

Date: Aug 07, 2000

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: Defining Targets/Initiators as Ports

Overview

In looking over the issue of how SCSI should deal with multiple port devices it appears a reasonable solution would be to change the notion of what a target and initiator is defined as. If you confine a target or initiator to only contain a service delivery port and allow a SCSI device to contain any number of targets or initiators then defining multi-port devices becomes an easier task. In many ways the current definitions and descriptions already support this idea.

What follows are suggested changes to SAM that would define targets and initiators as ports and allow multiples on a single SCSI device.

Multiple port target SCSI Device

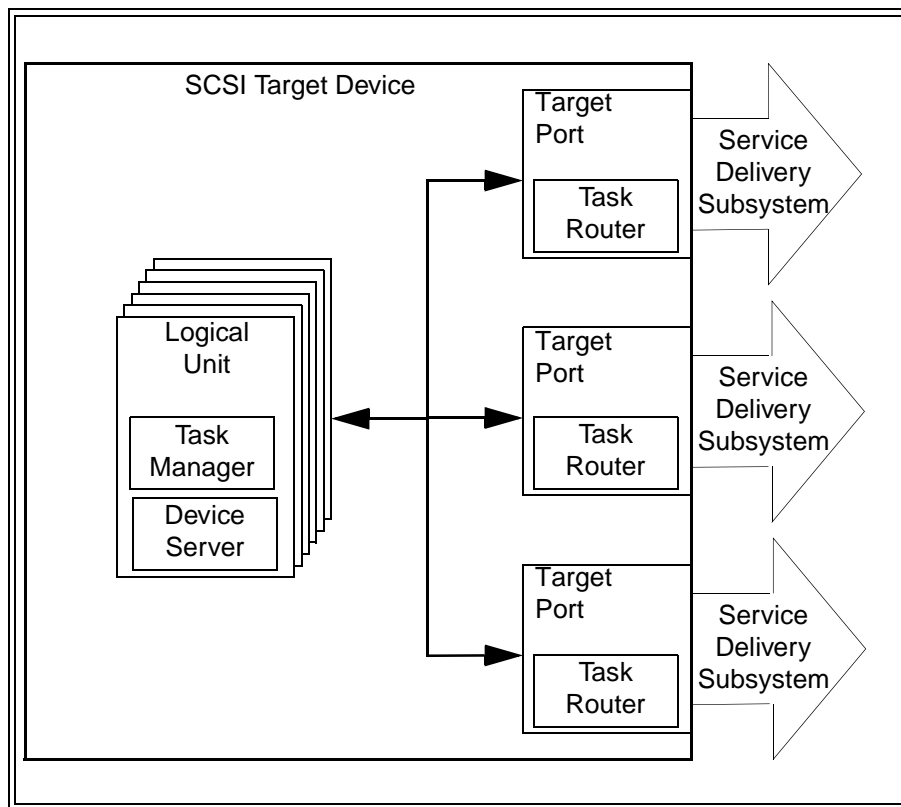


Figure 1 - Multiple port target SCSI device structure model

Multiple port initiator SCSI Device




Figure 2 - Multiple port SCSI initiator device structure model

Multiple port target/initiator SCSI Device

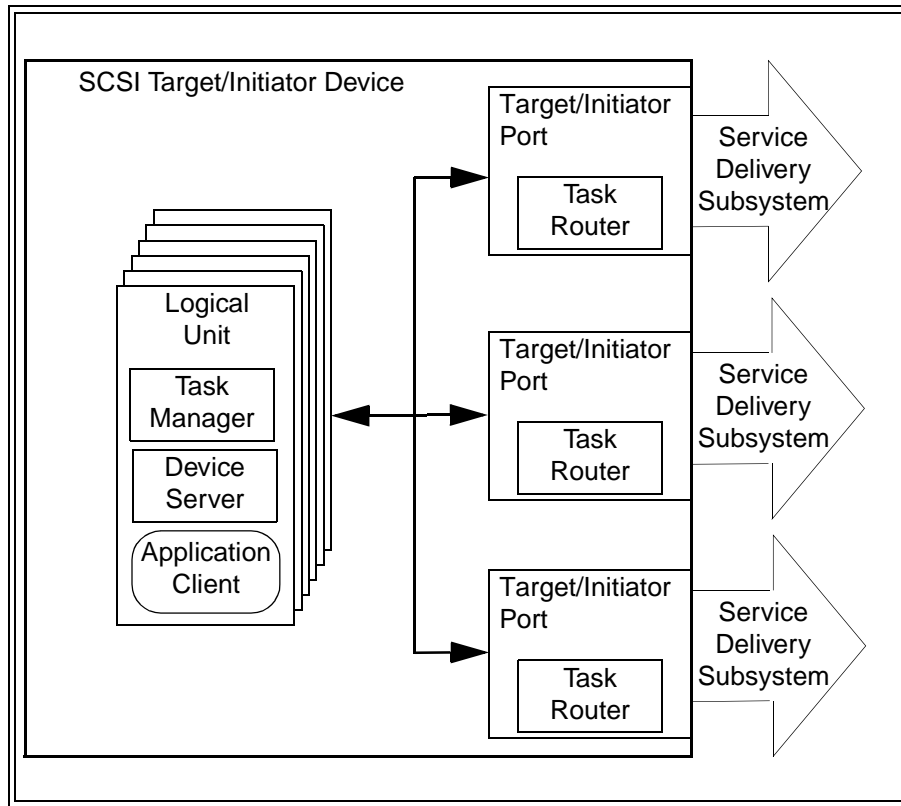


Figure 3 - Multiple port target/initiator SCSI device structure model

Multiple port SCSI Device model

The following show the changes to figures in section 4 of SAM-2 that would be required to change to support a port = target/initiator model.

SAM-2 Section 4.4 The SCSI structural model

Figure 4 shows the main functional components of the basic SCSI hierarchy. The following clauses define these components in greater detail.

