

To: X3T9.2 Membership

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

Subject: Previewing Linked Commands in SCSI-3

In discussions on the interaction between IO process queuing and linked commands, it has become clear that devices could optimize queued IO processes more effectively if they could preview linked commands. In SCSI-1 and SCSI-2, devices can only see the current command of each IO process. Thus they must make worst-case assumptions about subsequent linked commands, or in effect disable re-ordering whenever linked commands are used. If the device could preview the entire IO process, it could make more effective re-ordering and optimization decisions.

The following proposes a mechanism to allow SCSI-3 devices to preview all the commands in an IO process. The mechanism has two principal advantages. First, it is compatible with the present operation of SCSI-1 and SCSI-2. Second, a device may preview commands for many or all IO processes without necessarily having to store all the commands previewed.

An IO process consists of one or more ordered commands or CDBs. The last command in an IO process has the link bit clear in its CDB. All commands except the last in an IO process have the link bit set in their CDBs. The initiator shall not alter the CDBs in an IO process (including their sequence) after the IO process has been initiated.

The saved command pointer always points to the start of the CDB for the current command of the IO process. Whenever an IO process becomes active, the target may enter the Command Out bus phase to fetch the current command's CDB. If the link bit is set in that CDB, the target may optionally continue the Command Out bus phase to fetch the CDB of the next command in the IO process. The target may continue fetching CDBs until it obtains one whose link bit is clear, denoting the last command of the IO process.

A target may fetch the same CDB as many times as it wishes. So long as a command remains current, the saved command pointer will point to its CDB. Each time the IO process becomes active, or the target issues a RESTORE POINTERS message, a subsequent Command Out bus phase will retransfer the current command's CDB followed by subsequent CDBs in the IO process.

The saved command pointer is advanced to the next command's CDB whenever the current command completes. /tbd, should say whether this occurs on status transfer or Linked Command Complete message/.

Implementor's Note: when an IO process is first issued, devices might review all of the IO process's commands to obtain information for making optimization decisions. Rather than storing the entire command list internally, the device might instead discard them and re-fetch each command as it is executed.