



September 11, 1998

To: T10 Committee

From: Tom Coughlan
Compaq Computer Company
Mail Stop ZKO-3-4/U14
110 Spitbrook Road
Nashua, New Hampshire
Telephone: 603-884-0933
E-mail: tom.coughlan@digital.com

Subject: Changeable Device Identifiers

Summary

There is a requirement in some operating environments for the user to be able to associate a persistent identifier with a SCSI device. This proposal provides the required functionality by moving two command service actions from SCC-2 to SPC-2. These command service actions are defined for all device types. Implementation of these commands is optional.

Background

Many operating systems have traditionally assigned names to SCSI devices based on the path from the operating system (OS) to the device. These naming schemes are convenient because they provide names that are unique, and they convey some information about the logical or physical location of the device. Path-based names are problematic, however, when:

1. there are multiple paths from an OS instance to the device
2. there are multiple OS instances in the same naming domain with access to the same device (e.g. a cluster)
3. the interconnect employs path identifiers that are not as persistent as is required for OS device identification. Fibre Channel is one such interconnect.

One solution to these problems is to devise a device naming scheme that is based on Word Wide Identifiers (WWIDs). In such a scheme, each SCSI device is required to provide a persistent, path-independent, world-wide unique identifier in its Inquiry data. The operating system uses the WWID as the basis for device naming, thereby avoiding the problems with path-based naming listed above.

The WWID formats that are typically being implemented for SCSI are either 64-bit or 128-bit binary values. These identifiers are too long to be usable by humans as device

identifiers. To resolve this difficulty, the OS (or a group of OS instances in a cluster) will typically assign a short alias to each WWID, and will present this to the user as the device name. The OS is responsible for maintaining the alias-to-WWID mapping in such a way as to provide the device naming consistency and persistence that is expected by the users. For this discussion, we will assume that the OS implements a scheme for automatically assigning aliases to WWIDs, and that it implements a means for the user to modify aliases as desired to create meaningful device names.

The Problem

WWID-based naming does not work well in certain environments. For example:

1. In many system installations, the individuals who replace failed storage devices do not have access to the operating system. This presents a problem when WWID-based naming is in use, because the replacement storage device will, by default, have a different OS device name. The default OS name can be overridden, or OS parameters can be changed to compensate for the new name, but these actions require the involvement of personnel with different expertise, adding to the cost of the repair.
2. Producers of turn-key systems desire to ship and maintain identical copies of the OS on identically-configured hardware systems. This is not possible when WWID-based naming is in use, because the WWID-to-alias mapping on each system is necessarily different.
3. In some installations it is necessary to boot different instances of an operating system at different times. It is desirable for the device names in these environments to match. This is difficult to achieve in general with WWID-based naming, and may be impossible in an environment where there is a read-only (CD-ROM) system disk.
4. The default device name does not provide information about the location of the device. New tools are needed to provide the user with a mapping from the default WWID-based device name to the device's logical or physical location within the configuration. These tools may not be available in the same timeframe as when path-independent device names are required.

The Proposal

A variation on WWID-based naming is proposed to address the environments where the properties of WWID-based naming cause difficulties. In the proposed scheme, each SCSI device optionally implements a writable, persistent, device identifier.

The initial version of this proposal added the changable identifier to the Inquiry page 83 data. An advantage of this is that the identifier can be obtained without clearing unit attentions. Mode select and mode sense commands were proposed for setting and reporting the identifier. This proposal was rejected because of a reluctance to add changeable data to the Inquiry data.

The second proposal was to add two new commands to SPC for setting a reporting the device identifier. The difficulty here is that there is a scarcity of OPCODEs that can be assigned for all devices, at least for CDBs that are less than 16 bytes.

The OPCODE problem is avoided in the third revision of this proposal, by using two commands that are already defined in SCC. These commands have service actions to set and report a device identifier that can be used by the operating system for device naming. The proposal is to move those two service actions to SPC-2, and make them optional for all device types. One negative implication of this approach is that reading the identifier will clear unit attention conditions. The operating system must be prepared to deal with this.

Detailed Proposal

Two commands, each with one defined service action, are added to SPC-2. SCC-2 is unchanged.

0.1 Maintenance In command

The Maintenance In command (see table 1) is used to request that the logical unit send device identification information to the application client. This command is optional for all device types, except for SCC devices, and devices that set the SCCS bit in their Standard Inquiry data. The SCC-2 Standard defines additional service actions for this command, that apply to SCC devices, and devices that set the SCCS bit in their Standard Inquiry data (refer to SCC-2).

