To: **T10 Technical Committee**

From: **Cris Simpson, Intel Corp.** ***cris.simpson@intel.com***

Date: **11 April 2001**

Subject: **SRP MultiChannel Proposal**

Revision History

Revision 0: First Revision

Related Documents

T10/SRP-r03 – SCSI over RDMA protocol revision 3

T10/01-028r4 - SRP InfiniBand™ annex

Overview

This document proposes a method for providing an SRP Session over multiple channels. The ability to specify the set of channels over which the data transfer(s) shall occur is introduced. Also introduced is support for increased availability through a Control Channel Group, which allows control traffic to be carried over multiple channels.

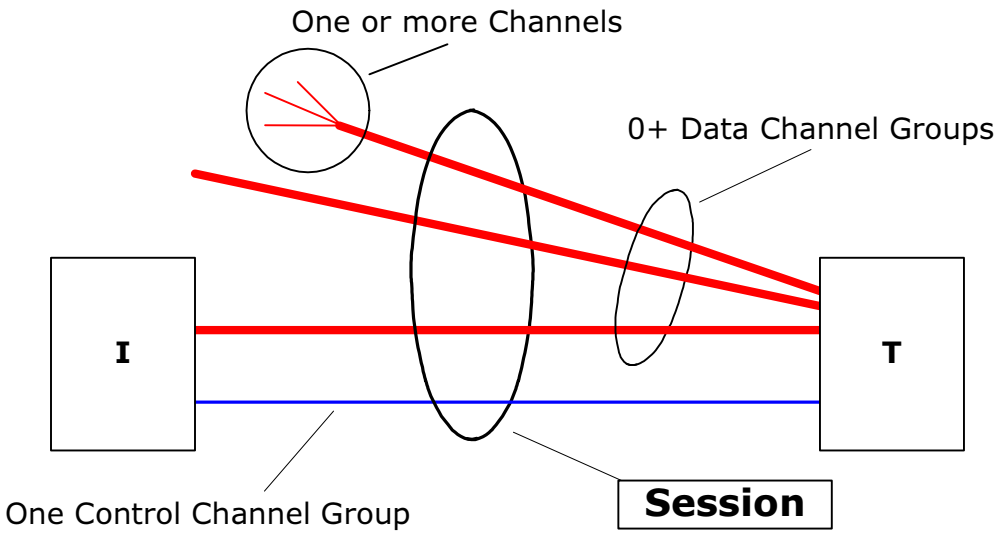

Figure 1 - SRP Session model

1 Terms

Channel	The association of two Queue Pairs for communication.
Channel Group	A set of Channels having identical memory access rights.
Channel Group Identifier	A numeric value that identifies a Channel Group within a Session.
Control Channel Group	The Channel Group that carries control traffic.
Data Channel Group	A Channel Group dedicated to carrying data traffic.
Session	One Control Channel Group and zero or more Data Channel Groups with a common Session Identifier.
Session Identifier	A numeric value that identifies a Session within a target.

2 Changes

2.1 SRP_LOGIN_REQ

Table 2 - SRP_LOGIN_REQ additions

Bit / Byte	7	6	5	4	3	2	1	0
N	SESSION ID							
N+1								
N+2	CHANNEL GROUP ID							
N+3								

SESSION ID

SESSION ID = 0xFFFF

This SRP_LOGIN_REQ is requesting the establishment of a new session. CHANNEL GROUP ID shall be set to 0, and the newly-established channel shall be a member of Channel Group 0.

Valid SESSION ID

This SRP_LOGIN_REQ is requesting that a channel be added to the session.

Invalid SESSION ID

The Status Response INVALID SESSION ID shall be returned.

Channel Group ID

CHANNEL GROUP ID = 0xFFFF

Requests the establishment of a new Channel Group, with the newly-established channel as a member.

VALID CHANNEL GROUP ID

Requests a new channel to be established as a member of the specified Channel Group.

INVALID CHANNEL GROUP ID

The Status Response INVALID CHANNEL GROUP ID shall be returned.

