The Fusion of Storage for Accelerated Archiving

Tom Coughlin
Coughlin Associates

The Fusion of Storage for Accelerated Archiving

Tom Coughlin
Coughlin Associates

Table of Contents

Introduction

The media and entertainment industry is undergoing big changes. These changes will redefine widely used content storage models. In order to provide superior quality and more immersive content, many content creators are generating high-resolution video with higher frame rates, greater color depth and advanced cameras. All of these factors increase both the size of individual media files as well as the total number of media files. The storage capacity needed to manage the raw content from 2K versus 8K is estimated to increase by a factor of 120X! **Table 1** illustrates how increasing resolution of digital cinema files are increasing data rate and capacity requirements[1].

These trends are driving media and entertainment content repositories towards high performance as well as high capacity digital storage. Higher performance storage devices are more expensive on a $/TB basis than lower performance storage devices. Thus there is a trade-off needed between the use of high performance and high capacity storage technologies, and as a consequence, high performance and high capacity storage are often used together.

[1] 2014 Digital Storage in Media an Entertainment Report, Coughlin Associates, http://tomcoughlin.com/techpapers.htm

Table 1. Example Resolution, Data Rates and Storage Capacity Requirements for Professional Media Content

Format	Resolution (width X height)	Frame Rate (fps)	Data Rates (MBps)	Storage Capacity/Hour (GB)
Digital Cinema 2K (4:2:4, 10-bit) YUV	2048 X 1080	24	199	716
Digital Cinema 4K (4:4:4, 12-bit) YUV	4096 X 2160	48	1,910	6,880
Digital Cinema 8K (4:4:4, 16 bit)[2]	7680 X 4320	120	23,890	86,000

This paper will explore the characteristics of various common storage devices used in media and entertainment content repositories and discuss how they can be fused together using intelligent software in order to create a cost-effective archive storage solution. With annual increases in media content of approximately 30% and typical budget increases of 8% or less, there is a need to economize on the storage costs of video content. By combining the best of storage technologies together along with intelligent management software into a unified storage environment, studios, post houses and media and entertainment professionals can meet their growing storage needs even with limited budgets.

Access – How fast and How Often

As content access is often critical to meet demanding deadlines, we will investigate how various technologies can meet this performance/cost tradeoff.

Touch Rate is the number of times every piece of content in a storage device or media is written or read in a given period of time.

Response time is how rapidly a storage system can retrieve content in normal operating conditions. The response time of a storage device depends on the size of the content. Response time, as used here, is the steady state response time in seconds for 100% random input/output operations in the storage system

Access time is the time for a user to reach the first byte of data after a request to reach that data in an idle system.

Most systems will access only a fraction of the bulk data stored in a year, such as 10%. This statement implies that the required touch rate is 0.1/Y.

Note that a touch rate (or file access) greater than one doesn't necessarily mean that the same data is accessed repeatedly, although it may. It can also mean that new

[2] "Super Hi-Vision" Video Parameters for Next Generation Television, SMPTE Motion Imaging Journal, May/June 2012, P. 63-68

data is coming in, which is also counted in the touch rate. What matters here is the amount of data accessed. The touch rate is thus a measure of how much of the total data in a system can be accessed during a given interval of time.

Touch rate vs. response time

These are important factors in optimizing needed storage system performance and cost. The response time measures the time to act on a single piece of content, while the touch rate relates to the time to act upon the data set as a whole. **Figure 1** shows such a chart that includes some indications of various sorts of applications (performance regions) and general trade-offs important to system design.

