

Date: June 20, 2001

To: T10 Committee (SCSI)

From: Jim Hafner (IBM) (hafner@almaden.ibm.com)

Subject: SAM-2 device and port names

## 1.0 Introduction

The multi-port model in the SCSI Architecture Model (SAM-2, see sam2r18) has carefully defined the constructs of “device” and “port”. Devices are naturally only addressed through ports, but there is a growing need to identify devices behind the ports. For example, it might be easier for a backup application intending to use EXTENDED COPY functions to reference the target device of a logical unit (by name in a target descriptor) and not have to specify the specific target port through which the copy manager must discover and communicate with the logical unit’s device server. Additionally, ports have traditionally only been identified by addresses, though even in FCP it was determined that port name constructs added value.

Furthermore, protocols such as iSCSI and SRP (and in retrospect FCP) have exposed the need for SAM-2 to contain the “name” concept for both SCSI ports and SCSI devices. FCP has used the WWPortName as a mechanism for reservations (in particular persistent reservations) to survive changes in the underlying FC addressing (such addressing changes should, by layering principles, be invisible to the SCSI layer). Additionally, FC has a notion of WWNodename that could have mapped nicely to SCSI device name, but it did not work out that way in practice; now FC is considering a “platform name” that might be used for this purpose. iSCSI’s addressing layer is both ephemeral and completely transparent to the SCSI layer; consequently, it needs to strongly leverage the iSCSI Name as the SCSI device name for the purposes of third party commands (a target descriptor can reference a logical unit and target device without specifying, necessarily, the target port through which that logical unit and device should be addressed). Similarly, SRP is providing a naming layer for ports that is above the network addressing layer; only the name layer is exposed directly to the SCSI layer.

This proposal adds the name concept to the SAM-2 model for both ports and devices (see 3.0 for a summary of the proposed changes); there are no limits to the size of names. The former 8-byte limit on identifiers seems to have disappeared in SAM-2 rev 18, so this proposal does not need to take that action.

With names introduced into the model, it would be possible to change other parts of SCSI in ways currently not available.

- a) provide generic language in reservation clause about “address” changes. In other words, we can have this clause say that the owner of a reservation is the port as identified by its name, not by its address. This removes the requirement to track address changes, or at least makes this more opaque to the SCSI layer.

- b) In the same reservation clause, we can redefine the terms “other initiators” so that the initiator that owns the port is identified by its name, independent of target port, provided it is clear to the device server that the named initiator is the same one (i.e., lives in the same SCSI domain). In that way, we get for free the property that a globally uniquely named port can access a logical unit through *any* target port with identical behavior.
- c) We could more easily extend the notion of reservation owner to the initiator device (by name) and make the behavior of all its ports identical with respect to the reservation.
- d) We could use the name of the device in third party contexts; this removes the dependence on particular port connectivity requirements. For example, if a copy manager has connectivity to a particular named target device through more than one of that target’s ports, it can chose which port is best. Today, only ports can be identified and that limits the options of the copy manager.
- e) We could use the name of the port or the device generically in access controls (right now, with the exception of iSCSI, we’re effectively only using the port name for TransportID - the AccessID was defined to give a higher level name/identifier than the port that was not available across all transports). The device name construct allows us to move in this direction.

This proposal stops well short of adding any of these potential changes to SCSI. It is concerned solely with adding the notion of name to ports and devices. Subsequent proposals may take on these additional tasks.

Section 2.0 of this document opens a broader debate about the fundamental notions of name and identifier, how they should/could be used and what they could mean. This is only opinion and is intended to spark debate. Section 3.0 briefly summarizes the proposed changes to SAM-2 (rev 19) and Section 4.0 details the changes.

