

Date: Feb 05,1994

X3T10/94-031 rev 1

To: X3T9.2 Committee (SCSI)

From: George Penokie (IBM)

Subject: SCSI-3 Logical Unit Addressing

1.0 Overview

SCSI-2 only allows eight Logical Units to be addressed under a single Target. SCSI-3 parallel has expanded this to 64 Logical Units by using three additional bits in the Identify message. However, 64 Logical Units is not enough to handle SCSI Disk Arrays.

SAM has defined a 64 bit address space for Logical Units and this address space is being used by the SCSI-3 serial interfaces. SCSI-3 parallel needs a way use the 64 bit address space. This proposal defines a method for SCSI-3 parallel to take advantage of the expanded address space. The method uses a five bit address space as a pointer into a translation table that contains 64 bit addresses.

This proposal also covers a proposed method of addressing the multiple layers which are created with SDA devices. This method uses the 64 bit address space to, not only address the device, but to also indicate the location of the device within the tree structure of a SDA device.

2.0 Base Device Address

All targets shall accept LUN 0 as an address. For targets that have more than one logical unit, logical unit zero shall be the logical unit that an application client addresses to determine information about the target and the logical units attached to it.

3.0 Translation Table Addressing

This method uses a translation table to address devices. The translation table is defined using Mode Sense/Select commands. When translation addressing is active logical unit addresses 1-31 of the Identify message shall each point to a specific eight byte location within a translation table. The eight byte location shall then be used by the target to determine to which bus/target/LUN to send the command.

This mode page shall only be used for SDA type devices.

The definition of this mode page follows:

.pa

Table x: LUN Mapping Page

Bit	7	6	5	4	3	2	1	0	
0	PS	Reserved	Page Code (xxh)						
1	Page Length (FAh)								
2	Reserved								
3	Reserved						Active		
4	(MSB)								
11	LUN 1 Mapping						(LSB)		
12	(MSB)								
19	LUN 2 Mapping						(LSB)		
244	(MSB)								
251	LUN 31 Mapping						(LSB)		

When translation addressing is active, the logical unit addresses 1-31 of an Identify message shall reference a specific eight byte location within a translation table. The eight byte location shall then be used by the target to determine to which bus/target/LUN to send the command.

The change of an address that is part of a current or pending task shall be delayed until the task is complete.

An Active bit of zero indicates the translation table shall not be used. An Active bit one indicate that the translation table shall be used to determine to which bus/target/LUN to send the command.

The LUN xx Address Translation fields contain the eight byte bus/target/LUN of the device to be addressed. The layout of this field is defined in section 5.0. A value of all zeros in this field shall indicate an undefined LUN.


Any attempt to address an undefined LUN xx Mapping field shall be terminated with a CHECK CONDITION status. In response to a REQUEST SENSE command the target shall return sense data. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED.

.pa

4.0 SDA Identify Message

For SCSI devices that are identified as SDAs bit 5 of the Identify message indicates the addressed SDA shall use the LUN as a V-LUI, P-LUI, or as a pointer into the LUN mapping mode page.

IDENTIFY message format



The IDENTIFY message is sent by either the initiator or the target to establish an I_T_L nexus. For dual port implementations: if the target disconnects from the bus during an I/O process, it shall reconnect through the same port when the I/O process is continued.

The identify bit shall be set to one to specify that this is an IDENTIFY message.

A disconnect privilege (DiscPriv) bit of one specifies that the initiator has granted the target the privilege of disconnecting. A DiscPriv bit of zero specifies that the target shall not disconnect. This bit is not defined and shall be set to zero when an IDENTIFY message is sent by a target.

The VolSel bit shall only be used for SDA peripheral device types. Other device types shall use this bit as an extension of the LUN field.

A VolSel bit of zero specifies that a SDA peripheral device type shall use the LUN as either a reference into the LUN mapping mode page (if the active bit within the LUN mapping mode page is set to one) or as the address of a P-LUI within a SDA peripheral device type (if the active bit within the LUN mapping mode page is set to zero or there is no LUN mapping mode page defined).

Note: All P-LUI addresses default to vendor specific values. V-LUI addresses may default to vendor specific values or may be defined by an Application Client during configuration.

A VolSel bit of one specifies that a SDA peripheral device type shall use the LUN as the address of a V-LUI controlled by the SDA.

The logical unit number (LUN) field specifies a logical unit number. For the target's response to an incorrect logical

unit selection see the SCSI Architecture Model.

Only one logical unit number shall be identified per I/O process. The initiator may send one or more IDENTIFY messages during a connection. A second IDENTIFY message with a different value in the LUN field shall not be issued before a BUS FREE phase has occurred; if a target receives a second IDENTIFY message with a different value in this field, it shall issue a Bus Free Request (see Unexpected Bus Free condition, 5.1.3). Thus an initiator may change the DiscPriv bit, but may not attempt to switch to another I/O process. (See the DTDC field of the disconnect-reconnect mode page in the SCSI-3 Command Set for additional controls over disconnection.)

An implied RESTORE POINTERS message shall be performed by the initiator prior issuing a message in response of a message in indication for an IDENTIFY message sent during reconnection.

