

To: T10 Technical Committee
 From: Chris Martin (chris.martin@hp.com) and Rob Elliott, HP (elliott@hp.com)
 Date: 26 February 2007
 Subject: 07-027r1 SAS-2 Enabling and disabling Transport Layer Retries

Revision history

Revision 0 (9 January 2007) First revision

Revision 0 (22 February 2007) Incorporated comments from January 2007 SAS protocol WG - simplified so the initiator just sets the new TLR CONTROL field to 01b or 11b blindly for all I_T_L nexuses unless it that specifically causes the frame to be rejected. Made the new field adjacent to the XFER_RDY RETRY DATA FRAMES bit rather than take it over. Removed VPD page and IDENTIFY address frame additions.

Related documents

sas2r06 - Serial Attached SCSI - 2 (SAS-2) revision 6

Overview

In SAS-1.1, transport layer retries are enabled or disabled via the Protocol-Specific Logical Unit mode page (18h/00h) TRANSPORT LAYER RETRIES bit (which has a mode page policy of shared or per-target port), controlled with the MODE SELECT command.

This has proven to be impractical for myriad reasons:

- a) If the mode page is shared, multiple initiators that don't agree on the setting interfere with each other. This is less of a problem if the mode page is per I_T nexus, but such implementations are rare.
- b) If the TRANSPORT LAYER RETRIES bit is set to one and saved, then the target exits hard reset with TLR enabled, before the initiator knows TLR is enabled. If it is not savable, then the target forgets the setting on a hard reset; the initiator has to constantly watch for unit attention conditions with an additional sense code of MODE PARAMETERS CHANGED (or an additional sense code for a reset condition that implies mode parameters changed) to realize that TLR has been turned off.
- c) In an initiator, the lower layers of software/firmware are the ones that need control TLR. However, sending a MODE SELECT command requires involvement by upper layers.
- d) MODE SENSE and MODE SELECT commands themselves may encounter errors and need to invoke TLR.
- e) MODE SENSE and MODE SELECT commands may take a long time to process
- f) An HBA may be confused if it receives an XFER_RDY with the RETRY DATA FRAMES bit set to one for a MODE SELECT command but it thinks TLR is disabled.

Transport layer retries means:

- a) Initiator: based on the incoming RETRY DATA FRAMES bit in the XFER_RDY frame, retry write DATA frames
- b) Target: based on the mode page, retry read DATA frames and XFER_RDY frames and set the RETRY DATA FRAMES bit to one in the XFER_RDY frame

Rather than having the target base its decision on the mode page, a TLR CONTROL field is proposed for the COMMAND frame header:

- a) If set to 01b, this enables TLR on the I_T_L_Q nexus (i.e., the command), overriding any mode page setting.
- b) If set to 10b, this disables TLR on the I_T_L_Q nexus, overriding any mode page setting.
- c) If set to 00b or 11b, the mode page is honored.

A SAS-2 initiator port would blindly set the TLR CONTROL field to 01b or 10b; the intent is not to actually change it per-I_T_L_Q nexus.

Because SAS does not mandate that reserved fields in the SSP frame header be checked or not checked, an initiator sending a COMMAND frame with the TLR CONTROL field set to non-zero to a SAS-1.1 target could get an error or be silently ignored. If a SAS-1.1 target ignores the field (which is also allowed), then the initiator just never receives an XFER_RDY with the RETRY DATA FRAMES bit set to one.

If a COMMAND frame is rejected with a non-zero value in its TLR CONTROL field, the initiator port sets the TLR CONTROL field to 00b on COMMAND frames for that particular I_T_L nexus and reverts to mode page control (or no TLR).

Table 1 shows the interactions between SAS-1.1 and SAS-2 ports.

