

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 25 February 2004
Subject: 04-073r0 SAM-3 SAS-1.1 SRP-2 Unit attention on transport layer frame discard

Revision history

Revision 0 (25 February 2004) First revision

Related documents

sam3r11 - SCSI Architecture Model - 3 revision 11
spc3r17 - SCSI Primary Commands - 3 revision 17
sas1r03 - Serial Attached SCSI 1.1 revision 3
srp2r00a - SCSI RDMA Protocol - 2 revision 00a

Overview

SAS ensures that commands are delivered in order at the link layer by interlocking COMMAND frame transmission. The sender must receive an ACK or NAK before proceeding. If the sender receives a NAK or detects an ACK/NAK timeout after sending a COMMAND frame, it can choose not to send the next command until it deals with the problem. This covers CRC errors and frames that are far too short or too long.

SCSI's ORDERED task attribute, autocontingent allegiance (ACA) feature, and unit attention interlock feature can be used to ensure that commands received correctly (i.e., ACKed by the link layer) are processed in order by the device server:

- a) Commands with ORDERED task attributes are processed in the order received;
- b) If a command fails and returns CHECK CONDITION status, ACA blocks all subsequent commands from being processed until the failure is cleaned up; and
- c) If a unit attention occurs (both for traditional unit attention reasons and also optionally whenever the logical unit returns BUSY, TASK SET FULL, or RESERVATION CONFLICT), the unit attention interlock blocks all subsequent commands from being processed until a REQUEST SENSE command is run.

However, SAS leaves one hole that is not covered by the link layer NAK. In certain circumstances the transport layer of the target port either:

- a) discards frames that were ACKed by the link layer; or
- b) returns a RESPONSE frame with RESPONSE_DATA indicating an error.

Reasons for discards include (see 9.2.5.1 SSP target port error handling and 9.2.6.3.2 ST_TFR state):

- a) COMMAND frame too short to contain a LUN field;
- b) COMMAND frame too short to contain a CDB;
- c) COMMAND frame with an ADDITIONAL CDB LENGTH field that doesn't match the frame size;
- d) COMMAND frame or TASK frame with TARGET PORT TRANSFER TAG field not set to FFFFh;
- e) TASK frame too short;
- f) TASK frame with tag already in use;
- g) TASK frame with unknown logical unit number;
- h) DATA frame with unknown tag (discard);
- i) DATA frame for which there is no XFER_RDY (discard);
- j) non-DATA frame arrives while ACK/NAK not balanced (discard);
- k) hashed source or destination SAS addresses don't match the connection (discard);
- l) XFER_RDY frame (discard);
- m) unsupported frame type (discard).

In these cases, the SCSI application layer doesn't realize anything has happened. If the application client has queued several ORDERED commands in an HBA and the first one gets a RESPONSE frame in reply (e.g. CDB length problem), the HBA could already have sent the next commands (after receiving the ACK for the first). If the first one is discarded (e.g. hashed address problem), the HBA does not know about this; subsequent commands are processed as if nothing happened.

To fill this hole, the device server should be allowed to create a unit attention if it discards a frame or returns a RESPONSE frame for any of these reasons. These cases all indicate that command ordering might be violated. This can be controlled with a bit to the Protocol-Specific Logical Unit mode page. It is optional. Software would probably choose to enable it only when using the ORDERED task attribute with the unit attention interlock feature also enabled.

If multiple logical units are present and the logical unit number for the bad frame is indeterminate, all logical units behind the target port need to create the unit attention just in case the application client had intended to send them that frame. In SAS, this includes hashed address mismatches (since the logical unit number in the frame is based on the target port, if the target port is suspect, so must be the logical unit number).

Other transport protocols might have this hole too.

FCP-3 has a FCP_CMND fields invalid (02h) response code to reject a frame in the same manner as SAS. However, if command references numbers (CRN) are being used (and anyone this concerned about ordering would be doing so), a CRN gap would be created for the rejected frame and the hole would be detected before subsequent commands are processed. The unit attention is *not* needed for FCP-3.

SRP-2 has a REQUEST FIELDS INVALID (02h) response code to reject a frame in the same manner as SAS. SRP-2 assumes it is running on a reliable, ordered (within an interconnect port nexus) network, and link layer errors result in the interconnect nexus being torn down. There are no sequence numbers to guarantee the transport layer didn't miss a frame with a higher level error, though. The unit attention is proposed for SRP-2.

The proposed additional sense code has a generic name so it can be used by any transport protocol.

Editor's Note 1: Commands with bad LUNs result in CHECK CONDITION/ABORTED COMMAND/LOGICAL UNIT NOT SUPPORTED per SAM-3. Which logical unit is assumed to return that and possibly enter ACA mode is not currently defined and needs to be discussed.

Suggested changes to SAM-3

Editor's Note 2: This section is being changed by 03-376 (Invalid LUN on Request Sense command), so no specific changes are included in this revision of the proposal

5.9.4 Incorrect logical unit selection

The SCSI target device's response to an incorrect logical unit number is described in this subclause.

