

T10/03-243 revision 0

Date: July 02, 2003

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: Optional INQUIRY command processing before entering task set

1 Overview

We have seen problems at customers involving failover, where one host is trying to take over the LUNs for a failed host. There were some LUNs on the failed host which had commands queued (some with ordered tags). The second host issued an Inquiry which didn't complete because of the proceeding ordered commands.

The second host could have gotten out of this situation by doing a LUN reset, but the software didn't want to do the LUN reset until it verified it was talking to the correct device (by comparing unit serial page).

Some SCSI devices handle this problem by treating Inquiry as a 'priority' command. What that means is that they will respond to Inquiry immediately upon receipt regardless of what is in the queue or how it is tagged. Those devices have done it that way since the first generation SCSI product.

To address this the following proposal would make it legal for a logical unit to process an INQUIRY command before it is entered into the task set without regard to it's tag or any tag for any command in the task set.

2 Proposal

In section 8.1 Introduction to task set management change the following:

The requirements for task set management only apply to a task after it has been entered into a task set. A task shall be entered into a task set unless:

- a) A condition exists that causes that task to be completed with a status of BUSY, RESERVATION CONFLICT, TASK SET FULL, or ACA ACTIVE;
- b) Detection of an overlapped command (see 5.9.3) causes that task to be completed with a CHECK CONDITION status; or
- c) SCSI transport protocol specific errors cause that task to be completed with a status other than GOOD.

to:

The requirements for task set management only apply to a task after it has been entered into a task set. A task shall be entered into a task set unless:

- a) A condition exists that causes that task to be completed with a status of BUSY, RESERVATION CONFLICT, TASK SET FULL, or ACA ACTIVE;
- b) Detection of an overlapped command (see 5.9.3) causes that task to be completed with a CHECK CONDITION status;
- c) SCSI transport protocol specific errors cause that task to be completed with a status other than GOOD; or
- d) it is an INQUIRY command in which case the device server may process the command without regard to the INQUIRY commands task attribute or the task attribute of any task currently in the task set.