5.0 8-byte Logical Unit Number structure

The logical unit number (LUN) 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.

This 8-byte LUI structure shall only be valid for SDA type devices.


If the LUN indicates that the command is to be passed to the next layer then the SCSI device shall use bytes 0-1 for the LUN field to determine the LUI to send the command. When the command is sent to the target the LUN value that was received shall be adjusted as follows:

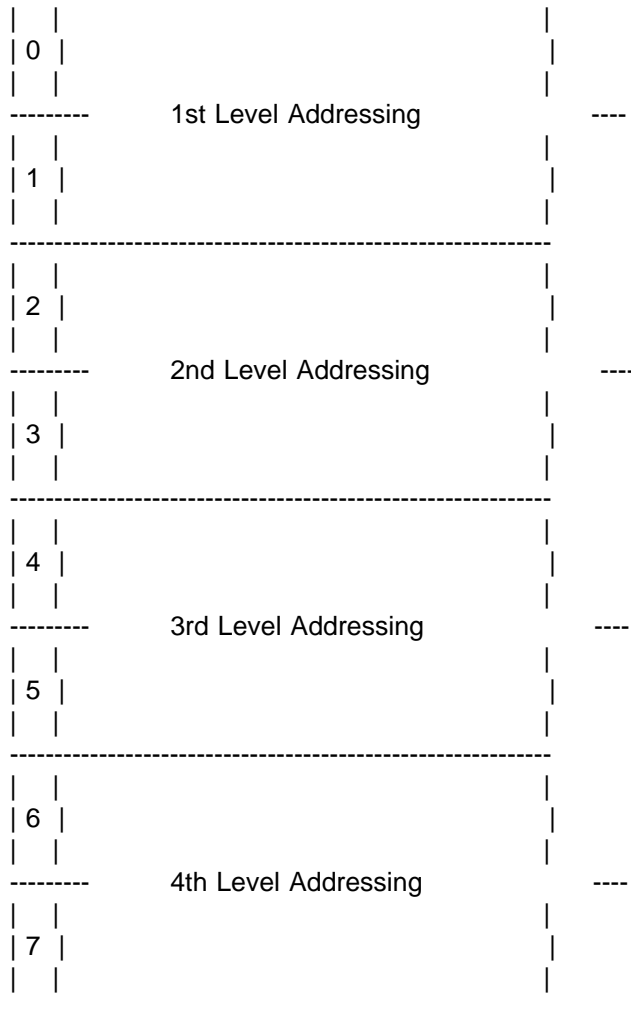
Byte Position		Old	New
0 - 1			Not Used
2 - 3	Moves to		0 - 1
4 - 5	Moves to		2 - 3
6 - 7	Moves to		4 - 5

Bytes six and seven of the new LUN value shall be set to zero.


Each SDA shall keep track of the necessary LUI information to allow reconnection to the correct I/O process during reselection.

.pa





.pa
The nth level addressing field is defined as follows:



The value of the address method field determines the contents of the remaining 14 bits of the nth level addressing field. The following defines the address methods:

Bit	Description
7 6	Physical Logical Units
0 1	Volume Logical Units
1 0	Reserved

1 1 Reserved

Physical logical units are either logical units that are located within the current level or logical units that require the received command to be sent to another target for execution. When this address method is indicated bits 5-0 of byte n shall indicate a bus number.

If the bus number is zero then byte n+1 shall be the LUN address of a logical unit located within the current level. If the bus number is greater than zero then byte n+1 shall indicate a target address. The bus number and target together indicate which SCSI device the command is to be sent for further processing.

Volume logical units are logical units that are controlled by the current level and require one or more accesses to other targets to execute the command. Volumes shall only be accepted on SDA device types. See SCSI Disk Array Model for a complete description of the rules for using volumes. When this address method is indicated bits 5-0 of byte n and 7-0 of byte n+1 shall be the LUN address. Bit 5 byte n shall be the MSB and Bit 0 byte n+1 shall be the LSB.

| Note: The value of the Target is defined by the individual
| protocol documents. (e.g. SIP defines Targets to be in the
| range 0-15.)

6.0 Addressing Examples for the 8-byte LUN structure

Several addressing examples follow. The conventions used within these examples are:

Layer 1 M:P:T or M:L or u
Layer 2 M:P:T or M:L or u
Layer 3 M:P:T or M:L or u
Layer 4 M:P:T or M:L or u

Where:

M is the Address Method (2 bit field)
P is the Bus Number (6 bit field)
T is the Target (8 bit field)
L is the Logical Unit Number (14 bit field)
u is means unused and set to zero (16 bit field)

Note that P is a value that starts at one, since the zero value is reserved for the SDA and devices that have no external path.

Note that T is a value that starts at zero and is limited to one less than the number of attachable SCSI devices, since the path initiator's address is also included in that address space.

Example 1:

Addressing the first layer SDA (for all control, creation, management functions and for identify)

Layer 1 0:0:0
Layer 2 u
Layer 3 u
Layer 4 u

Addresses will appear on the first level paths as required by the function.

