

To: T10 Technical Committee
From: Cris Simpson, Intel Corp. *cris.simpson@intel.com*
Date: 15 June 2001
Subject: SRP MultiChannel Proposal

Revision History

Revision 0: First Revision

Revision 1: Restricted CG0 to a single channel, eliminating sequencing issues.

Removed SRP_LOGIN_REJ, now 10-171r1

Related Documents

T10/SRP-r05 – SCSI over RDMA protocol revision 5

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.

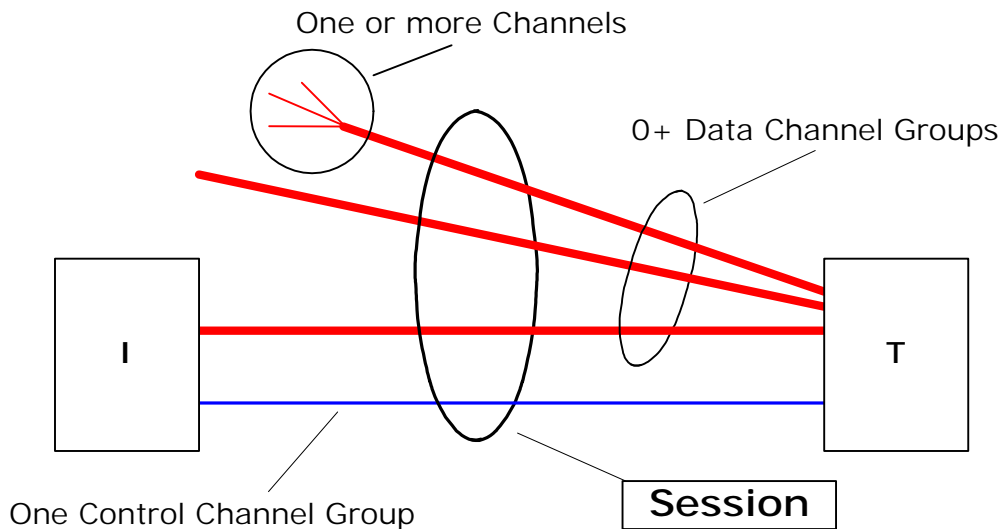


Figure 1 - SRP Session model

3 Definitions, symbols, abbreviations, and conventions

3.1 Definitions

3.1.1 channel group: A set of channels having identical memory access rights.

3.1.2 channel group identifier: A numeric value that identifies a channel group within a session. With respect to a target port, a channel group identifier is unique per initiator identifier. A target port may assign the same channel group identifier to different initiator ports.

3.1.3 control channel group: The single channel group of a session that carries control messages (i.e., requests and responses) and may carry RDMA data traffic. The control channel group always has the channel group ID of zero.

3.1.4 data channel group: A channel group that carries only RDMA data traffic.

3.1.5 session: One control channel group and zero or more data channel groups with a common session identifier.

3.1.6 session identifier: A numeric value that identifies a session within a target.

Changes

5.2 SRP_LOGIN_REQ

Table 1 - 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								

5.2.1 SESSION ID

The SESSION ID field contains the session identifier, and may be used to signal the initiator's request to create a new session.

If the SESSION ID field contains the value 0xFFFF, the target shall, if so able, establish a new session. The CHANNEL GROUP ID field shall be set to zero. Successful completion of this request shall result in a session consisting of a single-channel control channel group having a CHANNEL GROUP ID of zero.

A valid SESSION ID indicates that this SRP_LOGIN_REQ is requesting that a channel be added to the session. If the SESSION ID field contains an invalid session identifier, the target shall return an SRP_LOGIN_REJ IU with the REJECTION REASON field set to INVALID SESSION ID SPECIFIED.

5.2.2 Channel Group ID

The CHANNEL GROUP ID field contains the channel group identifier and may be used to signal the initiator's request for a new channel group.

If the CHANNEL GROUP ID field contains the value 0xFFFF, the target shall establish a new channel group with one channel. The channel group identifier shall be returned in the CHANNEL GROUP ID field of the SRP_LOGIN_RSP IU.