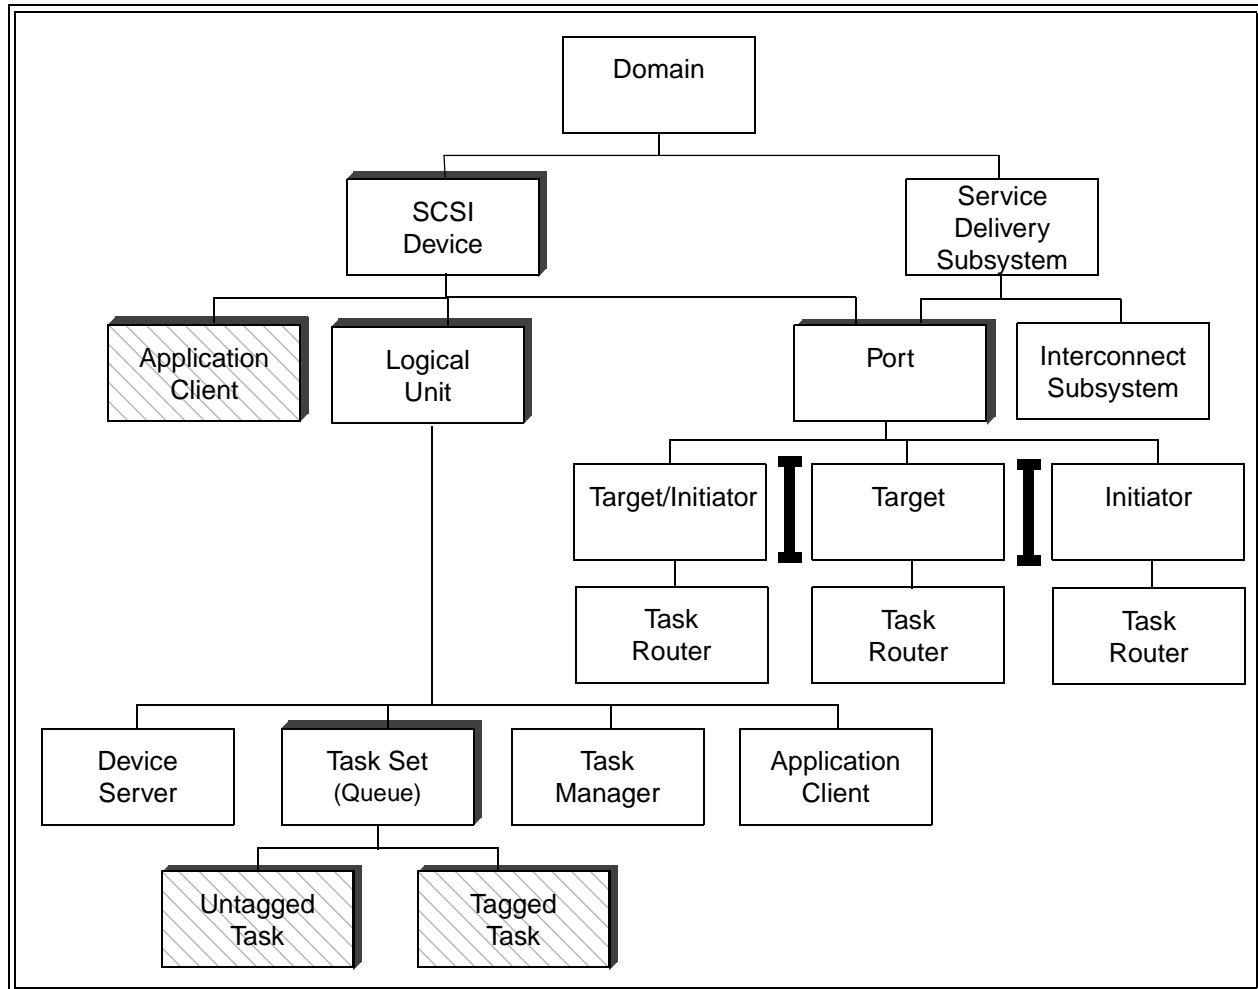


Figure 4 - Basic SCSI hierarchy

SAM-2 Section 4.5 SCSI domain

An example of some of the wording changes that would be required follow. In section 4.5 of SAM-2 the following paragraph would be modified as shown:

A SCSI device is an object that originates or services SCSI commands. As described in xxx, when a SCSI device originates a command that command goes through an initiator port or a target/initiator port; a SCSI device containing logical units that service commands receives commands through a target port or a target/initiator port. The service delivery subsystem connects all the ports in the SCSI domain, providing a subsystem through which application clients and device servers communicate (see 4.6). The boundaries of a SCSI domain are established by the system implementor, within the constraints of a specific SCSI protocol and interconnect standards.

SAM-2 Section 4.7 SCSI Device models

In the basic structural model, a SCSI device (see figure 5) is composed of one or more ports (see 4.6) each of which is a target, initiator, or target/initiator.

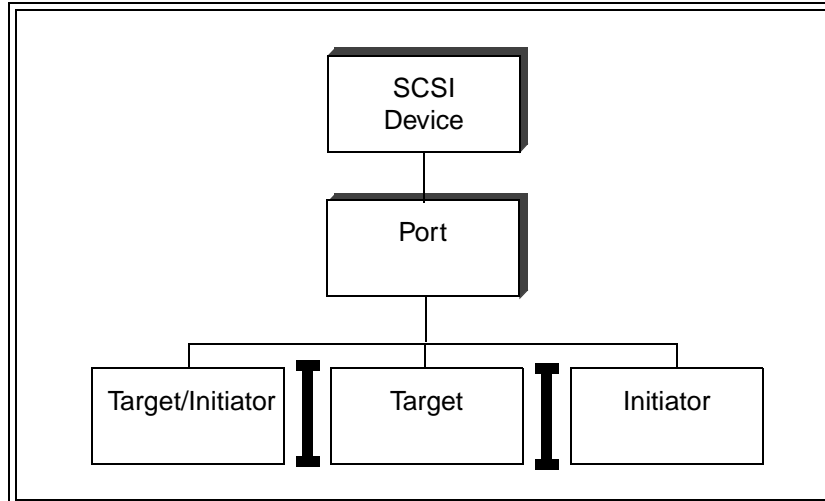


Figure 5 - SCSI Device hierarchy diagram

A SCSI initiator device is capable of originating SCSI commands and task management requests (see). A SCSI target device is capable of executing SCSI commands and task management requests (see). A SCSI target/initiator device is capable of originating and executing SCSI commands and task management requests. To be functional, an SCSI domain needs to contain a SCSI device that contains a target port or a target/initiator port and another SCSI device that contains an initiator port or target/initiator port.