Table 1 --- Maintenance In command

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation code (A3h)							
1	Reserved			Service Action (05h)				
2	Reserved							
3	Reserved							
4	LUN							
5								
6	Allocation Length							
7								
8								
9								
10	Reserved						PORCLU	Reserved
11	Control							

The Maintenance In command with a Service Action equal to 05h shall not be affected by reservations or persistent reservations.

The Service Action field shall equal 05h. The SCC-2 Standard defines additional Service Action values for SCC devices, and devices that set the SCCS bit in their Standard Inquiry data.

The LUN field is not defined for device types other than SCC devices, and devices that set the SCCS bit in their Standard Inquiry data (refer to SCC-2).

The Allocation Length field in the command descriptor block indicates how much space has been reserved for the returned parameter data. If the length is not sufficient to contain all the parameter data, the first portion of the data shall be returned. This shall not be considered an error. The actual length of the parameter data is available in the Identifier Length field in the parameter data. If the remainder of the parameter data is required, the application client should send a new Maintenance In command with an Allocation Length field large enough to contain all the data.

The report physical or component logical unit bit (PORCLU) is not defined for device types other than SCC devices, and devices that set the SCCS bit in their Standard Inquiry data (refer to SCC-2).

The Maintenance In parameter list (see table 2) contains a four-byte field that contains the length in bytes of the parameter list and the logical unit's Identifier.

Table 2 --- Maintenance In parameter list

Bit Byte	7	6	5	4	3	2	1	0
0	Identifier Length							
1								
2								
3								
4	Identifier							
n								

The Identifier Length field specifies the length in bytes of the Identifier field. If the Allocation Length field of the command descriptor block is too small to transfer all of the identifier, the length shall not be adjusted to reflect the truncation. The Identifier Length shall initially equal zero, and shall be changed only by a successful Maintenance Out command with a Service Action equal to 05h, and a non-zero List Length.

The Identifier field shall contain a vendor specific ASCII value. The value reported shall be the last value written by a successful Maintenance Out command with a Service Action equal to 05h.

The value of the Identifier shall be changed only by a Maintenance Out command with a Service Action equal to 06h. The Identifier value persists through resets, power cycles, media format operations, and media replacement.

The logical unit shall return the same Identifier to all initiators, on all of its ports.

The execution of a Maintenance In command with a Service Action equal to 05h may require the enabling of a nonvolatile memory within the logical unit. If the nonvolatile memory is not ready, the device server shall return CHECK CONDITION status, rather than wait for the device to become ready. The sense key shall be set to NOT READY and the additional sense data shall be set as described in the TEST UNIT READY command (see 7.24). This information should allow the application client to determine the action required to cause the device server to become ready.

0.0 Maintenance Out command

The Maintenance Out command (see table 3) requests that the device identifier information in the logical unit be set to the value received in the Maintenance Out parameter list. This command is optional for all device types. The SCC-2 Standard defines additional service actions for this command, that apply to SCC devices, and devices that set the SCCS bit in their Standard Inquiry data (refer to SCC-2).

On successful completion of this service action a unit attention shall be generated for all initiators except the one that issued the service action. When reporting the unit attention condition the additional sense code shall be set to DEVICE IDENTIFIER CHANGED.

Table 3 --- Maintenance Out Command

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation code (A4h)							
1	Reserved			Service Action (06h)				
2	Reserved							
3	Reserved							
4	LUN							
5								
6	List Length							
7								
8								
9								
10	Reserved					IDPORC	Reserved	
11	Control							

The Maintenance Out command with a Service Action equal to 06h shall be affected by reservations or persistent reservations in the same manner as the Mode Select command (refer to 98-164 for more specific information).

The Service Action field shall equal 06h. The SCC-2 Standard defines additional Service Action values for SCC devices, and devices that set the SCCS bit in their Standard Inquiry data.

The LUN field is not defined for device types other than SCC devices, and devices that set the SCCS bit in their Standard Inquiry data (refer to SCC-2).

The List Length field specifies the length in bytes of the Identifier that shall be transferred from the application client to the device server. A parameter list length of zero indicates that no data shall be transferred, and that subsequent Maintenance In commands with a service action of 05h shall return an Identifier length of zero

The set identification peripheral device or component device bit (IDPORC) is not defined for device types other than SCC devices, and devices that set the SCCS bit in their Standard Inquiry data (refer to SCC-2).

The Maintenance Out parameter list (see table 4) contains the identifier to be set by the addressed logical unit.

Table 4 --- Maintenance Out parameter list

Bit Byte	7	6	5	4	3	2	1	0
0	Identifier							
N								

The IDENTIFIER field shall be a vendor specific ASCII value, to be returned in subsequent Maintenance In commands that have a service action equal to 05h.