

To: T10 Technical Committee
 From: Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
 Date: 18-19 January 2001
 Subject: Minutes of the SRP WG – January 18-19, 2001 – Lake Buena Vista, FL

Revision History

Revision 0: first revision

Attendance

Name	S	Organization	Electronic Mail Address
Mr. William Lynn	V	Adaptec, Inc.	bill_lynn@corp.adaptec.com
Mr. Bill Galloway	P	BREA Technologies, Inc.	billg@breatech.com
Mr. Robert Snively	P	Brocade Comm. Systems	rsnively@brocade.com
Mr. David Peterson	P	Cisco Systems, Inc.	dap@cisco.com
Mr. Claudio Desanti	V	Cisco Systems, Inc.	cds@cisco.com
Mr. Robert C. Elliott	P	Compaq Computer Corp.	Rob_Elliott@compuserve.com
Mr. Wayne Bellamy	V	Compaq Computer Corp.	wayne.bellamy@compaq.com
Mr. Greg Pellegrino	V	Compaq Computer Corp.	greg.pellegrino@compaq.com
Mr. Ronald Stockford	P	Dell Computer Corp.	ron_stockford@dell.com
Mr. Robert H. Nixon	P	Emulex	bob.nixon@emulex.com
Mr. Ralph O. Weber	P	ENDL Texas	roweber@acm.org
Mr. Gaby Hecht	V	Gadzoox	ghecht@gadzoox.com
Mr. Donald Woelz	P	Genroco, Inc.	don@genroco.com
Mr. Randy Haagens	P	Hewlett Packard Co.	randy_haagens@hp.com
Mr. Cris Simpson	P	Intel Corp.	cris.simpson@intel.com
Mr. Dennis Moore	P	KnowledgeTek, Inc.	dmoore@ix.netcom.com
Mr. John Lohmeyer	P	LSI Logic Corp.	lohmeier@t10.org
Ms. Jie Fan	A	Madison Cable Corp.	jie.fan@madisoncable.com
Mr. Rob Haydt	A#	Microsoft Corp.	robhay@microsoft.com
Mr. Charles Monia	P	Nishan Systems Inc.	cmonia@nishansystems.com
Mr. Edward A. Gardner	P	Ophidian Designs	eag@ophidian.com
Mr. Mark Evans	P	Quantum Corp.	mark.evans@quantum.com
Mr. Martin Czekalski	V	Quantum Corp.	marty.czekalsky@quantum.com
Mr. Kenneth Moe	P	Sun Microsystems, Inc.	kenneth.moe@sun.com
Mr. William C. Terrell	P	Troika Networks, Inc.	terrell@TroikaNetworks.com
Mr. Ron Mathews	AV	UNISYS Corporation	ronald.mathews@unisys.com
Mr. Roger Cummings	AV	Veritas Software	roger.cummings@veritas.com

27 People Present

Status Key: P - Principal
 A,A# - Alternate
 AV - Advisory Member
 L - Liaison
 V - Visitor

Results of meeting

John Lohmeyer opened the SCSI over RDMA protocol meeting at 1:15 pm Thursday 18 January 2001 and thanked Adaptec for hosting. This protocol standard maps SCSI over InfiniBand™ Architecture, Virtual Interface (VI) Architecture, and similar transports supporting RDMA (remote direct memory access).

Agenda

The agenda was approved at the meeting.

1. Opening remarks and introduction
2. Attendance and membership
3. Approve agenda
4. Document distribution
5. Review minutes of previous meeting (01-024r0 by Rob Elliott)
6. Review old action items

7. Proposed InfiniBand annex (01-028r1 by Greg Pellegrino and Rob Elliott)
8. Internetwork addressing for SRP (01-034r0 by George Penokie)
9. Document review (SRP revision 2 by Ed Gardner)
10. Connecting to an IO Device (01-045r0 by Ed Gardner)
11. Review new action items
12. Meeting schedule
13. Adjournment

Topics

7. Proposed InfiniBand annex (01-028r1 by Greg Pellegrino and Rob Elliott)

[Thu and Fri]

Rob Haydt asked whether the annex should be normative or informative. The intention is that it be normative, i.e. there is no approved way to run SRP over InfiniBand that doesn't follow the rules in the annex.