## 2.0 Author’s Soapbox

### 2.1 “Identifiers”: addresses or what?

My personal preference is that the term “identifier” be converted to a more generic term that does not equate, as it does now, to the term “address”. I believe layering principles say that “address” is something that is within the scope of the service delivery subsystem. That subsystem can pass an identifier or handle to the SCSI port layer as a way to indicate a particular port, but that identifier need not map directly or indirectly to a subsystem address. For FCP, the N\_Port is a network address, but it has no intrinsic meaning to the SCSI layer; it is completely opaque to SCSI; it is only a pointer/handle shared between the SCSI and FC layers to indicate a port. The FC layer could easily create any other identifier (say permute the bits) and use that as the FC to SCSI layer identifier of a given port. That new identifier does not correspond to any address, but does identify a unique port. Additionally, there are situations (e.g., FC LIP) where the addresses change in ways that are supposed to be transparent to the SCSI layer (e.g., for persistent reser-

ventions). Narrowing the scope of the term “identifier” to mean “address” only forces the architecture to cross these layering boundaries in ways that are not required.

I have not chosen to write this proposal along these lines as I’ve run into resistance to this approach in private communications. Instead, this proposal concentrates on adding only the term “name” where I think it is needed and does not attempt to change the definition of “identifier”.

## 2.2 Names: mandatory or optional?

For names, my personal opinion is that all devices and ports should be required to have names (the names may be protocol-derived, but they can exist). This allows the model to rely on name constructs directly (e.g., in persistent reservations) rather than relying on the notion of identifier. Identifiers, as defined as addresses, can change on-the-fly (e.g., FC LIP) and are tightly bound to the protocol; names, though in practice are derived from the protocol are functionally independent and are more permanent. This provides a better foundation and better layering. Additionally, I believe there is an interpretation of current implementations of the model that supports this point of view in all protocols. Others disagree with this (see, for example, T10/01-084r2). Here’s how I see it.

SPI: in this case, it is possible (without breaking anything) to view every device as effectively single ported. Then we can equate the device name to the port name to the bus address. This has all the properties of uniqueness within the SCSI domain required for names. It is also the “degenerate” (as in simplest) interpretation of the model.

FCP: For the current view, we again assume that each device is single ported. We equate the device name with the port name (WWPortName). We equate the port identifier with the N\_Port (even though we don’t have to). This gives us uniqueness of identifiers only within a “local” FCP SCSI domain and uniqueness of names within the “global” FCP SCSI domain (or more precisely, independent of the local FCP SCSI domain). For FCP-3, it might be possible to leverage the platform name to go beyond the interpretation of single ported devices into truly multiported devices (because we’ll have a label/name for the device).

An argument against this point of view is that manufacturers are building multiported targets today (like dual ported drives or dual-headed multiported controllers). My response is a question: What is functionally different today between the interpretation of these “boxes” as multiple single-ported-devices that *share* one or more logical units and a single multi-ported-device containing one or more logical units. There may be different things happening with respect to caching, etc., but the logical unit’s view (actually the device server’s) simply sees some set of ports and it is effectively blind about the device boundary between those ports. A manufacturer may not think in the terms I’ve outlined, but I don’t think the model is fundamentally inconsistent with this point of view.

To complete the protocol picture:

SRP and iSCSI: In both of these protocols, names are being equated with identifiers (at the SCSI port level) and all network addressing is hidden from the SCSI layers. SRP has

not yet defined names for devices, and names for initiator ports are fairly dynamic things (derived from a root name). For iSCSI, names for devices are the root of the naming scheme, and names for ports are derived from them. So here, the concept of name is fundamental.

SBP: I don't know this well enough to comment.

## 2.3 Conclusions

In spite of the above perspective, this first approach to adding names to the model takes the position that names are optional (more precisely, they are required, optional or not as defined by the protocol standards). This is the "path of least resistance" to getting approval. If discussion concludes that names can be made mandatory, the author is more than willing to modify the proposal along those lines.

