What Price your Storage?

A paper explaining some of the real world considerations and pitfalls when buying storage systems for file-based media workflows.

Written and prepared by Graham Collins
April 2015

Tape replacement

20 years ago, it would be an easy technical decision when specifying media workflows. You'd choose your video tape format (there were only a few available) and go with the manufacturer of that format. Often you would talk with a trusted manufacturer and let them tell you what technology to invest in. And that was it, job done. You'd buy tape stock, record your material, and be pretty much guaranteed that you would always be able to play that material back on a similar deck. If you needed to do more editing, or record more lines, it was simple: you would simply buy more decks and install them. Your ambitions were only limited by the size of your budget.

Wind forwards to 10 years ago, and file-based media workflows relied on broadcast manufacturers providing bespoke systems that provided the high performance and features required for the broadcast industry, with suitably high price tags attached. I lost count of how many times I have heard people say "*well, nobody has ever been fired for buying brand X*" when asked why they made a particular technology choice. These systems performed well when used for what they were designed for, and gave years of good service

And then the industry woke up to the fact that commodity IT storage systems could be obtained for a fraction of the cost compared to buying from the more traditional broadcast equipment manufacturers. Storage is storage right? So then, how can these broadcast manufacturers justify many times price premium over the commodity products? I hope, in this paper to justify why that is so, at least to some extent.

Key requirements

I'm lucky enough to get to talk to both IT experts who spend their time successfully building datacentres, and broadcast workflow experts who build broadcast systems. You'd be forgiven for thinking they are talking two different languages! The former will talk about block transfers and byte level mirroring, RAID sets and spindle speeds. The later will talk about numbers of streams, latency and file formats. They still come from two different worlds, and really do see the 'challenges' of storage in two very different ways. Whoever you talk to for your storage, there are a number of key things you need to tell them in order to make sure they deliver what you need:

1. Total **usable** storage. It really doesn't matter what the raw storage is, you want to know how many files you can store. The raw storage (the actual capacity of all the discs or drives) will be

greatly reduced by things like how the files are protected (RAID levels will affect this most), numbers of 'hot spare' drives (see later about protecting your media), any storage used for system data such as file tables. Interestingly, most manufacturers from both sides of the fence will quote you raw storage (presumably because this makes the system look better value for money), so if it's not clear, ask them to tell you how much you can actually store. Beware that most storage is specified in Megabytes or MB, which is 1,000,000 bytes of data. Computers usually work in Mebibytes or MiB, which is 1,048,576 bytes. This may not seem a huge difference, but when you scale up the numbers to many 10's of Terabytes (or TiB accordingly), the numbers can become quite significant, and what you see available in your file system on your computer will likely a lot less than you expected. Make sure you are clear to your supplier which units you are working in

2. Total **usable** bandwidth. This is the practical speed that you can read or write data to your storage system at any time (on many systems there will be a value for read and a lower value for write). Some systems reserve bandwidth for housekeeping tasks such as drive rebuilds, others will simply drop the usable bandwidth when it needs to do this. It's always a good question to ask how much bandwidth a system will have when the storage is 95% full; many lower priced systems will start to perform less well when they reach this point. Ask for clarification on how this bandwidth can be achieved by whatever network connectivity is provided. For example if the system only supports 2 gigabit Ethernet connections, having 10 gigabits of internal bandwidth will be of little use to you
3. File transfer latency. This is all about how fast you can start reading a file on a storage system and how fast you can skip to the next part of the file. This is particularly important for editing systems, where the jog/shuttle performance will be critical. I have been shown systems which have great bandwidth and capacity performance, but when editing, a simple scrub up and down a time line proved the system to be useless. A lot will depend on what media file formats you are using, and how the edit client reads media, but some storage systems can perform clever buffering which will predict what you are going to want to read next and then pre-read that data. Ultimately I recommend that you perform some 'real-world' tests with your media and your editing systems. If it doesn't work for you in a small scale test, it's unlikely to work when scaled up to your enterprise
4. How safe do you need your data to be? Most storage systems will come with a number of choices on how you protect your data in case of a component failure (such as a hard drive failing). These range from complete replication of all data either at a block or file level, to parity protection using various RAID models, to complete backup of data to a separate storage system. Some systems make use of 'hot spare' drives which allow a system to automatically rebuild. Whatever systems are available and are recommended to you, I suggest asking the following questions:
    a. When there is a failure, will there be any downtime or loss of performance from the system?
    b. How long will it take to be protected again?
    c. What is the procedure for replacing a failed part? Is it something you can do yourself or does it need a specialist service engineer?
    d. What is the estimated time between failures on the system specified?
    e. What cost for preplacement parts, or is this included in a warranty or support agreement?

I suspect that a discussion around the above points will have just as big an impact on overall cost of a system as items 1,2 and 3 above.

Volumetric capture

Unless you know what you want to do, and how much of it and when, it's very unlikely that you'll end up with a storage system that provides the performance you need, without over specifying (and hence blowing the budget).

