Document:	T10/02-323r2	Date:	October 21, 2002
To:	T10 Committee Membership		
From:	Robert Sheffield (Robert.L.Sheffield@intel.com), Intel Corporatoin		
Subject:	SAS Data Corruption Problem		

Related Documents

- SAS-r02b
- T10/02-323r1 – SAS Data Corruption Problem [Ed Gardner, Basil Networks]

Overview of the Issue

As defined in SAS-r02b, there is a severe limitation in the detection and handling of frame transmission errors that will result in an unacceptably high frequency of corruption of data on media as compared to existing parallel-SCSI solutions. The frequency of data corruption results from a combination of the specified bit error rate (BER), the supported link rate, and the policy for handling frame errors detected during transmission through a structure of expanders. As currently defined, one might expect to see data corrupted on media at a frequency higher than once every minute.

BER and Link Rate

The bit-error rate is defined to be 10^{-12}, which says that some kind of transmission error will happen for every 100 Gbytes of data transferred. A SAS HBAs running 4 links full-duplex at 3 Gbps each, yields a common aggregate transfer rate of 2.4 Gbytes/sec. This means that such an HBA implemented according to the SAS working draft would *likely* experience a transmission error in one direction or the other once every 42 seconds. If actual implementations meet a much more stringent BER requirement such as 10^{-14}, transmission errors are still *likely* to occur once every 7 minutes.

Transmission Error Handling

A variety of error conditions, many related to transmission errors, result in a frame being discarded.

From SAS-r02b section 7.16.711:

> The frame (i.e., all the dwords between an SOF and EOF) shall be discarded if any of the following conditions are true:
>
> a) the number of bytes between the SOF and EOF is less than 28;
> b) the number of bytes after the SOF is greater than 1 052 bytes;
> c) the Rx Credit Status parameter from the SSP_RCM1:Rcv_Credit_Monitor state has an argument of Credit Exhausted;
> d) the SSP_R1:Receive state sent a DONE Received parameter.
>
> If consecutive SOF Received parameters are sent by the SSP_R1:Receive state without an intervening EOF Received parameter (i.e., SOF, data dwords, SOF, data dwords, and EOF instead of SOF, data dwords, EOF, SOF, data dwords, and EOF) then this state shall discard all dwords between the SOFs.

And from section 9.2.5.2:

> If a target port receives an XFER_RDY frame, it shall discard the frame.
>
> If a target port receives a frame with an unknown FRAME TYPE and an unknown TAG, it shall discard the frame.
>
> If a target port receives a COMMAND frame that is too short to contain a LUN field, it shall discard the frame.
>
> If a target port receives a DATA frame with an unknown TAG, it shall discard the frame.

Furthermore, there is nothing defining the disposition of a frame that incurs a CRC error except the condition that the recipient issues a NAK rather than an ACK. The connection remains active and subsequent data frames continue to arrive. There is no defined communication from the link

layer to the transport layer of the recipient of a frame with bad CRC that a frame is missing, and subsequent frames (now out of sequence) continue to be processed by the target as if they were in sequence.

This means there are several types of transmission errors that will cause frames to be dropped at the destination with no timely notification to the transport layer of the receiver.

In SPI similar consequences occur if REQ and ACK get out of sync, however many years of experience have proven that the probablity of REQ and ACK getting out of sync is so miniscule as to be negligible and for all practical purposes may be assumed to be zero. In contrast, the probability of dropped frames in SAS is clearly non-zero, and arguably quite frequent by comparison.

When SAS non-interlocked data frames are dropped, the recipient continues to receive subsequent data frames without knowledge that the data stream has been interrupted. This results in mis-identification of the data following the dropped data frame.

Consequences of a Dropped Data Frame to a Disk or Tape Target

Consider the situation where a transmission error occurs when an initiator is sending write data to a target. For the sake of argument assume that an SOF is corrupted, so that a frame is completely lost. However the result is substantially similar with other errors, e.g. coding violations or a CRC error.

Since data transfers are not interlocked, the initiator is not aware of the problem until the next time an interlocked frame is to be sent forcing the initiator to reconcile the number of ACKs recevied with the number of frames transmitted, thereby finding a mismatch. The target is not aware of the problem until after it times out the data's arrival or receives a TASK ABORT. Meanwhile the data frame in error is dropped from the middle of the data stream received at the target (even if a CRC error is detected, this is the behavior seen at the transport layer).

Disk and tape drives stream data from SCSI or Fibre Channel onto the media. Assuming SAS drives do the same, the corrupted (mis-identified) data stream (potentially very large quantities) will typically have already been written to the wrong location on the media by the time the error is discovered. This creates a substantial time window where a power failure, system crash or similar problem will leave the corrupted data on the media permanently. This is unacceptable for enterprise applications.

RAID is Severely Compromised

The *RAID*book published by the RAID Advisory Board (Sixth Edition, February 1997, Chapter-5, pp 120..121), describes the need for RAID-5 systems to implement protection against what's often referred to as "the write hole". It specifically describes the situation where, in the face of power-failure, incomplete writes to disk can leave data and parity inconsistent (one has old data and the other has new). A RAID device, to qualify for RAB certification, must meet the following requirement:

… each block of data specified in the unfinished write should appear after recovery as either:

- written completely or not written at all, (either is correct), or,
- clearly marked as incorrect or unrecoverable data.

