

To: T10 Technical Committee
From: Robert Sheffield (Robert.L.Sheffield@intel.com), Intel Corporation
Date: August 8, 2003
Subject: T10/03-273r0 SAS-1.1 Support for SATA Port Selector

Revision History

Revision 0 (August 8, 2003) first revision

Related Documents

SAS1-r00 – Working Draft Serial Attached SCSI –1.1 (SAS-1.1) revision 0.
[03-155r2, Serial ATA II Specification: Port Selector](#) (Revision 1.0RC 23 June, 2003)

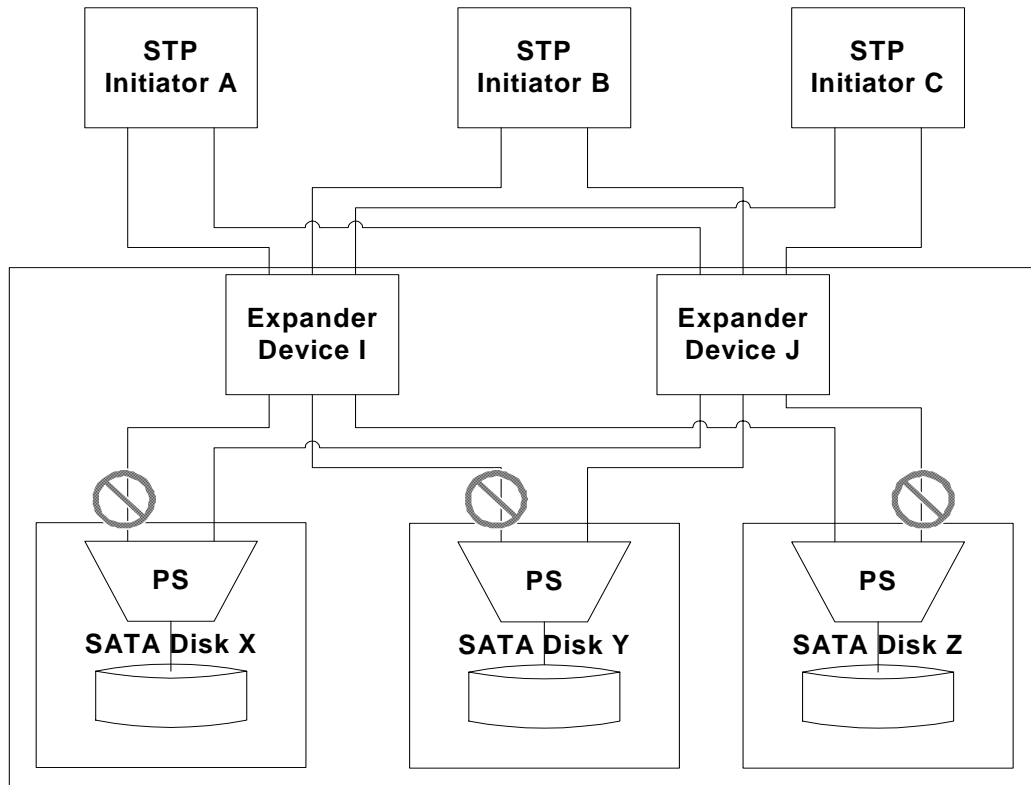
Problem Description

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

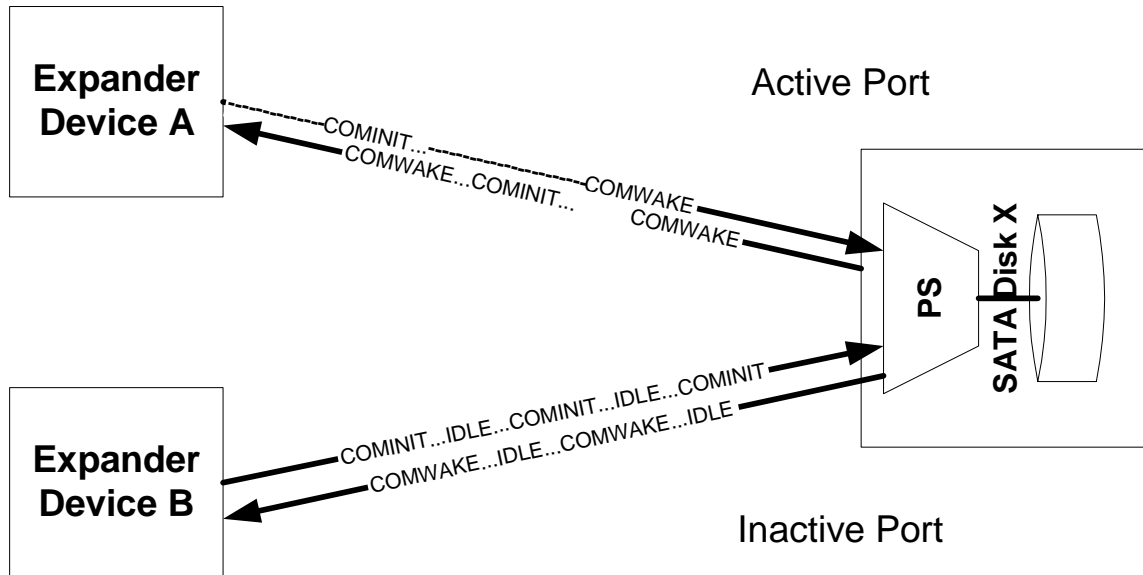
The latest version of the SATA port selector described in SATA-II defines elements that assist in the discovery and control of the SATA port selector in a SATA domain. These same attributes also enable effective management of the SATA port selector in a SAS domain. The following SATA port selector attributes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

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

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.