The first thing to start with is the file formats that you are likely to use. This may be one 'house format' at a given bit rate, but you may wish to store a number of different formats for production, distribution etc. A separate calculation will be needed for each format. When working out the stream bandwidth for each format, don't forget to include any audio. I remember once working with an engineer who had specified 50Mbit/s for the video, and ignored the audio as he said: "*it's insignificant compared to the video*". It turned out that the audio added a further 40Mbit/s for this particular customer, so almost doubled the payload, and more importantly halved the number of hours of content that could be stored.

Next, you'll need to work out how many hours of each format you'll need to store for each day, and how long it will need to stay on the storage before it is deleted. Some careful modelling is needed here in order to determine the peak storage requirements, and if the storage is likely to need to expand in the future, and if so, by how much over a number of years.

Calculating the peak bandwidth required is the next consideration. This is essentially how many streams of media and file transfers need to happen at any one time, and how fast. Moving a 50Mb/s video file from one server to another will happen at real-time if it is streamed at 50Mb/s, but if you need it to get there faster, then you'll need to assign more bandwidth for the transfer. If you intend to edit using files directly on your storage, then this can become a complex equation: you'll need to know how many streams each edit client will likely read at the same time (usually the number of tracks, but you'll also need to know how the audio is streamed, whether it is separate to the video or included); you'll need to know how jog/shuttle will affect the required stream, e.g. if you shuttle at 4x speed, does that require 4x the bandwidth for each track, or does the edit client do something 'smart'?; you'll need to know how the render back to storage (both for temporary files and for the final edit) will affect the write bandwidth. Of course your edit clients won't usually all be editing at the same time, but can you afford for the edit experience to deteriorate on those rare occasions when they do? The only realistic way to measure the sort of bandwidth profile needed for editing is to do some real world tests with your media, using your edit clients and your editing staff.

Expanding your business

Once you have a feel for the loads on your system both in terms of Terabytes to be stored and bandwidth required, you'll need to think about what you might need in 6 months' time and beyond. It always amazes me that no matter how large a storage system is, they are usually 95% full within a couple of weeks. If it's there and available, then users will fill it. Of course this can be mitigated if the

storage system allows quotas to be set up, but at some point, usually sooner than you would expect, demands on the system will go beyond your original plan.

Some storage systems allow you to expand both bandwidth and capacity by simply adding more storage chassis (some even without downtime), so all you need is spare rack space, some power and a bag of money. Other systems will allow you only to expand storage capacity, and some not expand at all. It's always worth checking what the expansion capabilities are of a storage system before you invest, and where the 'step upgrades' are: the point where you have to buy a lot more equipment in order to pass beyond a certain threshold.

Every couple of years (as Gordon Moore of Moore's law will tell you), IT capability roughly doubles, and this includes storage devices (both traditional hard drives and SSDs). There is therefore a corresponding drop in cost per Terabyte. This means that after only a few years, your new storage system may have been superseded by the latest technologies, meaning replacement may be more cost effective than upgrade. In this case, you will need to factor in any costs for migrating your data.

Connecting your clients and accessing your data

It may seem obvious, but connection of your storage system to your client workstations and other servers that consume and add media to your storage system is often overlooked. Many apparently low cost IT storage systems are designed to connect to a single server (or a cluster of servers providing one function) and not serve media to a number of clients simultaneously. In this case, you will need extra costly hardware to act as a gateway to your files, which will likely need specialist IT skills for both install and to maintain.

Mid-range storage systems usually come with some kind of gateway or 'head' which clients connect to, which in turn connect to the underlying storage system. Higher end storage will allow multiple heads which in turn will support more bandwidth and hence more streams of media. Some systems don't have a gateway or head at all, but rely on clients and other servers to reassemble data from discs, allowing the full bandwidth of the storage devices to be directly accessed, up to the speed that the network supports.

Most storage systems these days allow clients to connect via IP over Ethernet (usually multiple Gigabit or 10-Gigigabit connections to the storage). If care isn't taken in the network architecture, then this can easily become a bottleneck. Some storage system suppliers will be able to recommend suitable network infrastructure, but with others, you may need to hire experienced network design services, which is another cost that needs to be added to the overall budget.

Data transfer protocols used will greatly affect the performance of actual data transfers, and this will be determined by the operating systems you need to run on your clients and servers. Apple clients work well with NFS (Network File System) and AFS (Apple File System) transfers, where Microsoft Windows works reasonably well with SMB (Server Message Block {similar to, but not CIFS – that's a whole other paper!}) and well with iSCSI (Internet Small Computer Systems Interface), although iSCSI is designed for a point to point transfer and cannot work with multiple clients at the same time. Linux can support NFS,

SMB, and iSCSI, but will likely need special drivers to make this all work. FTP (File Transfer Protocol) can be used very effectively over short distances for file transfers, but won't allow you to stream unless some very clever tricks are played. Higher end storage systems will provide their own optimised transfer mechanisms, which will likely need a special driver on the client to make them work. Always check with the supplier if your OS (and version) is supported. Different storage systems will support a subset of all of these protocols, but the golden rule is, check that all of your clients will be able to access all of the storage in the way you need.

