grass valley
A BELDEN BRAND

Quasar File
Optimizing Performance

Version History

Date	Version	Release by	Reason for Changes
22/10/2014	1.0	J.Courtney	Document originated.
13/05/2015	1.5	J Courtney	Added information on PCI lanes to CPU section.
07/09/2015	1.6	J Metcalf	Edit for XF rebrand.
19/01/2016	1.7	G Emerson	Completed Rebrand.
01/03/2016	1.8	J Metcalf	4K timing graphs added.
19/09/2016	1.9	J Metcalf	New GPU added to timing graphs.
09/01/2017	2.0	J Metcalf	Updated for V3.0.0.14
14/06/2017	2.1	J Metcalf	Updated for V3.0.1.10
09/04/2017	2.2	J Metcalf	Updated for V4.0.0.0
27/11/2018	3.0	J Metcalf	Updated for V4.0.0.20 + GV rebrand.
29/04/2019	3.1	J Metcalf	RTX GPU added to timings graphs.

Table of Contents

List of Figures

List of Tables

Introduction

One of the most important design considerations of the GV File software framework was that it must be compatible with COTS (Commercial Off The Shelf) IT infrastructure. The versatile nature of the framework and its products, such as Quasar File, means that it can run on relatively low end, cost effective infrastructure through to state of the art high end enterprise solutions. Assuming a compatible GPU is present, the user will enjoy Grass Valley's unrivalled conversion quality. However the conversion duration is greatly affected by the IT specification.

For the majority of today's Broadcast and Post community time is money and speed is of the essence. Optimal performance and stability is crucial for maximizing the potential of products such as Quasar File. This document aims to provide Quasar File customers with the knowledge required to make an informed hardware and software purchasing decision, and ensure they have the appropriate performance for their requirements.

The GV File framework benefits from a Service Orientated Architecture (SOA) which has been engineered to support a range of deployments from a single node through to a cluster of nodes within a data center. In all cases the service responsible for the image processing is the **GV File Node** service.

There are many factors which will influence the conversion capability and speed of Quasar File. These are:

- **Compute power** – the specification of the hardware available to the GV File Node
- **Storage** – the bandwidth available for the GV File Node to access the media (read and write)
- **Source** – the resolution, frame rate, codec and wrapper of the source
- **User profile** – the resolution, frame rate, codec and wrapper of the output

Each will now be explored in more detail:

1. Compute

The performance of the compute hardware is one factor which will affect the conversion durations offered by the GV File Node.

The key components of the compute hardware are:

- **GPU**
- **CPU**
- **RAM**
- **Motherboard / PCI-E**

Typically, the GPU is the component that has the most significant impact on the conversion duration, however as we will discover later the other components play their part.

1.1 GPU

1.1.1 Introduction

It is essential a compatible GPU is present for successful operation of a GV File product.

Although single and double precision GPUs are available, for broadcast applications single precision arithmetic is more than sufficient.

The GPU/s of choice must offer at least:

- **2GB** of dedicated memory for SD and HD
- **4GB** of dedicated memory for 4K.

Not only is a valid GPU required but it must be configured to offer a **V1.1** or above **OpenCL** environment. Without this, no conversion can take place.

OpenCL (Open Compute Language) is supported by both AMD and NVIDIA and therefore the choice to select your preferred vendor. However some environmental limitations control when certain GPUs can be enabled to support OpenCL. These will be discussed later.

Within GV File, the GPU performs the render or conversion of the uncompressed video data. The CPU resource is responsible for all video decoding and encoding.

At the time of writing, most current GPU are capable of supporting HD conversions at real time speeds. However, multiple GPUs can be used within a single host machine to achieve real time speeds for higher than HD resolutions. **Real-time speed** simply means that the processing duration will match the file runtime. For example; a file with a duration of one hour will be processed in one hour.

It is important to appreciate that it is not recommended to mix and match different GPU models. When multiple GPU are employed in a GV File Node, all GPUs will run at the same speed of the lowest performing GPU.

