

Document: T10/00-357r0
To: T10 Committee Membership
From: Edward A. Gardner, Ophidian Designs
Subject: Identifying Devices

Date: September 13, 2000

When discussing SVP with the IBTA Application Working Group, several AWG members described a problem identifying devices in complex configurations. AWG requested that I bring this to the attention of T10 and seek a solution. Note that this problem is not specific to SVP, it exists with FCP and other protocols as well. It is not an architectural problem, rather a practical issue resulting from having to deal with existing devices.


Host operating systems need some way to identify devices or LUNs (specific physical LUNs or disk spindles). There must be some persistent identifier associated with each LUN that the host OS can use to determine whether it is talking to the same LUN. One manifestation of this is that the host OS needs to locate the same LUN following a power cycle or re-configuration. For example, it needs to locate its boot disk. Another manifestation is needing some identifier to recognize multiple paths to the same LUN.

The vital product data device identification page is intended to address this problem. The difficulty is that this page and use of device identifiers is optional. It is true that device identifiers are unnecessary in simple configurations. However, target devices that are intended for simple configurations may, through bridge controllers, be used in complex configurations.

The simplest approach would be to mandate device identification page support for all SCSI devices. This is impractical, if only because there are already so many existing devices that do not support that. A plausible solution is to define one or more algorithms by which bridge controllers can construct suitable device identifiers and emulate support for the device identification page. However, for this to work bridge controllers from different vendors would have to use the same algorithms. That is the reason this is being brought to T10.

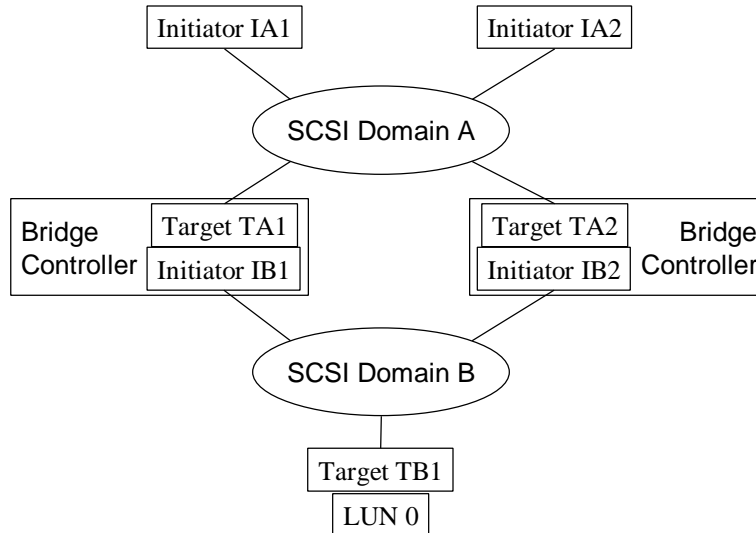
Note: FCP was silent on the subject of device identifiers. FCP-2 requires support for device identifiers (fcp2r04.pdf, clause 5.3, page 17, pdf page 33). How is this requirement met for LUNs that are not native to FCP-2? E.g., an FCP-2 JBOD controller with FCP or SPI disks behind it? This is the same issue, the same problem, that is addressed by this document.

Consider the following configuration:



For the sake of discussion assume that the SCSI Domain shown uses the FCP protocol, and that target T1 Lun 0 does not support the device identification page. Fibre Channel address or port identifiers are dynamic and may change following a reboot, reconfiguration, or a new phase of the moon. In the absence of device identification page support, the only reliable way for the initiators to identify target T1 is through its Fibre Channel port or node names (world-wide names). Therefore the LUN would be identified, within the host OS, by the tuple (T1.nodename, LUN 0). This has some problems, since there is no guarantee that all ports to the same target will return the same node name, but can be made to work with most devices.

Now consider the following configuration:



As before, the only reliable identifier for target TB1 is its Fibre Channel port and node names. Unfortunately, the bridge controller(s) hide that information, it is not accessible to initiators IA1 and IA2 (the host OS).

One approach to solving this would be to make the LUN field arbitrarily large, so the port or node names could be embedded in it. I don't believe that is practical. Another would require that the bridge controllers remember, in non-volatile memory, every port or node name they have ever seen so to always use the same LUN to target mapping. This also suffers from practicality problems, and provides no way to recognize that the same LUN is accessible through both bridge controllers.

A plausible approach is for the two bridge controllers to use the target's port or node names to construct device identifiers for target TB1, LUN 0. If initiators IA1 or IA2 requested the device identification page for a LUN that did not support that page (e.g. target TB1, LUN 0), the bridge controllers would return the constructed device identifiers. Since the device identifiers are constructed from information specific to target TB1, it is possible to choose an algorithm whereby both bridge controllers would construct the same device identifier(s) for each LUN.

My first question for T10 is whether this is the right technical approach for dealing with this issue. My second question is where it should be documented. For FCP devices, I suggest an informative annex to FCP-2, which might be justified as documentation for how to satisfy FCP-2's requirement that device identifiers be supported. For SPI devices there is no analog of the Fibre Channel port or node names. While we could document an algorithm that will work for a single whatever-to-SPI bridge controller, there is no way to construct device identifiers for multiple bridge controllers without resorting to vendor unique techniques.

I request direction from T10 on how to proceed.