Table 1 — Interactions between SAS-1.1 and SAS-2 ports

Initiator port	Target port ^a	Result
SAS-1.1 no TLR TLR CONTROL = 00b (use mode page)	SAS-1.1 no TLR	No TLR (mode page always has TLR disabled)
	SAS-2 no TLR	
	SAS-1.1 with TLR	Based on mode page; problem if mode page has TLR enabled.
	SAS-2 with TLR	Based on mode page; problem if mode page has TLR enabled
SAS-1.1 with TLR TLR CONTROL = 00b (use mode page)	SAS-1.1 no TLR	No TLR (mode page always has TLR disabled)
	SAS-2 no TLR	
	SAS-1.1 with TLR	Based on mode page; problem if mode page has TLR enabled and initiator not expecting it yet
	SAS-2 with TLR	Based on mode page; problem if mode page has TLR enabled and initiator not expecting it yet
SAS-2 no TLR TLR CONTROL = 10b (off)	SAS-1.1 no TLR permissive	No TLR (mode page always has TLR disabled)
	SAS-1.1 no TLR fussy	INVALID FRAME
	SAS-2 no TLR	No TLR
	SAS-1.1 with TLR permissive	Based on mode page; problem if saved mode page has TLR enabled
	SAS-1.1 with TLR fussy	INVALID FRAME
	SAS-2 with TLR	Based on mode page; problem if mode page has TLR enabled
SAS-2 with TLR TLR CONTROL = 01b (on)	SAS-1.1 no TLR permissive	No TLR (mode page always has TLR disabled)
	SAS-1.1 no TLR fussy	INVALID FRAME
	SAS-2 no TLR	No TLR (RETRY DATA FRAMES always 0)
	SAS-1.1 with TLR permissive	Based on mode page. Initiator always ready to receive TLR. RETRY DATA FRAMES specifies if it shall do TLR itself.
	SAS-1.1 with TLR fussy	INVALID FRAME
	SAS-2 with TLR	TLR enabled
^a "permissive" means it does not check reserved fields; "fussy" means it checks reserved fields		

Suggested changes to SAS-2

9.2.1 SSP frame format

Table 134 defines the SSP frame format.

Table 134 — SSP frame format

Byte\Bit	7	6	5	4	3	2	1	0
0	FRAME TYPE							
1	(MSB)	HASHED DESTINATION SAS ADDRESS						(LSB)
3								
4	Reserved							
5	(MSB)	HASHED SOURCE SAS ADDRESS						(LSB)
7								
8	Reserved							
9	Reserved							
10	Reserved			Reserved TLR CONTROL		RETRY DATA FRAMES	RETRANSMIT	CHANGING DATA POINTER
11	Reserved					NUMBER OF FILL BYTES		
12	Reserved							
13								
15	Reserved							
16	(MSB)	TAG						(LSB)
17								
18	(MSB)	TARGET PORT TRANSFER TAG						(LSB)
19								
20	(MSB)	DATA OFFSET						(LSB)
23								
24	INFORMATION UNIT							
m	(e.g., see table 136, table 138, table 140, table 141, or table 142)							
Fill bytes, if needed								
n - 3	(MSB)	CRC						(LSB)
n								

Table 135 defines the FRAME TYPE field, which defines the format of the INFORMATION UNIT field.

Table 135 — FRAME TYPE field

Code	Name of frame	Type of information unit	Originator	Information unit size (bytes)	Reference
01h	DATA frame (i.e., write DATA frame or read DATA frame)	Data information unit (i.e., write Data information unit or read Data information unit)	SSP initiator port or SSP target port	1 to 1 024	9.2.2.4
05h	XFER_RDY frame	Transfer Ready information unit	SSP target port	12	9.2.2.3
06h	COMMAND frame	Command information unit	SSP initiator port	28 to 284	9.2.2.1
07h	RESPONSE frame	Response information unit	SSP target port	24 to 1 024	9.2.2.5
16h	TASK frame	Task Management Function information unit	SSP initiator port	28	9.2.2.2
F0h - FFh	Vendor specific				
All others	Reserved				

The HASHED DESTINATION SAS ADDRESS field contains the hashed value of the destination SAS address (see 4.2.3). See 9.2.6.2.2 and 9.2.6.3.2 for transport layer requirements on checking this field.

The HASHED SOURCE SAS ADDRESS field contains the hashed value of the source SAS address (see 4.2.3). See 9.2.6.2.2 and 9.2.6.3.2 for transport layer requirements on checking this field.

If an initiator port supports transport layer retries, it shall set the TLR CONTROL field to 01b in each COMMAND frame it sends unless it has determined that the I T L nexus does not support the TLR CONTROL field.

If an initiator port does not support transport layer retries, it shall set the TLR CONTROL field to 10b in each COMMAND frame it sends unless it has determined that the I T L nexus does not support the TLR CONTROL field.