Mixing AMD and NVIDIA GPU's will not work as the GPU drivers will conflict and prevent the GV File Node from successfully starting.

It is highly recommended that this is not attempted!

1.1.2 Consumer or Professional GPU?

One of the initial decisions a user will be faced with when selecting a GPU is whether to use/purchase Consumer/Commodity or Professional/Enterprise GPU/s.

There are many considerations which will influence this decision but generally it is fair to say the professional path will provides a greater chance of a smooth and efficient trouble-free operation.

Professional cards differ from consumer cards in the following ways:

- **Certified Solutions:** The major IT vendors such as Supermicro, Dell and HP work closely with both AMD and NVIDIA to create certified solutions (rack mounted servers or workstations). Certification ensures that all physical issues have been managed, there is sufficient power and cabling for the specified GPU/s and that the solution will operate 24/7 within the safe recommended thermal limits.

- **Quality and stability**: The professional cards have a huge amount of effort invested into the quality and stability of both the hardware and the associated software drivers.

- **Lifetime and roadmaps**: Consumer cards are frequently replaced as they are at the forefront of the latest and fastest technology. Professional cards enjoy longer development periods and are replaced less often. This helps large organisations standardize and assists with stability.

- **Support:** The support infrastructure offered by AMD and NVIDIA for professional cards is very comprehensive. This is important for businesses as down time impacts revenue. The consumer cards are generally based on the chipsets provided by AMD and NVIDIA but are manufactured and supported by 3rd party vendors. Professional GPU's generally offer longer warranty periods.

Table 1: Summary Professional or Consumer GPU

Professional	Consumer
Guaranteed certified solutions (24/7 operation)	More cost effective
Higher quality and stability	User responsible for physical compatibility
Better support infrastructure	User responsible for thermal management
Increased lifetime and roadmaps	User responsible for power management

Table 2: GPU, OS Compatibility with GV File Node

GPU	OS	Node Install
Consumer AMD (R9, R7, HD6000, HD7000)	Windows Server 2012	Service or application
Consumer AMD (R9, R7, HD6000, HD7000)	Windows 7,8, 10	Service or application
Consumer AMD (R9, R7, HD6000, HD7000)	Redhat/CentOS	Application only
Consumer NVIDIA (GeForce GT, GeForce GTX)	Windows Server 2012	Service or application
Consumer NVIDIA (GeForce GT, GeForce GTX)	Windows 7,8, 10	Service or application
Consumer NVIDIA (GeForce GT, GeForce GTX)	Redhat/CentOS	Application only
Professional AMD (FirePro)	Windows Server 2012	Service or application
Professional AMD (FirePro)	Redhat/CentOS	Application only
Professional NVIDIA (Quadro, Tesla)	Windows Server 2012	Service or application
Professional NVIDIA (Quadro, Tesla)	Redhat/CentOS	Application only

1.1.3 GPU Performance

These charts show the relative performance between various single GPUs with Quasar File. These results may not reflect the relative performance of other applications running on the same hardware. The data has been derived from benchmarking GV File on our reference hardware with dual SSD's (see Appendix A).