Figure 1. Touch rate versus response time indicating various types of uses

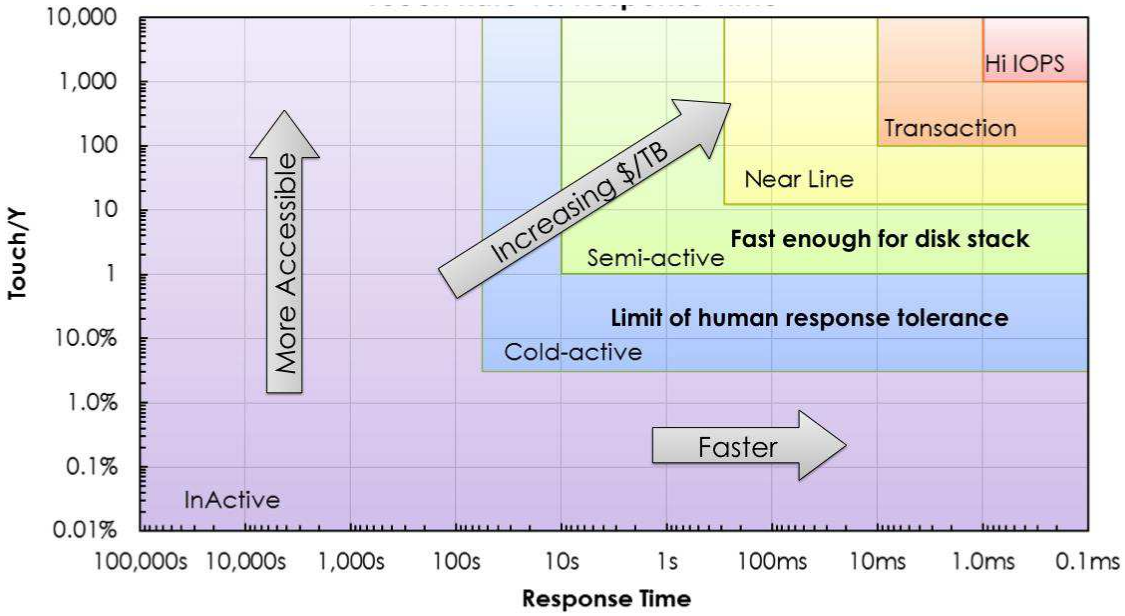

The chart shows annual touch rate on the vertical axis, and response time on the horizontal axis. Shorter response time means data can be accessed more quickly, increasing its value in fast-paced workflows. Higher touch rate means more of the total stored data can be processed in a given time period, increasing the value that can be extracted. In a perfect world, all of our content would be rapidly accessible. But the reality is that system costs tend to increase going to the upper right.

Taking into account the increasing costs as we move to the upper right, organizations that need all of their data rapidly available will see an increase in costs for doing so. But, intelligent blending of storage technologies provides a path to optimized performance and economics. Workflow applications have minimum performance requirements that can dictate storage choices.

The nearline region has performance limits matching the characteristics of high capacity hard disk drives, with access around 1-2 times each month. An example of

this would be broadcast content form several seasons ago that may only require access for a special event or programming.

The semi-active region is defined as having a response time fast enough to not have time-outs when used with an HDD SW stack (typically about 10 seconds). A system with requirements in this region can use HDDs although it would not be using them at their full performance.

The Cold-Active archive region is characteristic of on-line archives. Here, the response time limit is 60 seconds, which is about the human tolerance for data requests. If access takes longer than this, most humans would assume the request failed. This region is appropriate for non-active older content that is infrequently accessed in the workflow but still has value.

For transactional data, such as active production content, fast access is critical. But, this transactional data is a small fraction of the total content a business might own. Many studies indicate that less than 10% of an organizations content is accessed 90 days after it's been created and for long term professional media and entertainment libraries that number is often less than 1%.

That's where the nearline and archive scalability are so important. Unfortunately, traditional archives are mostly offline, cumbersome and don't work with today's fast-paced workflows. As a consequence many people keep everything on more expensive storage because they don't have a good way to manage data movement so they can easily access less used content when they need it.

Understanding Storage Technology

Flash, HDD and tape are used to address workloads in different regions as shown in **Figure 2**.

These storage technology regions have a hockey stick shape due to the way performance limits impact the system performance. On the right side of a region, the performance for small objects is dominated by the access time (time to get to the first byte of data). On the left side, the performance for larger objects is dominated by the data transfer time. In media and entertainment applications, many video files are large objects that can span multiple gigabytes.

