

Document: T10/02-324r0

Date: September 8, 2002

To: T10 Committee Membership

From: Edward A. Gardner, Basil Networks

Subject: SAS Disconnect-Reconnect Mode Parameters

The SAS SCSI layer carries over many Disconnect-Reconnect mode page parameters from parallel SCSI. With parallel SCSI only the target can directly control the bus. These parameters were necessary as a means for the initiator to instruct the target how the bus should be controlled.

With SAS the initiator also has direct control of the bus. It can withhold R_RDYs, initiate closing a connection, etc. The Disconnect-Reconnect mode page parameters are unnecessary and, in some cases, contradict other areas of the spec. SAS should define the entire contents of the Disconnect-Reconnect mode page as reserved.

The remainder of this proposal discusses each parameter individually.

1 BUS INACTIVITY TIME LIMIT

Various sections of SAS specify when a device may or shall close a connection. Depending on circumstances this is either after a 1 millisecond timeout or after a vendor specific time. The definition of the BUS INACTIVITY TIME LIMIT contradicts those.

An initiator has a better method to instruct a target to close an inactive connection, specifically sending both CREDIT BLOCKED and DONE.

This field should be made reserved.

2 MAXIMUM CONNECT TIME LIMIT

This parameter provides a way for an initiator to instruct a target to not "hog" the bus. That is appropriate for a shared bus such as parallel SCSI.

SAS is considerably different. With wide links, there is much less incentive to terminate a connection simply because it has persisted (and is doing useful work) for a long time.

An initiator can terminate a connection whenever it chooses, such as by sending both CREDIT BLOCKED and DONE. This field should be made reserved.

3 MAXIMUM BURST SIZE

This is much the same as the MAXIMUM CONNECT TIME LIMIT. With parallel SCSI, it was important to have a way to limit data transfer burst sizes, since the only way an initiator could interrupt a data transfer burst was with a bus reset. With SAS an initiator can interrupt data transfers at frame boundaries by sending CREDIT BLOCKED or by simply not sending frames.

This field should be made reserved.

4 FIRST BURST SIZE

The sole merit of this field is with high latency interconnects – SCSI protocols and configurations where the latency to fetch write data is comparable to command execution time. This does not apply to SAS. SAS has limited, short distance connections. It would be straightforward to design hardware acceleration that would allow a WRITE command, XFER_RDY and DATA to all be transferred in a single connection. Noone is doing so because expected SAS latency is so low as to make it not worth the bother.

This is not worth the complexity and should be made reserved.