

To: T10 Technical Committee
 From: Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
 Date: 3 May 2001
 Subject: Minutes of the SRP WG – May 3, 2001 – Nashua, NH

Revision History

Revision 0: (3 May 2001) first revision

Revision 1: (10 May 2001) a few corrections (changes highlighted)

Agenda

~~Agenda – SCSI SRP Working Group
 May 3, 2001 – Dallas, TX~~

1. Opening remarks and introduction
2. Attendance and membership
3. Approve agenda
4. Review minutes from previous meeting (01-092r0 by Greg Pellegrino)
5. Review old action items
6. SRP document status (John Lohmeyer)
7. SRP-2 project proposal (01-~~xxxr0~~-160r0 by Rob Elliott)
8. SRP InfiniBand™ annex (01-028r4 by Greg Pellegrino and Rob Elliott)
9. Names, addresses, and identifiers, Oh My! (01-084r1 by George Penokie)
10. Reduce SRP_CMD IU types to one (01-164r0 by Bill Galloway)
11. Flow control question (Cris Simpson)
12. SRP Multichannel proposal (01-085r0 by Cris Simpson)
13. SRP Document review (srp-r04 by Ed Gardner)
14. Review new action items
15. Meeting schedule
16. Adjournment

Topics

1 Opening Remarks and introduction

Ralph Weber opened the SCSI over RDMA protocol meeting at 1:00 pm Thursday 3 May 2001 and thanked Hitachi for hosting. This protocol standard maps SCSI over InfiniBand™ Architecture, Virtual Interface (VI) Architecture, and similar transports supporting RDMA (remote direct memory access).

2 Attendance and membership

Name	S	Organization	Electronic Mail Address
Mr. Bill Galloway	P	BREA Technologies, Inc.	billg@breatech.com
Mr. Robert C. Elliott	P	Compaq Computer Corp.	Robert.Elliott@compaq.com
Mr. John Tyndall	P	Crossroads Systems, Inc.	jtyndall@crossroads.com
Mr. Robert H. Nixon	P	Emulex	bob.nixon@emulex.com
Mr. Ralph O. Weber	P	ENDL Texas	roweber@acm.org
Mr. Joe Breher	P	Exabyte Corp.	joebre@exabyte.com
Mr. Randy Haagens	P	Hewlett Packard Co.	randy_haagens@hp.com
Mr. George O. Penokie	P	IBM / Tivoli Systems	gpenokie@tivoli.com
Mr. John P. Scheible	A	IBM Corp.	Scheible@vnet.ibm.com
Dr. Jim Hafner	V	IBM Corp.	hafner@almaden.ibm.com
Mr. Dan Colegrove	V	IBM Corp.	dcolegro@us.ibm.com
Mr. Cris Simpson	P	Intel Corp.	cris.simpson@intel.com
Mr. John Lohmeyer	P	LSI Logic Corp.	lohmeier@t10.org
Mr. Mark Evans	P	Maxtor Corp.	mark.evans@quantum.com

Mr. Rob Haydt	A# Microsoft Corp.	robhay@microsoft.com
Mr. Gerald Maurer	V QLogic Corp.	gerald.maurer@qlogic.com
Mr. Kenneth Moe	P Sun Microsystems, Inc.	kenneth.moe@sun.com
Mr. Phil Shelton	V UNISYS Corporation	phil.shelton@unisys.com

18 People Present

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

3 Approve agenda

The agenda was approved as listed above.

4 Review minutes from previous meeting (01-092r0 by Greg Pellegrino)

The minutes from the March meeting, 01-092r0, were approved.

5 Review old action items

No action items were created at the previous meeting.

6 SRP Document status (John Lohmeyer)

John reviewed the status of SRP revision 4. On Tuesday, Ed Gardner emailed that he was sick and unable to attend this week. Revision 4, incorporating all of the SRP work of the working groups since November, has been expected since January.

John will ask Ed to deliver revision 4 by 10:30am Thursday 10 May 2001 or whatever source files exist on that date.

7 SRP-2 Project proposal (01-~~xxxr0~~-160r0 by Rob Elliott)

In the interest of getting SRP published in a timely matter, the multichannel proposal by Cris Simpson of Intel may best be deferred until SRP-2. This project proposal lists that as the only specific work item (in addition to the usual "anything else" category).

The group reviewed the proposal, recommending a target date for public review is December 2002.

A document number will be obtained for the proposal as presented and revised in the meeting.

8 SRP InfiniBand™ annex (01-028r4 by Greg Pellegrino and Rob Elliott)

The group reviewed the annex. The group discussed the ordering implications of multiple channels between the same initiator port and target port. The March meeting contained a discussion of that topic. The minutes 01-092r0 reflect this conclusion:

An SRP target port receiving a login that communicates an initiator port identifier already in use at that target port could have one of two reactions:

- disconnect existing connections and abort any outstanding tasks on those connections
- accept the new connection and treat it as independent, except for checking reservation conflicts.

