



White Paper

Crossroads StrongBox to Bring Unmatched Cost Savings
and Flexibility to Media and Entertainment

Broadcasters say StrongBox could “future proof” their digital media vaults

By Debasmita Roychowdhury

Crossroads StrongBox to Bring Unmatched Cost Savings and Flexibility to Media and Entertainment

They may need to access it in five minutes or five decades, but one thing is certain – they have to keep it. Digital media storage managers at broadcast and entertainment giants around the globe say they're continually building a mountain of media files that are steadily increasing the size of their archive. The question is: how do they know their data storage will remain cost-effective and that their media files will always be accessible, usable and portable? Additionally, how do they know their media files aren't forever linked to proprietary archive software? Right now, those questions are unanswered, but not for long.

The media and entertainment (M&E) industry knows their content IS their business. Without it, their business would crumble. Whether it's a sporting event that gets broadcasted or news coverage that requires old footage, the media files that hold all this content are crucial to their day-to-day business. Thus, flexible media file storage and access is paramount. In many ways, the demand for M&E firms to store media files goes far beyond even the rigors of compliance requirements that other industries face. It's not because M&E firms are required to keep their content, it's because their customers demand and expect access to it as part of a stellar entertainment service. Furthermore, for firms that own the rights to content, that content is a company asset with tangible balance sheet impact, as the "asset value" of a studio is based on its archive.

This paper examines the benefits of StrongBox by Crossroads Systems, a new network-attached storage solution that leverages the new, open source (non-proprietary), Linear Tape File System (LTFS) technology in order to address the M&E industry's need for low cost and highly reliable long-term media file preservation. This paper also discusses the storage challenges that M&E firms face and includes interviews with M&E industry insiders, including Media Technology Market Partners LLC (MTMP), MLB Network and NBCUniversal.



StrongBox from Crossroads Systems

StrongBox serves as an online active data vault for media files. It presents itself to the user as familiar hard disk storage, but combines hard disk storage (for fast data storage and retrieval) with high-capacity data LTFS tape storage (for cost-effective, long-term preservation of data). StrongBox shatters the perception that media file storage on data-tape will forever remain proprietary, inflexible, high maintenance and expensive.

Designed to address the need for low cost and highly reliable long-term media file preservation, StrongBox integrates a LTO-5 tape library, LTFS, and HSM-like software with a front-end NAS for standard NAS (CIFS/NFS) mounts. StrongBox presents itself as traditional network-attached storage (NAS) while leveraging LTFS formatted LTO tape systems.

With LTFS and StrongBox, users can randomly access their media files from the tape media. No application modifications are required and a persistent view for all files on tapes stored in the library is finally possible.

A brief summary of StrongBox attributes includes:

- Plug-n-Play CIFS/NFS file share(s)
- Non-proprietary data vault
- Persistent file view
- High-performance system write and read capacity
- Fully portable, non-vendor lock in
- Policy-based data management
- Self healing
- Auto migration to future LTO generations
- Seamlessly scalable to 5 billion files

By incorporating caching and buffering capabilities as well as data management policies, StrongBox delivers a new level of performance for tape systems, including:

Caching and Buffering

- **The Ingest Buffer** – As files are delivered to the system (ingested), they are first written to disk. The files remain on disk until all data movement policies have been applied. Files are always moved to tape. The number of tape copies made of the file, either one or two, is specified by policy. A replication policy may also need to be applied. Upon completion of all data movement policies for a given file, the file's data storage will be reduced to only that required for the Read Cache.
- **Read Cache** – After a file is moved to tape, the first portion of the file is retained on disk. When a user or application requests a file the initial reads are satisfied from disk, allowing time for the tape to be loaded, after which reads are done from the tape. Currently Read Cache size is fixed at 512KBs per file from the factory (tunable in future product releases).
- **File Cache** – Files read from tape are kept on disk on the assumption there may well be repeated reads of the same file. The size of the File Cache is a system configuration option and set at 2.6 TB from the factory. Files are aged out of the File Cache using a least-recently-used algorithm.
- **Read Buffer** – Files being read from tape are read in at full tape speed and stored in memory to be delivered to the client at the speed of the client's application and hardware (over the 10 Gbit link).

Crossroads StrongBox to Bring Unmatched Cost Savings and Flexibility to Communications, Media and Entertainment Archives

This frees up the tape drive for other use sooner than would be the case if the tape were only read at the speed of the client's requests. The Read Buffer is sized at 60 GBytes, which will handle multiple parallel read operations, allowing tape operations to occur even if the network link is saturated. The File Cache (disk) will be written to in parallel as a secondary process; however, the client data path takes priority and will not be delayed (due to potential disk performance limitations).

