Document: T10/02-064r4 Date: 15 February 2002
To: T10 Committee Membership
From: Edward A. Gardner, Ophidian Designs
Subject: SRP normal and solicited message reception (comment OD006)

Revision 0 of this document proposed adding a bit to every SRP information unit, as suggested by the text of comment OD006. Spelling out the detailed changes made it clear (to both the author and the 17 January 2002 SRP working group) that that was not the best approach. Revision 1 of this proposal limits changes to only the relevant information units.

Revision 2 reflects comments from the January 25, 2002 teleconference.

Revision 3 reflects an editorial change requested at the February 1, 2002 teleconference.

Revision 4 changes the position of the bits added to SRP_LOGIN_REQ and SRP_LOGIN_RSP, as requested by Cris Simpson and discussed at the March 8, 2002 teleconference. That change is shown in blue.

A related editorial issue was discovered while preparing this proposal. SRP revision 10 does not consistently describe how SRP information units are transferred using an RDMA communication service. In some cases the description is inconsistent, in others the description is missing. This proposal also addresses those editorial inconsistencies, rather than making them a separate proposal.

This proposal has four parts. Part 1 discusses how an initiator discovers whether a target supports solicited events. Part 2 discusses commands and task management functions. Part 3 discusses target requests. Part 4 discusses the editorial corrections described in the previous paragraph.

Part 1 Discovery

A target port is responsible for determining whether it is capable of supporting solicited events (e.g. whether it's TCA supports solicited events). The target reports that to the initiator during login. If a target does not support solicited events, an initiator may still request them, but the target ignores the request.

Add the following bit definition to subclause 6.3, SRP_LOGIN_RSP:

Table 5 -

Bit / Byte	7	6	5	4	3	2	1	0
26	Reserved			SOLNTSUP	Reserved		MULTI-CHANNEL RESULT	

The solicited notification supported bit (SOLNTSUP) indicates whether the SRP target port supports solicited message reception notification for messages sent from the SRP target port to an SRP initiator port (see 4.3). If the SOLNTSUP bit is one, the SRP target port supports solicited message reception notification. If the SOLNTSUP bit is zero, the SRP target port only supports normal message reception notification.

Part 2 Commands and Task Management Requests

The main purpose of this proposal is to allow an SRP initiator port to specify whether a command's response should use normal or solicited message reception notification (i.e. whether the Solicited Event flag is set in the InfiniBand message containing the response). One can easily envision wanting different notifications depending on whether a command completes successfully or encounters an exception. The simple and flexible solution is two bits, one controlling notification for successful completion and the other for unsuccessful completion. Successful completion means return of GOOD status, unsuccessful means any other status.

Task management requests are very similar to commands and share the same response. Therefore they should share the same normal or solicited notification controls. It's arguable whether any initiator would care about successful vs. unsuccessful completion of a task management request. But since they have so much in common with commands, and a response code indicating success or otherwise, it is simpler to exactly parallel commands.

Add the following bit definitions to subclause 6.7, SRP_TSK_MGMT, and to subclause 6.8, SRP_CMD:

Table 6 -

Bit / Byte	7	6	5	4	3	2	1	0
0	TYPE (01h)							
1	Reserved					UCSOLNT	SCSOLNT	Reserved
2	Reserved							
• • •								
7								

The unsuccessful completion solicited notification bit (UCSOLNT) specifies whether an SRP_RSP response reporting unsuccessful completion of the task management request should use normal or solicited message reception notification. This bit shall be set to one to request solicited notification, or set to zero to request normal notification. See 6.9.

The successful completion solicited notification bit (SCSOLNT) specifies whether an SRP_RSP response reporting successful completion of the task management request should use normal or solicited message reception notification. This bit shall be set to one to request solicited notification, or set to zero to request normal notification. See 6.9.

Add the following bit definition to subclause 6.9, SRP_RSP:

Table 7 -

Bit / Byte	7	6	5	4	3	2	1	0
0	TYPE (C1h)							
1	Reserved							SOLNT
2	Reserved							
3								

The solicited notification (SOLNT) bit indicates whether the SRP initiator port specified normal or solicited message reception notification for this response. If the STATUS field is non-zero or if the RSP_CODE field is present and non-zero, then the SOLNT bit shall contain the value that was specified in the UCSOLNT bit of the corresponding SRP_CMD or SRP_TSK_MGMT request. Otherwise the SOLNT bit shall contain the value that was specified in the SCSOLNT bit of the corresponding SRP_CMD or SRP_TSK_MGMT request.

If the solicited notification (SOLNT) bit is one and the SRP target port supports solicited message reception notification (see 6.3), the SRP target port shall send the SRP_RSP response with solicited message reception notification (see 4.3). Otherwise the SRP target port should send the SRP_RSP response with normal message reception notification. An SRP initiator port shall not validate the SOLNT bit against whether an SRP_RSP response was actually received with normal or solicited message reception notification.

Part 3 Other Target to Initiator Messages

While the primary benefit derives from controlling normal vs. solicited notification for SRP_RSP, we also need to specify the behavior of other messages sent from the target to the initiator. There are three of these, SRP_T_LOGOUT, SRP_CRED_REQ and SRP_AER_REQ. A flexible solution is to specify the notification type for each of these during login.

Add the following bit definitions to subclause 6.2, SRP_LOGIN_REQ:

Table 8 -

Bit / Byte	7	6	5	4	3	2	1	0
26	Reserved	AESOLNT	CRSOLNT	LOSOLNT	Reserved		MULTI-CHANNEL ACTION	

The asynchronous event solicited notification bit (AESOLNT) specifies whether an SRP_AER_REQ request should use normal or solicited message reception notification. This bit shall be set to one to request solicited notification, or set to zero to request normal notification. See 6.12.

