Document:	T10/00-359r0	Date:	September 13, 2000
To:	T10 Committee Membership		
From:	Edward A. Gardner, Ophidian Designs		
Subject:	Unit Attention Issue		

When discussing SVP with the IBTA Application Working Group, I was asked to explain how Unit Attention worked in non-interlocked protocols. After some investigation I determined that it did not work the way I had expected, and might be considered by some to be broken. This document describes one awkward scenario to ask how it ought to be handled.

Consider the configuration:

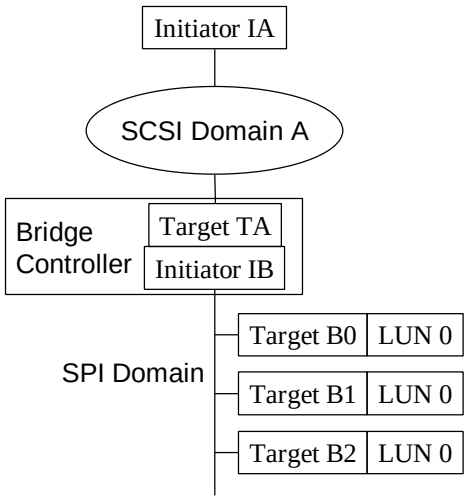

SCSI Domain A uses a non-interlocked protocol such as FCP or FCP-2. Targets B0, B1, etc. are normal parallel SCSI disk drives. The bridge controller is a parallel SCSI to Fibre Channel adapter (JBOD), presenting the devices behind it as individual LUNs. For example, it might present target B0 LUN 0 as its LUN 0, target B1 LUN 0 as its LUN 1, etc.

Assume that all commands specify NACA=0. Since support for NACA=1 is optional, one expects NACA=0 to be the common usage.

Consider the following sequence of events:

1. Initiator IA (the host OS) has mounted and is accessing the LUN corresponding to target B1.
2. During a period of inactivity (no outstanding commands) an operator removes target B1 and replaces it with a different drive (hot-swap). The reason for the hot-swap is irrelevant; perhaps the operator was supposed to swap a different drive but chose the wrong one.
3. Shortly after the hot-swap, initiator IA issues a burst of commands to the LUN corresponding to target B1. SCSI Domain A is FCP or some equivalent, multiple commands may be issued and in-flight before any response is received.
4. Per the definition of Unit Attention conditions and how they are cleared, the first command of the burst completes immediately, returning CHECK CONDITION and auto-sense data describing the Unit Attention condition (power-on or the like). This clears the Unit Attention condition, allowing the remaining commands of the burst to be executed against the wrong drive.

In this scenario initiator IA is informed that the LUN has been changed, although too late to prevent commands from being executed improperly. Another interesting scenario is when the response IU for the first command is lost, since that is the only notification to the initiator that a hot-swap has occurred.

My question is how should such scenarios be handled or avoided? I've come up with the following possibilities:

1. Don't use command queuing. This is impractical for performance reasons.
2. Specify NACA=1. My impression was that most current targets and initiators do not support this. If this is the preferred solution, then support for NACA=1 should be mandatory.
3. Prohibit hot-swap of LUNs. This is impractical to say the least.
4. Require that bridge controllers manufacture a new, unique LUN whenever a hot-swap occurs. In the example described above, the bridge controller would permanently retire the LUN assigned to the old drive, then use a new LUN for the new drive. I question whether this is practical and don't believe existing bridge controllers operate this way.

How do existing bridge controllers (also called JBOD adapters or controllers) and existing host software deal with this? Do they just ignore it? How ought we to deal with it?