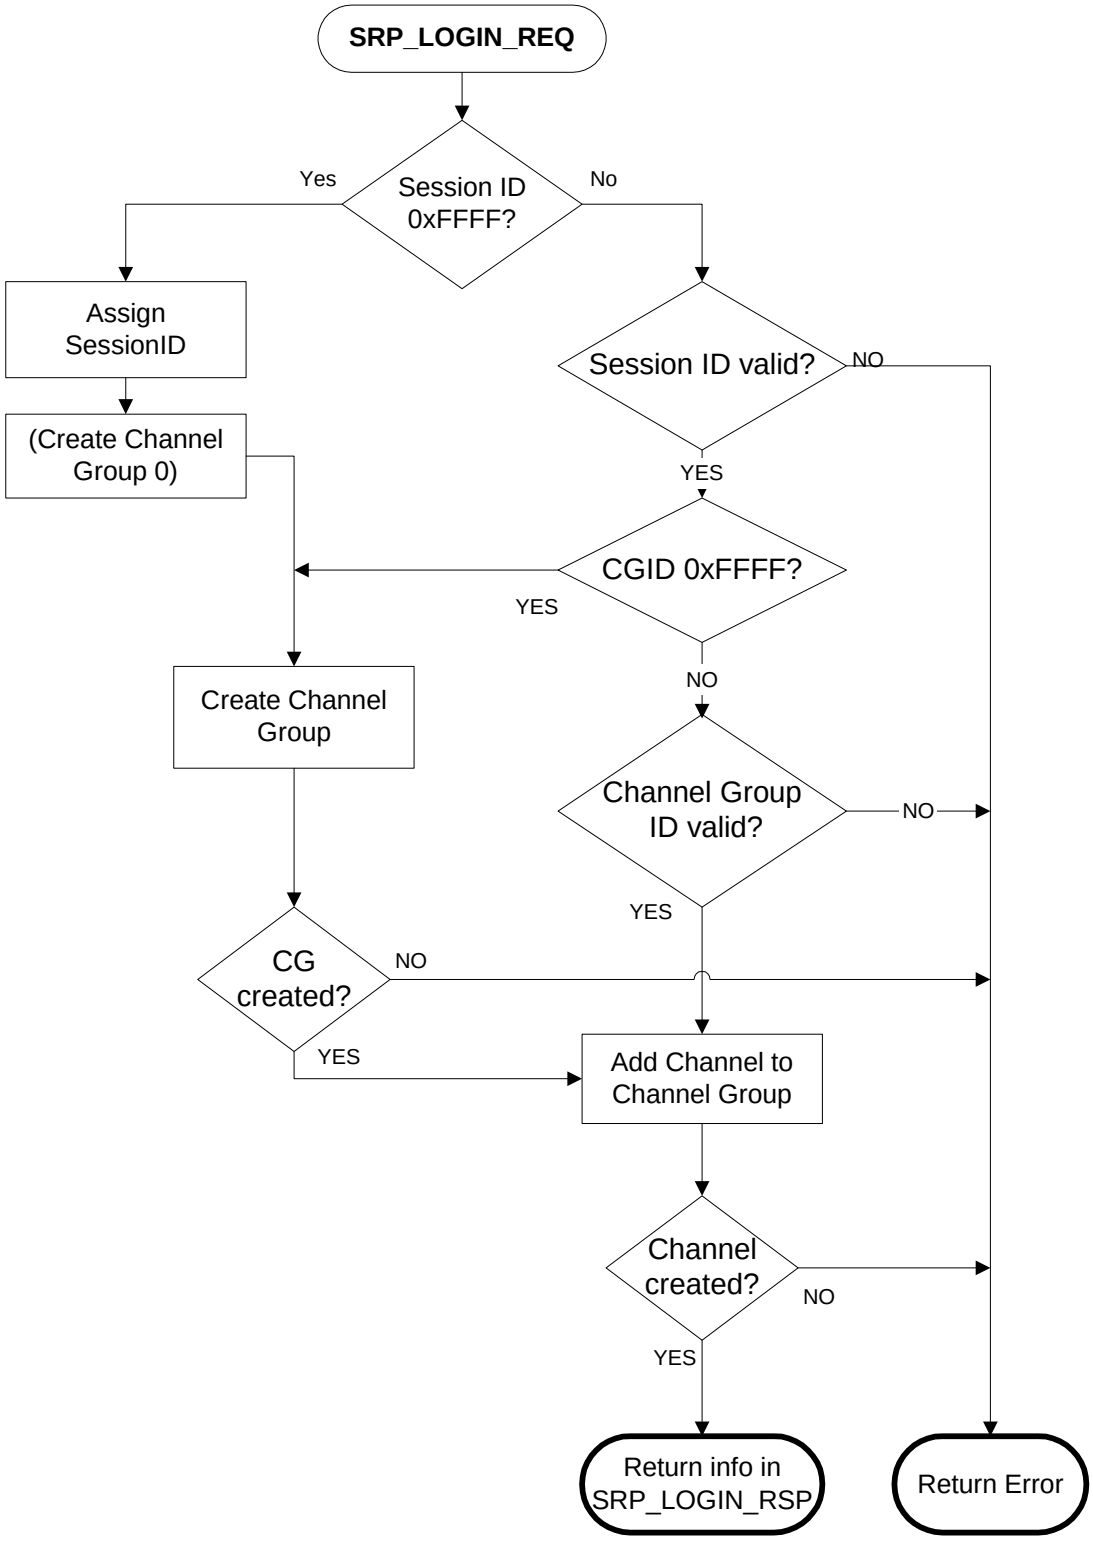

Figure 2 - SRP_LOGIN_REQ processing

2.2 SRP_LOGIN_RSP

SRP_LOGIN_RSP is returned if SRP_LOGIN_REQ is successful, SRP_LOGIN_REJ if not successful.

Table 3 - SRP_LOGIN_RSP additions

Bit / Byte	7	6	5	4	3	2	1	0
N	SESSION ID							
N+1								
N+2	CHANNEL GROUP ID							
N+3								

SESSION ID identifies the Session associated with the SRP_LOGIN_REQ. If the SRP_LOGIN_REQ was to create a new session, SESSION ID is represents the value assigned to the new session.

CHANNEL GROUP ID identifies the Channel Group to which the new channel was added.

2.3 SRP_LOGIN_REJ

SRP_LOGIN_REJ is returned when SRP_LOGIN_REQ is not successful.

Table 4 - SRP_LOGIN_REJ

Bit / Byte	7	6	5	4	3	2	1	0
N	REJECTION REASON							
N+1	RESERVED							
N+2	RESERVED							
N+3	REERVED							

REJECTION Reason codes:

0x01	INVALID SESSION ID
0x02	INVALID CHANNEL GROUP ID
0x03	UNABLE TO CREATE SESSION
0x04	UNABLE TO CREATE CHANNEL GROUP
0x05	UNABLE TO ADD CHANNEL TO CHANNEL GROUP
0x06	UNABLE TO CREATE CHANNEL

NOTE 6 - These may overlap with some of the IB-specific codes CM:Reject codes, such as 'No QP Available'. Policy?

