

To: T10 Technical Committee
From: Bob Sheffield (Robert.L.Sheffield@Intel.com), Intel Corporation
Date: March 11, 2003
Subject: T10/03-131r0, SAS Model for use of Affiliation

Revision History

Revision 0 (March 11, 2003) first revision

Related Documents

SAS-r03e – Serial Attached SCSI – revision 3e

Overview

There has been some confusion over the expected use of Affiliations for STP connections. Affiliations are intended to provide a means to ensure the atomicity of commands issued to SATA devices in an environment where multiple STP initiator ports may access a SATA device in a SAS domain. Of particular concern is the misinterpretation that Affiliation might be intended as a substitute for static allocation of SATA or STP devices among STP initiator ports in a SAS domain coordinated through side-band communication. This proposal is to add a description of the likely use of Affiliations (i.e. to avoid possible collisions when a SATA device is accessed by more than one initiator port, for example during discovery) and to clarify that it is not intended as a substitute for static mapping of devices to initiator ports that would normally be established using side-band communication between initiator ports.

Suggested Changes

Insert sub clause F.3 as in informative annex:

F.3 Affiliation Policies

SATA, like ATA, is based on a model that assumes a device is controlled by a single host, and does not address the notion of multiple hosts having the ability to access any given SATA device. Through the definition of the STP protocol, SATA devices are cast into an environment where multiple STP initiator ports (assuming the role of hosts in a SATA domain) have access to the same SATA or native STP devices in the SAS domain. SATA protocol carried across STP connections does not account for the possibility that more than one STP initiator port might be vying for access to the device. SAS affiliation provides a way to ensure contention for a SATA device does not result in incoherent access to the device when commands from different STP initiator ports collide at the SATA device.

To prevent a SATA device from confusing commands from one STP initiator port with commands from another, an STP initiator port needs a means to maintain exclusive access to a SATA or STP device for the duration of the execution of a command. For example, consider the case where an STP initiator port establishes a connection to send a command (perhaps read), and then closes the connection while the disk performs a seek operation to the track containing the data. If, after the connection is closed, another STP initiator port is allowed to establish a connection to send another command, the SATA device would no longer have a means to determine which STP initiator port should receive the data when the device requests the connection to send the data for the first command. This is because, unlike SCSI devices, SATA devices have no notion of multiple initiator port environments. The consequences are worse for write commands since it could result in the wrong data being written to media with the original data being overwritten and lost forever.

Affiliation provides a means for an STP initiator port to establish atomic access to a SATA device across the execution a command or series of commands to the SATA device without requiring the STP initiator port to maintain a connection open to the device for the duration.

F.3.1 Affiliation Policy for Static SATA device – STP initiator port mapping

Affiliations should not be used to enforce policies establishing fixed associations between SATA devices and STP initiator port ports.

F.3.2 Affiliation Policy with SATA queued commands and multiple STP Initiator ports

STP initiator ports using queued commands when other STP initiator ports may be accessing the same device should periodically allow commands to complete and release the affiliation to allow other STP initiator ports access to the device.

F.3.3 Applicability of Affiliation for STP target devices

Affiliation may or may not be necessary for target devices that support STP protocol directly depending on whether the device tracks the STP Initiator port address on each command received from each initiator port establishing a connection with the device. If the STP target port has the means to manage and track commands from each STP initiator port independently, then affiliations are not necessary because the device can associate each information transfer on the interface with the appropriate initiator port, and can establish connections with the appropriate STP initiator port when sending information back to the STP initiator port.

An STP target port that behaves the same as a SATA device in that it maintains only a single ATA task file register image shared among STP initiator ports supports affiliations in order to provide a way for initiator ports to maintain exclusive access to the STP target port while executing commands remain outstanding. In this model an STP target port has the ability to establish connections to an initiator port, but like an expander with a SATA device attached, can only remember the SAS address of the last STP initiator port to establish a connection, and therefore can only request a connection back to that same STP initiator port.

An STP target indicates that it can maintain coherency among commands from multiple initiator ports outstanding simultaneously by setting the AFFILIATIONS bit in the REPORT PHY SATA RESPONSE (Table 161) to a value of zero. An STP target indicates that it cannot maintain coherency among commands from multiple initiator ports outstanding simultaneously by setting the AFFILIATIONS bit in the REPORT PHY SATA RESPONSE (Table 161) to a value of one, indicating that the initiator port must rely on affiliations to maintain coherency.