If you have a mixture of operating systems that are going to share the same storage, and need to access the same media files, make sure that this is going to be possible, and that there are not going to be any incompatibilities with things like directory and file naming.

A tiered approach

When balancing functional, space, power and heat requirements, it is often found that different storage systems will be needed for different purposes. For example, storage needed for high-end editing may need to serve many 100s of megabits per second with very low latency to many clients at the same time and is likely to be expensive; archive material may only need to be accessed infrequently by a couple of clients, but will require many 100s of terabytes of storage. This will have a much lower cost per terabyte than the editing storage.

It is therefore often the case that two or more 'tiers' of storage are used to provide for these different purposes. This tiered approach is sometimes used to provide a lower cost per Terabyte backup of a main storage system, however caution should be exercised. A high performing top tier storage system may allow many 10s of Terabytes of data be written to it each day. If this is to be copied to a less performant tier, it may turn out that this lower tier does not support sufficient bandwidth to allow for this amount of data to be copied in the same period of time. Bandwidth will also be required on the top tier (to copy the day from), potentially increasing the overall demand on this system too. The overall cost may therefore be higher than a larger, high performance single tier, and that's before we even think about the cost of file mirroring management applications

R&D and added value

More traditional broadcast manufacturers who supply storage systems will undoubtedly understand your requirements better than a generic IT company.

> *I recall a conversation between a broadcaster and a global IT cloud storage solution company, where it was explained that the files were as large as they were as they were encoded at 220Megabits per second. The IT company couldn't understand why anyone would want anything above 5Megabits per second for HD.*

The broadcast manufacturers will also likely have spent a considerable amount of time testing real-world broadcast workflows as well as optimising their systems for the demands they have specifically

been designed for. If you choose generic IT storage, you'll likely need to bear this testing cost yourself, and will not necessarily be in a position to influence any development to improve the system if needed.

Most storage systems will provide additional 'added value' tools useful to your storage management needs. These include everything from self-monitoring (alerting users when there is a problem), to automated backup (allowing mirroring to other storage systems), to full MAM (Media Asset Management) functionality. The usefulness of these tools will be determined by what you want to use the storage system for, but it is more likely that the broadcast specific manufacturers' tools will provide tools more closely aligned with your broadcast business needs

Build and support

When choosing a storage system, it's important to consider the overall cost of the project and running costs, and not just the capital cost of the equipment. On-site build services range from nothing except an office hours number to call if there is a problem (the office may be on the other side of the world of course!) to full commission and testing. Be clear when ordering services what you really expect. I have been caught out before when ordering a storage system, I was promised a 'full working system'. When the system was handed over to me, I could find no way of getting files in and out, only to be told "the installation of the file system is the customers' responsibility". It took a further week of very expensive specialist knowledge to get the system in a state where I could write a single file, a cost that blew the budget of that job. Ultimately you will want a storage system that performs a certain task, in a certain way. Unless it can be proven at the end of the install process that it does what it requires, then in my opinion, the job isn't complete.

On-going support is another critical differentiator between manufacturers. All systems eventually fail in one way or another, and certainly will need regular maintenance to keep them running properly. The down time in case of failure and the cost (and availability) of replacement parts should be factored around your business requirements. Having an experienced engineer in your city with a case of spare parts will likely cost more than remote telephone support and a three day delivery time of replacement parts that you have to fit yourself. You'll also need to be able to trust the support engineer to look after your valuable media assets and get you back in use with your broadcast workflows as soon as possible.

So…what price your storage, and what I haven't told you

From the above discussion, I hope it is clear that there is no easy answer when deciding which storage system to choose. The price of a system is far more than just the capital cost per Terabyte, but should include performance, upgrade, functionality, connectivity (including network), installation and support costs. Until you have a clear cost on all of these factors, it is impossible to make a good comparison. Beware the IT storage system that looks good on paper when it comes to bandwidth and capacity. This may only be available in certain circumstances and not achievable in your media workflow, leading to a system that wouldn't work in your specific use-cases. Knowing what you need the system to do is important work to do before you even approach storage vendors; the more information you can give them, the better chance you have at assessing their responses.

I've not really mentioned different types of storage technology. SSDs (Solid State Drives) are falling in price and offer great performance at low power; spinning hard drives are available at 6Terabytes or more, so offer good storage 'density' and continue to expand while dropping in cost each year; data tape technology gives zero power consumption for very good cost per terabyte. All of this is important when making your decisions, but only if it supports your requirements and budgets. Storage technology will continue to evolve, and inevitably perform better and better at lower and lower costs. Make sure whatever you choose, that it really is what you need, and that you do a fair comparison with other systems, and don't get too blinded by 'new and shiny' is best!

Advanced Broadcast Integration is a supplier of technology consultancy, design, implementation and training services. With over 25 years of broadcast experience, the company's founder, Graham Collins has the experience to help you with your broadcast workflow projects, including specifying, assessing and implementing your storage systems. Please visit www.abcintegration.co.uk and complete our contact form to get in touch