Figure 1: GPU Timings, HD De-interlace and Rescale

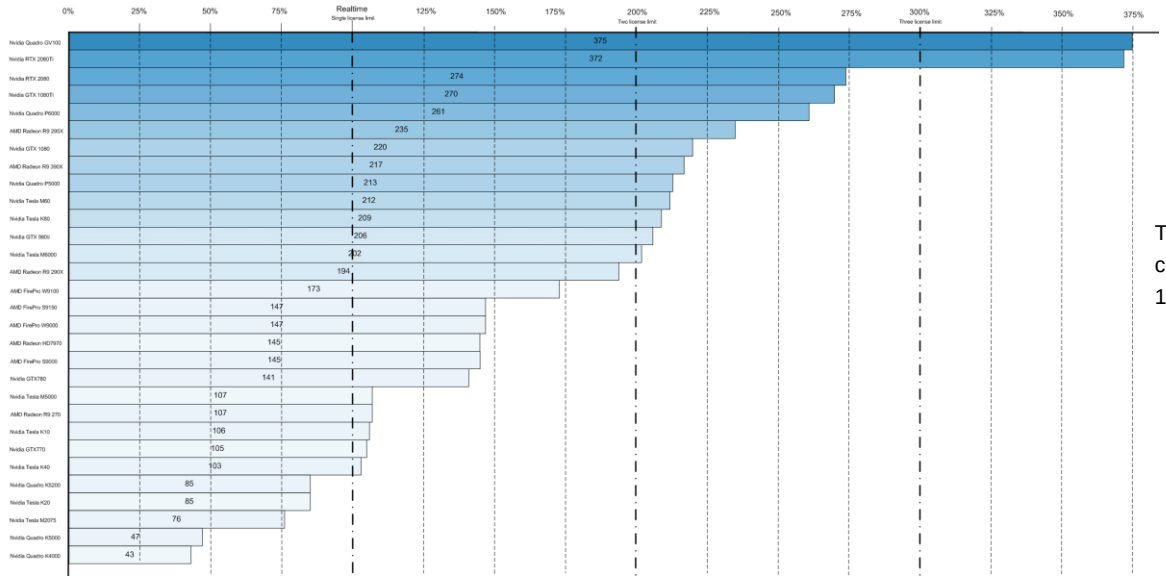

Timings based on conversions of 1080 50i to 720 50p

Figure 2: GPU Timings, low frame rate

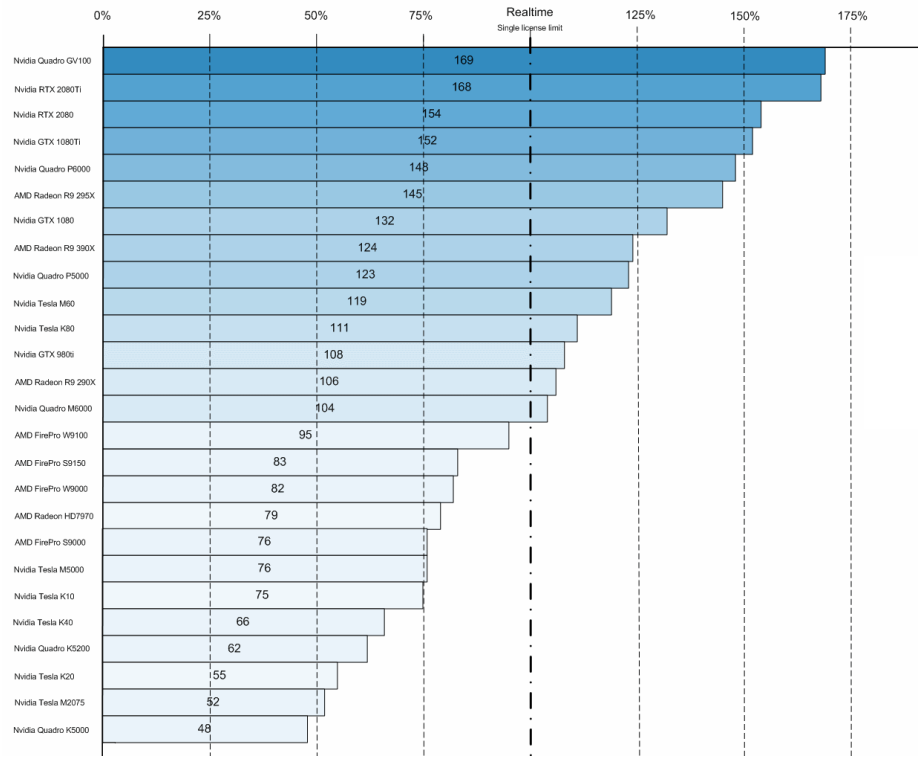

Timings based on conversions of:

1080 25p to UHD 25p

Figure 3: GPU Timings, high frame rate up-conversion HD to UHD.