The initiator could indicate in login which of the two options the target should do. The login response can indicate whether anyone else was/is logged in. The group agreed Ed should pursue that approach.

Case b) may be too vague to create a usable system, even with software that doesn't expect ordering in the fabric.

A picture showing a target in a processor node and an initiator in an IO Unit will be added.

The term "port" will be changed to 'InfiniBand port' wherever appropriate to keep it from being confused with SCSI port.

Cris Simpson will propose a Login Reject IU that can be carried as PrivateData with the CM_REJ message.

The group discussed whether a CM disconnect can destroy the queue pair before the application has processed all outstanding IOs on that queue pair. If so, a Logout Acknowledge IU would be needed to ensure clean disconnections. Cris Simpson will research the issue.

The alias proposal 00-425r2 has changed its terms to "designations." The annex needs to be updated to match the new terminology.

Cris will work with the InfiniBand Trade Association to fix the IOController attributes page in AWG (perhaps replacing it with a new page).

Rob will prepare Revision 5 in Frame format.

9 Names, Addresses, and Identifiers, Oh My! (01-084r1 by George Penokie)

The group reviewed the SRP columns of 01-084r1 and George made several corrections.

The group agreed that a note should be added to SRP indicating that linked commands are not supported.

The group agreed that SRP needs to clarify that the initiator port identifier is also a name that needs to be persistent. It serves as both the initiator port identifier and the "initiator port name" (a generic replacement for "Worldwide_Name" used in SPC-2's persistent reservations clause).

Jim Hafner will propose eliminating the initiator port identifier and target port identifier sizes from SAM-2 in an upcoming "port name" proposal.

10 Reduce SRP_CMD IU types to one (01-164r0 by Bill Galloway)

Bill proposed combining SRP_CMD_16, SRP_CMD_32, and SRP_CMD_LONG back into a single SRP_CMD_IU. The target must accept the whole packet and check its CRC before parsing it, so there is no real advantage in fixed length formats.

The group recommended that 01-164r0 be incorporated into SRP by a vote of 9-0-0.

11 Flow control question (Cris Simpson)

Cris asked why InfiniBand flow control is not sufficient and SRP implements its own. The November 2000 SRP minutes (01-009r0) contain the explanation, repeated here:

1. This protocol is not designed solely for InfiniBand. Generic VI transports do not include flow control – if a queue is overflowed, the entire connection is closed.
2. IB flow control may not be sufficient. IB flow control stalls the whole connection, preventing task management functions from being issued. SRP potentially allows prioritization.

12 SRP Multichannel proposal (01-~~086r0~~ 085r0 by Cris Simpson)

Cris reviewed his proposal for supporting multiple channels. Questions were asked about ordering implications if multiple control channels are allowed.

Cris agreed to make sure everything defaults to SRP-1 compatible behavior.

13 SRP Document review (srp-r04 by Ed Gardner)

This topic was deferred until the next meeting.

14 Review new action items

Cris Simpson will propose a Login Reject IU that can be carried as PrivateData with the CM_REJ message.

Cris Simpson will research whether a Login Acknowledge IU is needed.

Cris Simpson will work with the InfiniBand Trade Association to fix the IOController attributes page in AWG (perhaps replacing it with a new page).

The SRP editor should add a note should be added to SRP indicating that linked commands are not supported.

The SRP editor should clarify that the initiator port identifier is also a name that needs to be persistent.

The SRP editor should incorporate 01-064r0.

Jim Hafner will propose eliminating the initiator port identifier and target port identifier sizes from SAM-2 in an upcoming "port name" proposal.

15 Meeting schedule

The next T10 week in Colorado Springs, CO at the Wyndham Colorado Springs Hotel, hosted by LSI Logic, will have special SRP meetings in addition to the usual CAP meeting.

10 May 2001 Thu 10:30am Mountain time: Ed's due date

11 May 2001 Fri 1pm-2pm Pacific time: conference call hosted by Cris Simpson, Intel.

Topic: what to do (please RSVP to Cris)

additional conference calls TBD.

letter ballot date TBD.

18 June 2001 Mon Redmond, WA 9am – 7pm (RSVP recommended – Rob Haydt, Microsoft)

19 June 2001 Tue Redmond, WA 9am – 1pm

18 July 2001 Wed Colorado Springs CAP

19 July 2001 Thu Colorado Springs SRP (begins 1 hour after T10 plenary, usually around 1:00 pm

20 July 2001 Fri Colorado Springs SRP

16 Adjournment

The meeting adjourned at 6pm Thursday.