2.4 SRP_CMD IU

Add Control Sequence Number

Table 5 - SRP_CMD addition

Bit / Byte	7	6	5	4	3	2	1	0
N	CONTROL SEQUENCE NUMBER							
N+1								
N+2	RESERVED							
N+3	REERVED							

A CONTROL SEQUENCE NUMBER of 0x0 indicates that the target may process the IU immediately. A target may process IUs with non-zero values only when the IU with the preceding CONTROL SEQUENCE NUMBER (modulo 2^{16}, excepting 0) has been processed.

NOTE 7 - 'Process' here means 'consider placing into task set' or 'begin executing task management command'. Is there a better word? 'Deliver to Task Manager?'

Add Channel Group Identifier to each data direction

(after CDB, before Buffer descriptors)

Table 6 - SRP_CMD information unit

Bit / Byte	7	6	5	4	3	2	1	0
N	DATA OUT CHANNEL GROUP ID							
N+1								
N+2	DATA IN CHANNEL GROUP ID							
N+3								

NOTE 8 - I'd prefer we combine all the data transfer info for each direction (Indirect bit, Descriptor count, Channel Group ID) into a contiguous block.

The DATA OUT CHANNEL GROUP ID field indicates the Channel Group to be used for transferring the data specified in the DATA OUT BUFFER DESCRIPTOR.