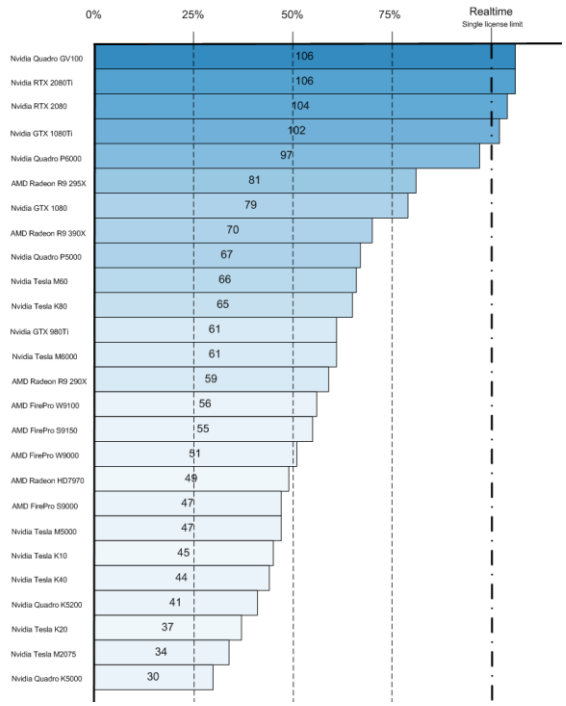

Timings based on conversions of 1080 50i to UHD 50p

Figure 4: GPU Timings, low frame rate, down convert UHD to HD.

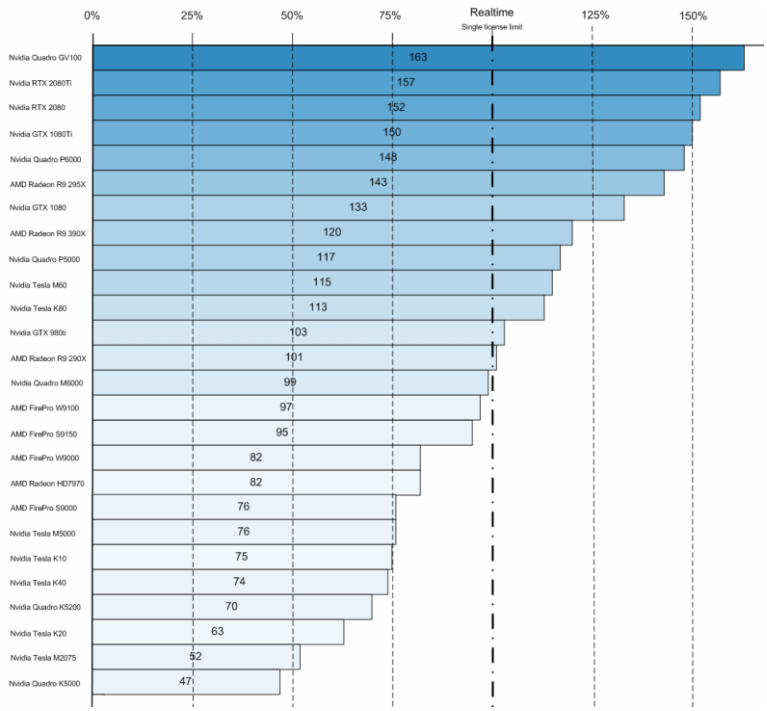

Timings based on conversions of UHD 25p to 1080/25p

Figure 5: GPU Timings, high frame rate, down convert UHD to HD.