Rob Haydt asked that "verbs" be left out of all figures and text.

Ed Gardner asked that reference to Volume 2 concepts like module and chassis be left out of the annex. Although they may be useful for discussion, they will probably not be needed in the final text.

Ed Gardner asked that if GUID is identified as EUI-64, a reference to the EUI-64 specification should be included. Otherwise, just refer to GUID and let the InfiniBand specs refer to EUI-64.

Rob Haydt expressed dissatisfaction with the "platform GUID" concept. Ed Gardner asked that GUID not be used in the name, as that implies an EUI-64. Ed plans to add a generic initiator identifier to the SRP login procedure, as the issue is not InfiniBand-specific. Jim Hafner noted that iSCSI has a similar feature. Ed felt that the annex should not attempt to describe how to create this unique identifier. The group did not discuss if this is only needed for the initiator or whether a target identifier is also needed. Cris Simpson noted that InfiniBand channel construction involves exchanging 32-bit connection IDs.

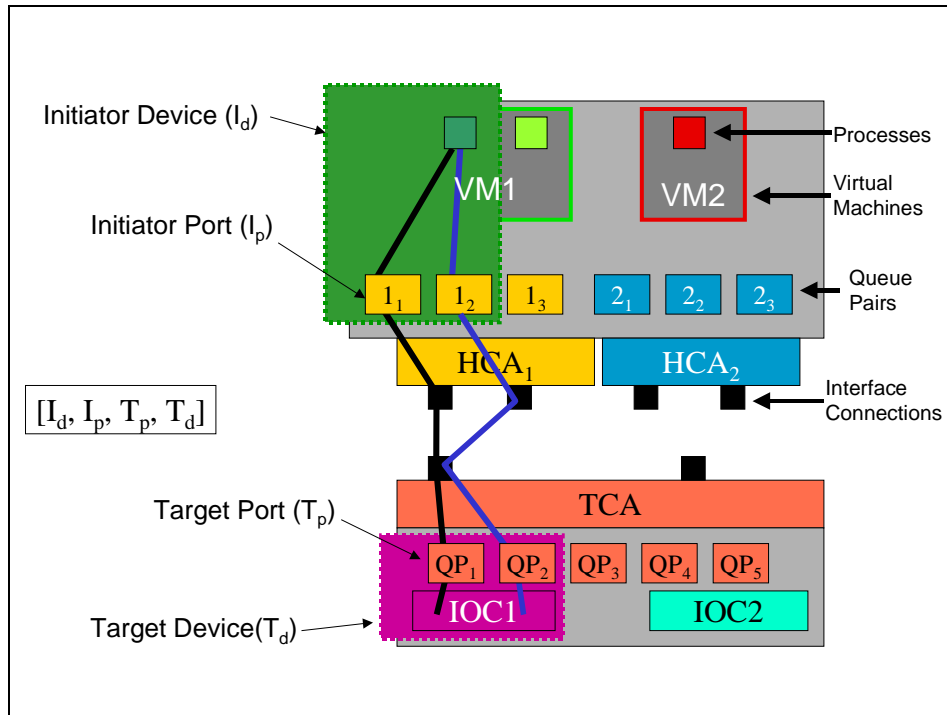
Cris Simpson noted that InfiniBand did not intend to use the term "node" and would prefer that "channel adapter" be used. Until a correction to Volume 1 is published with new terminology, the annex proposal will continue to use "node."

The group agreed to these definitions:

- "Initiator port identifier" is the target's QPN. This is the only identifier that target SRP software needs to use to identify the initiator once a connection has been established.
- "Target port identifier" is the initiator's QPN.
- "Initiator device identifier" is the initiator's GUID. This is a new identifier to the SCSI architecture, identifying the "initiator device" in 00-268r3.
- "Target device identifier" is the target's GUID.

The first two are the identifiers used when referring to an I_T nexus (IP_TP nexus?). The ID and TD identifiers are not currently used in SCSI.

Cris Simpson provided a picture with processes, virtual machines, and queue pairs, showing the location of initiator port, initiator device, target port, and target device. This picture will be added to the annex to help future discussions.



