

To: X3T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: Addressing Model for SAM-2

1 Overview

1.1 Example of hierarchical system

A system that is composed of many layers appears as a tree. For example, a driver may connect to multiple HBAs, which in turn may connect to multiple SCSI devices, etc. See figure 1 for an example of a system that consists of:

- a) One initiator that has three SCSI devices attached on a single SCSI bus that is not expandable. One of the SCSI devices is a dual ported SCSI bridge controller.
- b) One initiator has two SCSI devices attached on a single SCSI bus that is expandable. One of the SCSI devices contains a dual ported SCSI bridge controller.
- c) The SCSI bridge controller has three SCSI buses with SCSI devices attached and is capable of driving more SCSI buses.
 - a) Two of the SCSI buses contain two SCSI devices each and these SCSI buses are not expandable. One of the SCSI devices contains a SCSI bridge controller.
 - b) One of the SCSI buses contains two SCSI devices and is expandable.
 - c) The SCSI bridge controller has three SCSI buses with SCSI devices attached and is capable of driving more SCSI buses.
 - a) Two of the SCSI buses contain two SCSI devices each and these SCSI buses are not expandable.
 - b) One of the SCSI buses contains two SCSI devices and is expandable.

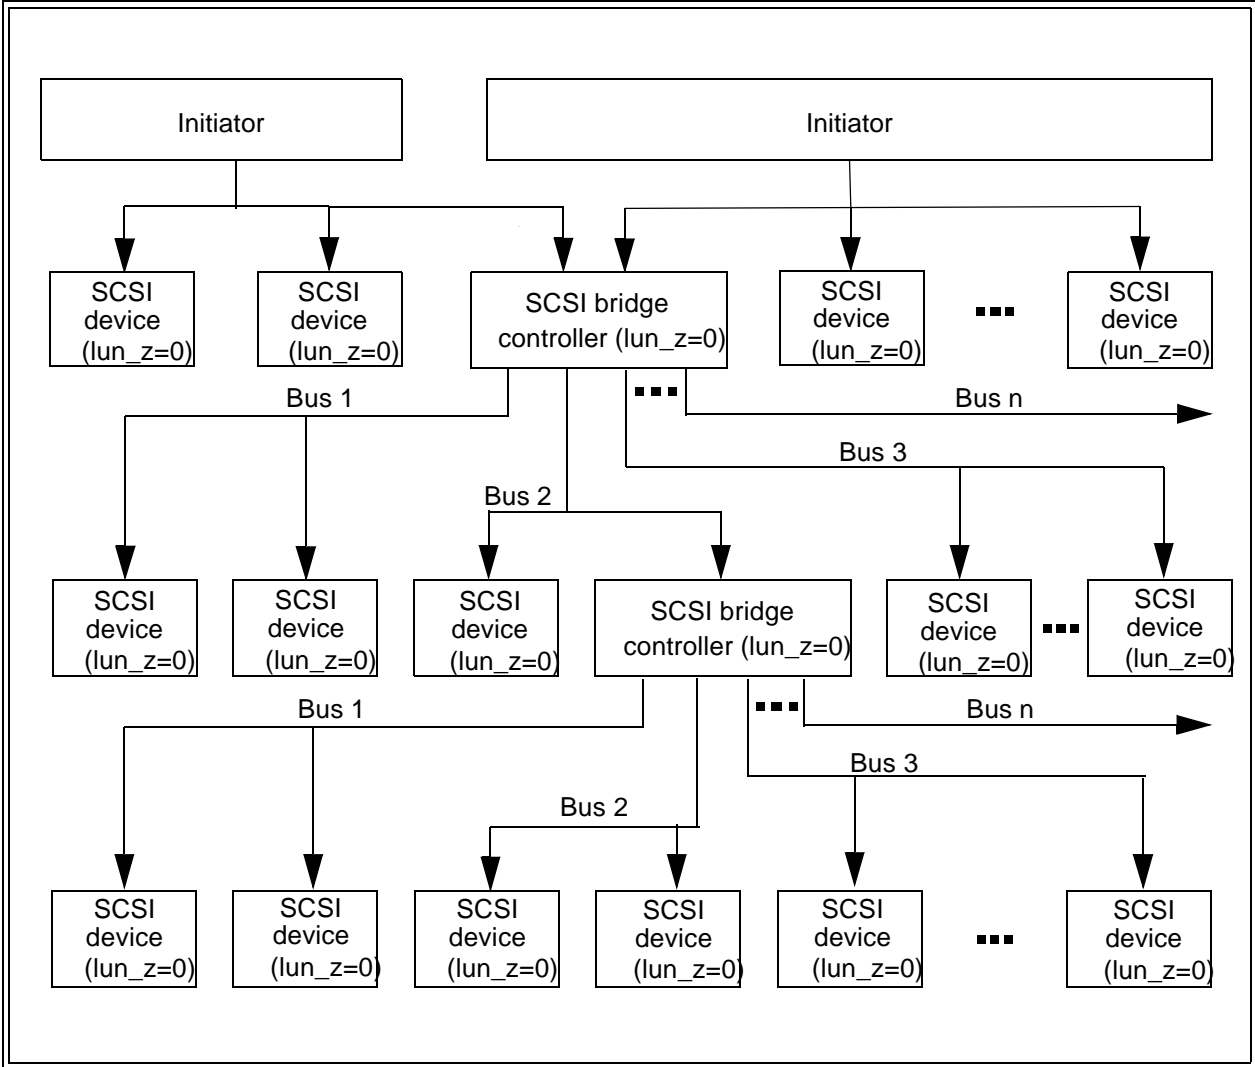


Figure 1 - Example of hierarchical system diagram

2 Addressing Model

All peripheral device addresses, except LUN_Z, default to vendor specific values. All addressable objects may default to vendor specific values or may be defined by an application client during configuration.

2.1 LUN_Z address

A LUN_Z address shall be equivalent to addressing an SCSI device's LUN of zero.

All SCSI devices shall accept LUN_Z as a valid address. For SCSI devices that support the hierarchical addressing model the LUN_Z shall be the logical unit that an application client addresses to determine information about the target and the logical units contained within the target.

To address the LUN_Z of an SCSI device the peripheral device address method shall be used.

2.2 Buses

Within the hierarchical system there may be target devices that have multiple logical units that are

connected through separate physical interconnects. Within the addressing model these physical interconnects are referred to as buses. A target device that has logical units attached to buses shall assign numbers, other than zero, to those buses. The bus numbers shall be used when assigning LUN's to the logical units attached to those buses.