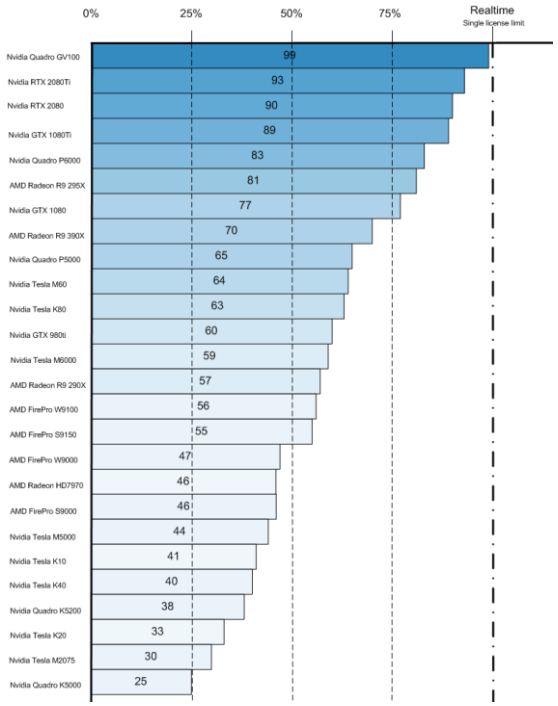

Timings based on conversions of UHD 50p to 1080/50i

A single GV File license offers the conversion of one media asset at a time, running up to real time speed. Licenses may be aggregated together to allow faster than real time conversion. For example 2 license can reduce conversion times by halve, but only if there is sufficient resource to process at the higher speed.

In our experience multiple GPUs offer near linear performance improvement, unless a different system bottleneck is reached. For example, for a conversion 1080/50i to UHD 50p, 1x AMD FirePro w9000 will offer approx. 50% real-time speed. It would be reasonable to expect to 2x AMD FirePro w9000, would offer real-time conversion. However, in practice, 2x AMD FirePro w9000 will actually perform the conversion in approximately 90% of real-time.

Note, some GPU models are actually ‘dual’ GPU mounted in one GPU solution. Examples of this include the NVIDIA K10, the NVIDIA K80, and the AMD R9 295X. These should be counted as two GPU cards, and will influence the quantity of system RAM required (see section 1.3 below).

A typical 4K conversion has far more data to process, compared to HD and conversion times will be significantly longer.

1.1.4 Driving a Display and Running GV File Simultaneously

It should be considered that although multiple GPU can be employed in a single GV File Node, that if GPUs of differing performance are installed, the performance of the higher spec GPU will throttle back to match the performance of the lower spec GPU. It is for this reason it is not recommended to install different models of GPU in a single Node environment (although if the relative performance of two different GPU are similar, then using different GPU’s isn’t a problem).

In the scenario, employing a lower specification GPU in the host machine for display purposes, it is recommended this GPU is disabled within the configuration of the Node (see GV File Installation Guide, section 7.3, for details how to disable). If this isn’t done and it is compatible with an Open CL environment, it will cause the performance of all other GPU in the Node to throttle back, which would have a very detrimental effect on performance of the Node.

1.1.5 GPU Addressable Memory

If the Compute system has GPUs with more than 4 GB of addressable on-board memory, then the motherboard BIOS needs to have the PCI-E/PCI configuration changed so they can address above 4 Gb.

1.1.6 Consumer GPU design

Modern gaming GPU can be extremely fast and offer very good value for money. However, the actually hardware may not be manufactured by either AMD or Nvidia. For example, there are lots of manufacturers making Nvidia GTX1080 GPU and although the silicon that the GPU is based on is essentially the same, hardly any two GPU look alike. There are differences in the approach that different manufacturers design their GPU cooling systems. Some just mount fans to the side of the GPU and simply circulate air within the server enclosure. Others actually design the GPU so that the exhausted hot air is routed out of the server enclosure. This latter approach is much better in terms of cooling efficiency and is much the preferred option in Quasar File installation, and is essential for multi GPU solutions.

GPU employing simple fan cooling. Not recommended for Quasar File installations but probably alright for single GPU solution.

GPU employing external exhaust cooling. Recommended for Quasar File installations that employ multi-ple GPU solutions

1.2 CPU

Introduction

The CPU primarily manages the host system, data transfers and decoding and encoding of compressed video.

Our benchmarking has focussed on two types of host machine. One based on Enterprise components and the other Consumer. To achieve real time decoding and encoding of 1080i input/output the host CPU requirements would be:

- 2 x Xeon 2650 V2 (Enterprise)
- 1 x i7 4790 or 1 x i7 5930 (Consumer/commodity)

In both of these examples it would be beneficial to have the appropriate GPU resources to offer real time speeds as this would provide the optimum system.