There are several models for implementing ports in an SCSI device, some are shown in figure 6. An SCSI device may contain only target ports, only initiator ports, only target/initiator ports or any combination there of.

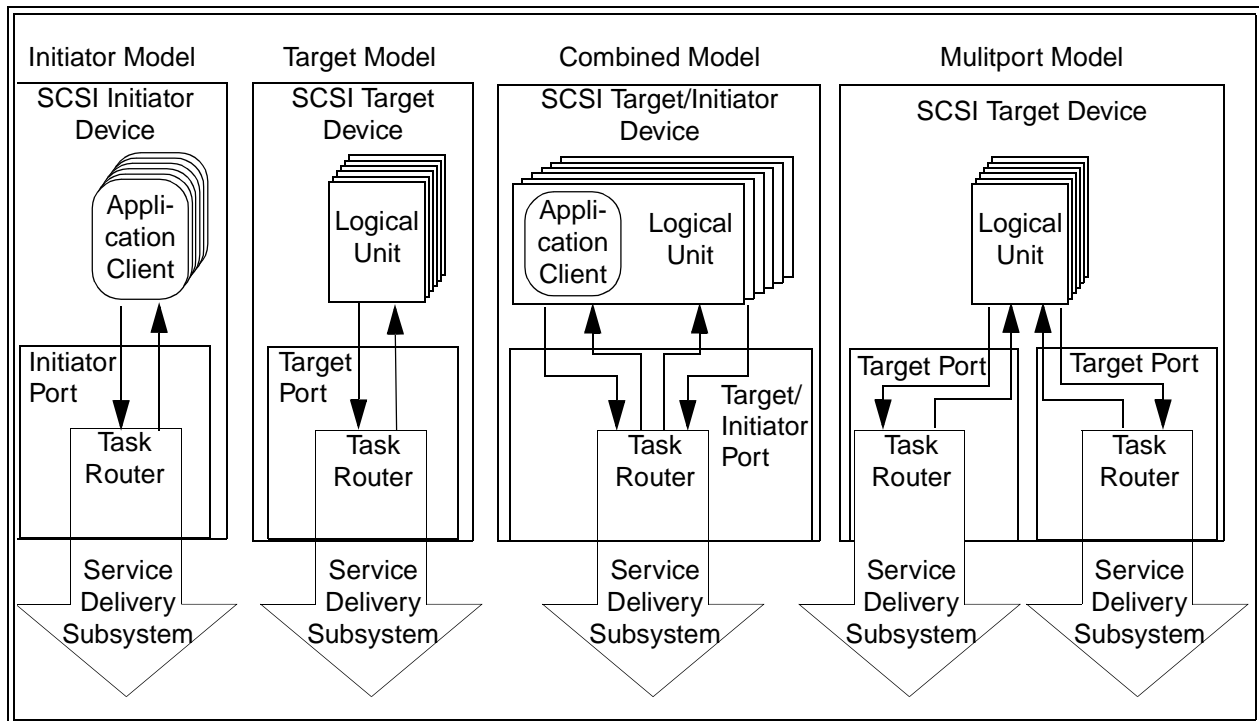


Figure 6 - SCSI device functional models

A target/initiator SCSI device is referred to by the role its port takes when it participates in an I/O operation. That is, when a SCSI target/initiator device receives SCSI commands or task management functions the SCSI target/initiator device takes on the characteristics of a SCSI target device, or when a SCSI target/initiator device issues SCSI commands or task management functions the SCSI target/initiator device takes on the characteristics of a SCSI initiator device.

SAM-2 section 4.7.1 SCSI initiator model

A SCSI initiator device (see figure 7) is composed of a task router and an Initiator Identifier for each initiator port and zero or more application clients.

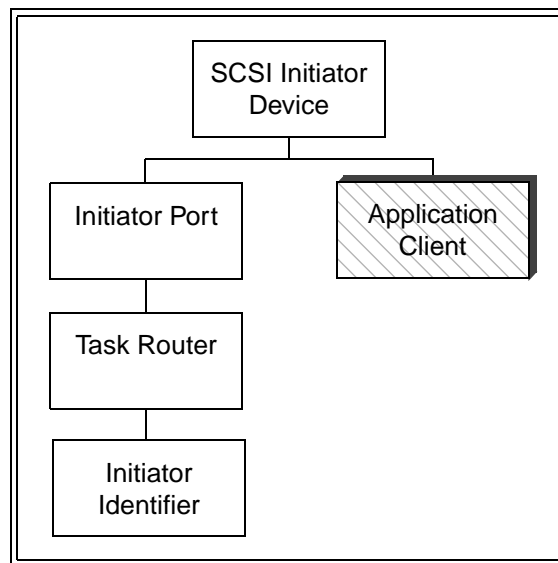


Figure 7 - Initiator hierarchy diagram

An Initiator Identifier is a field containing up to 64 bits that is a SCSI device identifier for the initiator port. An Application Client is the source of commands and task management functions. In this model, there is one application client for each pending command or task management function. A task router moves the commands and task management functions between the service delivery subsystem and the appropriate application client.

SAM-2 section 4.7.2 SCSI target model

A SCSI target device (see figure 8) is composed of a task router and target identifier for each target port and one or more logical units composed of a task manager and a device server.