Target devices shall assign a bus number of zero to all the target devices logical units not connected through a separate physical interconnect.

2.3 Eight byte LUN structure

The eight byte LUN structure (see table 2) allows up to four levels of devices to be addressed under a single target. Each level shall use bytes 0-1 to define the address and/or location of the SCSI device to be addressed on that level.

If the LUN indicates that the command is to be passed to the next layer then the current layer shall use bytes 0-1 of the eight byte LUN structure to determine the address of the device to which the command is to be sent. When the command is sent to the target the eight byte LUN structure that was received shall be adjusted to create a new eight byte LUN structure (see table 1).

Devices shall keep track of the necessary addressing information to allow reconnection to the correct task during reselection.

Table 1 - Eight byte LUN structure adjustments

Byte position		
Old		New
0 - 1	Moves to	Not used
2 - 3	Moves to	0 - 1
4 - 5	Moves to	2 - 3
6 - 7	Moves to	4 - 5

Bytes six and seven of each new eight byte LUN structure shall be set to zero.

Table 2 - Eight byte LUN structure

Bit Byte	7	6	5	4	3	2	1	0
0	FIRST LEVEL ADDRESSING							
1								
2	SECOND LEVEL ADDRESSING							
3								
4	THIRD LEVEL ADDRESSING							
5								
6	FOURTH LEVEL ADDRESSING							
7								

The FIRST LEVEL ADDRESSING field indicates the first level address of a device. See table 3 for a definition of the FIRST LEVEL ADDRESSING field.

The SECOND LEVEL ADDRESSING field indicates the second level address of a device. See table 3 for a definition of the SECOND LEVEL ADDRESSING field.

The THIRD LEVEL ADDRESSING field indicates the third level address of a device. See table 3 for a definition of the THIRD LEVEL ADDRESSING field.

The FOURTH LEVEL ADDRESSING field indicates the fourth level address of a device. See table 3 for a definition of the FOURTH LEVEL ADDRESSING field.

The device pointed to in the FIRST LEVEL ADDRESSING field, SECOND LEVEL ADDRESSING field, THIRD LEVEL ADDRESSING field, and FOURTH LEVEL ADDRESSING fields may be any physical or logical device addressable by an application client.

Table 3 - FIRST LEVEL ADDRESSING field, SECOND LEVEL ADDRESSING field, THIRD LEVEL ADDRESSING field, and FOURTH LEVEL ADDRESSING field

Bit Byte	7	6	5	4	3	2	1	0
n-1	address method		ADDRESS METHOD SPECIFIC					
n								

The ADDRESS METHOD field defines the contents of the ADDRESS METHOD SPECIFIC field. See table 4 for the defined address methods.