If the CPU resource was significantly lower there would be little benefit in having enough GPU for real time conversion speeds, as the CPU would become the bottleneck. Equally, it could be argued, it would not be beneficial to invest in a capable host machine if it is only planned to run one low end GPU.

1.2.1 PCI-E Lanes & Link Speed

PCI-E lane width and link speed is important as this will dictate the bandwidth available to each GPU in the system and influence the overall conversion duration.

- PCI-E lanes allow the host CPU to pass data to and from the GPU/s
- PCI-E lanes can be aggregated together to increase the width of the bus to offer a greater bandwidth
- PCI-E width is typically described as x2, x4, x8, x16
- PCI-E revision dictates the link speed of each lane
- The width together with the link speed will dictate the overall bandwidth available
- CPU/Motherboard analyses PCI-E on power-up, creates a traffic map and negotiates link speed

Table 3: PCI-Express Revisions

Revision	Speed per lane	X16 lane speed	Transfers speed
PCI Express 1.0	250 MB/s	4GB/s	2.5 GT/s
PCI Express 2.0	500 MB/s	8GB/s	5 GT/s
PCI Express 3.0	985 MB/s	15.75GB/s	8 GT/s

Different CPUs can support a different number of PCI lanes.

Table 4: CPU Information Table

CPU	Max PCI-E Lanes	PCI-E Revision
I7 4790K	16	3.0
I7 5820K	28	3.0
I7 5930K	40	3.0
Xeon 2650 V2	40	3.0

The vast majority of GPUs now support PCI-e Revision 3 with a width of 16 lanes. However some older GPU, or cheaper consumer GPU, may be PCI-e Revision 2 and may only support up to 8 lanes.

Table 5: GPU Information Table

GPU	Max PCI-E Lanes	PCI-E Revision
NVIDIA K20	16	2.0
AMD RX560	8	3.0
AMD R9 290X	16	3.0
AMD S9000	16	3.0
Nvidia GTX 950	8	3.0
Nvidia Tesla C2075	16	2.0
Nvidia GTX 980Ti	16	3.0
AMD R9 390X	16	3.0
Nvidia GTX 1080Ti	16	3.0

It is recommended that each GPU installed can operate at a PCI-E width of at least 8 lanes and a link speed of 8GT/s (Gen 3.0)

Slower configuration will work but will have a detrimental effect on conversion durations.

All GPUs installed should meet this recommendation for optimal performance.

Below is a table illustrating some recommended/typical PCI-E configurations.

Table 6: Recommended CPU / PCI-E width Configurations

CPU	Chipset	PCI Revision	Quantity of CPU	Recommended Quantity of GPU	Lanes width per GPU
I7 4790K	Z97	3.0	1	2	x8 x8
I7 4790K	Z97	3.0	1	1	x16
I7 5930K x1	X99	3.0	1	3	x8 x8 x8
I7 5930K	X99	3.0	1	2	x16 x16
Xeon 2650 V2	C602	3.0	2	4	x16 x16 x16 x16

1.2.2 CPU Power Management

Typically, motherboards will offer a CPU power management profile in the BIOS. In some cases this BIOS setting must be set in addition to the operating system power mode. Failure to do so will result in sub optimal hardware performance.

Here is an example of a typical power management BIOS page:

- **Power Technology** set to **Max Performance**
- **Energy/Performance Bias** set to **Performance**