Data Management Policies

- Grouping by share helps to group all data from one share into a tape or a set of tapes for security and fast data retrieval.
- Pre-fetch policy enables the retrieval of files from tape in advance on disk to ensure quick and fast access to the data.
- File retention policy retains files for a user-configured time on the disk for fast retrieval. These files, like all StrongBox ingested files, will also be copied to tape for archiving and data protection, even though they remain on disk for a specified period of time.
- Optimization of data writes to tape prevents the thrashing of drives and efficiently utilizes the tape drive resources.

ME Media File Storage: Today and Tomorrow

Not only are M&E businesses charged with storing media they create today, they're also in the midst of digitizing and managing older non-digital and non-file media as well as the media that will be shot with new, ever-increasing resolution digital image capture devices. David Trumbo, principal at Hollywood, CA-based Media Technology Market Partners LLC, explains that if specific content has revenue generation potential, it's going to be protected and kept accessible. Even the content that is designated as having only potential value needs to be kept and made easily accessible.

"That's what makes the M&E sector so unique – they have to keep their content," Trumbo said. "In the digital media realm there really is no other solution besides data storage. Everything that has already been created and everything that's currently being created has to become some type of storable, and hopefully managed, data file."

With the ever-increasing amount of content, cost becomes a larger and larger concern. Trumbo continued, "The industry is very sensitive to storage costs with high storage demands that will only keep climbing at monstrous rates. Tape storage is the only affordable option. An all disk solution is out of the question."

With these large storage amounts and extremely large file sizes, a master plan that helps drive an organization's methods and processes for storing data is a must. But proprietary management software and different systems within one firm can quickly allow consistency to fly out the window. Yet, consistency is the first element of a true master plan that ensures all media files are safe and usable.

"You'll frequently find that each department and even each sub department have their own system," Trumbo said. "However, in the last 10 years, we've started to see media-specific asset management and practices come to light." But even as consistent systems for storage are adopted within an organization, Trumbo said M&E firms will still retain the same top four concerns:

Greatest Concerns Surrounding Media File Storage in the M&E Market

1. What kind of value are they getting from their investment in tape and storage infrastructure?
2. What is the ease of operating and deploying the system, and is employee training difficult?
3. What will throughput and performance be like?
4. What will happen if I don't use the latest version of LTO tape?

Media Storage Requirements Soaring

Jim Miles, project manager for the Digital Media Systems Group at NBCUniversal, manages a vault of content that is always online and is expected to have its media files available at a moment's notice. The only exception is content that is not yet been digitized. Therefore, nearly all media files that NBCUniversal creates and stores need to be continually available.

"I always tell people that it's like we're building a mountain," Miles said. "Every year the mountain gets bigger and bigger. Right now, I might have 10 years of content stored, and five years from now it is 15 years worth of content. It's really important to archive only what you need, but what we need is very large."

NBCUniversal's content is minimally compressed with a typical yield of 1.9 Terabytes on an LTO-5 tape. And the size of NBCUniversal's media files are enormous:

Crossroads StrongBox to Bring Unmatched Cost Savings and Flexibility to Communications, Media and Entertainment Archives

- Studio recordings can easily reach 100 Gigabytes
- 10 minutes of cut stories and B-roll footage are typically 10 Gigabytes and larger
- Long-form content varies greatly from 10 Gigabytes and larger

NBCUniversal employs a team of more than 24 playback and storage specialists to manage its media files. Miles said their system is almost exclusively tape and that speed is their largest concern. He also expects read/write access rates to run at native device speeds. "For us, the speed at which we can obtain footage is crucial," Miles said.

For NBCUniversal to keep all of its content on disk would require an enormous data center that would be far too expensive to operate. With data tape as the only viable option for storing such large amounts of data, Miles said flexibility is key. "It's really important for us to future proof our use of data tape," he said. "When we are forced to make changes or go through some sort of migration, we want to be good to go."

StrongBox Benefits in the M&E Market

The vast majority of current data tape archive systems used in the M&E industry wrap multiple files of content together typically using a TAR function (which can cause huge latencies to ingest or access files to and from tape) or some proprietary vendor format locking access to the content through that vendor's system. StrongBox offers fast and random access to any specific file on data tape and does it in a non-proprietary manner using open source LTF5 technology. It's unwrapped. Users don't need a copy of any proprietary software whatsoever to utilize the files managed and stored in StrongBox.

Trumbo said, "StrongBox has a built-in disaster recovery component. Let's say you're in a large facility with a big network and one day something goes wrong. If their content is sitting in a proprietary system, they may be in trouble. And at a minimum, they're locked into something they may not want to continue using. With StrongBox, they don't have to be stuck buying this and that."

If the organization had deployed StrongBox, they would know exactly what tape contains the files they needed and find that tape. After all, Strong Box is not the tool manipulating files. It securely stores the files for some tool or function that does manipulate and add value them. It allows someone to put files where colleagues can get to them quickly – all without any third-party software that locks users into something they may need to change down the road.

