

To: T10 Technical Committee
 From: Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
 Date: 8-9 March 2001
 Subject: Minutes of the SRP WG – March 8-9, 2001 – Dallas, TX

Revision History

Revision 0: (8-9 March 2001) first revision

Agenda

Agenda -- SCSI SRP Working Group
 March 8-9, 2001 -- Dallas, TX

1. Opening remarks and introduction
2. Attendance and membership
3. Approve agenda
4. Review minutes from previous meeting (01-073r0 by Greg Pellegrino)
5. Review old action items
6. SRP InfiniBand™ Target Identification (01-079r0 by Ed Gardner)
7. SRP Initiator Identification (01-080r0 by Ed Gardner)
8. SRP Multichannel proposal (01-085r0 by Cris Simpson)
9. SRP Transfer Length proposal (01-086r0 by Keith Holt)
10. SRP Dropping TARGET RESET support (01-108r0 by Rob Elliott)
11. Names, addresses, and identifiers, Oh My! (01-084r0 by George Penokie)
12. SRP InfiniBand™ annex (01-028r3 by Greg Pellegrino and Rob Elliott)
13. Microsoft document review comments (Rob Haydt)
14. SRP Document review (srp-r04 by Ed Gardner)
15. Review new action items
16. Meeting schedule
17. Adjournment

Topics

1 Opening Remarks and introduction

Ralph Weber opened the SCSI over RDMA protocol meeting at 1:35 pm Thursday 8 March 2001 and thanked Texas Instruments for hosting. This protocol standard maps SCSI over InfiniBand™ Architecture, Virtual Interface (V) Architecture, and similar transports supporting RDMA (remote direct memory access).

2 Attendance and membership

Name	S	Organization	Electronic Mail Address
Mr. Robert C. Elliott	P	Compaq Computer Corp.	Robert.Elliott@compaq.com
Mr. Robert H. Nixon	P	Emulex	bob.nixon@emulex.com
Mr. Ralph O. Weber	P	ENDL Texas	roweber@acm.org
Mr. Nathan Hastad	P	General Dynamics	nathan.j.hastad@gd-is.com
Mr. Donald Woelz	P	Genroco, Inc.	don@genroco.com
Mr. George O. Penokie	P	IBM / Tivoli Systems	gpenokie@tivoli.com
Dr. Jim Hafner	V	IBM Corp.	hafner@almaden.ibm.com
Mr. Cris Simpson	P	Intel Corp.	cris.simpson@intel.com
Mr. John Lohmeyer	P	LSI Logic Corp.	lohmeier@t10.org
Mr. Keith Holt	V	LSI Logic Corp.	keith.holt@lsil.com
Mr. Rob Haydt	A#	Microsoft Corp.	robhay@microsoft.com
Mr. Edward A. Gardner	P	Ophidian Designs	eag@ophidian.com

Mr. Mark Evans	P	Quantum Corp.	mark.evans@quantum.com
Mr. Gerald Houlder	P	Seagate Technology	gerry_houlder@seagate.com
Mr. William C. Terrell	P	Troika Networks, Inc.	terrell@TroikaNetworks.com
Mr. Rick Casaly	A	Troika Networks, Inc.	rcasaly@troikanetworks.com
Mr. Ron Mathews	AV	UNISYS Corporation	ronald.mathews@unisys.com

17 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 (01-073r0 by Greg Pellegrino)

The minutes from the February meeting, 01-073r0, were approved.

5 Review old action items

Bill Galloway will post a message to the T10 reflector suggesting removing the multiple command IUs. Bill did not do that; this action item will be removed.

Ed Gardner will remove the IU Size fields from the PAR IUs. Ed is doing so in revision 4; this action item will be removed.

6 SRP Target Identifiers (01-079r0 by Ed Gardner)

Ed proposed using an EUI-64 as the target port identifier, communicated during login. It is 64 bits and can be set to the IOC GUID in InfiniBand. It can be used for other VI transports too. Targets implemented in processor nodes would have to provide management datagram support and offer the IOUnit, IOController, and ServiceEntries attribute pages.

Rob Elliott asked that a larger identifier (e.g. 128 bits) be used so the target port identifier can be constructed from channel adapter GUIDs for targets implemented in processor nodes (e.g. where multiple operating systems are running on the same hardware and presenting separate target ports). It would be difficult to create a worldwide unique 64-bit identifier from another 64-bit identifier; it is easy to create a 128-bit identifier.

Rob Haydt suggested that different service IDs could be used to differentiate the different target ports. The channel adapter GUID could be shared among targets.

Bob Nixon suggested that target software own a range of GUIDs so it can provide a unique target port identifier for every target that it generates. It can generate virtual IOCs and IOC GUIDs (EUI-64s) as needed. This preserves the concept that IOC can equal target port. It turns GUID assignment into a licensing issue. The group agreed with this reasoning.

Bob expressed concern about using this identifier for non-InfiniBand VI transports. FC-VI uses IPv4 addresses and does not have EUI-64s available, for example. Ed felt that software could generate them as described for InfiniBand processor nodes.

7 SRP Initiator Identifiers (01-080r0 by Ed Gardner)

Ed proposed communicating an initiator port identifier during login.