A given technology is most cost effective at the upper right boundary (the knee of the storage technology curve), where its full performance can be extracted. It becomes less cost effective as it is used more to the lower left of this knee. At some point, as the requirements move to the left, a lower performing technology is usually more cost effective. This is why tape can be an effective storage solution for large storage objects, such as video files, especially if the time to first byte can be kept at an acceptable level.

Figure 2. Digital storage technologies regions overlaid on the Touch Rate/Response Time chart

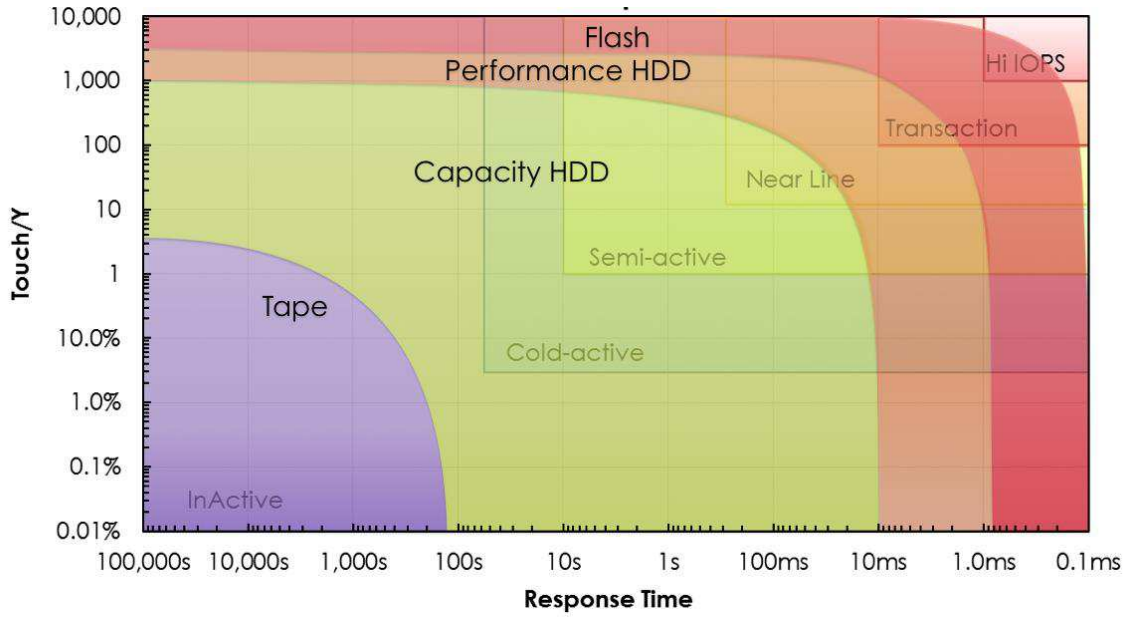

What 4K and 8K content means for file access and storage devices

The file size used by an application influences both the touch rate and the response time. Thus, we will get a curve for the touch rate vs. response time as a function of the object size for various storage technologies. The object size is indicated along this curve.

Figure 3 shows a comparison of data object size vs. response times available for capacity HDDs and a magnetic tape library (here we assume LTO 6 tape with 1 drive per PB of data). This chart shows the differences in our ability to access a given amount of data in a year as well as the difference in available response times for each technology and object size.

The HDD storage has a response time as high as 10s milliseconds while the tape response time is over 100 seconds. As shown in this figure, with one drive per PB of media it takes 2 hours to migrate a 1 TB object (closer to the size of a season of broadcast content) the tape storage to the HDD storage with a tape touch rate of 5. For 1 GB objects (closer to the size of a shorter video clip) it takes 2.5 minutes to migrate the data from tape to the HDD with a tape touch rate of 0.20.