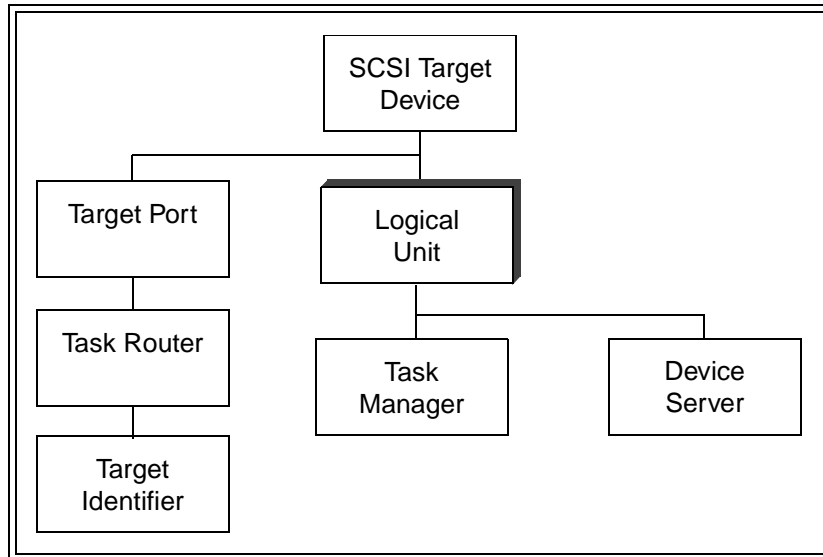


Figure 8 - Target hierarchy diagram

A Target Identifier is a field containing up to 64 bits that is a SCSI device identifier for the SCSI target device. A task router moves the commands and task management functions between the service delivery subsystem and the appropriate logical unit. A logical unit (see 4.8) is the object to which SCSI commands are addressed. One of the logical units composing within the SCSI target device shall be addressed using the logical unit number zero. A Task Manager is a server that controls one or more tasks in response to task management requests (see 4.7.4). A device server carries out the received commands.

SCSI target/initiator model

A SCSI target/initiator device (see figure 8) is composed of a task router and target identifier or initiator identifier for each target/initiator port and one or more logical units composed of a task manager, application client, and a device server.

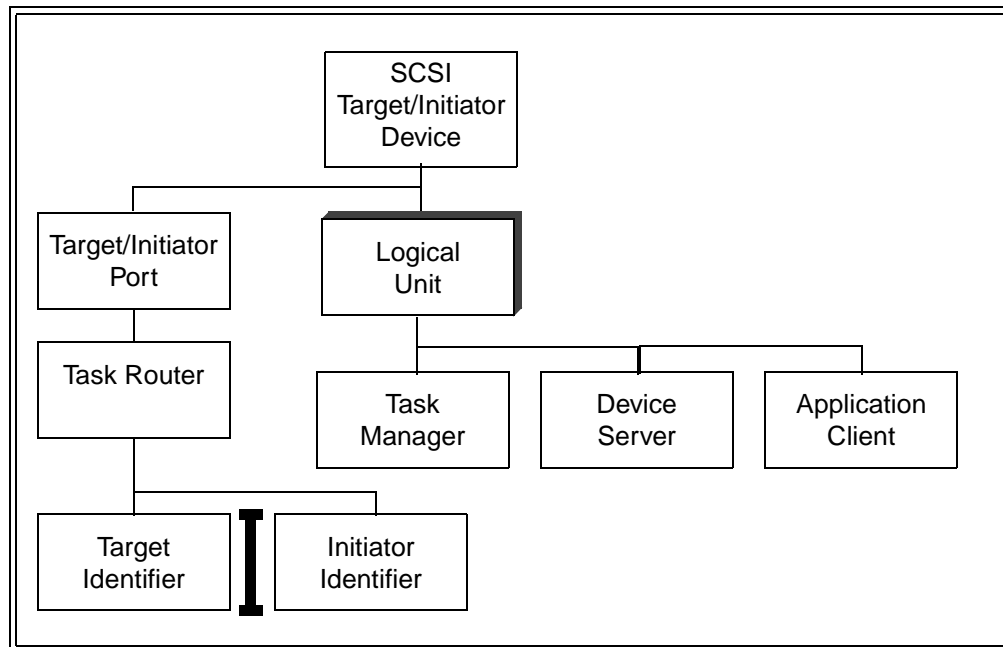


Figure 9 - Target hierarchy diagram

When the SCSI target/initiator device is operating as a SCSI target device the Target Identifier is a field containing up to 64 bits that is a SCSI device identifier for the SCSI target/initiator device. A task router moves the commands and task management functions between the service delivery subsystem and the appropriate logical unit. A logical unit (see 4.8) is the object to which SCSI commands are addressed. One of the logical units composing within the SCSI target/initiator device shall be addressed using the logical unit number zero. A Task Manager is a server that controls one or more tasks in response to task management requests (see 4.7.4). A device server carries out the received commands.

When the SCSI target/initiator device is operating as a SCSI initiator device the Initiator Identifier is a field containing up to 64 bits that is a SCSI device identifier for the initiator port. An Application Client is the source of commands and task management functions. In this model, there is one application client for each pending command or task management function. A task router moves the commands and task management functions between the service delivery subsystem and the appropriate application client.

SAM-2 section 4.9.3 Identification of target tasks

A SCSI target device and a SCSI target/initiator device identifies a target task with a Task Identifier. The Task Identifier object represents either a Tagged Task Identifier or an Untagged Task Identifier. A Tagged Task Identifier is composed of an Initiator Identifier (see), a Logical Unit Identifier (see 4.8) and a Tag (see 4.9.1). An Untagged Task Identifier is composed of an Initiator Identifier and a Logical Unit Identifier.

If a SCSI target device and a SCSI target/initiator device contains multiple target ports (see), then the Task Identifier objects contain one additional object (beyond those mentioned above). For a multiple port SCSI target device and a SCSI target/initiator device, a Tagged Task Identifier is composed of a target Port Identifier (see), an Initiator Identifier (see), a Logical Unit Identifier (see 4.8) and a Tag (see 4.9.1). An Untagged Task Identifier is composed of a target Port Identifier, an Initiator Identifier, and a Logical Unit Identifier.

SAM-2 section 4.9.3 Identification of initiator tasks

A SCSI initiator device identifies an initiator task to a target task using a Task Address. The Task Address object represents either a Tagged Task Address or an Untagged Task Address. A Tagged Task Address is composed of a Logical Unit Identifier (see 4.8) and a Tag (see 4.9.1). An Untagged Task Address is composed of a Logical Unit Identifier.

If a SCSI initiator device contains multiple initiator ports (see), then it may enhance the definition of a Task Address with initiator port identification information.