The logical unit number may be incorrect because:

- a) The SCSI target device does not support the logical unit (e.g., some SCSI target devices support only one peripheral device).

In response to any other command except REQUEST SENSE and INQUIRY, the SCSI target device shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED;

- b) The SCSI target device supports the logical unit, but the peripheral device is not currently attached to the SCSI target device.

In response to an INQUIRY command the SCSI target device shall return the INQUIRY data with the peripheral qualifier set to the value required in SPC-2. In response to a REQUEST SENSE command, the SCSI target device shall return GOOD status and parameter data that contains sense data. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED.

In response to any other command except REQUEST SENSE and INQUIRY, the SCSI target device shall terminate the command with CHECK CONDITION status. The sense key shall be set to

ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED;

- c) The SCSI target device supports the logical unit and the peripheral device is attached, but not operational.

In response to an INQUIRY command the SCSI target device shall return the INQUIRY data with the peripheral qualifier set to the value required in SPC-2. In response to REQUEST SENSE, the SCSI target device shall return GOOD status and parameter data that contains sense data appropriate to the condition

that is making the logical unit not operational.

The SCSI target device's response to any command other than INQUIRY and REQUEST SENSE is vendor specific; or

- d) The SCSI target device supports the logical unit but is incapable of determining if the peripheral device is attached or is not operational when it is not ready.

In response to an INQUIRY command the SCSI target device shall return the INQUIRY data with the peripheral qualifier set to the value specified in SPC-2. In response to a REQUEST SENSE command the SCSI target device shall return GOOD status and parameter data that contains sense data with a sense key of NO SENSE.

The SCSI target device's response to any other command is vendor specific.

Suggested changes to SPC-3

Add a new additional sense code next to the PREVIOUS BUSY/TASK SET FULL/RESERVATION CONFLICT additional sense codes used by the unit attention interlock feature:

2Ch/0Bh TRANSPORT PROTOCOL DISCARDED FRAME

Suggested changes to SAS-1.1

10.2.6.3 Protocol-Specific Logical Unit mode page

10.2.6.3.1 Protocol-Specific Logical Unit mode page overview

The Protocol-Specific Logical Unit mode page (see SPC-3) contains parameters that affect SSP target port operation on behalf of the logical unit. If the mode page is implemented, the mode page policy shall be shared (i.e., there is one copy of the mode page shared by all SSP initiator ports)(see SPC-3).

If a SAS target device has multiple SSP target ports, changes in the short page parameters for one SSP target port should not affect other SSP target ports.

Table 1 defines the subpages of this mode page.

Table 1 — Protocol-Specific Logical Unit mode page subpages

Subpage	Description	Reference
Short page	Short format	10.2.6.3.2
Long page 00h	Not allowed	
Long page E0h - FEh	Vendor specific	
Long page FFh	Return all subpages for the Protocol-Specific Logical Unit mode page	SPC-3
All others	Reserved	

10.2.6.3.2 Protocol-Specific Logical Unit mode page - short format

Parameters in this page shall affect all phys in the SSP target port, and may affect all SSP target ports in the SAS target device.

Table 132 defines the format of the page for SAS SSP.

Table 2 — Protocol-Specific Logical Unit mode page for SAS SSP - short format

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (18h)					
1	PAGE LENGTH (06h)							
2	Reserved		TRANSPORT LAYER UNIT ATTENTION	TRANSPORT LAYER RETRIES	PROTOCOL IDENTIFIER (6h)			
3	Reserved							
4	Reserved							
7	Reserved							

The PARAMETERS SAVEABLE (PS) bit is defined in SPC-3.

The SPF field shall be set to zero for access to the short format mode page.

The PAGE CODE field shall be set to 18h.

The PAGE LENGTH field shall be set to 06h.

The PROTOCOL IDENTIFIER field shall be set to 6h indicating this is a SAS SSP specific mode page.

[A TRANSPORT LAYER UNIT ATTENTION bit set to one specifies that the logical unit shall create a unit attention condition with an additional sense code set to TRANSPORT PROTOCOL DISCARDED FRAME after:](#)

- a) [the transport layer of the target port \(see 9.2.5 and 9.2.6.3.2\) discards a frame that was addressed to the logical unit;](#)
- b) [the transport layer of the target port returns a RESPONSE frame with the DATAPRES field set to RESPONSE DATA and the RESPONSE CODE field set to INVALID FRAME after receiving a frame addressed to the logical unit;](#)
- c) [the transport layer of the target port returns a RESPONSE frame with the DATAPRES field set to RESPONSE DATA and the RESPONSE CODE field set to INVALID LOGICAL UNIT NUMBER; or](#)
- d) [any other logical unit in the SCSI target device returns a CHECK CONDITION status with a sense key set to ABORTED COMMAND and an additional sense code to LOGICAL UNIT NOT SUPPORTED.](#)

[A TRANSPORT LAYER UNIT ATTENTION bit set to zero specifies that the logical unit shall not create a unit attention condition for those reasons.](#)

[Editor's Note 3: Item d\) might belong in SAM-3.](#)

A TRANSPORT LAYER RETRIES bit set to one specifies that the target port shall support transport layer retries for XFER_RDY and DATA frames for the logical unit as described in 9.2.4. A TRANSPORT LAYER RETRIES bit set to zero specifies that transport layer retries shall not be used.

