

Date: Feb 16, 2001

To: T10 Committee (SCSI)

From: George Penokie (Tivoli)

Subject: Transport of Initiator Identification Across Large Network

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).

At this point it has been determined that a talk-through model is not workable. However, a related issue of identification of the actual source of a task needs to be resolved. In SCSI the source of a command for purposes of reservation (and possibly access controls) is the I in the I_T_L nexus. For a task to be accepted by a logical unit that has, for example a reservation, the I_T_L nexus must be the same for all tasks. It is not easy for gateways to keep track of the initiator identifier of the I_T_L nexus when the task is moving across networks. A better solution would be to have the initiator port that originates a task place an initiator identifier in the command information unit.

That initiator identifier would then be carried with the command information unit all the way to the destination logical unit which could then determine if the task is allowed to be executed.

This proposal defines a way to do this for protocols that support placing the initiator identifier into the command information unit.

2 Terminology

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

2.0.2 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.

3 Overview on inter-network initiator identification

If any SCSI operations are used to restrict access to a logical unit (e.g., reservations, persistent reservations, access controls) then an initiator identifier shall be placed into the command information unit by the initiator port through which the command information unit is routed. If access to the logical unit is restricted the target shall use that initiator identifier to determine if the task is to be accepted or rejected. If access is not restricted to the logical unit then the initiator identifier shall be ignored.

4 Additions to command information units

The format of command information unit with the INITIATOR IDENTIFIER field follows.

Table 1 - Additions to SRP command

Bit Byte	7	6	5	4	3	2	1	0
n	(MSB) _____							
	INITIATOR IDENTIFIER							
n+7	_____ (LSB)							

The INITIATOR IDENTIFIER field contains the identification of the initiator port from which the command information unit originated.

Editors Note 1 - GOP: There are some issues that need to be looked at: Should the size of the initiator identifier be expanded to 16 bytes to allow for IPv6 address?

Editors Note 2 - GOP: Does the initiator identifier be placed into the task information unit?
