

Tech Note - HDD based recorders

Document no.: GEC-US-0040

Author: Hugh Tarver

Revision: 1.0.0

Date: June 15, 2018

Preface

Conventions

IEC prefixes are used to separate between decimal and binary multiples.

- Decimal prefixes k (kilo), M (mega), G (giga), and T (tera) denote 10^3 , 10^6 , 10^9 , and 10^{12} respectively.
- Binary prefixes Ki (kibi), Mi (mebi), Gi (gibi), and Ti (tebi) denote 2^{10} , 2^{20} , 2^{30} , and 2^{40} respectively.

Dates printed in numerical format are in the order day month year. E.g. 24.01.10 (or 24.01.2010) denotes January 24th, 2010.

Copyright

Copyright © 2018 Galleon Embedded Computing AS. All rights reserved.

Limited Liability

Galleon Embedded Computing AS assumes no liability arising from any use of information provided within this document.

Revision History

Rev.	Date	Changes
0.1.0	8/6/18	First draft
1.0.0	15/6/18	First release

Note: Revision numbers are on the form X.Y.Z, where X indicates the major revision number, Y indicates minor revision number, and Z indicates revision number for editorial changes that do not alter the technical content of this document.

Contents

Preface.....	ii
Revision History	ii
Contents	iii
1 HDD RECORDING	4
1.1 HDD bandwidth problems.....	4
1.1.1 Seek times	4
1.1.2 Fragmentation of the file system	4
1.1.3 Sustained bandwidth	4
1.2 Mitigation	5
1.2.1 Memory buffers	5
1.2.2 Formatting.....	5
1.2.3 Avoid simultaneous activities	5
1.3 Impact.....	5
2 SUMMARY	6

1 HDD recording

For rugged environments, SSDs are used for all recording applications. HDDs are not reliable enough to be sensible competition to the SSDs in these circumstances.

However, in commercial environment applications, HDDs could be used, although still not ideal (as explained later in this document).

When storage requirements are large, then SSDs are extremely expensive, so HDDs are certainly worthy of serious consideration.

1.1 HDD bandwidth problems

Hard Disk Drives (HDDs) display extreme variations in write bandwidth to the disk. This is in stark contrast to Solid State Device based disks (SSDs) which have a predictable worst-case bandwidth.

1.1.1 Seek times

HDD bandwidths can be very high while streaming data to one location on the disk.

However, sustained bandwidth (for a recording application) are affected to the need for the write head to seek a new location when required. This seek time hugely reduces the effective write bandwidth of the disk. Also, the RAID controller will be delayed until all of the disks (4 or 8, depending on the configuration of the recorder) have completed the write cycle. i.e. if just one of the disks has a seek time delay, all of the RAID array will be affected.

When recording data, the recording software needs to write to more than 1 file. In the simplest case, this could be a raw data file, and an index file. Writing to multiple files leads to more seek time and less time spent writing to the disk.

The problem is exacerbated if multiple different channels are being recorded simultaneously (each with two or more files).

To further complicate the access pattern, the file system will occasionally read from the drives to update file allocation tables or journals. These intermittent read accesses will also severely affect the disk access speed.

1.1.2 Fragmentation of the file system

Fragmentation of the files spread across the disk forces the write process to pick new locations for writing to more frequently. This also results in more seek time, and less time spent writing to the disk.

1.1.3 Sustained bandwidth

Whereas with SSDs, recording bandwidths to HDDs cannot be guaranteed. It is possible (although unlikely) that the disk bandwidth could reduce to much less than normal. This would only happen if certain factors combine at the same time, causing the disk to spend almost all its time seeking and hardly any time writing to the disk.

Also, HDDs keep track of bad blocks, which gradually develop (and are detected) over the lifetime of the disks. Managing these bad blocks (and needing to avoid using those locations on the disk), will also impact the capability for the HDD to sustain any particular bandwidth over the full duration of its life.

HDD based recorders

Therefore, it really isn't possible to give a 100% guaranteed bandwidth which can be sustained into the disk.

1.2 Mitigation

It is possible to mitigate the variability of the HDD write bandwidth. The techniques described below will reduce the likelihood of problems, but they cannot completely remove the issue.

1.2.1 Memory buffers

Using memory in the computer to buffer data going to disk helps to reduce the impact of the variability of bandwidth to the disk.

Galleon recording software all uses a main memory buffer.

1.2.2 Formatting

If the disk is formatted before recording is started, then the likelihood of problems occurring is reduced significantly. Formatting removes fragmentation issues, especially if the disk is not filled during a recording mission.

1.2.3 Avoid simultaneous activities

Recording the data into a single file (or set of files) will help to mitigate the issue.

Reducing the number of files being written during the recorder operation will also help. For instance, recording only video data, rather than the full SDI data set (these are options for the Galleon video recording application).

Similarly, avoiding playback of the data at the same time as recording will reduce seek times.

1.3 Impact

If the HDD bandwidth does drop down below the incoming data rate, and the memory buffers fill up, then some data will be dropped.

In many cases, the Galleon recording software will deal with the issue in a clean way.

For instance, with the Galleon video recording software, whenever possible whole video frames will be dropped, rather than storing corrupted video data on the disk. This behaviour cannot be guaranteed, though.

So, for all cases, it is possible that some corrupted data will end up being stored onto the disk.

If this happens, some post processing would be needed to identify the corrupted data and remove it from the recorded file.

Note: Galleon does not provide any tools for this type of post processing.

2 Summary

If a recording requirement cannot afford to lose any data, then use SSDs as the recording medium.

But if dropping data (and/or corrupted data) on rare occasions is acceptable, then a large storage recorder based on HDDs would be a good option (because it will be much cheaper than with SSDs).