SAM-2 section 4.10 SCSI devices with multiple ports

The model for a SCSI device with multiple ports is a single SCSI target device, SCSI initiator device, or SCSI target/initiator device (see) with multiple ports. Similarly, a single SCSI target device, SCSI initiator device, or SCSI target/initiator device may respond to multiple SCSI device identifiers. The model for such a SCSI device also is one of multiple SCSI target devices, SCSI initiator devices, or SCSI target/initiator devices, one for each SCSI device identifier.

SCSI devices with multiple ports do not fit in the two dimensional hierarchy shown in 4.4. Instead, they combine basic objects from the SCSI device model in a way that produces a different structural plane. Figure 10 also shows that SCSI devices from different domains may be connected to a single SCSI device with multiple ports.

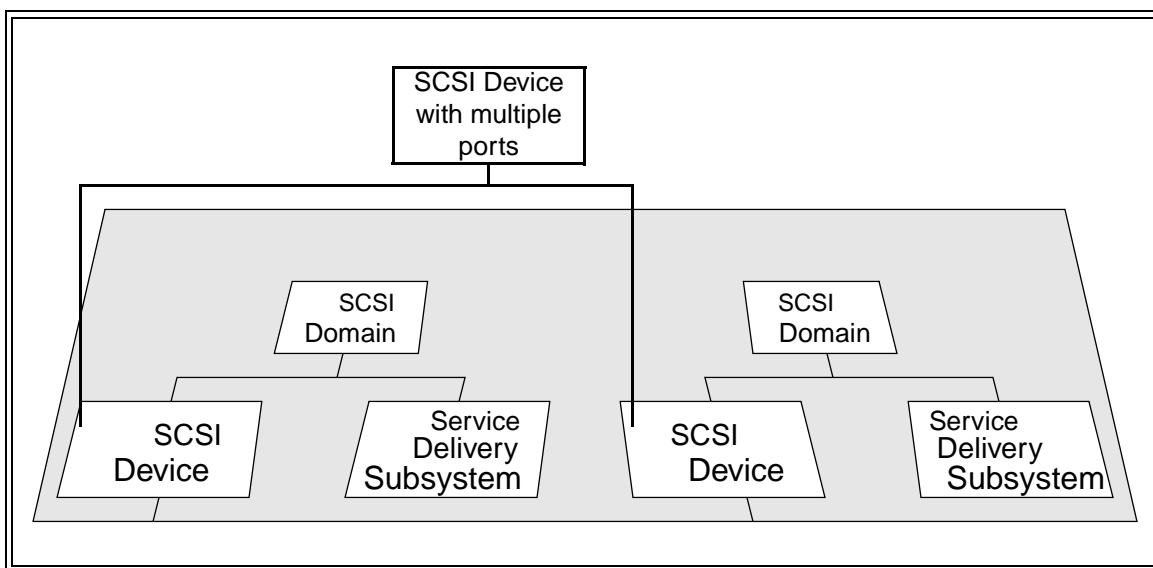


Figure 10 - A SCSI multi-port unit and multiple SCSI domains

The multiple SCSI device identifiers representing the ports shall meet the requirements for Initiator Identifiers (see) or Target Identifiers (see) or both. SCSI target/initiator devices with multiple ports implement both target and initiator models and combine the target/initiator port structures in vendor-specific ways that meet product requirements while maintaining the multi-port model for the target and initiator functions performed by the product. How a multiple port SCSI device is viewed by counterpart SCSI devices in the SCSI domain also depends on whether an initiator port is examining a target port or target/initiator port, or a target port is servicing an initiator port or target/initiator port. The structures and views of SCSI devices are asymmetric for target ports, initiator ports, and target/initiator ports. The subclauses that follow discuss the principle cases described above.

SAM-2 section 4.10.1 Multiple port target SCSI device structure

Figure 10 shows the structure of a SCSI target device with multiple target ports. Each target port consists of a task router and is shared by a collection of logical units. Each logical unit contains a single task manager and a device server.

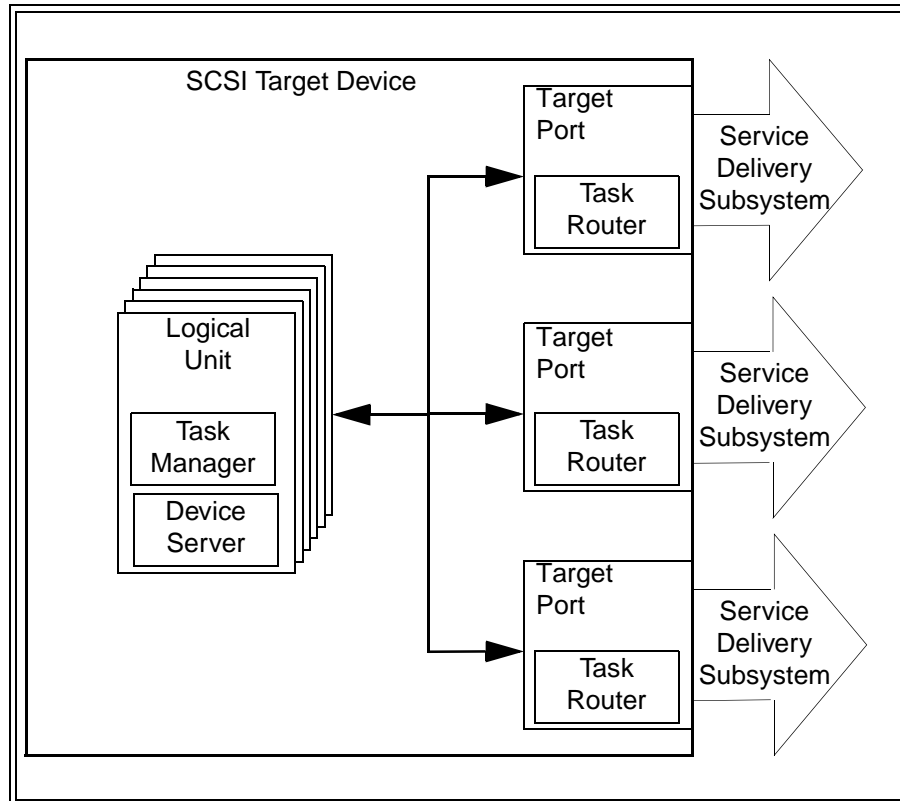


Figure 11 - Multiple port target SCSI device structure model

Two-way communications shall be possible between all logical units and all target ports, however, communications between any logical unit and any target port may occasionally be inactive. Two-way communications shall be available between each task manager and all task routers. Each target port shall accept commands addressed to LUN 0 and the task router shall forward them to a device server for processing. The REPORT LUNS commands (see SPC-2) shall be accepted by logical unit 0 from any target port and shall return the logical unit inventory available via that target port. The availability of a the same logical unit through multiple target ports is discovered by matching Device Identifier values in the INQUIRY command Vital Product Data page (see SPC-2).

