

Date: Jan 12, 2001

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: Inter-network addressing for SRP

1 Overview

There have been several discussions on defining a method for transmitting SCSI commands and task functions over multiple networks. This is currently being focused on by the SRP working group. The SRP working group is looking onto SCSI commands and task functions that originate in an application client that is in communication with a SCSI Initiator device that attaches to an InfiniBand (TM) network. The desire is that the application clients SCSI commands and task management functions be transmitted through the InfiniBand network to a converter that would either act as an end-point (i.e., talk-to) or transmit the SCSI information to an end-point (i.e., talk-through) contained within in a non-InfiniBand network (e.g., Fibre Channel, Ethernet).

This proposal defines a way to do this for operations that originate on a InfiniBand network. It does not cover a way to allow an application client in communication with a SCSI initiator device that attaches to a non-InfiniBand network to issue SCSI commands or task management functions across the non-InfiniBand network to a converter and have that converter use talk-through to transmit the SCSI information to an InfiniBand end-point (i.e., the converter must be the end-point). This could be changed if other protocols adopt the some of the changes purposed below.

This proposal also introduces the idea of a WW name to be associated with a SCSI target device. This is not the same as the WW name that is currently defined for ports.

2 Terminology

2.0.1 network converter: A device that allows communication between a InfiniBand network and a non-InfiniBand network. It may support LUN addressing, target port routing, or both.

2.0.2 SCSI target functions: Specific codes within the EXTENDED ADDRESS METHOD field of the 2 byte addressing field in a LUN that cause the indicated function to be carried out by the SCSI target device that contains the addressed target port.

2.0.3 target port routing: The ability of a device to use one of it's ports as a source port to route SCSI information to a destination port. The address of the ports is contained within a SCSI command or task management function.

3 Overview on how inter-network address works

If the network converter supports the talk-through function then the addresses of the source and destination ports must be placed into any information unit (IU) that currently contains a LUN. This information is used by the network converter to route the SCSI information to the correct end-point.

To find the ports a SCSI target function is issued to the network converter that causes a report of the source port addresses the network converter contains and the destination ports that can be reached from each of those source ports. In addition to the addresses of the ports the type of network each source port supports is returned.

The same SCSI target function is used to determine if the network converter supports the talk-to function by reporting the addresses of any logical units has access to.

4 LUN structure

SCSI target functions allow an application client to issue requests to a SCSI target. An application client builds the selected SCSI target function into a LUN and then sends that LUN to a specific target port. See xxx for the LUN format and xxx for a list of SCSI target functions.

Table 1 — Format of addressing fields

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

The ADDRESS METHOD field defines the contents of the ADDRESS METHOD SPECIFIC field. See table 2 for the address methods defined for the ADDRESS METHOD field. The ADDRESS METHOD field only defines address methods for entities that are directly addressable by an application client.

Table 2 — ADDRESS METHOD field values

Code	Description	Clause
10b	Logical unit addressing method	xxx
00b	Peripheral device addressing method	xxx
01b	Device type specific	
11b	Extended addressing	xxx

4.1 Extended addressing

Extended addressing allows for more address methods to be defined for the 8 byte LUN. These may include additional 2 byte formats or formats that use more than 2 bytes. See xxx for the format of the 2 byte EXTENDED ADDRESS METHOD field and xxx for the format of the 8 byte EXTENDED ADDRESS METHOD field.

Table 3 — Format of 2 byte extended addressing fields

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

Table 4 — Format of 2 byte extended addressing fields

Bit Byte	7	6	5	4	3	2	1	0
0	1	1	EXTENDED ADDRESS METHOD					
1	EXTENDED ADDRESS METHOD SPECIFIC							
7								

The extended ADDRESS METHOD field defines the contents of the EXTENDED ADDRESS METHOD SPECIFIC field. See table 2 for the address methods defined for the EXTENDED ADDRESS METHOD field. The EXTENDED ADDRESS METHOD field only defines address methods for entities that are directly addressable by an application client.