When reading the touch rate chart, the object size chosen (for instance a video project file) should be suitable for the application being considered. Similarly, when comparing different storage technologies for a given application, the touch rate and response time values should be compared at the same object size (such as 1 TB and 1 GB in Figure 3).

Figure 3. Touch/Y and response time comparison for magnetic tape and a capacity HDD

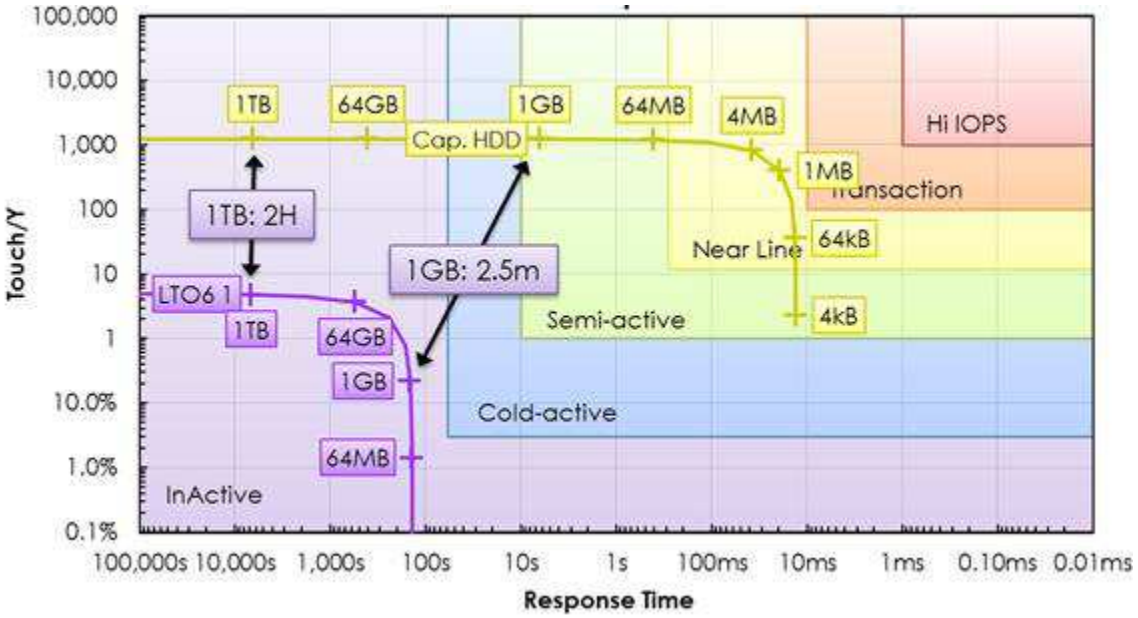

The Fusion of Tape and HDD

As shown in Figure 3, an HDD NAS provides higher performance storage while magnetic tape provides cost effective, although lower performing, storage. In a typical media environment data may go from being more frequently used (hot) to less frequently used (cold) and than back again to more frequent use.

Since the long term costs of tape storage is are at least three times less than HDD storage, using tape for the bulk of the video content storage can result in much lower total storage system costs than a HDD-only archive system. A storage system that seamlessly moves data from the hot storage to colder storage and back again as needed, provides an optimized combination of performance and lower storage costs.

If data is originally put on tape with some cached content moved to HDDs or SSDs to provide fast initial access to that content, then the time to first byte can be kept low. Software management allows fusing the tape and HDD storage into a seamless storage system and avoids some of the issues of working directly with magnetic tape libraries. In addition this software management can provide automatic data protection and disaster recovery since files are always protected on the tape layer.

A New Approach: Automated, Intelligent Data Management with a Hybrid Storage Architecture

The ideal storage software to manage content access should tune performance to match needs in an active archive. It must provide near line HDD storage when needed and lower cost tape archive for long term data retention.