SAM-2 section 4.10.2 Target port Task Identifiers

In addition to the Task Identifier constituent objects described in xxx, Target ports shall provide a Port Identifier object to contribute to the construction of Task Identifiers. The exact nature of the Port Identifier is vendor specific. The complete contents of the Task Identifier object, including the Port Identifier, is defined in xxx.

SAM-2 section 4.10.3 Multiple port initiator SCSI device structure

Figure 11 shows the structure of a SCSI initiator device with multiple initiator ports. Each initiator port consists of a task router and is shared by a collection of application clients.

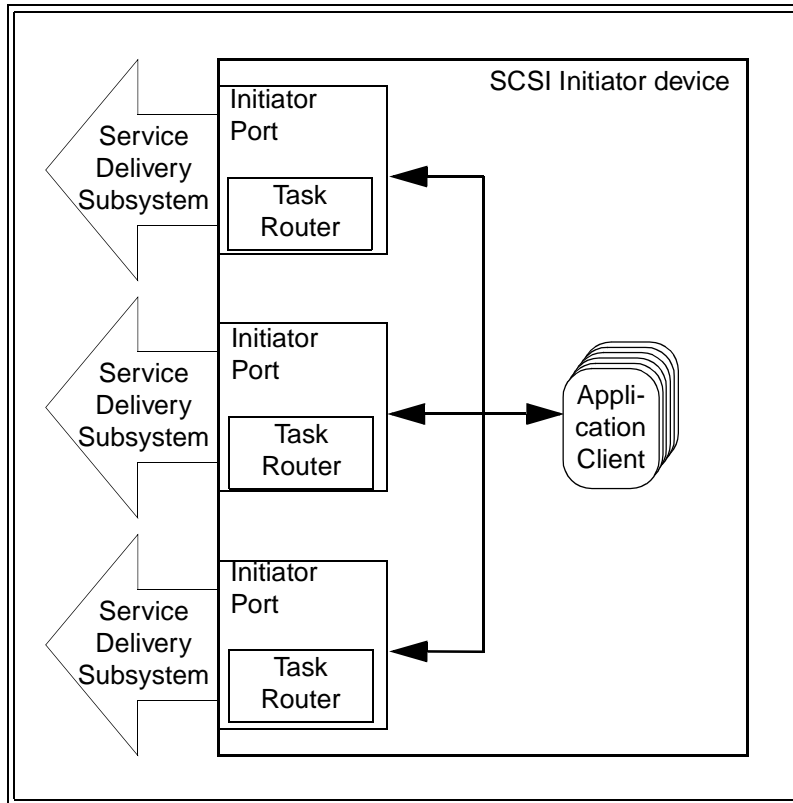


Figure 12 - Multiple port SCSI initiator device structure model

Two-way communications shall be possible between all application clients and all task routers, however, communications between any application client and any task router may occasionally be inactive. Mechanisms by which a SCSI target device would have the ability to discover that it is communicating with an initiator port are beyond the scope of any standards in the SCSI family of standards.

Multiple port target/initiator SCSI device structure

Figure 13 shows the structure of a SCSI target/initiator device with multiple target/initiator ports. Each target/initiator port consists of a task router and is shared by a collection of logical units. Each logical unit contains a single task manager, a device server, and an application client.

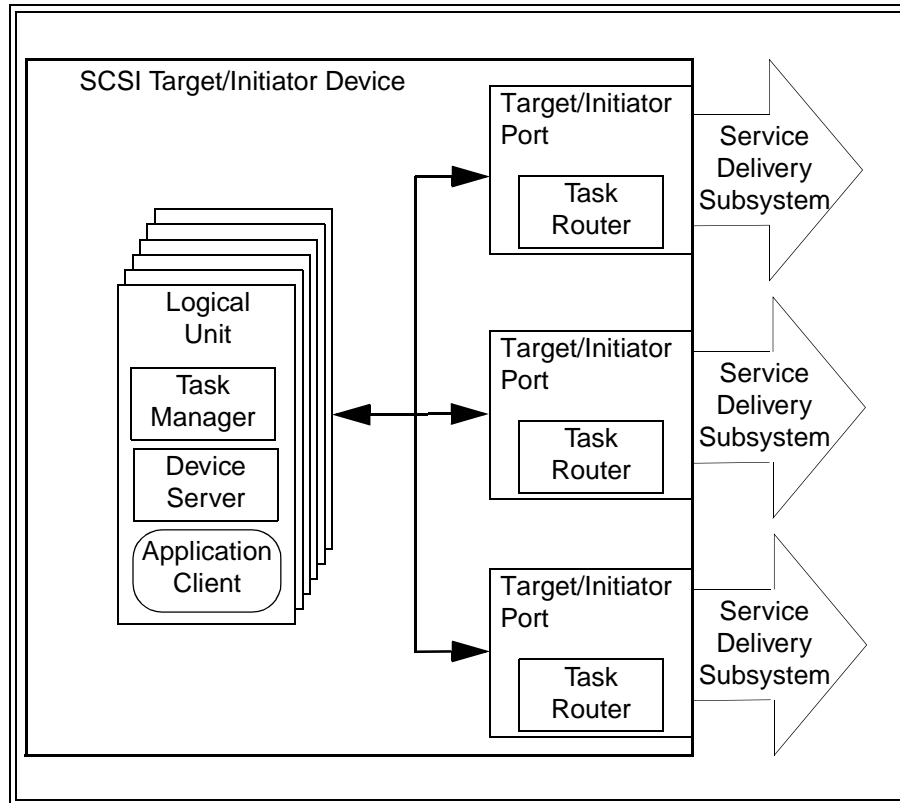


Figure 13 - Multiple port target/initiator SCSI device structure model

Two-way communications shall be possible between all logical units and all target/initiator ports, however, communications between any logical unit and any target/initiator port may occasionally be inactive. Two-way communications shall be available between each task manager or application client and all task routers. Each target/initiator port shall accept commands addressed to LUN 0 and the task router shall forward them to a device server for processing. The REPORT LUNS commands (see SPC-2) shall be accepted by logical unit 0 from any target port and shall return the logical unit inventory available via that target port. The availability of a the same logical unit through multiple target/initiator ports is discovered by matching Device Identifier values in the INQUIRY command Vital Product Data page (see SPC-2).