The text further points out that failure to address the write-hole not only puts the data written at risk, but will prevent proper reconstruction of "bystander data" in the same RAID stripe in case of subsequent failure of another member disk.

So in a RAID system using SAS disks that occasionally overwrite previously written data with the wrong data stream, in the face of power-failure it not only results in the loss of the data-blocks

being written at the time, but it also loses the redundancy required to reconstruct bystander data unrelated to the transaction being processed.

SAS, as currently defined, exacerbates the problem in that not only are we concerned about inconsistency between data and parity; now we're also concerned about the general integrity of the primary data and parity. Standard techniques that cover the write-hole by logging incomplete operations in non-volatile memory and then rebuild parity after a power cycle, fail in this scenario, leaving not only the blocks being written corrupted, but also leaving a window open to corrupt innocent bystander data.

It's imporant to remember at this point that SAS systems designed according to the standard may incur transport errors on average somewhere in the range from once every 42 seconds to once every 7 minutes, so the probability that such an error will coincide with a power-failure is extremely high, resulting in permanent loss of application data.

In short, a RAID system that easily meets the RAB requirements for data integrity using a parallel-SCSI transport to disk will not meet the same requirement when the parallel SCSI interface is replaced with SAS.

In this respect SAS, as defined, is deficient as compared to existing parallel SCSI solutions in its ability to deal with data integrity standards long established by the RAID Advisory Board. This clearly does not meet the industry's expectations of the next generation SCSI transport intended to offer whole-scale replacement of parallel SCSI.

SAS must provide a mechanism to allow traditional RAID systems to protect user data in the face of these types of failures.

Solution

Fibre Channel avoids this problem by defining two header fields for error checking. The one most commonly implemented is SEQ_CNT, a sequence count used to check that all data frames arrived. Less often used is RELATIVE OFFSET, which allows the recipient to check that all data bytes arrived.

Parallel SCSI bit-errors in themeselves are detected between CRC and parity algorithms and do not cause misidentification of data packets. A similar problem results if REQ/ACK streams lose syncronization, however on high-speed (U320) SCSI links using packetized transfers, these are comparatively low-frequency signals which, experience has proven, virtually never lose syncronization.

The overhead for inserting a CRC is negligible. In contrast, the only way a SAS target can validate a write data transfer is to break it up into many smaller transfers, each with its own XFER_RDY. The bus round-trip required to issue a new XFER_RDY essentially guarantees that the connection will be closed and a new connection will have to be opened. That is not a low overhead operation.

The requirement is for a low overhead way for a target to validate a write data transfer at granularity comparable to the disk block size or frame size. This proposal suggests a per connection counter that counts frames or information units.

Each link maintains a frames-transmitted counter that is initialized whenever a connection is opened. The counter increments for every frame or information unit sent on the connection. The counter is included in the information unit header. Likewise the receiver maintains a count of frames received during the course of the connection. If the received value does not match the expected value, it implies that one or more frames have been lost. The recipient transport discards that information unit and all subsequent information units received on the connection, similar to discarding frames with incorrect address hashes. Note that if a frame is lost, the sender

will detect an ACK/NAK timeout when it next interlocks the connection. This guarantees that all frames that are discarded are for the same (non-interlocked) data transfer as the frame that was lost. Note also that the existing limit of at most 255 outstanding R_RDYs ensures that an eight-bit counter is sufficient.

This proposal defines a flag to indicate the presence or absence of the counter.

Recommended Changes

Clause 9.2.1, SSP frame format, and table 59
Define bit 2 of byte 11 in table 59 to be CFCE. The corresponding bit in Fibre Channel headers is defined as “Reserved for Exchange reassembly”, for use in exchanges that span many source or destination ports. Define byte 15 in table 59 to be CONNECTION FRAME COUNTER. The corresponding byte in Fibre Channel headers contains the low byte of the sequence count. Add the following definitions to the text following table 59:

> The connection frame counter enable (CFCE) bit is set to one to indicate that the CONNECTION FRAME COUNTER field is valid. The CFCE bit is set to zero to indicate that the CONNECTION FRAME COUNTER field is reserved.
>
> The CONNECTION FRAME COUNTER field counts frames with the CFCE bit set to one that are sent on a single connection. SAS devices shall set the CONNECTION FRAME COUNTER field to one in the first frame with the CFCE bit set to one that is sent on a connection. The contents of the CONNECTION FRAME COUNTER field shall increment in subsequent frames with the CFCE bit set to one that are sent on the same connection. The counter value FFh shall wrap around to 00h when incremented.
>
> The recipient may compare the CONNECTION FRAME COUNTER field to the expected value. If it is checked and the CONNECTION FRAME COUNTER field does not match the expected value, that frame and all subsequent frames received on the same connection shall be discarded by the transport layer.