

Status of SCSI Error Recovery Proposal (X3T10/95-352R3)

Charles Monia

Digital Equipment Corporation

March 12, 1997

T10/97-151R0

Problem Addressed by Proposal

- For a non-interlocked bus, multiple commands can be in transit to a device that supports command queuing.
- The failure of any command in the pipeline may result in subsequent commands being executed 'out of order' relative to the failed command.

Proposed Solution

- Add a generic, device-resident interlock (similar to ACA)
 - Automatically set on any command error;
 - Prevents the device from accepting commands;
 - Is cleared by the host after all ‘in flight’ commands are flushed from the pipeline.
- Feature is an implementation option

Assumptions (all open to question)

- The order of command delivery should be preserved for the following reasons
 - Certain aspects of the full queuing model (ordered and head-of-queue commands, the ability to preserve order on disk commands whose LBN ranges overlap) require that there be some way for the drive to reconstruct the order in commands were sent from a given host.

Assumptions (cont.)

- Ordering is required for sequential devices that support command queuing.
- As seen by the “Class Driver”, interlocked error handling emulates the driver’s view of error handling in parallel SCSI.

Counterarguments

- Aspects of the queing model that require ordering either don't work or are not useful in practice.
- Sequential devices don't require command queing for performance.
 - Large amounts of write-behind/read-ahead cache are available to sustain throughput.

Counterarguments (cont.)

- Sequential and Block devices are the only SCSI device types of interest on a non interlocked bus.
 - Other types will either never be implemented in this environment or will not support command queuing.

Proposal for Closure

- The need for block device command ordering must be justified by host implementations.
 - Especially given the reduced complexity and cost of the Basic queing model as an alternative.
- Systems integrators and tape vendors must justify the need for queing in stream devices.

Closure (con't)

- Issues should be reopened on the Serial Solutions reflector.
 - Need to justify cost and complexity.
- Need to reach closure on sequential device model.
 - I'll work with the tape and systems vendors to understand costs and benefits.