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:

- a) disconnect existing connections and abort any outstanding tasks on those connections
- b) 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.

Ed will change “connection” to “channel” in the future.

For the format of the initiator port identifier, Ed proposed EUI-64. Cris Simpson noted the Wired for Management specification has an algorithm to generate 128 bit globally unique IDs. The group agreed to increase it to 128 bits with an EUI-64 as the base.

Cris later provided these links on generating identifiers:

- <http://www.alternic.org/drafts/drafts-l-m/draft-leach-uuids-guids-01.html>
- <http://archives.neohapsis.com/archives/crypto/2000-q4/0128.html>
- <http://www.lyra.org/pipermail/dav-dev/1999-May/000245.html>
- http://msdn.microsoft.com/library/psdk/midl/mi-laref_2g10.htm

The group debated how the format should be documented. Rob Haydt disagreed with mandating that EUI-64 be used; he’d prefer that it “should” be used. Rob Elliott requested that EUI-64 be mandated to avoid interoperability problems between initiators that follow the recommendation and those that do not. The group agreed to phrase the requirement as “shall be globally unique and should be based on EUI-64.”

8 SRP Transfer Length proposal (01-086r0 by Keith Holt)

The last WG agreed to add a total data length field when indirect lists are used. This proposal documents the specific format. This is mainly for IB to FC bridges, since FC packets need the Data Length field upfront.

The target is not required to check this length compared to that computed by adding up the lengths in the SG list. Ed asked if the target happens to learn that a mismatch occurs, can it report an error? How?

Rob Haydt felt that if the SG list were longer than the total transfer length, there is no problem and the target should be silent. If the SG list were shorter, the target should definitely report an error.

There are several lengths involved:

- a) the total transfer length field;
- b) the sum of the SG list entry lengths;
- c) the length indicated in the command itself; and,
- d) the actual amount of data transferred.

Ed asserted that the total transfer length must equal the SG list entry lengths by definition, so the underrun calculation can use either value. The group agreed that each descriptor will define what constitutes the “data length” used for under/overrun calculations.

9 SRP Multichannel proposal (01-086r0 by Cris Simpson)

This topic was deferred to the next meeting.

10 SRP Dropping Target reset support (01-108r0 by Rob Elliott)

Based on T10's acceptance of 01-015r2, this proposes that SRP not support the TARGET RESET task management function.

Rob Elliott moved and George Penokie seconded that 01-108r0 be incorporated into SRP. The group recommended this by a vote of 9-1-2.

11 Names, Addresses, and Identifiers, Oh My! (01-084r0 by George Penokie)

The group filled in the SRP columns of this document, presented the previous day at the CAP working group.

- Initiator port identifier size = 16 bytes (EUI-64 + identifier extension)
- Target port identifier size = 8 bytes (EUI-64)
- Initiator port name size = 16 bytes (EUI-64 + identifier extension)
- Target port name size = 8 bytes (EUI-64)

The port name and port identifier are the same.

The device identifier is not defined (same as all other protocols so far).

Persistent reservations and reservations use the initiator port identifier (or name). Extended Copy uses the target port identifier (plus other information).

12 Microsoft document review comments (Rob Haydt)

Rob had two general concerns:

1. Not much reflection of multi-hosted nature of InfiniBand (e.g. multiple hosts through single IB to FC bridge to storage). Jim Hafner noted that access controls and reservations can be implemented on the target. InfiniBand partitioning can be used to hide devices from each other (outside the scope of SRP).
2. Overloading functionality on top of an InfiniBand connection establishment. Including the SRP LOGIN in PrivateData seems to violate protocol layering and may lead to problems in the future (due to the limited length available). Cris Simpson replied that SRP was following the InfiniBand trade association's intent by putting additional login-related information there. It allows the recipient to refuse to create an InfiniBand connection based on the data, possibly saving queue pairs.

13 SRP InfiniBand™ annex (01-028r3 by Greg Pellegrino and Rob Elliott)

The initiator port identifier and target port identifier formats were chosen earlier in the meeting.

The third-party target descriptors were discussed. The group agreed that LID never needs to be specified; given a GID, the copy manager can query the subnet manager to find the LID and other PathRecord information. The format indicating just the target port identifier is preferred, but requires a "name server" (not the Subnet Manager... something like the Configuration Manager) to relate port identifiers to GIDs.

Revision 4 will incorporate changes discussed at this meeting.

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

This topic was deferred until the next meeting.

15 Review new action items

No new action items were created.

16 Meeting schedule

The next T10 week in Nashua, NH at the Sheraton Nashua Hotel, hosted by Hitachi Cable Manchester, will have special SRP meetings in addition to the usual CAP meeting:

2 May 2001 Wed Nashua CAP

3 May 2001 Thu Nashua SRP (begins 1 hour after T10 plenary, usually around 1:30 pm)
multichannel proposal and other tech topics

4 May 2001 Fri Nashua SRP 9am – 2pm
mostly document review

There will be no interim meetings. A conference call may be held for Cris Simpson's multichannel proposal (and any other topics like the SAM mapping annex) – Cris Simpson will host.

17 Adjournment

The meeting adjourned at 6:45 pm Friday.