The credit request solicited notification bit (CRSOLNT) specifies whether an SRP_CRED_REQ request should use normal or solicited message reception notification. This bit shall be set to one to request solicited notification, or set to zero to request normal notification. See 6.10.

The logout solicited notification bit (LOSOLNT) specifies whether an SRP_T_LOGOUT request should use normal or solicited message reception notification. This bit shall be set to one to request solicited notification, or set to zero to request normal notification. See 6.6.

Add the following bit definition to subclause 6.6, SRP_T_LOGOUT, subclause 6.10, SRP_CRED_REQ, and subclause 6.12, SRP_AER_REQ:

Table 9 -

Bit / Byte	7	6	5	4	3	2	1	0
0	TYPE (C1h)							
1	Reserved							SOLNT
2	Reserved							
3								

The solicited notification (SOLNT) bit indicates whether the SRP initiator port specified normal or solicited message reception notification for SRP_T_LOGOUT requests during login (see 6.2). The SOLNT bit shall contain the value that was specified in the LOSOLNT bit of the SRP_LOGIN_REQ request.

If the solicited notification (SOLNT) bit is one and the SRP target port supports solicited message reception notification (see 6.3), the SRP target port shall send the SRP_T_LOGOUT response with solicited message reception notification (see 4.3). Otherwise the SRP target port should send the SRP_T_LOGOUT response with normal message reception notification. An SRP initiator port shall not validate the SOLNT bit against whether an SRP_RSP response was actually received with normal or solicited message reception notification.

Part 4 Editorial Clarifications

This part provides consistent wording for each SRP information unit describing how that information unit is transported. The affected text is the first paragraph of each information unit subclause. These are purely editorial changes, there are no technical changes.

Subclause 6.2 SRP_LOGIN_REQ request:

An SRP_LOGIN_REQ request (see table 9) conveys SRP protocol login parameters from an SRP initiator port to an SRP target port. ~~The~~ An SRP_LOGIN_REQ request shall ~~only~~ be sent as login data during RDMA channel establishment (see 4.2).

Subclause 6.3 SRP_LOGIN_RSP response:

An SRP_LOGIN_RSP response (see table 11) indicates successful RDMA channel establishment and conveys SRP protocol login parameters from an SRP target port to an SRP initiator port. ~~The~~ An SRP_LOGIN_RSP response shall ~~only~~ be sent as accept data during ~~to indicate successful~~ RDMA channel establishment (see 4.2).

Subclause 6.4 SRP_LOGIN_REJ response:

An SRP_LOGIN_REJ response (see table 13) indicates that ~~is sent by a SRP target port to notify the SRP initiator port that~~ an RDMA channel could not be established. An SRP_LOGIN_RSP response shall be sent as reject data (see 4.2).

Subclause 6.5 SRP_I_LOGOUT request:

An SRP_I_LOGOUT request (see table 15) is sent by an SRP initiator port to notify the SRP target port that the SRP initiator port is disconnecting the RDMA channel. An SRP_I_LOGOUT request may also be used to notify the SRP target port that an RDMA channel has failed, rendering it non-operational. An SRP_I_LOGOUT request shall be sent as a 16 byte message with normal message reception notification (see 4.3).

Subclause 6.6 SRP_T_LOGOUT request:

An SRP_T_LOGOUT request (see table 16) is sent by a SRP target port to notify the SRP initiator port that the SRP target port is disconnecting the RDMA channel. An SRP_T_LOGOUT request may also be used to notify the SRP initiator port that an RDMA channel has failed, rendering it non-operational. An SRP_T_LOGOUT request shall be sent as a 16 byte message.

Subclause 6.7 SRP_TSK_MGMT request:

An SRP_TSK_MGMT request conveys a SCSI task management request (table 18). An SRP_TSK_MGMT request shall be sent as a 48 byte message with normal message reception notification (see 4.3).

Subclause 6.8 SRP_CMD request:

An SRP_CMD request conveys a SCSI command (see table 20). An SRP_CMD request shall be sent as a message whose length is 48 bytes plus the lengths of the ADDITIONAL CDB, DATA-OUT BUFFER DESCRIPTOR and DATA-IN BUFFER DESCRIPTOR fields. An SRP_CMD request shall be sent with normal message reception notification (see 4.3).

Subclause 6.10 SRP_CRED_REQ request:

An SRP target port may use SRP_CRED_REQ requests (see table 25) to adjust an SRP initiator port's REQUEST LIMIT value (see 5.3). All SRP initiator ports shall support receiving SRP_CRED_REQ requests. An SRP_CRED_REQ request~~s~~ shall be sent as a 16 byte message.

Subclause 6.11 SRP_CRED_RSP response:

An SRP_CRED_RSP response (see table 26) is the response to an SRP_CRED_REQ request (see 6.10) received by an SRP initiator port. All SRP initiator ports shall support generating SRP_CRED_RSP responses. SRP_CRED_RSP responses shall be sent as a 16 byte message with normal message reception notification (see 4.3).

Subclause 6.12 SRP_AER_REQ request:

An SRP_AER_REQ request (see table 27) conveys a SRP target port request to report an asynchronous event. An SRP_AER_REQ request~~s~~ shall be sent as the minimum length message capable of carrying the fields. ~~All~~ SRP initiator ports shall support receiving SRP_AER_REQ requests and ~~all~~ SRP target ports shall support generating SRP_AER_REQ requests.

Subclause 6.13 SRP_AER_RSP response:

An SRP_AER_RSP response (see table 28) conveys an SRP initiator port's SRP response to an SRP_AER_REQ request (see 6.12). An SRP_AER_RSP response shall be sent as a 16 byte message with normal message reception notification (see 4.3).