

Document: T10/02-319r0
To: T10 Committee Membership
From: Edward A. Gardner, Basil Networks
Subject: SAS Relative Offset

Date: August 20, 2002

SAS does not presently allow random access to data buffers. Targets must transfer data sequentially for each command. While this is a reasonable restriction for disk and tape drives, other environments have other needs.

We are designing a RAID controller for use with SAS disk drives. RAID controllers commonly implement caching. Many applications require redundant RAID controllers with critical data duplicated in both controllers' caches.

SAS is an attractive way to communicate critical data between redundant controllers. It appears to have the requisite characteristics, and it reuses hardware that we will already have implemented. However, using SAS for this purpose requires random access to data. While we could define a vendor unique way to implement random access, we would prefer not to do so. For this purpose we only require random access for data returned by READ commands.

Other environments that may require random access to data include copy managers that implement the EXTENDED COPY command and similar commands that use the data buffer to hold a large, complex data structure.

This proposal defines random access completely, including both read and write data and a way to enable it for individual commands. I have made this a complete proposal because I assumed it was necessary to get this approved. The only issue we actually care about is the changes to DATA information units sent by the target (read data). That is the changes to clause 9.2.3.4 and any new fields in table 60. These are the changes that would affect our hardware. If it were plausible to approve those while deferring other changes to SAS-2, we would be happy to do so.

I have also assumed that proposal 02-318r0 is accepted, and therefore that data offsets should be required to be a multiple of 1024 bytes.

Clause 9.2.1, SSP frame format, and table 60. Define a bit named RANDOM. I suggest bit 0 of byte 10. Add the following definition:

The RANDOM bit may be set to one in COMMAND information units to indicate that the target may use random access to the command's data buffer(s). The RANDOM bit may be set to one in DATA information units sent by a target to indicate that random data-in buffer access is being used. The RANDOM bit may be set to one in XFER_RDY information units to indicate that random data-out buffer access is being used. If the RANDOM bit is set to zero in a COMMAND information unit, it shall be set to zero in all DATA and XFER_RDY information units sent for that command. The RANDOM bit shall be set to zero in TASK information units, RESPONSE information units and DATA information units sent by an initiator.

Clause 9.2.1, SSP frame format, and table 60. Define a four byte field named READ DATA OFFSET. I suggest bytes 20-23. Add the following definition:

The READ DATA OFFSET field is used in DATA information units to implement random data-in buffer access (see 9.2.3.4). It is reserved in DATA information units when the RANDOM bit is set to zero and in all other information unit types.

Clause 9.2.3.3, XFER_RDY information unit and table 66. Rename bytes 0-3 to be a four byte WRITE DATA OFFSET field. Add the following:

If the RANDOM bit (see table 60) is set to zero in the first XFER_RDY information unit for a command, the initiator shall transfer data from the beginning of the data-out buffer. If the RANDOM bit is set to zero in a subsequent XFER_RDY information unit for a command, the initiator shall transfer data from the data-out buffer immediately following the data transferred by the previous XFER_RDY information unit for the same command.

If the RANDOM bit (see table 60) is set to one in a XFER_RDY information unit, the initiator shall transfer data from the data-out buffer starting at the byte offset specified by the WRITE DATA OFFSET field. The value in the WRITE DATA OFFSET field shall be divisible by 1024.

Clause 9.2.3.4, DATA information unit. Add the following:

If the RANDOM bit (see table 60) is set to zero in the first DATA information unit sent by a target for a command, the DATA field shall be transferred to the beginning of the data-in buffer. If the RANDOM bit is set to zero in a subsequent DATA information unit sent by a target for a command, the DATA field shall be transferred to the data-in buffer immediately following the data transferred by the previous DATA information unit sent by the target for the same command.

If the RANDOM bit (see table 60) is set to one in a DATA information unit sent by a target, the DATA field shall be transferred to the data-in buffer at the offset specified by the READ DATA OFFSET field (see table 60). The value in the READ DATA OFFSET field shall be divisible by 1024.

The RANDOM bit (see table 60) shall be set to zero in DATA information units sent by an initiator.