

To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@Intel.com), Intel Corporation
Date: October 23, 2002
Subject: T10/02-430r0, SAS Simplified Support for Multiple STP Initiators

Revision History

Revision 0 (October 8, 2002) first revision

Related Documents

SAS-r02a – Serial Attached SCSI revision 02a.

Overview

This proposal offers a few very simple modifications to the current SAS working draft that greatly simplify the way expanders handle multiple STP initiator access to SATA targets in a SAS domain without imposing any additional requirements on SATA targets. The effect of this proposal is to shift the greater burden of arbitrating access to shared SATA targets from the expander to the STP initiator, and to limit the role of the expander to detecting and reporting collisions accessing shared SATA ports.

The SAS working draft defines support of SATA targets attached to expanders in a SAS domain with text in several places indicating multiple STP initiators 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 initiators falls upon the expander 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 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 must provide to enable access to SATA targets from multiple STP initiators:

- For each SATA port, the expander must replicate a shadow ATA task file register set for each STP initiator that may access the attached SATA device.
- The expander must translate Queue Tags between any STP initiator and a SATA target to avoid collision of identical Queue Tags issued by different STP initiators (note: queue tags are undefined in SATA 1.0, and the SAS standard cannot reference material not yet in the public domain).

SAS Simplified Support for Multiple STP Initiators

- The expander 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 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.

Suggested Changes

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

After power-on reset or link reset and subsequent initialization affecting an expander link with a SATA target attached, the expander may accept a connection request to the SATA target port from any STP initiator. Once the expander has accepted a connection request for that SATA target port from a given STP initiator, the expander shall reject all subsequent connection requests for that SATA target port from other STP initiators, responding with OPEN_REJECT (STP RESOURCES BUSY). The expander shall continue to reject connection requests from other STP initiators to that SATA target port even if the connection with the original STP initiator is closed. In this state, the expander port is said to maintain an *affiliation* between the STP initiator and the SATA target port. This affiliation shall persist until any of the following occurs:

- A link reset sequence is performed affecting Phy of the affiliated SATA target port for any reason (e.g. power cycle, or receipt of a Phy Control Request specifying a Phy Operation of LINK_RESET, HARD_RESET, or ENABLE).
- The STP initiator with which the expander port maintains an affiliation issues a Phy Control Request specifying a Phy Operation of CLEAR_AFFILIATION).

If an STP initiator other than the one in an existing affiliation with the SATA expander port issues a Phy Control Request specifying a Phy Operation of CLEAR_AFFILIATION, the expander shall not change the existing affiliation and shall return a PHY CONTROL response specifying a FUNCTION RESULT of 11h – Operation specified by PHY OPERATION failed.

Add the following row to Table 114 – Phy operation:

Code	Operation	Description
07h	CLEAR AFFILIATION	Remove the affiliation between the specified link and the last STP initiator to open a connection with the corresponding SATA target port, allowing the corresponding expander port to accept the next connection request from any STP initiator.

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

Destination device exists but has an affiliation with another STP initiator.