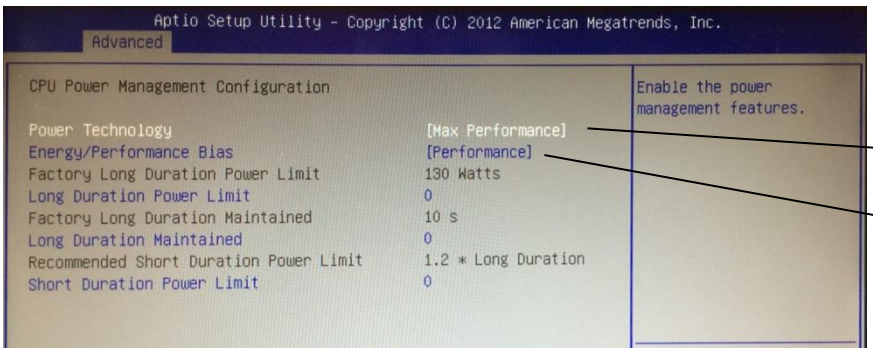

1.3 RAM

The amount of system RAM required is proportional to the number of GPU fitted. Please refer to Table 7.

In circumstances where there is not enough RAM resource provided, the system may resort to substituting memory with HDD resources. Read and Write times to a hard disk are significantly slower relative to accessing RAM and conversion duration will be adversely affected.

The RAM specification should also be considered. High specification RAM allows faster reading and writing times and so effects conversion duration. The faster the better!

Table 7: Memory Requirements

Format	No of GPUs	Minimum Required Memory	Recommended Memory
SD/HD	1	8GB	16GB
SD/HD	2	16GB	32GB
SD/HD	3	24GB	32GB
SD/HD	4	32GB	48GB
SD/HD	5	40GB	64GB
SD/HD	6	48GB	64GB
4K	1	16GB	32GB
4K	2	32GB	64GB
4K	3	48GB	64GB
4K	4	64GB	96GB
4K	5	80GB	128GB
4K	6	96GB	128GB

Minimum RAM recommendations maybe applied when the host machine is being utilised solely for GV File services. In circumstances where the host machine is supporting other applications and processes in addition to GV File services, then recommended RAM figures should be applied. Please note running other applications may affect the performance of GV File.

It should also be considered that the recommended RAM resources allow for future-proofing of the host machine.

Note 1 – Some GPU models are actually ‘dual’ GPU mounted in one GPU solution. Examples of this include the NVIDIA K10, the Nvidia K80, the Nvidia M60 and the AMD Radian R9 295X. It is important to understand that these should be counted as two GPU. This is important when considering the amount of system RAM.

Note 2 – RAM utilization can be affected by data transport speeds to and from disc. A system where fast reading of source files, combined with slow writing of output files represents the worst case scenario. Where read and write times are similar, the utilization of RAM can be significantly less.

1.4 Storage

1.4.1 Introduction

The type of storage employed can have a significant effect on conversion duration. If the storage medium is slow, the compute node will simply slow down and process data at the speed it can access the storage. Typically, single large files such MXF or MOV files, place a smaller demand on the storage as the access is constant and sequential. Uncompressed DPX files on the other hand put a huge demand on storage due to the concurrent transactional nature of having separate, but reasonably large files for each frame. All forms of storage will slow down when multiple concurrent accesses are applied unless the necessary infrastructure is used.

Real-time delivery of 4k DPX would generally require bandwidths of up to 2GB/s. A Single compressed file at real time might be more in the order of 100MB/s in the real world.

1.4.2 Internal / Local Storage

SSD

For optimal performance SSD hard drives are superior to SATA, SCSI, or IDE. The read and write speeds of modern SSD drives make these the preferred choice for optimum GV File performance.

SSD's are relatively more expensive, but considering the benefits they bring in terms of performance, they are often the best choice.

SATA Drive

SATA drives offer good performance and are relatively cheap compared to SSD drives. Where there is a requirement for a large capacity of internal storage, SATA drives offer a practical solution for a more cost effective system.

Several SATA drives maybe RAID'ed to increase read and write times, but typically due to the mechanical nature of hard disk drives concurrent performance is poor.

1.4.3 External Storage

NAS and SAN drives

External Storage devices such as NAS (Network Attached Storage) and SAN (Storage Area Network) offer a mechanism to have a centralized storage system available to many users. Whilst enterprise IT infrastructure offers some impressive fast connectivity (10G Ethernet, 40G Infiniband, 8G / 16G Fibre channel) measures need to be put in place to ensure multiple users can achieve concurrent fast access.