If CHANNEL GROUP ID contains the identifier of a valid control group, the target, if so able, shall establish a new channel as a member of the specified Channel Group. If the CHANNEL GROUP ID field contains an invalid session identifier, the target shall return an SRP_LOGIN_REJ IU with the REJECTION REASON field set to INVALID CHANNEL GROUP ID SPECIFIED.

The channel group having a CHANNEL GROUP ID of zero is limited to a single channel. If a SRP_LOGIN_REQ is received with the CHANNEL GROUP ID equal to zero, the target shall return an SRP_LOGIN_REJ IU with the REJECTION REASON field set to ADDITIONAL CHANNELS NOT ALLOWED ON THIS CHANNEL GROUP.

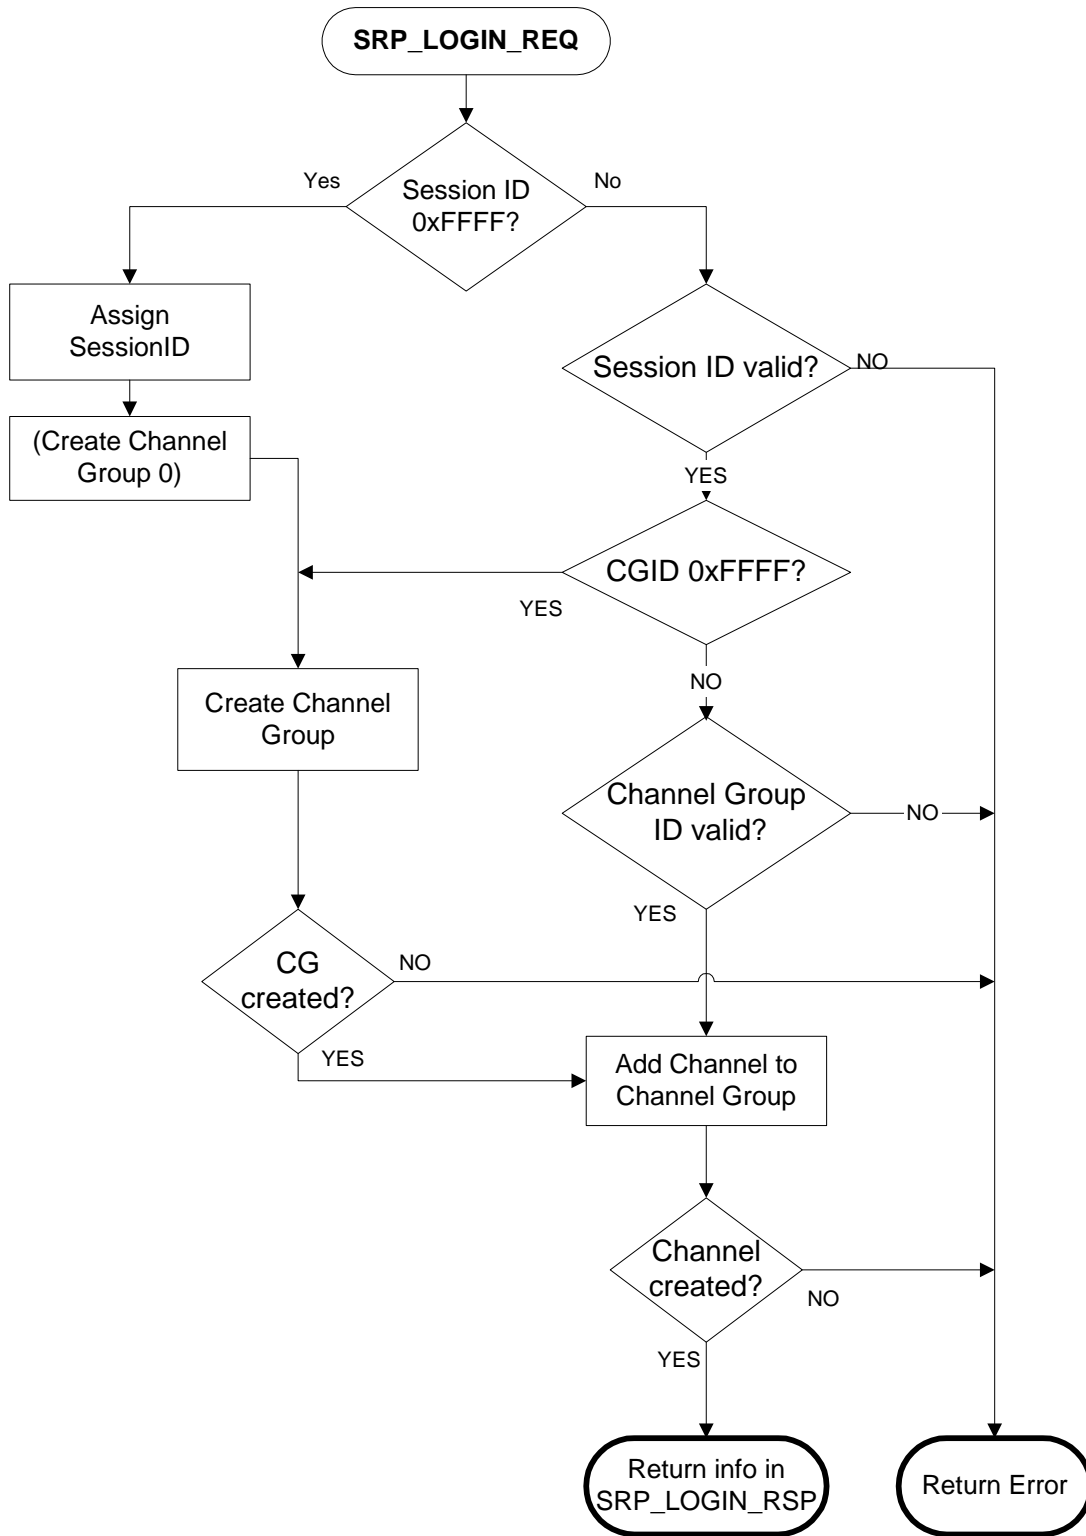


Figure 2 - SRP_LOGIN_REQ processing

5.3 SRP_LOGIN_RSP

The target shall return SRP_LOGIN_RSP IU if SRP_LOGIN_REQ is successful, SRP_LOGIN_REJ if not successful.

Table 2 - 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								

The SESSION ID field identifies the session associated with the SRP_LOGIN_REQ. If the SRP_LOGIN_REQ requested a new session, the SESSION ID field contains the session identifier assigned to the new session.

The CHANNEL GROUP ID field contains the channel group identifier of the of the newly-added channel.

5.4 SRP_CMD IU

Add four reserved bytes for future Control Sequence Number

Table 3 - SRP_CMD addition

Bit Byte	7	6	5	4	3	2	1	0
N	RESERVED							
N+1								
N+2								
N+3								

Add Channel Group Identifier to each data direction

(after CDB, before Buffer descriptors)

Table 4 - SRP_CMD information unit

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

NOTE 6 - 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 SPECIFIED shall be returned. Targets should check the CHANNEL GROUP ID(s) before entering a command into the task set.

5.5 SRP_RSP

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

Table 5 - 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	RESERVED							
N+3	RESERVED							
N+4	RESERVED							
N+5	RESPONSEDATA							
N+6	RESPONSEDATA							
N+7	RESPONSEDATA							

TABLE 6 - NEW RESPONSE CODES

Code	NAME	Description
0x07	INVALID CHANNEL GROUP ID SPECIFIED	
0x08	INVALID SESSION ID SPECIFIED	

6 Behavior

6.1 Channel Group Management

6.1.1 Creation

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

6.1.2 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.

6.1.3 Using Channel Groups

Control messages shall be sent over the Control Channel Group (i.e., CG 0) only. Initiators may specify CG0 for RDMA data transfer, and if CG0 is the only channel group, must specify it for commands that transfer data.

6.1.4 Removing Channels

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

6.1.5 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 7 - 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.