If the initiator port sends a COMMAND frame with the TLR CONTROL field set to 01b or 10b and receives a RESPONSE frame with the DATAPRES field set to RESPONSE_DATA and the RESPONSE CODE field set to 02h (i.e., INVALID FRAME), it shall set the TLR CONTROL field to 00b for subsequent COMMAND frames for that I T L nexus.

NOTE 1 - Initiator ports compliant with previous versions of this standard always set the TLR CONTROL field to 00b.

A TLR CONTROL field set to 00b or 11b in a COMMAND frame specifies that the SSP target port shall use the TRANSPORT LAYER RETRIES bit in the Protocol-Specific Logical Unit mode page (see 10.2.7.3) to enable or disable transport layer retries for this command:

- a) if the TRANSPORT LAYER RETRIES bit is set to one, the target port shall set the RETRY DATA FRAMES bit to one in any XFER_RDY frames that it transmits for this command; and
- b) if the TRANSPORT LAYER RETRIES bit is set to zero, the target port shall set the RETRY DATA FRAMES bit to zero in any XFER_RDY frames that it transmits for this command.

A TLR CONTROL field set to 01b in a COMMAND frame specifies that the SSP target port may enable transport layer retries for this command:

- a) if it enables transport layer retries, the target port shall set the RETRY DATA FRAMES bit to one in any XFER_RDY frames that it transmits for this command; and
- b) if it does not enable transport layer retries, the target port shall set the RETRY DATA FRAMES bit to zero in any XFER_RDY frames that it transmits for this command.

A TLR CONTROL field set to 10b in a COMMAND frame specifies that the SSP target port shall:

- a) disable transport layer retries for this command; and
- b) set the RETRY DATA FRAMES bit to zero in any XFER_RDY frames that it transmits for this command.

The TLR CONTROL field is reserved for frames other than COMMAND frames.

~~The RETRY DATA FRAMES bit is set to one for XFER_RDY frames under the conditions defined in 9.2.4 and shall be set to zero for all other frame types. When set to one this bit specifies that the SSP initiator port may retry write DATA frames that fail.~~

A target port sets the RETRY DATA FRAMES bit in an XFER_RDY frame based on the TLR CONTROL field received in the COMMAND frame for the command and the TRANSPORT LAYER RETRIES bit in the Protocol-Specific Logical Unit mode page (see 10.2.7.3).

A RETRY DATA FRAMES bit set to one in an XFER_RDY frame specifies that the SSP initiator port shall enable transport layer retries for write DATA transfers related to this XFER_RDY.

A RETRY DATA FRAMES bit set to zero in an XFER_RDY frame specifies that the SSP initiator port shall disable transport layer retries for write DATA transfers related to this XFER_RDY.

The RETRY DATA FRAMES bit is reserved for frames other than XFER_RDY frames.

The RETRANSMIT bit is set to one for TASK frames, RESPONSE frames, and XFER_RDY frames under the conditions defined in 9.2.4 and shall be set to zero for all other frame types. This bit specifies that the frame is a retransmission after the SSP port failed in its previous attempt to transmit the frame.

The CHANGING DATA POINTER bit is set to one for DATA frames under the conditions defined in 9.2.4 and shall be set to zero for all other frame types. When set to one this bit specifies that the frame is a retransmission after the SSP target port failed in its previous attempt to transmit the frame or a subsequent frame and the DATA OFFSET field of the frame may not be sequentially increased from that of the previous frame.

The NUMBER OF FILL BYTES field specifies the number of fill bytes between the INFORMATION UNIT field and the CRC field. The NUMBER OF FILL BYTES field shall be set to zero for all frame types except DATA frames as specified in 9.2.2.4 and RESPONSE frames as specified in 9.2.2.5 (i.e., all other frame types are already four-byte aligned).

The TAG field contains a value that allows the SSP initiator port to establish a context for commands and task management functions.

For COMMAND frames and TASK frames, the SSP initiator port shall set the TAG field to a value that is unique for the I_T nexus established by the connection (see 7.12). An SSP initiator port shall not reuse the same tag when transmitting COMMAND frames or TASK frames to different LUNs in the same SSP target port. An SSP initiator port may reuse a tag when transmitting frames to different SSP target ports. An SSP initiator port does not reuse a tag until it receives indication from the SSP target port that the tag is no longer in use (see 9.2.4, 9.2.5, and 10.2.2).