Table 4 - ADDRESS METHOD

Codes	Description
10b	Logical unit addressing method
00b	Peripheral device addressing method
01b	Virtual device addressing method
11b	Reserved

See 2.3.1, 2.3.2, and 2.3.3 for the definitions of the ADDRESS METHOD SPECIFIC field.

2.3.1 Logical unit address method

All SCSI commands are allowed when the logical unit address method is selected, however logical units are only required to support mandatory SCSI commands. Devices are not required to honor pass-through requests from the application client. Any command that is not supported or passed-through shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID COMMAND OPERATION CODE.

If the logical unit addressing method is selected the device shall relay the received command, if supported, to the addressed logical unit. See table 5 for the definition of the ADDRESS METHOD SPECIFIC field used when the logical unit addressing method is selected.

Table 5 - Logical unit addressing

Bit Byte	7	6	5	4	3	2	1	0
n-1	1	0	target					
n	BUS			LUN				

The TARGET field indicates the target address of the device to which the received command shall be relayed. The TARGET field indicates the address of the target on the bus indicated by the BUS NUMBER field that the received command shall be relayed.

NOTE 1 - The value of targets within the TARGET field are defined by individual standards. (e.g., SCSI-3 Parallel Interface Standard defines targets to be in the range 0-7, 0-15, and 0-31).

The BUS NUMBER field indicates the location of the bus that shall be used to relay the received command.

The LUN field indicates the address of the logical unit within the target indicated by the TARGET field that the received command shall be relayed to.

2.3.2 Peripheral device address method

All SCSI commands are allowed when the peripheral device address method is selected, however peripheral devices are only required to support mandatory SCSI commands. Devices are not required to honor pass-through requests from the application client. Any command that is not supported or passed-through shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID COMMAND OPERATION CODE.

If the peripheral device addressing method is selected the device shall relay the received command, if supported, to the addressed peripheral device. See table 6 for the definition of the ADDRESS METHOD SPECIFIC field used when the peripheral device addressing method is selected.

Table 6 - Peripheral device addressing

Bit Byte	7	6	5	4	3	2	1	0
n-1	0	0	bus number					
n	TARGET/LUN							

The BUS NUMBER field indicates the location of the bus that shall be used to relay the received command.

The TARGET field indicates the target address of the device to which the received command shall be relayed. The TARGET field indicates the address of the target on the bus indicated by the BUS NUMBER field that the received command shall be relayed.

The TARGET/LUN field indicates the address of the target or LUN to which the received command shall be relayed. If the BUS NUMBER field is set to zero the TARGET/LUN field indicates the address of the logical unit to relay the received command to, is located within the current level. If the bus number field is not zero the TARGET/LUN field indicates the address of the target on the bus indicated by the bus number field that the received command shall be relayed to. The logical unit to relay the received command to shall be LUN_Z.

The LUN_Z located within the current level shall be addressed by a BUS NUMBER field and a TARGET/LUN field of all zeros.

NOTE 2 - The value of targets within the TARGET/LUN field are defined by individual standards. (e.g., SCSI-3 Parallel Interface Standard defines targets to be in the range 0-7, 0-15, and 0-31).

2.3.3 Virtual device address method

All SCSI commands are allowed when the virtual device address method is used, however virtual devices are not required to support all SCSI commands. Any command that is not supported shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID COMMAND OPERATION CODE.

In the response to an INQUIRY command the addressed virtual device shall return a valid SCSI peripheral device type.(e.g., direct access device, streaming device, etc.)

When the virtual device addressing method is selected the device at the current level addresses peripheral devices as required to execute the received command. See table 7 for the definition of the ADDRESS METHOD SPECIFIC field used when the volume set addressing method is selected.

Table 7 - Virtual device addressing

Bit Byte	7	6	5	4	3	2	1	0
n-1	0	1	(MSB)					
n	LUN							(LSB)

The LUN field indicates the address of the virtual device the current level shall direct the received command to.

3 SPC additions

3.1 Inquiry command additions

A bit needs to be added to the inquiry data to indicate when a device uses the addressing model to assigned LUNs to logical units. The bit would be as follows:

A hierarchical support (HISUPPORT) bit of zero indicates the target does not use the hierarchical addressing model to assign LUNs to logical units. A HISUPPORT bit of one indicates the target uses the hierarchical addressing model to assign LUNs to logical units.

3.2 Command requirements

For SCSI devices that support hierarchical addressing the REPORT LUNS command shall be mandatory.