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);
 - Add the CLEAR ATTACHED SATA PORT SELECTOR subcommand to clear the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response².

² The CLEAR ATTACHED SATA PORT SELECTOR subcommand is added for debug support. As such, it is provided through the SMP commands to allow an STP initiator to use in whatever way the vendor sees best fit. Use of this to clear the ATTACHED SATA PORT SELECTOR bit is not a necessary element of the protocol to manage the SATA port selector.

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 the SATA-II Port Selector Specification (T10 reference 03-155r2) 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 elements³ needed to use the SATA port selector in a SAS domain.

*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 COMWAKE Detected message and the phy supports attachment to SATA devices (i.e. the phy is attached to an STP/SATA bridge); 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:

- o **SP1:OOB_AwaitCOMX** sub clause – 6.7.3.3.1 STATE description
- o **SP2:OOB_NoCOMSASTimeout** sub clause – 6.7.3.4.1 STATE description
- o **SP3:OOB_AwaitCOMINITSent** sub clause – 6.7.3.5.1 STATE description
- o **SP3:OOB_COMSAS** sub clause – 6.7.3.6.1 STATE description

If this state receives COMWAKE Detected message and the phy supports attachment to SATA devices (i.e. the phy is attached to an STP/SATA bridge); it shall cause the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response to be set to one.

Table 147 – Discover Response in section 10.4.3.5 DISCOVER function, add bit-4 in byte-15: ATTACHED SATA PORT SELECTOR.

Add accompanying descriptive text:

The ATTACHED SATA PORT SELECTOR bit set to one indicates the SATA port selector is attached. Another device may or may not be attached behind the SATA port selector. If no other attached device bits are set to one, then the SATA port selector phy attached to the expander device is inactive, or it may be active with no SATA device attached. If the ATTACHED SATA DEVICE bit is set to one, then the SATA port selector phy attached to the expander is active.

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

³ Perhaps not all the protocol elements are defined in this standard. Additional coordination between STP initiators may be needed to avoid conflicting usage of the SATA port selector between STP initiators.

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

ATTACHED SATA		Description
PORT SELECTOR	DEVICE ⁴	
zero	zero	Neither a SATA port selector nor a SATA device is attached and ready on the selected phy.
zero	one	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. No SATA port selector is present and a SATA device is attached to the selected phy.
one	zero	The SP state machine detected COMWAKE while waiting for COMINIT. This means 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 (i.e. THE PORT SELECTOR IS present and not ATTACHED SATA DEVICE does not necessarily mean Inactive Path).
one	one	The SP state machine detected COMWAKE while waiting for COMINT, 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 BROADCAST(CHANGE) shall be generated anytime the ATTACHED SATA PORT SELECTOR bit changes state.

Add a new PHY CONTROL PHY OPERATIONS: TRANSMIT SATA PORT SELECTION SIGNAL and CLEAR ATTACHED SATA PORT SELECTOR

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, then this command shall cause the phy to transmit the SATA port selection signal defined in 03-155r2 subclause 3.1.1 that causes the SATA port selector to engage the attached phy as the active phy (and consequently render the alternate phy inactive). 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 UNKNOWN PHY OPERATION. 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.

⁴ A SATA port selector is not considered a SATA device.

Code	Operation	Description
08h	CLEAR ATTACHED SATA PORT SELECTOR	This function clears the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response. This function is provided for diagnostic purposes and its use is not required for correct execution of the protocol.

Add the following text and figures yy and zz 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 issue the SATA port selection signal as specified in T10 document 03-155r2 subclause 3.1.1 Port Selection Signal Definition.

During transmission of the SATA port selection signal, the SP (phy layer) state machine shall suppress the COMINIT Detected and COMWAKE Detected messages. Immediately following transmission of the SATA port selection signal, the SP state machine shall transition to the SP0:OOB_COMINIT 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.

Editor's note: should subsequent link initialization be held off until the next Hot-Plug timeout? Some SATA port selectors may pass-through the device-bound COMINITs constituting the fail-over OOB signal as well as the subsequent COMINIT response from the device, which more or less represents the start of link initialization anyway. The proposal doesn't require the STP/SATA bridge to recognize it as such however, and would allow to the expander to process link initialization at the next hot-plug timeout if that's more convenient.

Modify sub clause 10.3 ATA application layer as follows:

10.3 ATA application layer

~~No SAS-specific ATA features are defined.~~ For systems that support SATA port selectors in the SAS domain, STP initiators should maintain information about the topology including which other STP initiators may access SATA devices through the SATA port selectors, and the rules that govern access to the SATA port selectors in the specific environment. The mechanism by which STP initiators obtain this information and coordinate with each other is outside the scope of this standard.

STP initiators in systems that support SATA port selectors should recognize DISCOVER response elements indicating the presence of a SATA port selector, and should send the SMP TRANSMIT SATA PORT SELECTION SIGNAL Phy Control command to effect fail-over when doing so does not conflict with other STP initiators that may access the common SATA device attached to the SATA port selector.