The storage software must support data the ability to automate second copies of data for onsite storage and offsite protection. As mentioned before, caching of content on HDDs and SSDs to reduce the initial time to data is an important software function. It must also be able to expand as needed in terms of the tape storage and front end HDD storage and be able to work between as many HDD and tape vendors equipment as possible.

These capabilities will have enormous benefits to media and entertainment professionals. They can control their costs for long term asset retention while providing the fast access to data required in an active media archive. They can keep remote copies of their content to protect that data in the event that the first copy of the data is incapacitated or destroyed. They can also save on expansion costs by working with most NAS, object and tape storage environments. Furthermore accessible APIs must be able to link the backend tape storage with many applications for ubiquitous access to the archive content.

Sidebar on StrongBox and Dternity

One such solution to intelligently manage data in a hybrid storage archive environment is the combination of Fujifilm's Dternity NAS solution, powered by SDS StrongBox, which provide a cost effective way to create a useful fusion of high performance HDD-based and even SSD-based storage NAS or object store with a cost effective tape archive. Dternity uses intelligent management to meet customer workflow requirements and manages the balance between active and inactive data.

The resulting active archive provides the required performance for modern workflows with cost effective tape storage, in order to reduce the costs of long-term retention of media and entertainment content at the same time that this content is available on demand. By working with any type of application via CIFS, NFS or in an object storage environment, such as S3, this solution is a simple addition to an existing storage environment.

Likewise, the ability to connect with any vendor's tape library via block-bases SAS or Fibre Channel gives users the flexibility to choose hardware solutions that fit professional media users budget and data needs.

The Dternity allows almost limitless capacity increase over time using heterogeneous NAS storage and tape libraries. **Figure 4** shows how the Dternity NAS can be used as an intermediary device moving data from any external HDD or

SSD storage front end to an LTFS tape library backend, providing a central point of management for data and leveraging LTFS tape for economic scalability fused with HDD/SSD for rapid data access.

Figure 4. Dternity NAS System Architecture

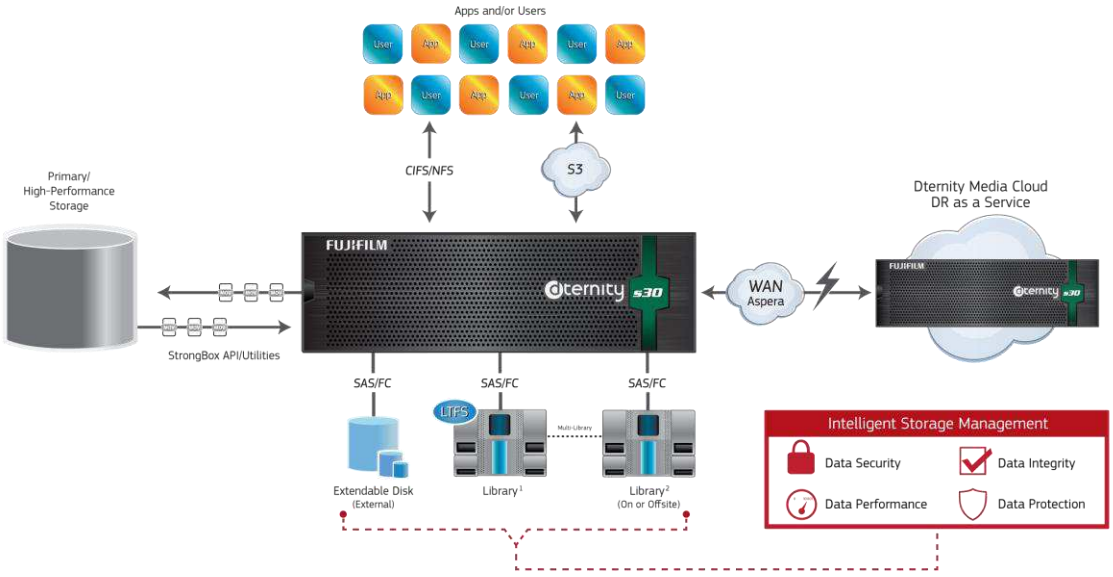

There is no change required to file formats or attributes or applications. Dternity provides built in data integrity, data reliability and data redundancy. It also provides grouping and controlled content access, active directory system access and API and utilities to provide fast file access, pre-fetching and application integration with media and entertainment companies such as AVID.