The DATA IN CHANNEL GROUP ID field indicates the Channel Group to be used for transferring the data specified in the DATA IN BUFFER DESCRIPTOR.

If a non-existent or invalid Channel Group ID is specified, the Status code INVALID CHANNEL GROUP ID shall be returned. Targets should check the Channel Group ID(s) before entering a command into the task set.

2.5 SRP_RSP

Modify RESPONSE DATA to allow returning data (e.g. missed CSN), similar to Sense Key/ Sense Code.

Table 7 - SRP_RSP RESPONSE_DATA change

Bit / Byte	7	6	5	4	3	2	1	0
N	RSP_CODE							
N+1	RESERVED							RESPONSE DATA VALID
N+2	RESPONSEDATA							
N+3								

New RESPONSE Codes

0x07 INVALID CHANNEL GROUP ID

0x08 INVALID SESSION ID

0x09 MISSING CONTROL SEQUENCE NUMBER

A RESPONSEDATA VALID bit of one indicates that the RESPONSEDATA field contains valid data, zero indicates the field does not contain valid data.

Targets must return valid data for the Response Code 0x09, MISSING CONTROL SEQUENCE NUMBER. For other Response Codes which are due to an invalid value in the IU, the target should return the invalid value.

3 Behavior

3.1 Channel Group Management

3.1.1 Creation

A Channel Group is created through an SRP_LOGIN_REQ IU with CHANNEL GROUP ID = 0xFF.

3.1.2 Adding Channels

A Channel is added to a Channel Group through the SRP_LOGIN_REQ IU. The first channel is added implicitly when the Channel Group is created.

3.1.3 Removing Channels

A Channel may be removed from a Channel Group through the transport-specific disconnection protocol.

3.1.4 Verifying Channels

As data transfer is target-driven, the target must initiate any operations intended to verify that Channels are operating correctly.

3.1.5 Errors

An error that causes a Queue Pair to be placed in the Error state affects only the associated channel, not the Channel Group.

When all the channels of a Channel Group have been torn down, the Channel Group ID becomes invalid.

> NOTE 9 - It seems desirable to give the initiator an opportunity to save the session even if all the channels of CG0 die. SBP has a timer for a similar case.

Targets should assign Channel Group IDs on a least-recently used basis.

3.2 Message sequencing

The SRP_CMD CONTROL SEQUENCE NUMBER allows an initiator to use multiple channels of the control Channel Group to send IUs to a target. The target uses the CONTROL SEQUENCE NUMBER to order the commands before processing.

As the channels of a Channel Group may have differing characteristics, an IU may arrive before another IU with a lower CSN arrives. In Figure 3, the IUs with CSNs 6,8, and 9 have arrived, but the IUs with CSNs 5 and 7 have not. IU-6 shall not be forwarded to the task manager until IU-5 has been forwarded. The size of a target's marshalling area determines how great a CSN skew can be accomodated.

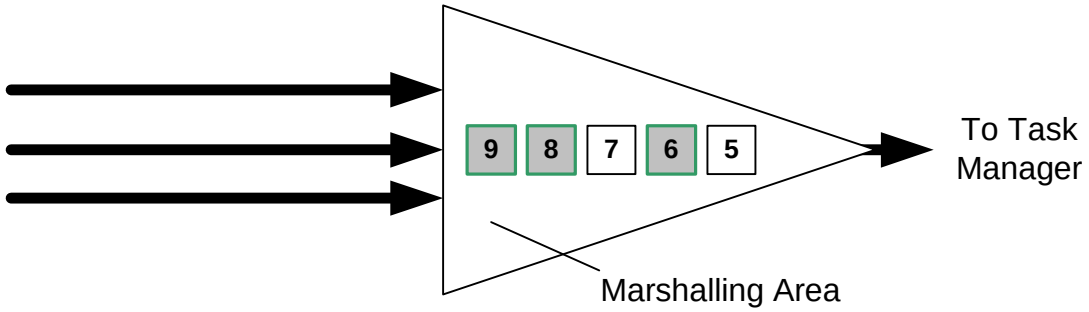

Figure 3 - Control Channel Group message sequencing