

Document: T10/01-289r0
To: T10 Committee Membership
From: Edward A. Gardner, Ophidian Designs
Subject: SRP Teleconference Minutes

Date: August 3, 2001

Minutes of SRP teleconference held Monday, September 24, from 11:00am to 2:00pm CST. Thanks to Greg Pellegrino of Compaq for hosting the teleconference. Thanks to Bob Snively for taking the notes from which these minutes are derived.

Attendance:

John Tyndall	Crossroads
George Penokie	Tivoli/IBM
Greg Pellegrino	Compaq
Bob Griswold	Crossroads
Cris Simpson	Intel
Edward A. Gardner	Ophidian Designs
Nathan Ober	Microsoft
Rob Haydt	Microsoft
Bob Snively	Brocade

Cris Simpson was called away a few minutes into the teleconference and did not return.

Time and contact information for next SRP teleconference:

Friday, September 28
11:00am-2:00pm CST
call in number: 281-518-9999
pass code number: 9655712#
Hosted by Rob Elliott at 281-518-5037.

Agenda and action items for next meeting:

1. Cris Simpson / Rob Haydt: IOControllerProfile attribute extensions to include GUID extension. See heading 1 below.
2. Choice of IB reference: v1.0, v1.0a, both, neither.
3. John Tyndall / Rob Haydt: Third party behavior. See heading 3 below.
4. Removal of processor unit concept. See heading 4 below.
5. Explicitly addressing redirection. See heading 5 below.
6. Any comments provided by any reviewer.

1. Addressing structure and IB/SRP device discovery

Rob Haydt and Cris Simpson discussed the addressing structure.

An SRP target port identifier consists of a 64-bit GUID plus a 64-bit GUID extension. The GUID extension can be assigned by a vendor to facilitate creating multiple target ports from a single hardware GUID. The issue is how an SRP initiator determines the proper GUID extension value.

Cris Simpson has prepared a proposal describing one way this might be done. Various people including Rob Haydt have criticized it as being too complex. Neither the proposal nor the criticism have been made publicly available. Cris Simpson and Rob Haydt were the only people present who had seen either. No decision was possible at this time.

Both were given action items to provide a proposal, the past discussion, or anything else that T10 could discuss in time for the Friday teleconference.

2. Document comment review

See 01-289r0.fdf.

3. Third-party operation

John Tyndall questioned whether third party operations (i.e. extended copy) were sufficiently specified. He was concerned that some additional information might be needed to establish the third-party connection, which might result in "back door" interfaces or proprietary solutions.

Various people present stated that third party connections should be established the same as normal connections. That the copy controller should use a target port identifier to locate and connect to the target port in the same manner as a normal initiator.

Rob Haydt expressed concern that that would not be sufficient. For example a fabric might be configured on an assumption that backup traffic flowed through a specific path. The current provision to specify a specific GID may not be sufficient to ensure the proper path is chosen.

John Tyndall volunteered to look into this, as Crossroads has experience with third-party operation. Rob Haydt offered to work with him.

4. Brocade comment on processor unit

Text of Brocade comment:

What is the point of this "processor units" vs. "IO units" exercise? "Processor unit" is NOT a concept defined in InfiniBand. IOUs and IOCs are indeed defined by IB in the Device Management context, but they are to IB what logical units are to SCSI---abstractions for purposes of representation that have no a priori relationship to any specific underlying implementation. If a device is going to operate in the role of a target, it must present the appearance of being an IOU with IOCs regardless of the device's intrinsic function or implementation---in other words, it must look and act like an "IO unit" (and the text so states, though not so succinctly). Similarly, if a device intends to operate as an initiator, the IOU/IOC representation is irrelevant, and there's no reason to care whether the initiator is an IO unit or a processor unit.

If the point of the exercise (and I'm speculating here) is to suggest that a target that also acts as an initiator should keep the same port identifier in either role, well, it probably ought to say that.

The bottom line here is that figure A.9 should be kept-as-is (except that the "processor unit" labels need be elided), and Figure A.10 should be deleted. The text needs to be similarly adjusted."

Ed Gardner expressed agreement with the comment. Others expressed similar sentiment. In the absence of Cris Simpson and Rob Elliott, further discussion (and any decision) was postponed until Friday.

5. Brocade comment on redirection

Text of Brocade comment:

p. 75: Section C.6.3: would it be worthwhile adding something to the following effect? "A target MAY generate CM:Reject messages for the purposes of redirecting a connection by setting the CM:Reject Reason field to 24 or 25, and that an initiator SHOULD recognize such rejections as redirections, and retry the connection requests as redirected."

This functionality is already in the IB Connection Manager spec, but feels like the kind of thing that a lot of implementers might mindlessly blow off without appropriate encouragement. Maybe it's a non-issue, but this redirection is shaping up as an important function."

Ed Gardner agreed with the sentiment that something should be said about support for redirection. In the absence of Cris Simpson and Rob Elliott, further discussion (and any decision) was postponed until Friday.

6. Brocade comment on memory handle definition for IB

Text of Brocade comment:

As defined in Table 1, a memory descriptor contains a virtual address, a memory handle, and a data length. InfiniBand does not make use of a memory handle as such---it's implicitly defined by the QP to which an RDMA operation applies.

InfiniBand does, however, define the notion of an "Remote Key", or R_Key, that's associated with each memory area made available for RDMA access. The R_Key must be present in the RDMA Extended Transport Header (RETH) found in each RDMA---related IB packet, and perforce must be communicated from the initiator to the target as part of each memory descriptor. The fairly obvious choice is to stick the R_Key value into each memory descriptor's memory handle field; but it seems like this ought to be explicitly stated."

It was agreed that a mapping from the SRP memory descriptor to IB RDMA operations needs to be included in Annex C.