

To: T10 Technical Committee  
From: Bob Sheffield ([Robert.L.Sheffield@Intel.com](mailto:Robert.L.Sheffield@Intel.com)), Intel Corporation  
Date: October 31, 2002  
Subject: T10/02-430r1, SAS Simplified Handling of Multiple STP Initiators

### **Revision History**

Revision 0 (October 24, 2002) first revision

Revision 1 (October 25, 2002) Modified so as not to preclude the complex Multiple STP initiator port model supported by extensive SATA Target Port capabilities in the expander. Added an SMP command extensions to query current initiator affiliations. Added an SMP command to clear the affiliation for specified initiators. Added an option for the expander to broadcast an AFFILIATION\_CLEARED primitive to all STP initiator ports whenever an affiliation is cleared.

### **Related Documents**

SAS-r02a – Serial Attached SCSI revision 02a.

### **Overview**

This proposal offers a means to greatly simplify the way expander devices handle multiple STP initiator port access to SATA targets in a SAS domain without imposing any additional requirements on SATA targets. This proposal offers a way to shift the greater burden of arbitrating access to shared SATA targets from the expander device to the STP initiator port, and to limit the role of the expander device to detecting and reporting collisions accessing shared SATA ports.

The SAS working draft defines support of SATA targets attached to expander devices in a SAS domain with text in several places indicating multiple STP initiator ports can access any given SATA target. Because SATA targets have no intrinsic capability to support access by multiple initiators, the requirement to arbitrate access to a SATA disk among multiple STP initiator ports falls upon the expander device supporting attachment of SATA targets.

SAS supports control of SATA targets in a SAS domain in order to address cost-sensitive markets where SSP targets may prove too expensive. To that end, the incremental cost required to implement an expander device supporting attachment of SATA targets should be negligible as compared to the cost of an expander that supports SSP and SMP only.

As currently written, the SAS draft standard specifies a number of capabilities an expander device must provide to enable access to SATA targets from multiple STP initiator ports:

## SAS Simplified Support for Multiple STP Initiators

- For each SATA port, the expander device must replicate a shadow ATA task file register set for each STP initiator port that may access the attached SATA device.
- The expander device must translate Queue Tags between any STP initiator port and a SATA target to avoid collision of identical Queue Tags issued by different STP initiator ports.
- The expander device must detect when all available target queue-tags are in use and queue subsequent commands until outstanding commands complete, freeing up target queue tags.

The required capabilities for each expander port supporting a SATA target roughly correspond to the capability of a complete SATA host bus adapter, replicated for each STP initiator port that might communicate with the SATA target. Clearly this is at odds with the objective to provide a cost-effective interconnect for SATA devices and is far beyond the intended usage of SAS or SATA.

A decided simplification of the protocol to support SATA devices is essential before SAS can effectively support SATA target devices in a multi-initiator SAS domain. This proposal defines a simplified alternative to the complex model that provides limited multiple STP initiator port access to shared SATA disks without ~~imposing the same~~ requirements on the expander device.

#### Suggested Changes

Modify Table 40 – Primitives not specific to the type of connection to replace the “CHANGE (RESERVED 2) primitive as follows:

Primitive	Use <sup>a</sup>	From <sup>b</sup>			To <sup>b</sup>			Primitive sequence type <sup>c</sup>
		I	E	T	I	E	T	
AFFILIATION CLEARED	SAS			T	I			Redundant

Modify Table 43 – Primitive encoding for primitives not specific to type of connection as follows:

Primitive	Character			
	1 <sup>st</sup>	2 <sup>nd</sup>	3 <sup>rd</sup>	4 <sup>th</sup>
AFFILIATION CLEARED	K28.5	D02.0	D29.7	D16.7

Modify paragraph 7.1.4.4 to read as follows:

#### **7.1.4.4 CHANGE**

CHANGE is sent by an expander device to notify initiator ports and other expander devices that a configuration change has occurred. CHANGE shall only be sent outside of connections.

CHANGE (RESERVED 0) and CHANGE (RESERVED 1) are reserved, and shall be broadcast by expander devices the same as CHANGE.

See 7.11 for details on domain changes.

**7.1.4.4 AFFILIATION CLEARED**

AFFILIATION CLEARED is sent by an STP Port (or an expander when acting on behalf of a SATA target) to notify all STP initiator ports that an affiliation maintained by an STP target port with an STP initiator port has been cleared because the STP target port received a PORT\_CONTROL request specifying a function of CLEAR\_AFFILIATION. Expanders shall broadcast this primitive to all end ports in a SAS domain, but all port types except STP initiator port shall process this primitive as an ALIGN. An STP initiator port may use this primitive as a signal to retry connection requests to STP target ports that were previously rejected with OPEN\_REJECT(STP RESOURCES BUSY).

Add the following text under the heading 9.3.2 tunneling for multiple initiator ports:

After power-on reset or initial link reset and subsequent initialization affecting an expander phy with a SATA target device attached, the expander device may accept a connection request to the SATA target port from any STP initiator port. Once the expander has accepted a connection request for that SATA target port from N distinct STP initiator ports, where N is the NUMBER OF INITIATORS supported by the SATA target port as reported in the REPORT SATA CAPABILITIES SMP request, the expander shall reject all subsequent connection requests for that SATA target port from other STP initiator ports, with OPEN\_REJECT (STP RESOURCES BUSY). The expander shall continue to reject connection requests from other STP initiator ports to that SATA target port even if the connection with one of the the original STP initiator ports is closed. In this state, the expander port is said to maintain an affiliation between the STP initiator port and the SATA target port. This affiliation shall persist until any of the following occurs:

- A HARD\_RESET sequence is performed affecting phy of of the affiliated SATA target.