By leveraging the 'open' nature of LTO technology, and the open source, self describing format of LTF5 technology,

StrongBox delivers a truly enterprise-class data vault that writes and reads data in an open, non-proprietary and fully portable manner. It provides total freedom from any OS or application with no specialized agents required, resulting in the ability to share data among disparate systems, departments and even businesses. By not altering the file format, attributes, applications, or adding proprietary software layers, StrongBox operates as a fully vendor neutral archive.

Miles at NBCUniversal said StrongBox would open up many doors for his team. "StrongBox has tremendous value," he said. "It offers alternatives in a large archive infrastructure and allows the smaller workflows to work the same as the larger ones. With flexibility to grow, it's very applicable for our secondary workflows. For example, one of our sister units has smaller content and smaller teams, but they still need to store things like we do right now, and they don't. With StrongBox, this would be easy."

Although StrongBox makes data portable, Miles said that's not the only major benefit. For his organization, he said, it's the ability to have another vendor read their media files later on, opening up future opportunities for him.

Tab Butler, director of media management and post production at MLB Network, also reiterated that the flexibility StrongBox would bring is of great interest to his company. As one of the largest consumers of LTO-4 tape, MLB Network, in partnership with Major League Baseball Productions, share a broadcast facility that telecasts, licenses and archives the footage of every Major League Baseball game at 30 fields throughout the United States.

"The sheer number of games and the volume of media files that we're dealing with becomes astronomical," Butler said. "There can be up to seven different recordings of one single game, as we record and log every shot of every game, with every pitch being catalogued. In 2010, we placed 60,000 hours of content in the archive, and this year we archived over 70,000 hours. With this much content, a storage system that scales and reduces cost is key."

Prior to MLB Network's founding in 2008, Major League Baseball Productions' archives contained more than 150,000 hours of film and video tape recordings. That content is currently in the process of being transferred to digital media files. Today, in an average week, MLB Network records more than 2,500 hours of content. Fifteen to 20 Terabytes of video content is added daily to the archive, which consumes 20 to 25 LTO cartridges. Each LTO-4 cartridge holds between 34-36 hours of content. MLB Network and MLB Productions archive system has more than 10,000 LTO-4 tapes in the archive, with each holding approximately 800 Gigabytes.

“Technology at MLB Network is a snap shot in time,” Butler said. “I have hardware and software that was available in 2008 that I would not have bought in 2010. And this year I know I’ll say the same about what I’m buying now. I want to keep advancing the infrastructure and our management demand more and more. As a result, flexibility and open systems are paramount. All of our infrastructures have to be flexible and modular in such a manner that we can easily migrate to the next best thing. StrongBox would make that possible.”

Butler continued, “Our current archive and proxy systems will eventually become outdated. Right now, it costs service, support, power and air conditioning. It has a high level of cost. But what do I do? I see StrongBox as being the type of tool that regardless of where and when I migrate, it has the potential to fill that void. StrongBox could help us further strengthen our foundation and continue to turn data on a dime when we’re asked to. StrongBox would make it easy to retrieve any content we need, at any time, for any platform.”

Bottom Line

The M&E market stores huge amounts of content and needs nearly all of it to be affordably accessible for the foreseeable future. Today, migration from one storage system to another can be an expensive and massive headache, if not impossible. The proprietary road blocks that an organization can face are often overwhelming and cost prohibitive. StrongBox changes this completely by providing a significantly less expensive system that puts the user, not their software vendors, in true control of their data vault.

About The Author

Debasmita Roychowdhury is Senior Product Manager at Crossroads Systems, Inc. For more information on StrongBox or to schedule a demonstration, contact Crossroads Systems at 512-349-0300 or info@crossroads.com.

About Crossroads Systems

Crossroads Systems, Inc. (NASDAQ: CRDS), is a global provider of solutions and services that ensure stored data is proactively protected and reliably recovered. Crossroads offers organizations powerful data protection, proactive data security, intelligent storage connectivity, unmatched performance, and significant cost savings. Founded in 1996 and headquartered in Austin, Texas, Crossroads holds more than 100 patents granted and pending and has been honored with numerous industry awards for innovation in data protection and storage. Visit www.crossroads.com.

www.crossroads.com/products/strongbox



11000 North MoPac Expwy. Ste. 100 Austin, Texas 78759 866.289.2737 512.349.0300 sales@crossroads.com

© 2012 Crossroads Systems, Inc. - Crossroads and StrongBox are registered trademarks of Crossroads Systems, Inc. Crossroads is a trademark of Crossroads Systems, Inc. All other trademarks are the property of their respective owners.

SB_WP_122711