA Port Selectors

Not all the protocol elements for STP initiator ports to manage a SATA port selector in a SAS domain are defined in this standard. Additional coordination between STP initiator ports may be needed to avoid conflicting usage of the SATA port selector between STP initiator ports. Such additional coordination is outside the scope of this standard.