Mechanisms by which a SCSI target device would have the ability to discover that it is communicating with a target/initiator port are beyond the scope of any standards in the SCSI family of standards.

Target/initiator port Task Identifiers

In addition to the Task Identifier constituent objects described in xxx, Target/initiator ports, when functioning as a target port, shall provide a Port Identifier object to contribute to the construction of Task Identifiers. The exact nature of the Port Identifier is vendor specific. The complete contents of the Task Identifier object, including the Port Identifier, is defined in xxx.

SAM-2 section 4.10.4 SCSI initiator device view of a multiple port SCSI target device

In all cases when the term SCSI target device is used it refers to SCSI target devices and SCSI target/initiator devices that are using the target/initiator port as a target port. In all cases when the term target port is used it refers to a target port and a target/initiator port operation as a target.

In all cases when the term SCSI initiator device is used it refers to SCSI initiator devices and SCSI target/initiator devices that are using the target/initiator port as an initiator port. In all cases when the term initiator port is used it refers to an initiator port and a target/initiator port operation as an initiator.

A SCSI target device may be connected to multiple independent service delivery subsystems through target ports in such a way that no single initiator port can communicate with a logical unit using two or more of the target ports in the SCSI target device. In this case, the application clients in the SCSI initiator device cannot determine if a SCSI target device has multiple ports.

However, SCSI target devices with multiple ports may be configured in many different ways where application clients have the ability to discover that one or more logical units are accessible via multiple target ports. Figure 14 and figure 15 show two examples of such configurations.

Figure 14 shows a SCSI target device with multiple target ports participating in a single SCSI Domain with two SCSI initiator devices. There are three SCSI devices, one of which has two target ports, one with one initiator port, and one with one initiator port. There are two Target Identifiers and two Initiator Identifiers in this SCSI Domain. Using the INQUIRY command Vital Product Data page as described in xxx, the application clients in each of the SCSI initiator devices that have the ability to discover the logical units in the SCSI target devices are accessible via multiple Target Identifiers (target ports) and map the configuration of the SCSI target devices.

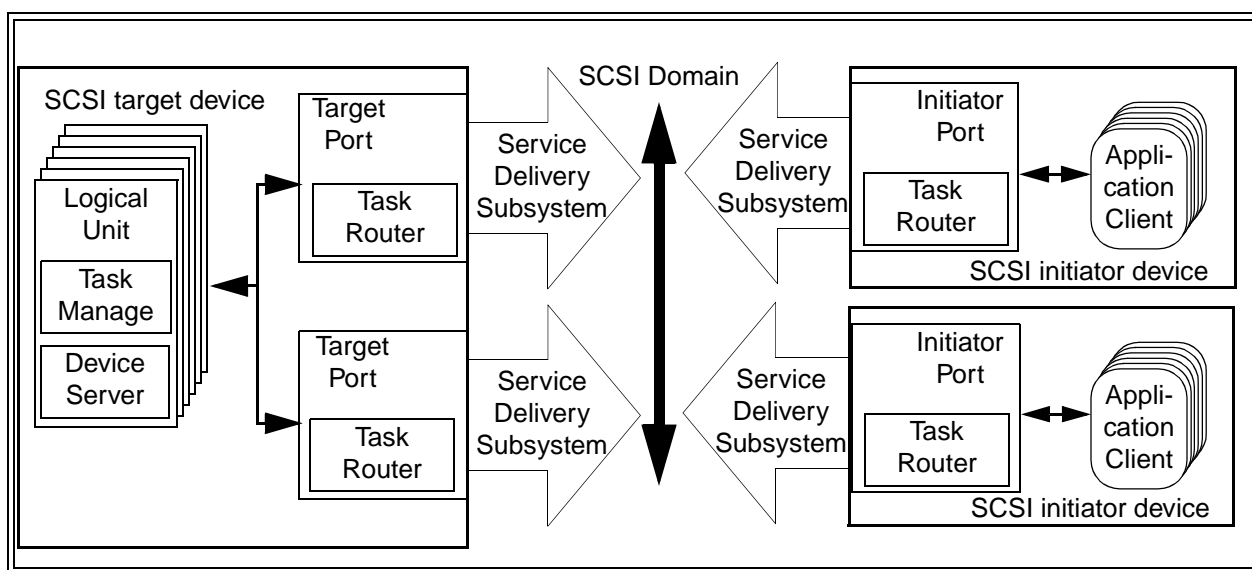


Figure 14 - SCSI target device configured in a single SCSI Domain

Figure 15 shows a SCSI target device with multiple target ports participating in two SCSI Domains and a SCSI initiator device with multiple initiator ports participating in the same two SCSI Domains. There is one SCSI target device two target ports and the a SCSI initiator device with two initiator ports. There is one Target Identifier and one Initiator Identifier in each of the two SCSI Domains. Using the INQUIRY command Vital Product Data page as described in xxx, the application clients in the SCSI initiator device have the ability to discover that logical units in the SCSI target device are accessible via multiple ports and map the configuration. However, the methods available to application clients to distinguish between the configuration shown in figure 15 and the configuration shown in figure 22 may be beyond the scope of the SCSI family of standards.