The TAG field in a COMMAND frame contains the task tag defined in SAM-4. The TAG field in a TASK frame does not correspond to a SAM-4 task tag, but corresponds to an SAM-4 association (see 10.2.1). The tag space used in the TAG fields is shared across COMMAND frames and TASK frames (e.g., if a tag is used for a COMMAND frame, it is not also used for a concurrent TASK frame).

For DATA, XFER_RDY, and RESPONSE frames, the SSP target port shall set the TAG field to the tag of the command or task management function to which the frame pertains.

The TARGET PORT TRANSFER TAG field provides an optional method for an SSP target port to establish the write data context when receiving a write DATA frame (i.e., determine the command to which the write data

corresponds). Unlike the TAG field, which was assigned by the SSP initiator port, the TARGET PORT TRANSFER TAG field in a write DATA frame contains a value assigned by the SSP target port that was delivered to the SSP initiator port in the XFER_RDY frame requesting the write data.

NOTE 2 - The TARGET PORT TRANSFER TAG field may be useful when the SSP target port has more than one XFER_RDY frame outstanding (i.e., the SSP target port has transmitted an XFER_RDY frame for each of two or more commands and has not yet received all the write data for them).

SSP target ports may set the TARGET PORT TRANSFER TAG field to any value when transmitting any SSP frame. SSP target ports that use this field should set the TARGET PORT TRANSFER TAG field in every XFER_RDY frame to a value that is unique for the L_Q portion of the I_T_L_Q nexus (i.e., that is unique for every XFER_RDY that is outstanding from the SSP target port).

SSP initiator ports shall set the TARGET PORT TRANSFER TAG field as follows:

- a) For each write DATA frame that is sent in response to an XFER_RDY frame, the SSP initiator port shall set the TARGET PORT TRANSFER TAG field to the value that was in the corresponding XFER_RDY frame;
- b) For each write DATA frame that is sent containing first burst data (see 9.2.2.4), the SSP initiator port shall set the TARGET PORT TRANSFER TAG field to FFFFh; and
- c) For frames other than write DATA frames, the SSP initiator port shall set the TARGET PORT TRANSFER TAG field to FFFFh.

For DATA frames, the DATA OFFSET field is described in 9.2.2.4. For all other frame types, the DATA OFFSET field shall be ignored.

The INFORMATION UNIT field contains the information unit, the format of which is defined by the FRAME TYPE field (see table 135). The maximum size of the INFORMATION UNIT field is 1 024 bytes, making the maximum size of the frame 1 052 bytes (1 024 bytes of data + 24 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the INFORMATION UNIT field so the CRC field is aligned on a four byte boundary. The number of fill bytes are specified by the NUMBER OF FILL BYTES field. The contents of the fill bytes are vendor specific.

The CRC field contains a CRC value (see 7.5) that is computed over the entire SSP frame prior to the CRC field including the fill bytes (i.e., all data dwords between the SOF and EOF). The CRC field is checked by the link layer (see 7.16), not the transport layer.

10.2.7.3.2 Protocol-Specific Logical Unit mode page - short format

The mode page policy (see SPC-4) for the Protocol-Specific Logical Unit mode page short format subpage shall be either shared or per target port. If a SAS target device has multiple SSP target ports, the mode page policy should be per target port. Parameters in this page shall affect all phys in the SSP target port if the mode page policy is per target port, and shall affect all SSP target ports in the SAS target device if the mode page policy is shared.

Table 136 defines the format of the page for SAS SSP.

Table 136 — 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 RETRIES	PROTOCOL IDENTIFIER (6h)			
3	Reserved							
7	Reserved							

The PARAMETERS SAVEABLE (PS) bit is defined in SPC-4.

The SUBPAGE FORMAT (SPF) bit 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 the number of bytes in the page after the PAGE LENGTH field (i.e., 06h).

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

A TRANSPORT LAYER RETRIES bit set to one specifies that, [for COMMAND frames received with the TLR CONTROL field set to 00b or 11b](#), the target port shall support transport layer retries for XFER_RDY and DATA frames for the logical unit as described in 9.2.4 (i.e., transport layer retries are enabled). A TRANSPORT LAYER RETRIES bit set to zero specifies that, [for COMMAND frames received with the TLR CONTROL field set to 00b or 11b](#), transport layer retries shall not be used (i.e., transport layer retries are disabled).

[NOTE 3 - The TRANSPORT LAYER RETRIES bit may become obsolete in a future version of this standard.](#)