

To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@Intel.com), Intel Corporation
Date: October 24, 2002
Subject: T10/02-435r0, SAS STP Buffering

Revision History

Revision 0 (24 October 2002) first revision

Related Documents

SAS-r02b – Serial Attached SCSI revision 02b

Overview

The requirement for an expander that supports attachment of a SATA target to provide additional buffering to compensate for the latency it adds in propagating SATA flow-control primitives is not adequately bounded to establish a proper basis for interoperability. Only the expander device that actually has a SATA disk attached is aware of the need to buffer data. But when the SAS domain consists of multiple levels of expander devices, including fan-out expander devices and edge expander device sets composed of edge expander devices, the expander device with the SATA target attached has no means to determine the added latency between the STP initiator and SATA target, and therefore has no means to determine the amount of buffering required. Even if there were a means to determine dynamically the amount of buffering required, an expander must be built with a fixed amount of buffering capability and cannot dynamically allocate more based on the topology it finds itself in.

A means to solve this problem is to establish bounds on the topology for expander interconnection that support SATA devices, and to place limits on the latency introduced by each expander in the pathway between an STP initiator and a SATA target. It is then possible to determine the minimum amount of buffering required to implement an expander that supports attachment of SATA targets.

The changes to the working draft SAS-r02a proposed in this document establish the limits on expander topology and transport latency required to make SATA incremental buffering requirements deterministic.

Suggested Changes

The following restrictions on expander topology and introduced latency should be added to section 9.3.1 for SAS domains that support attachment of SATA targets:

- A pathway through an edge expander set attached to a SATA target or STP initiator shall traverse no more than 4 links from the target or initiator to the subtractive decode port of the expander set, including the link to the initiator

or target and the link connected to the subtractive decode port. Behavior of expander topologies that exceed this limit is not defined in this standard.

- The latency introduced by any expander between any two successive links along the pathway between an STP initiator and SATA target shall introduce no more than 500 nS of latency in the transmission of dwords from one link to the next along the pathway between STP initiator and SATA target. Behavior of expanders that permit attachment of SATA targets but that exceed this latency is not defined in this standard.
- Each expander port that supports connection of a SATA target shall provide a minimum of 100 dwords of additional buffering to hold dwords transmitted by the sender during the time interval from receipt of a HOLD request from the receiver and the time the sender recognizes the HOLD and suspends data transmission.