Rob Haydt stated that Microsoft would like to manage everything with IPv6 addresses (the GID in InfiniBand).

Additional comments will be incorporated in 01-028r2.

8. Internetwork addressing for SRP (01-034r0 by George Penokie)

[Thu]

George proposed:

- 1) Using the LUN with byte 0 bits 7:6 set to 11b as a way to address the target device rather than a logical unit. This would provide an alternative to overloading LUN 0 with target-based functionality (e.g. REPORT LUNS).
- 2) Defining a set of target commands encoded in the rest of the LUN field. These commands would help describe the LUN structure.
- 3) Adding source and destination port identifiers to the SRP packet structure. These would be used when bridging from one SCSI domain to another (e.g. InfiniBand SRP to Fibre Channel FCP gateway).

George agreed to use standard SCSI commands rather than encode them in the LUN field.

The group discussed the usefulness of source and destination port identifiers at length. In an IB to FC gateway the source port ID would be the ID of the IB initiator and the destination port ID would be the ID of the FC target.

This doesn't help unless protocols on both sides of the gateway support the new fields. Otherwise, the gateway has to resolve all multiple initiator issues itself (e.g. two IB hosts trying to use reservations on one FC target). The source/destination fields only work through one level, too. The alternative is a gateway that makes remote devices appear as native devices, so the FC devices appear as IB SRP devices on the IB side and IB devices appear as FC devices on the FC side. This requires no protocol modifications, requires no LUN mapping, requires no access control and reservation handling code in the gateway, and supports initiators and targets on either side of the gateway.

George will revise the proposal.

9. Document review

[Fri]

Bill Galloway raised a few issues with SRP revision 2.

Bill Galloway proposed redefining the CDB IU to move the Additional CDB Length to byte 1 and eliminating the multiple command IUs (16-byte, 32-byte, and variable-length). George Penokie protested moving the Additional CDB Length to a different place from other protocols. The people previously arguing for multiple command IUs were not present, so Bill will post a message on the issue on the T10 reflector.

Bill asked why memory descriptors (buffer descriptors) can be included in the command IU when they are already available via the indirect pointer. The group replied that it is an optimization, and the target is not obligated to use them.

Bill noted section 4.5 (of revision 2) says "The target may send one or more requests before returning the target login response IU" while section 5.3 says "as the second message (first sent by the target) on a newly established connection." Ed Gardner replied that the 5.3 behavior is intended and section 4.5 will be corrected in revision 3.

Bill asked why the devices would ever want to renegotiate the MaxCmdIUSize and MaxRspiIUSize while operating. Ed replied that a recipient is allowed to reject such a change by always returning the same value. The group discussed whether requiring a new login is preferable and the PAR IUs could be removed. The PAR IU does contain the RequestLimitDelta field which is needed, and may also be used as a ping/no-op message. The group unanimously recommended Ed remove the IU Size fields from the messages.

10. Connecting to an IO device (01-045r0 by Ed Gardner)

[Fri]

Ed presented a slide set on finding SRP devices on InfiniBand he gave at the InfiniBand 1.0 event in Las Vegas.

11. Review new action items

Bill Galloway will post a message to the T10 reflector suggesting removing the multiple command IUs.

Ed Gardner will remove the IU Size fields from the PAR IUs.

12. Meeting schedule

InfiniBand Application Working Group members are meeting in San Francisco on these dates:

24 January 2001 Wed IBTA AWG

25 January 2001 Thu IBTA AWG

26 January 2001 Fri IBTA AWG

An interim SRP meeting will be held in conjunction with the InfiniBand Trade Association Applications Working Group in Denver:

20 February 2001 Tue Denver SRP 1pm – 7pm with AWG

21 February 2001 Wed Denver SRP 9am – 6pm with AWG

The next T10 week in Dallas, TX at the Crowne Plaza Suites, hosted by Adaptec, will have special SRP meetings in addition to the usual CAP meeting:

7 March 2001 Wed Dallas CAP

8 March 2001 Thu Dallas SRP (begins 1 hour after T10 plenary, usually around 2pm)

9 March 2001 Fri Dallas SRP 9am – 4pm

13. Adjournment

The meeting adjourned at 11:15 am Friday.

.