If demand increases, the bandwidth available might drop. As stated in the introduction, if the bandwidth of file access falls, the node will slow down and process at the speed the storage medium can service.

1.5 Hardware Enclosure

The choice of Enclosure is of critical importance when deciding upon hardware to run a GV File solution on. In particular the choice of GPU and the quantity of GPU will have a significant bearing on choice of Enclosure.

Enclosures come in all shapes and sizes.

- **1U rackmount** – generally supports passively cooled GPUs only (relies on internal chassis cooling).
- **2U rackmount** – generally supports either type of GPU.
- **4U rackmount** – very versatile chassis size. Can generally support either type of GPU.
- **4U workstation** – very versatile chassis size. Can generally support either type of GPU.
- **Gaming / consumer enclosures** of various sizes – generally actively cooled GPUs only.

1.5.1 Enclosure considerations

- **Quantity of GPUs**
- **GPU type (actively cooled or passively cooled)**
- **Power**
- **Quantity of Storage**
- **Motherboard type (single socket / dual socket)**

2. Source

The source resolution and frame rate dictates the quantity of data to be processed. This, combined with the codec, will govern the extent of the demand on the Compute node. A 4096x2048 spatial resolution is significantly more pixels than 720x576 spatial resolution. 60Hz is a much higher temporal rate than 24Hz. AVC-I 100Mbps places a greater demand on the CPU for decoding than AVC-I 50mbps.

3. User Profile

The User Profile and its associated conversion parameters will define the resolution, frame rate, codec and wrapper of the output file.

As with the source, the higher the uncompressed data rate the greater the demand on the GV File Node. The load on the host is also affected by the codec profile selected.

4. Operating Systems

The GV File framework is cross platform and can be deployed on both Windows and Linux operating systems.

4.1 Windows

Snell Advanced Media recommends the use of 64bit Windows Server 2012 as this is the version all product validation is performed on. Other 64bit operating systems such as Windows 7, 8, 10, Server 2008 R2 and Server 2016, will all work but GV does not actively validate on them

4.2 Linux

Grass Valley recommends the use of 64bit Redhat or CentOS operating systems. We cannot guarantee the successful operation of other Linux based operating systems.

5. FAQS

1. **What is the maximum temperature a GPU should reach during peak performance?**

   Typically no more than 95-98 deg C would be recommended, however you should always consult the datasheet for the exact GPU in use.

2. **Is it possible to damage a GPU through overheating?**

   No. The GPU drivers will reduce the effort level in an attempt to manage the temperature.

3. **Can I use consumer NVIDIA cards with a Linux OS?**

   Yes. Consumer NVIDIA cards will operate successfully in Linux.

4. **How can I achieve faster than real time conversions?**

   Yes, it is now possible to aggregate GV File licenses together to achieve faster than real time conversions.

5. **Why are some consumer GPUs faster than professional GPUs?**

   Consumer GPUs utilize the very latest and fastest technology and are generally available before their Professional counterparts. Also, high end professional GPU utilize double precision technology, which does not offer any benefit to GV File conversions, but can have an effect on the peak processing performance of the GPU. Consumer GPU do not offer double-precision technology and so are not similarly hampered.

6. **If I had two similar specification GPUs which are the same brand could I use them in the same host machine?**

   Yes. It is just not recommended to mix low specification GPUs with high specifications.

7. **Why can't I purchase a passively cooled Tesla online?**

   Passively cooled NVIDIA Tesla cards are only sold as part of a certified system. This ensures they are only used in certified solutions where the appropriate thermal management is present.

8. **When monitoring my GPU, its performance does not seem to be optimised.**

   This is most likely caused because the system is running in power save mode. Please refer to section 1.2.2 on page 14, and check you system is set to Performance Mode.

Appendix A. Reference System

- Microsoft Windows Server 2012
- Single socket i7 5930K or dual socket Xeon 2650V2
- PCI-E GEN 3.0
- 64 GB, DDR4 System RAM

Note: Internal timing tests were performed with two SSD (reading from one, writing to the other).