

To: T10 Membership
From: Paul Suhler, Seagate Technology
Date: 27 October 2000
Subject: Target Handling of Out-of-Order
Commands, T10/00-405r0



This document outlines the need for specification at some level of how a target should respond when a command is received whose CRN is two or more greater than the previous command with a non-zero CRN. The intent is for the Joint T10/T11.3 Working Group to choose a means for handling this, and for me to then develop a detailed proposal.

FCP-2, rev. 4a, section 4.3, paragraph e) specifies that when precise delivery control is enabled, a device server shall not accept a command with non-zero CRN into the task set until all commands with a previous CRN have been received. There are a number of options for initiator-driven recovery, but not for target-driven recovery.

In particular, for non-queued operation if the target chooses to reject the command, how should it do so? The choices seem to be:

1. Handle it at the transport level by reporting a (probably-new) RSP_CODE.
2. Handle it at a higher level by reporting a CHECK CONDITION with appropriate SK/ASC/ASCQ.
3. Report BUSY.

In what standards and/or profiles should this be addressed?

In queued operation, there are more options, though we may not want to specify what to do. Options are:

1. Reject each subsequent command as above.
2. Hang onto subsequent commands, hoping that the missing one will arrive, filling the inside straight and allowing the subsequent commands to continue. If we run out of resources, return either BUSY, CHECK CONDITION, or the RSP_CODE for later commands.

Should this behavior be addressed in FCP-2? In a profile as part of error recovery?