More on Dternity

Fujifilm's Dternity, powered by SDS StrongBox, comes in two product configurations: the S10 with an MSRP of $22,555 and the S30 with an MSRP of $37,520. Combined with a StrongBox Vseries tape library it can scale from 6U up to 42U of rack space providing from 184 TB to 1.3 PB. This means $125,370 can provide 1.1 PB of native tape storage capacity that can be combined with a NAS or object storage front end.

The cost of Dternity is estimated at $0.003/GB/month at todays HDD capacities (which will be increasing with time). SDS estimates that over 5 years for the same storage capacity the Dternity will cost $180,872 compared to a cost on Amazon Glacier of $625,000. Clearly the Dternity is a cost effective way to create an active archive for a busy media environment. In fact, Dternity works with AVID, Adobe, CATDV, Vitec and other popular media and entertainment platforms.

Case Study: LA Kings Video Archive

The LA Kings is a heavy user of video content, in the arena during a game, for marketing, for online content and of course for TV. With the introduction of HD and 4K the amount of data that they generated has exploded. While they used to shoot just 100 GB per season they are now shooting 30-35 TB of new content per season that needs to be archived. The team's former archive involved buying more HDDs as needed or working with single LTO tapes. This consumed too many hours and too much money on what was a very slow workflow.

Since moving to Dternity the LA Kings have found that they can scale their archive storage as needed, it requires less management and provides much faster workflows that before. In addition it was easy to install. They started on Friday and by Monday the system was ready to go. According to Aaron Brenner, Director of Production for the LA Kings:

"The Dternity has been a huge improvement to our workflow. If an editor's looking for a file or a producer's looking for a specific interview bite, it doesn't matter if it was five years ago or a month ago. We know without a shadow of any doubt that they're able to find that file and just start working with it.

It was important for us to find a solution that was open source, non-proprietary. LTFS is that. If the hardware doesn't exist in 10 years, I know I'll still be able to open this tape.

Not only are we archived but what we're doing is we're taking those tapes out of the Dternity and we're vaulting them in the Dternity Media Cloud. And what that allows us to do is know that "Hey, we've got a copy here and it's safe, and we're able to access it anytime, but, God forbid anything happens, we still have a copy off-site."

As we're looking ahead, as we're looking into the future, the way that we're archiving our content now is important. The 2012 Stanley Cup Championship, 100 years from now we need to be able to access those tapes. We need to be able to retrieve all that data, byte for byte.

Sources

[1] Touch Rate, Steve Hetzler and Tom Coughlin, White Paper, Coughlin Associates, 2015, http://www.tomcoughlin.com/techpapers.htm

About the Author

Tom Coughlin is a respected storage analyst and consultant. He has over 30 years in the data storage industry with multiple engineering and management positions at high profile companies.

Dr. Coughlin has many publications and six patents to his credit. Tom is also the author of Digital Storage in Consumer Electronics: The Essential Guide, which was published by Newnes Press. Coughlin Associates provides market and technology analysis as well as Data Storage Technical Consulting services. Tom publishes the *Digital Storage Technology Newsletter, The Media and Entertainment Storage Report*, and other industry reports

Tom is active with SMPTE, SNIA, the IEEE and other professional organizations. Tom is the founder and organizer of the Annual Storage Visions Conference (www.storagevisions.com), a partner to the International Consumer Electronics Show, as well as the Creative Storage Conference (www.creativestorage.org). He is the general chairman of the annual Flash Memory Summit. He is a Senior member of the IEEE, Leader in the Gerson Lehrman Group Councils of Advisors and a member of the Consultants Network of Silicon Valley (CNSV). For more information on Tom Coughlin and his publications. go to www.tomcoughlin.com.