

To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@Intel.com), Intel Corporation
Date: November 4, 2002
Subject: T10/02-430r2, 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.

Revision 2 (November 4, 2002) Modified to specify removal of task file register shadowing in expander devices, specifies that HARD_RESET clears affiliation and causes broadcast of CHANGE, specifies that link reset sequence does not clear affiliation but still causes broadcast of CHANGE, does not specify an AFFILIATION_CLEARED primitive, specifies an SMP PHY CONTROL function for an STP initiator to clear it's own affiliation, defines an extension to the REPORT_PHY_SATA SMP command to list the STP Initiator SAS address for a single active affiliation and corresponding valid bit to handle the case where there is no affiliation, defines an optional CLOSE variant (CLOSE AND CLEAR AFFILIATION) that specifies to clear the affiliation when processing of the CLOSE.

Related Documents

SAS-r02c – Serial Attached SCSI revision 02c.

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 Simplified Support for Multiple STP Initiators

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:

- 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 severe requirements on the expander device.

The proposed changes are summarized as follows:

- A successful connection between an STP initiator and a SATA port establishes an explicit affiliation.
- Delete task-file register shadowing in expanders
- Any SMP PHY CONTROL HARD RESET operation clears affiliation & causes change primitive,
- SMP PHY CONTROL LINK RESET does not clear affiliation but causes change primitive,
- There is no need for AFFILIATION CLEARED primitive,
- Define a new SMP PHY CONTROL request phy operation to clear an STP Initiator port's own affiliation. If the sender doesn't have an affiliation, the expander device responds with PHY OPERATION FAILED status.
- SMP REPORT_PHY_SATA reports the current affiliation with a single valid bit.

SAS Simplified Support for Multiple STP Initiators

- A new CLOSE primitive is defined to allow an STP initiator to clear its affiliation upon close. If this appears anywhere other than in an STP connection, it is interpreted as a normal CLOSE. This is optional for STP initiators (can choose to use SMP method). It is required for expander devices that support SATA targets to use the CLOSE and CLEAR affiliation primitive to clear an affiliation.

Suggested Changes

Modify Table 42 – Primitives not specific to the type of connection to add the “CLOSE AND CLEAR AFFILIATION” primitive sequence:

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
CLOSE AND CLEAR AFFILIATION	STP	I					T	Triple

Modify Table 45 – Primitive encoding for primitives not specific to type of connection as follows:

Primitive	Character			
	1 st	2 nd	3 rd	4 th
CLOSE AND CLEAR AFFILIATION	K28.5	D02.0	D07.3	D04.7

Add section 7.1.4.6 CLOSE AND CLEAR AFFILIATION:

CLOSE AND CLEAR AFFILIATION shall be processed in the same manner as CLOSE except that an STP initiator port may issue CLOSE AND CLEAR AFFILIATION in place of a normal CLOSE to cause the expander device on the far end of the STP connection to clear the affiliation between the STP initiator and the SATA phy at the other end of the STP connection, leaving the phy available for another STP initiator to establish a connection. This behavior is in addition to the normal effect of closing the connection. If an expander that supports attachment of a SATA target receives CLOSE AND CLEAR AFFILIATION, the expander shall clear the affiliation between the STP initiator that sent the CLOSE AND CLEAR AFFILIATION and the phy at the other end of the connection being closed.

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

After power-on reset or a HARD RESET issued through the PHY CONTROL request, 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 device has accepted a connection request for that SATA target port from an STP initiator port, the expander device shall reject all subsequent connection requests for that SATA target port from other STP initiator ports with OPEN REJECT (STP RESOURCES BUSY). The expander device shall continue to reject connection requests from other STP initiator ports to that SATA target port even if the connection with the original STP initiator port is closed. In this state, the expander device

SAS Simplified Support for Multiple STP Initiators

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:

- Another power-on condition occurs,
- A PHY CONTROL request specifying a PHY OPERATION of HARD RESET from any initiator specifying the phy with an affiliation is received,
- An STP initiator port issues an SMP PHY CONTROL request specifying a PHY OPERATION of CLEAR AFFILIATION specifying the phy which has an affiliation with that STP initiator,
- An STP initiator closes a connection to the affected phy with CLOSE AND CLEAR AFFILIATION.