We have also opted for devices to have zero or more names (i.e., even more than one). This isn't the best modelling principle, but it seems to be necessary in practice for the following reasons. Some devices may have interfaces to different transport protocols at the same time. But different protocols are defining device naming conventions in different ways (see, for example, iSCSI and FCP). Other protocols (as noted above) may not define a name at all. The "zero or more device names" in this proposal accommodates these concerns.

## 3.0 Summary of proposed changes

The following summarizes the proposed changes to SAM-2:

- a) define the terms "device name" and "port name" (and related terms) and add them to the model clause appropriately (including the figures) glossary
- b) add the above terms to the glossary.

Names have the following properties:

- a) names are unique within the SCSI domain in which they are presented (this is specified both for device names that might be presented during some login phase of the transport protocol and for port names);
- b) names may be worldwide unique, i.e., unique across all SCSI domains
- c) both device and port names are optional in the model, however protocols may define a requirement for device or port names;
- d) a port may have at most one name (so zero or one);
- e) a device may have more than one name (so zero or more); see 2.3 for the rationale.

## 4.0 Detailed description of proposed changes to SAM-2 rev 19

### 4.1 Changes for clause “3.1 Definitions”

Add the following entries:

**3.1.aa initiator device name:** a SCSI device name of a SCSI initiator device or of a SCSI target/initiator device when operating as a SCSI initiator device (see 4.7.1).

**3.1.bb initiator port name:** a SCSI port name of a SCSI initiator port or of a SCSI target/initiator port when operating as a SCSI initiator port (see 4.7.1).

**3.1.cc SCSI device name:** a name of a SCSI device that is unique within the SCSI domain in which the device has SCSI ports (see 4.7.xx). The name may be presented to other devices or ports in that domain in protocol specific ways.

**3.1.dd SCSI port name:** a name of a SCSI port that is unique within the SCSI domain of that port (see 4.7.yy). The name may be presented to other devices or ports in that domain in protocol specific ways.

**3.1.ee target device name:** a SCSI device name of a SCSI target device or of a SCSI target/initiator device when operating as a SCSI target device (see 4.7.2).

**3.1.ff target port name:** a SCSI port name of a SCSI target port or of a SCSI target/initiator port when operating as a SCSI target port (see 4.7.2).

### 4.2 Changes to model clause “4.7 SCSI devices”

[AUTHOR'S NOTE: The description for changes to figures is textual because the author is “graphically impaired”!]

#### 4.2.1 Changes to clause “4.7.1 SCSI initiator device”

Change the text at the beginning of the clause as indicated:

A SCSI initiator device (see figure 11) contains:

- a) one or more SCSI initiator ports each containing an initiator port identifier\_ and an optional initiator port name; and
- b) zero or more initiator device names; and
- c) zero or more application clients.

A protocol standard may place additional requirements for initiator device names (see 4.7.xx) and initiator port names (see 4.7.yy).

Change Figure 11 to contain the following:

- a) a zero-or-more box on a level with the “SCSI Initiator Port” and “Application Client” boxes and connected to the top “SCSI Initiator Device” box that contains the term “Initiator Device Name”
- b) a zero-or-one box on a level with the “Initiator Port Identifier” box and connected with the “SCSI Initiator Port” box that contains the term “Initiator Port Name”.

Add the following text after the paragraph that begins “An initiator port identifier...”

An initiator device name is a name that is a SCSI device name (see 4.7.xx) for the initiator device.

The initiator port name is a name that is the SCSI port name (see 4.7.4) for the initiator port.

#### **4.2.2 Changes to clause “4.7.2 SCSI target device”**

Change the text at the beginning of the clause as indicated:

A SCSI target device (see figure 12) contains:

- a) one or more SCSI target ports each containing a task router and SCSI target port identifier and an optional target port name; and
- b) zero or more target device names; and
- c) zero or more application clients.

A protocol standard may place additional requirements for target device names (see 4.7.xx) and target port names (see 4.7.yy).