[Editor's Note 4: Changes to the SAS transport layer state machine \(9.2.6.3.2\) and error-handling overview 9.2.5\) are also needed. Since 03-245 \(Terminate Data Transfer protocol service\) is modifying the state machine text, no specific changes for those sections are included in this revision of this proposal.](#)

Suggested changes to SRP-2**6.9 SRP_RSP response**

An SRP_RSP response (see table 22) conveys an SRP response to an SRP_TSK_MGMT request (see 6.7) or an SRP_CMD request (see 6.8) received by a SRP target port. SRP_RSP responses that contain neither RESPONSE DATA nor SENSE DATA shall be sent as a 36 byte message. SRP_RSP responses that contain either RESPONSE DATA or SENSE DATA shall be sent as the minimum length message containing those fields.

The RESPONSE DATA field (see table 23) contains information describing protocol failures detected during processing of an SRP request received by the SRP target port. The RESPONSE DATA field shall be present if the SRP target port detects any of the conditions described by a non-zero RSP_CODE value (see table 24).

...

The RSP_CODE field is defined in table 24.

...

If the RSP_CODE field is set to 02h (i.e., REQUEST FIELDS INVALID), the device server may establish a unit attention condition with an additional sense code set to TRANSPORT LAYER DISCARDED FRAME.

7.3 Protocol specific ~~LUN~~ Logical unit mode page

~~The Protocol Specific LUN page shall not be implemented by SRP target ports.~~

7.3.1 Protocol-Specific Logical Unit mode page overview

The Protocol-Specific Logical Unit mode page (see SPC-3) contains parameters that affect SRP-2 target port operation on behalf of the logical unit. If the mode page is implemented, the mode page policy shall be shared (i.e., there is one copy of the mode page shared by all SRP-2 initiator ports)(see SPC-3).

If a SCSI target device has multiple SRP-2 target ports, changes in the short page parameters for one SRP-2 target port should not affect other SRP-2 target ports.

Table 1 defines the subpages of this mode page.

Table 3 — Protocol-Specific Logical Unit mode page subpages

Subpage	Description	Reference
Short page	Short format	7.3.2
Long page 00h	Not allowed	
Long page E0h - FEh	Vendor specific	
Long page FFh	Return all subpages for the Protocol-Specific Logical Unit mode page	SPC-3
All others	Reserved	

7.3.2 Protocol-Specific Logical Unit mode page - short format

Parameters in this page may affect all SRP-2 target ports in the SCSI target device.

[Table 132](#) defines the format of the page for SRP-2.

Table 4 — Protocol-Specific Logical Unit mode page for SRP-2 - short format

Byte\Bit	7	6	5	4	3	2	1	0
<u>0</u>	<u>PS</u>	<u>SPF (0b)</u>	<u>PAGE CODE (18h)</u>					
<u>1</u>	<u>PAGE LENGTH (06h)</u>							
<u>2</u>	<u>Reserved</u>		<u>TRANSPORT LAYER UNIT ATTENTION</u>	<u>Reserved</u>	<u>PROTOCOL IDENTIFIER (4h)</u>			
<u>3</u>	<u>Reserved</u>							
<u>4</u>	<u>Reserved</u>							
<u>7</u>	<u>Reserved</u>							

[The PARAMETERS SAVEABLE \(PS\) bit is defined in SPC-3.](#)

[The SPF field shall be set to zero for access to the short format mode page.](#)

[The PAGE CODE field shall be set to 18h.](#)

[The PAGE LENGTH field shall be set to 06h.](#)

[The PROTOCOL IDENTIFIER field shall be set to 4h indicating this is a SRP-2 specific mode page.](#)

[A TRANSPORT LAYER UNIT ATTENTION bit set to one specifies that the logical unit shall create a unit attention condition with an additional sense code set to TRANSPORT PROTOCOL DISCARDED FRAME if the target port returns a SRP_RSP response with the RSPVALID field set to one and the RSP_CODE field set to REQUEST_FIELDS_INVALID after](#)

- a) [receiving an SRP_CMD or SRP_TSK_MGMT request addressed to the logical unit;](#)
- b) [receiving an SRP_TSK_MGMT request addressed to an unknown logical unit; or](#)
- c) [any other logical unit in the SCSI target device returns a CHECK_CONDITION status with a sense key set to ABORTED_COMMAND and an additional sense code to LOGICAL_UNIT_NOT_SUPPORTED.](#)

[A TRANSPORT LAYER UNIT ATTENTION bit set to zero specifies that the logical unit shall not create a unit attention condition for those reasons.](#)

Editor's Note 5: Item c) might belong in SAM-3.