Any condition causing a link reset sequence other than power-on or PHY CONTROL request specifying a PHY OPERATION of HARD RESET shall not clear an affiliation held between an expander phy connected to a SATA target port and an STP initiator port.

Modify the description of OPEN_REJECT (STP RESOURCES BUSY) in

Table 50 – OPEN_REJECT primitives as follows:

Destination device exists but the corresponding phy has an active affiliation with another STP initiator port.

Modify Table 105 – REPORT SATA CAPABILITIES response as follows:

Table 105 – REPORT SATA CAPABILITIES response

Bit Byte	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (01h)							
2	FUNCTION RESULT							
3	Reserved							
4	Ignored							
7	Ignored							
8	Reserved							
9	Ignored							
10	Reserved							
11	Reserved							
12	Reserved (Editor's note: removed ATA QUEUEING CAPABLE)							SATA CAPABLE
13	Reserved (Editor's note: removed NUMBER OF INITIATORS)							
16	STP SAS ADDRESS							
16	(MSB)							
17	Reserved (Editor's note: deleted SATA VERSION)							
18	Reserved							
19	Reserved							
20	(MSB)	CRC						
23								(LSB)

Delete descriptive text following table 105 for NUMBER OF INITIATORS, SATA VERSION, and ATA QUEUEING CAPABLE.

Add the following descriptive text for the AFFILIATED STP INITIATOR SAS ADDRESS:

The AFFILIATED STP INITIATOR SAS ADDRESS in bytes 48 through 55 is the SAS address of the STP Initiator port that currently has an affiliation with the SATA target port indicated by PHY IDENTIFIER. The AFF VALID bit shall be set to a one when the AFFILIATED STP INITIATOR SAS ADDRESS is valid indicating the expander maintains an active affiliation between the indicated phy and the reported STP initiator SAS address. The AFF VALID bit shall be cleared to a value of zero when no STP initiator has an affiliation with the specified phy to indicate the AFFILIATED STP INITIATOR SAS ADDRESS field is not valid.

Modify Table 119 – 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							
4	Ignored							
7	Ignored							
8	Reserved							
9	PHY IDENTIFIER							
10	Ignored							
11	Reserved							
12	Reserved							
15	Reserved							
16	(MSB)	STP SAS ADDRESS						(LSB)
23								
24	(MSB)	REGISTER DEVICE TO HOST FIS						(LSB)
43								
44	Reserved							
46	Reserved							
47	Reserved							AFF VALID
48	(MSB)	AFFILIATED STP INITIATOR SAS ADDRESS						(LSB)
55								
56	(MSB)	CRC						(LSB)
59								

Modify the description of the LINK RESET operation and HARD_RESET operations in Table 128 – Phy operations as follows:

Code	Operation	Description
------	-----------	-------------

Code	Operation	Description
01h	LINK RESET	Perform a link reset sequence (see 4.4) on the specified phy. Any affiliation active at the time the LINK RESET operation is issued shall continue to be active following the LINK RESET sequence executed as a result of receiving this request. Execution of this request shall cause the expander to broadcast CHANGE.
02h	HARD RESET	Perform a link reset sequence (see 4.4) on the specified phy. If the attached phy is a SAS phy, the link reset sequence shall indicate an identification sequence. If the attached phy is a SATA phy, and a hard reset (see 4.4.2) on the specified phy. If the specified phy is attached to a SATA phy, only a link reset sequence (see 4.4) shall be performed. FIX THIS Execution of this request shall clear any active affiliation maintained by the expander between the specified phy and an STP initiator and shall cause broadcast of CHANGE.

Add the following Phy Operation to table 128 – Phy operation:

Table 128 – Phy operation

Code	Operation	Description
07h	CLEAR AFFILIATION	This request, directed to an expander device and affecting the phy specified in the PHY IDENTIFIER, clears the affiliation held by the STP initiator port issuing the request. If the STP initiator issuing this request does not currently have an affiliation with the specified phy, the expander shall issue a response of PHY OPERATION FAILED. If clear affiliation is received by an expander for a phy that does not have a SATA device attached the expander shall fail the command with UNKNOWN PHY OPERATION.