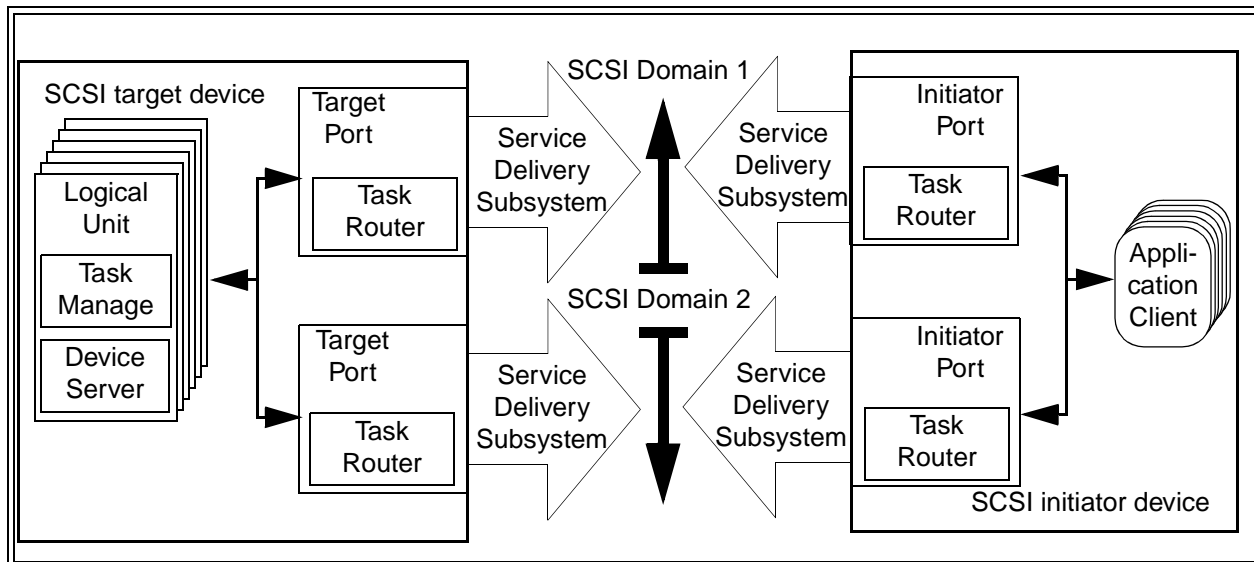


Figure 15 - SCSI target device configured in multiple SCSI Domains

Figure 16 shows the same configuration as figure 15 except that the two SCSI Domains have been replaced by a single SCSI Domain.

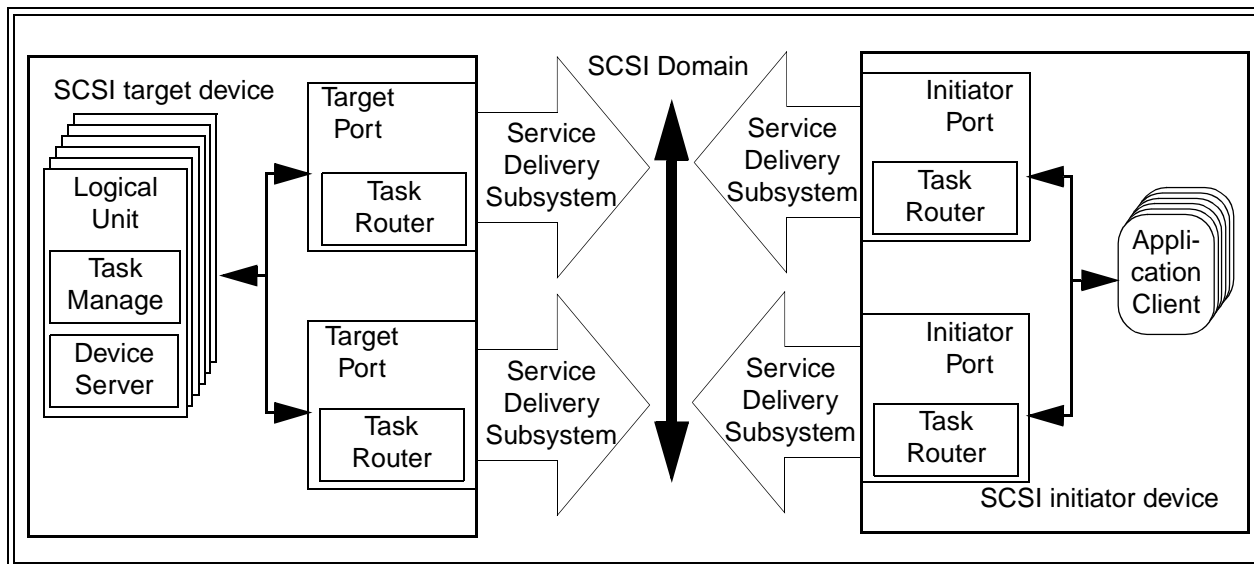


Figure 16 - SCSI target device and SCSI initiator device configured in a single SCSI Domain

This model for application client determination of multiple target port configurations relies on information that is available only to the application clients via SCSI commands. The initiator ports in the SCSI initiator devices (figure 14) or SCSI initiator device (figure 15 and figure 16) are unable to distinguish multiple target ports from individual target ports in two separate SCSI target devices.

SAM-2 section 4.10.5 SCSI target device view of a multiple port SCSI initiator device

A SCSI target device does not have the ability to detect the presence of a SCSI initiator device with multiple initiator ports. Therefore, a SCSI target device handles a SCSI initiator device with multiple initiator ports exactly as it would handle multiple separate SCSI initiator devices each with a single initiator port.

For example, a SCSI target device handles the configurations shown in figure 15 and figure 16 in exactly the same way it handles the configuration show in figure 14.

NOTE 1 The implications of this view of a SCSI initiator device are more far reaching than are immediately apparent. For example, if a SCSI initiator device with multiple initiator ports makes an exclusive access reservation via one initiator port, then access is denied to the other initiator port(s) on that same SCSI initiator device.

SAM-2 section 4.10.6 SMU considerations for task management functions

Although, the task manager for each logical unit handles all the target ports and target/initiator ports in a SCSI target device, the task manager in each logical unit shall observe the requirements described in the following clauses in addition to the requirements placed on a task set manager by the SCSI architecture basic structural model.