Table 5 — ADDRESS METHOD field values

Code	Description	Clause
000000b	Reserved	
000001b	SCSI target function	xxx
000010b-111111b	Reserved	

4.2 SCSI target functions

The SCSI target function defines a number of operations that can be carried out by a SCSI target device. A SCSI target device may support zero or more of the target functions however a single SCSI target device shall only support one instance of each supported SCSI target function regardless of the number of target ports allow entry into the SCSI target device.

See table 6 for the definition of the EXTENDED ADDRESS METHOD SPECIFIC field used when the SCSI target method is selected.

Table 6 — SCSI target function

Bit Byte	7	6	5	4	3	2	1	0
n-1	1	1	SCSI TARGET FUNCTION (000001b)					
n	FUNCTION							

The FUNCTION field identifies SCSI target function to be carried out by the SCSI target device. See xxx for a list of SCSI target functions.

Table 7 — FUNCTION field values

Code	Description	Clause
00h	REPORT ADDRESSABLE OBJECTS	xxx
01h	REPORT ADDRESSABLE OBJECTS ASSOCIATIONS	xxx
02h	REPORT SCSI TARGET DEVICE IDENTIFICATION	xxx
03h-FFh	Reserved	

4.3 REPORT ADDRESSABLE OBJECTS

Reports the address of all source target ports and the destination target ports that are accessed through each source target port. In addition to the address each source port includes information on the type of

network it is attached to (i.e., FC NL_Port, FC N_Port, IB, Ethernet, Parallel SCSI).

Editors Note 1 - GOP: This will be expanded if this technique is agreed to.

4.4 REPORT ADDRESSABLE OBJECTS ASSOCIATIONS

Reports the address of all the target ports and the LUNs of the logical units that are accessed through each target port contained within the SCSI target device.

Editors Note 2 - GOP: This will be expanded if this technique is agreed to.

4.5 REPORT SCSI TARGET DEVICE IDENTIFICATION

Reports a unique world wide name for the addressed SCSI target device. Other information may be returned such as a list of all the target ports and all the logical units contained within the SCSI target device.

Editors Note 3 - GOP: This will be expanded if this technique is agreed to.

4.6 Other target function issues

There is further information which would need to be added. Such as:

If a SCSI target device receives a LUN that contains a SCSI target function and it has no such LUN then it should handle in the same manner as it does today when it receives a command that is address to an unknown LUN.

There should also be a method defined to allow notification of a change to any of the reported information. This would be protocol specific as some protocols do not have a good way to do this.

5 Discovery

To discover the port addresses and/or the LUNs a REPORT ADDRESSABLE OBJECTS function is issued by a application client to the network converter. The network converter shall support the REPORT ADDRESSABLE OBJECTS SCSI target function. A network converter shall report source and destination port addresses or LUNs and may report both.

Editors Note 4 - GOP: How does the IB application client know the device is a network converter?

6 Additions to SRP task management and command information units

To make the jump from InfiniBand to non-InfiniBand network information on the non-InfiniBand network must be placed in the task management and command information units. This information is the source port address and the destination port address that was reported as a result of the REPORT ADDRESSABLE OBJECTS function. The format of this information follows.

Table 8 - Additions to SRP command and task IUs

Bit Byte	7	6	5	4	3	2	1	0
n	(MSB) _____							
n+7	SOURCE PORT ADDRESS _____ (LSB)							
n+8	(MSB) _____							
n+15	DESTINATION PORT ADDRESS _____ (LSB)							

The SOURCE PORT ADDRESS field contains the address of the port on the network converter that the IB message transmitted from.

The DESTINATION PORT ADDRESS field contains the address of the port that the network converter shall send the IB to.

Editors Note 5 - GOP: There are some issues that need to be looked at: Should the size of the port address be expanded to 16 bytes to allow for IPv6 address? Should we allow multiple pairs of source/destination ports which would allow multiple hops in on a single operation?