Example 2:

Addressing a fan at address 7 within the first layer SDA (a SCSI device not physically on an identifiable SCSI path, or if you prefer, a LUN of the SDA controller itself.)

Layer 1 0:0:7
Layer 2 u
Layer 3 u
Layer 4 u

No first level path will be used.

Example 3:

Addressing a local P-LUI

The address of the second drive on the third path would be:

Layer 1 0:3:1
Layer 2 0:0:0
Layer 3 u
Layer 4 u

The second level path would use path 3 to access target 1. The LUN value of 0 would be taken from the 0:0:0.

Example 4:

Addressing any V-LUI controlled by the first SDA (including V-LUI's constructed from PS-extents defined by lower SDA's)

Layer 1 1:L
Layer 2 u
Layer 3 u
Layer 4 u

Addresses will appear on the first level paths as required by the function.

Example 5:

Addressing an SDA at the second layer of the hierarchy. Note that the second layer SDA is also a physical target on the first layer path. In this example the SDA is on the fourth path and third target address.

Layer 1 1:4:2
Layer 2 0:0:0
Layer 3 u
Layer 4 u

A LUN address of the following form would be emitted on the fourth path to target address 2 (following the shift the address rules) to get at the SDA.

Layer 1 0:0:0
Layer 2 u
Layer 3 u
Layer 4 u

Example 6:

Addressing a P-LUI of the above second level SDA (the second drive on the third path of that second level SDA).

Layer 1 0:4:2
Layer 2 0:3:1
Layer 3 0:0:0
Layer 4 u

The LUN address emitted on the fourth path to target address 2 (following the shift the address rules). This P-LUI could be used as a component of a V-LUI defined by the first level SDA or the second level SDA.

Layer 1 0:3:1
Layer 2 0:0:0
Layer 3 u
Layer 4 u

The LUN address emitted on the third path of the second level SDA to target address 1 would be 0, taken from the 0:0:0 value.

Example 7:

Addressing any V-LUI of the above second level SDA. Note that the second level SDA's entry path is being addressed directly. This V-LUI could be used as a component of a V-LUI defined by the first layer SDA.

Layer 1 0:4:2
Layer 2 1:L
Layer 3 u
Layer 4 u

Address emitted on the fourth path to target address 2 (following your shift the address rules). This is a standard V-LUI address.

Layer 1 1:L
Layer 2 u
Layer 3 u
Layer 4 u

Addresses will appear on the second level paths as required by the function.

Example 8:

Addressing a P-LUI that has real LUN's behind a standard target. As an example, LUN 4 behind the second target on the third path of the first level SDA.

Layer 1 0:3:1
Layer 2 0:0:4
Layer 3 u
Layer 4 u

Note that the LUN address emitted on the third first level path to target 4 is taken from the same place as it normally would be taken. The value would be 0/0/4 (or 4). This is analogous to the fan case, above.

Example 9:

Note that the fourth layer devices must be single LUN devices. As a P-LUI addressing example, if the first layer was B=3,T=7; the second layer was B=4, T=6; the third layer was B=1, T=5; and the fourth layer was B=7, T=2 then the address would be:

Layer 1 0:3:7
Layer 2 0:4:6
Layer 3 0:1:5
Layer 4 0:7:2
0:0:0

the fifth layer must be understood to have only single LUN devices in it. As we shift down we see:

The LUN issued on the first layer path 3, target 7 would be:

Layer 1 0:4:6
Layer 2 0:1:5
Layer 3 0:7:2
Layer 4 0:0:0

The LUN issued on the second layer path 4, target 6 would be:

Layer 1 0:1:5

Layer 2 0:7:2
Layer 3 0:0:0
Layer 4 u

The LUN issued on the third layer path 1, target 5 would be:

Layer 1 0:7:2
Layer 2 0:0:0
Layer 3 u
Layer 4 u

The LUN issued on the fourth layer path 7, target 2 would necessarily be:

Layer 1 0:0:0
Layer 2 u
Layer 3 u
Layer 4 u

7.0 Addressing Examples for the 6-bit LUN structure

The following examples assume the LUN Mapping Page is not active.

Several addressing examples follow. The conventions used within these examples are:

Layer 1 V:L

Where:

V is the Volume Select (1 bit field)

L is the Logical Unit Number (5 bit field)

Note: Information on logical units at layers below layer 1 is not available to the application client when the 6-bit LUN structure is used. Logical units below layer 1 may only be addressed directly using a vendor specific translation of the 6-bit LUN.

Example 1:

Addressing the first layer SDA (for all control, creation, management functions and for identify)

Layer 1 0:0

Addresses will appear on the first level paths as required by the function.

Example 2:

Addressing a fan at address 7 within the SDA.

Layer 1 0:7

| No first level path will be used.

| Example 3:

| Addressing any P-LUI.

| The address of a drive in the SDA:

| Layer 1 0:L

| Addresses will appear on the first level paths as required by
| the function.

| Example 4:

| Addressing any V-LUI:

| Layer 1 1:L

| Addresses will appear on the first level paths as required by
| the function.