

To: X3T9.2 Membership

From: Edward A. Gardner, Digital Equipment Corporation
gardner@ssag.enet.dec.com

Subject: The Meaning of Linked Commands in SCSI-3

At December's meeting in San Diego there was a lively discussion of what linked commands might mean or imply in SCSI-3, possibly incorporating the IPI queuing model. Without judging the many IPI queuing modes, I suggest we adopt the following principles:

1. The meaning of or guarantees provided by individual commands should not be implicitly altered by the command being linked to others. For example, assume that a disk device implements an operating mode with unrestricted re-ordering and multiple commands executing concurrently. Two multi-block writes to the same disk blocks might result in the blocks containing a mixture of the data from the two writes. If this were the case for unlinked commands, simply making the writes linked commands (in different IO processes) should not alter anything. (Note that this is just an example, I am not advocating that SCSI encourage such an operating mode).
2. When we do want to provide special guarantees for a command or command sequence, it should be explicitly specified using a queue tag message or similar mechanism. For example, an indivisible read-modify-write sequence is desirable in some file systems. If we choose to add such a feature to SCSI, it should be explicitly requested by the protocol, most likely by a newly defined "indivisible action" queue tag message. Note that the same queue tag message used with the concurrent overlapping writes of the previous example would ensure that one write's data or the other's would result, not a mixture.
3. RESERVATION CONFLICT status shall only be returned for conflicts induced by use of the RESERVE or RESERVE UNIT commands. Execution requirements imposed by queue tag messages or similar mechanisms will only defer or alter the execution scheduling of an IO process, they do not cause other commands or IO processes to be rejected.