Change Figure 12 to contain the following:

- a) a zero-or-more box on a level with the “SCSI Target Port” and “Logical Unit” boxes and connected to the top “SCSI Target Device” box that contains the term “Target Device Name”
- b) a zero-or-one box on a level with the “Target Port Identifier” and “Task Router” boxes and connected with the “SCSI Target Port” box that contains the term “Target Port Name”.

Add the following text after the paragraph that begins “A SCSI target port identifier...”

A target device name is a name that is a SCSI device name (see 4.7.xx) for the target device.

The target port name is a name that is the SCSI port name (see 4.7.4) for the target port.

### 4.2.3 Changes to clause “4.7.3 SCSI target/initiator device”

Change the text at the beginning of the clause as indicated:

A SCSI target/initiator device (see figure 13) contains:

- a) one or more SCSI target/initiator ports each containing a task router, target port identifier, ~~and an initiator port identifier,~~ an optional target port name and an optional initiator port name; and
- b) zero or more target device names; and
- c) zero or more initiator device names; and
- d) zero or more application clients.

A protocol standard may place additional requirements for initiator device names and target device names (see 4.7.xx), and for initiator port names and target port names (see 4.7.yy).

Change Figure 13 to contain the following:

- a) a zero-or-more box on a level with the “SCSI Target/Initiator Port”, “Logical Unit” and “Application Client” boxes and connected to the top “SCSI Target/Initiator Device” box that contains the term “Target Device Name”
- b) a zero-or-more box on a level with the “SCSI Target/Initiator Port”, “Logical Unit” and “Application Client” boxes and connected to the top “SCSI Target/Initiator Device” box that contains the term “Initiator Device Name”
- c) a zero-or-one box on a level with the “Task Router”, “Target Port Identifier” and “Initiator Port Identifier” boxes and connected with the “SCSI Target/Initiator Port” box that contains the term “Target Port Name
- d) a zero-or-one box on a level with the “Task Router”, “Target Port Identifier” and “Initiator Port Identifier” boxes and connected with the “SCSI Target/Initiator Port” box that contains the term “Initiator Port Name”.

Add the following text after the paragraph that begins “The target port identifier...”

A target device name and an initiator device name are names that are SCSI device names (see 4.7.xx) for the target/initiator device when operating as a target device and initiator device, respectively. A target device name and an initiator device name may or may not be identical.

The target port name and the initiator port name are names that are the SCSI port name (see 4.7.4) for the target/initiator port when operating as a target port and initiator port, respectively. The target port name and the initiator port name may or may not be identical.

#### 4.2.4 Additional clauses to “4.7 SCSI devices”

Add the following clause after “4.7.4 SCSI port identifier”.

##### **4.7.xx SCSI device name**

A SCSI device name is an optional label of a SCSI device that is unique within a SCSI domain in which the device has SCSI ports. A SCSI device may have more than one name if that device has SCSI ports in different domains. A name shall never change and may be world wide unique in a way that is independent of the SCSI domain. A SCSI device name is used to persistently identify a SCSI device in contexts where specific references to port names or port identifiers is not required.

A protocol standard may require that a SCSI device include a SCSI device name if the device has SCSI ports in a domain of that protocol. Additionally, a protocol standard may require that a SCSI device name be world wide unique. The name may be presented to other devices or ports in a given SCSI domain in protocol specific ways.

##### **4.7.yy SCSI port name**

A SCSI port name is an optional label of a SCSI port that is unique within the SCSI domain of that port. A SCSI port may have at most one name. A name shall never change and may be world wide unique in a way that is independent of the SCSI domain. A SCSI port name may be used to persistently identify a SCSI initiator port or SCSI target port in contexts similar to that of SCSI port identifier (see 4.7.4).

A protocol standard may require that a SCSI port include a SCSI port name if the port is in a domain of that protocol. Additionally, a protocol standard may require that a SCSI port name be world wide unique. The name may be presented to other devices or ports in the given SCSI domain in protocol specific ways.