An STP initiator port issues an SMP Port Control request specifying a phy operation that clears one or more affiliations maintained by the expander between the specified phy and STP initiator ports.

Modify the description of OPEN\_REJECT (STP RESOURCES BUSY) in Table 40 – OPEN\_REJECT primitives as follows:

Destination device exists but has already established the maximum number of affiliations supported with other STP initiator ports.

Modify Table 114 – REPORT PHY SATA response as follows:

Table 114 – REPORT PHY SATA response

Bit Byte	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (12h)							
2	FUNCTION RESULT							
3	Reserved							

Bit Byte	7	6	5	4	3	2	1	0
4	Ignored							
7	Ignored							
8	Reserved							
9	PHY IDENTIFIER							
10	Ignored							
11	VALID AFFILIATIONS BIT MASK							
12	Reserved							
15	Reserved							
16	(MSB)	STP SAS ADDRESS						(LSB)
23	(LSB)							
24	(MSB)	REGISTER DEVICE TO HOST FIS						(LSB)
43	(LSB)							
44	(MSB)	AFFILIATED INITIATOR 0 SAS ADDRESS						(LSB)
51	(LSB)							
52	(MSB)	AFFILIATED INITIATOR1 SAS ADDRESS						(LSB)
59	(LSB)							
60	(MSB)	AFFILIATED INITIATOR 2 SAS ADDRESS						(LSB)
67	(LSB)							
68	(MSB)	AFFILIATED INITIATOR 3 SAS ADDRESS						(LSB)
75	(LSB)							
76	(MSB)	AFFILIATED INITIATOR 4 SAS ADDRESS						(LSB)
83	(LSB)							
84	(MSB)	AFFILIATED INITIATOR 5 SAS ADDRESS						(LSB)
91	(LSB)							
92	(MSB)	AFFILIATED INITIATOR 6 SAS ADDRESS						(LSB)
99	(LSB)							
100	(MSB)	AFFILIATED INITIATOR 7 SAS ADDRESS						(LSB)
107	(LSB)							
108	(MSB)	CRC						(LSB)
111	(LSB)							

Bytes 44 through 107 are a list of SAS addresses of all STP Initiator ports that currently share an affiliation with the SATA target port indicated by PHY IDENTIFIER. Valid affiliations bit mask indicates, by bit number, which of the AFFILIATED INITIATOR\_SAS ADDRESSES is valid. The mapping for an affiliation shall not change for the duration of the affiliation.

Modify Row 11 of table 125 – PHY CONTROL request as follows:

Bit Byte	7	6	5	4	3	2	1	0
11	AFFILIATED STP INITIATOR_ BIT MASK							

Add the text below describing the AFFILIATED STP INITIATOR\_ BIT MASK:

This bit mask is used for the phy operations that clear STP port affiliations. Each bit set to one corresponds to the number of the AFFILIATED\_INITIATOR\_SAS\_ADDRESS as reported in the REPORT AFFILIATIONS SMP request for that SATA target port. A bit set to one indicates the affiliation is to be cleared for the corresponding STP initiator port.

Modify the description of the LINK RESET operation in Table 126 – Phy operations as follows:

Code	Operation	Description
01h	LINK RESET	Perform a link reset sequence (see 4.4) on the specified phy. The expander device shall not clear any affiliations with STP initiator ports existing as a result of performing the link reset sequence in response to this request.

Add the following Phy Operations to table 126 – Phy operation:

Table 126 – Phy operation

Code	Operation	Description
07h	CLEAR OWN AFFILIATION	This request, directed to an expander device and affects the phy specified in the PHY IDENTIFIER, clears the affiliation held by the STP initiator port issuing the request. The phy may broadcast AFFILIATION_CLEARED to signal other STP initiator ports to retry connection requests.
08h	CLEAR OWN AFFILIATION AND RESET	This request operates the same way as the CLEAR OWN AFFILIATION function except that the link shall be reset and a CHANGE shall be broadcast instead of AFFILIATION_CLEARED.
09h	CLEAR COMPANION AFFILIATION	This request, directed to an expander device and affects the phy specified in the PHY IDENTIFIER, clears affiliations maintained by the expander device, for the SATA target port associated with the designated phy. The expander uses the AFFILIATED_INITIATOR_BIT_MASK to determine which affiliations to clear. The phy may broadcast AFFILIATION_CLEARED to signal STP initiator ports to retry connection requests.
10h	CLEAR COMPANION AFFILIATION AND RESET	This request operates the same way as the CLEAR COMPANION AFFILIATION function except that the link shall be reset and a CHANGE shall be broadcast instead of AFFILIATION_CLEARED.

Modify the NUMBER OF INITIATORS description under paragraph 10.3.1.3 REPORT SATA CAPABILITIES function to read as follows:

The NUMBER OF INITIATORS field shall indicate how many initiator ports the expander device is capable of allowing to share access to a SATA target port. This may be restricted in expander devices which support the ATA queued feature set. A SATA target port shall report a number no greater than 08h in this field. Connection requests exceeding limit reported in this field shall result in OPEN\_REJECT (STP RESOURCES BUSY).