

Active Archive: Front-End File System for Tape and Disk



April 2010

Contents

Active Archive: Tape with a File System Front-End	3
Tape and Disk in Parallel	3
Reliability.....	3
Performance/Speed	4
Capacity and Affordability.....	5
Active Archive Implementation	6
Active Archive Advantages.....	7
Tape Libraries' Role in Active Archive.....	7
Next Steps	8


Active Archive: Tape with a File System Front-End

Tape has come back into its own, with the advent of new hardware and software that lets tape act like disk. This completes the cycle, with tape initially used like disk for mainframes. In that context, tape served as storage of data to be acted on, and was loaded into the computer only for manipulation. In the 1960s disk became an affordable technology that was used in place of tape for storing data. Tape moved to the role of data backup medium given its long archival life, low cost, and low maintenance, low-energy use. Then disk was used as tape by way of VTLs. Now the cycle is complete, through the use of active archive applications that put a file system in front of tape.

The latest change in the cycle of disk/tape usage started in about 2004, when VTLs entered into wider use for backup. Disk-as-tape moved disk into the spotlight. Disk steadily continued its role in data storage, with less expensive SATA disk providing relatively cost-effective short-term storage.

Then disk vendors started pushing disk as the only storage medium organizations require, given its relative simplicity once it was installed on a site. But the

economy, disk's true high cost in terms of initial purchase price and over the life of the product, the intensive power use disk demands and associated costs, and the green IT movement with an emphasis on reducing power use all helped to focus attention back to tape.



Tape and Disk in Parallel

The disk incursion into the storage space continued during a period of enormous advancements in tape that improve tape's reliability to match and exceed that of disk, and also exceed disk's capacity and speed.

Reliability

Tape's improvement of 700% over the last ten years¹ in terms of reliability is now coupled with additional features that ensure tape usability. Contemporary drives already perform automatic read-after-write data verification along with error checking as the data itself is written to tape. Some backup software applications run a second read-after-write verification. These three levels of confirmation of data integrity are just the start. The undetected error rate on LTO media, for example, is 10^{-15} , while

¹ Beech, Debbie. "Best Practices for backup and long-term data retention" Sylvatica Whitepaper. The evolving role of disk and tape in the data center. June 2009.


CERN testing found that disk “error rates are at the 10^{-7} level.” The report goes on to say that “the data loss rate would be much higher for compressed files.”²

This confirms that LTO tape technology is orders of magnitude more reliable than disk. The archivists at the Library of Congress continue to use tape as the most reliable and longest-lived medium available today.³

Library automation has shored up tape’s already high reliability by providing new tools. For example, the Spectra T-Series libraries provide additional layers of assurance concerning data’s integrity on tape.

Among the tools now available through the library interface:

- **Media Lifecycle Management** tracks tape health, letting IT staff identify failing or ailing tapes before data is at risk. Benefit: substantially increase backup/recovery effectiveness.
- **Drive Lifecycle Management** tracks drive health, so administrators identify, early in the process, drives that are having issues. This data lets administrators address these issues early, for example through a cleaning or replacement.
- **Global spare tape drive** parallels the spare disk in a RAID set, by providing a tape drive that is available as a spare to be used in case a tape drive is unavailable for whatever reasons.
- **ScanTape** parallels the scan-disk function, scanning a tape to ensure it is in good shape and the data is healthy and will be readable when needed potentially years into the future.



Performance/Speed

In active archive, determine how fast you need your data. Putting heavily/frequently used data on disk just makes sense. For millisecond retrieval, use the highest end disk. For data that is less used, SATA/SAS disk is a good fit, with retrieval times of one or two seconds. That leaves the data that is rarely used—and that, prior to active archive, didn’t support real-time retrieval. This is where the active archive file-system on tape pays off: you can get at the data—it requires a delay of a minute or two on average, but that’s often tolerable for data that in every other environment can’t be accessed without a much longer wait.

Note that storage media performance does not scale directly. Accessing or writing one file takes advantage of disks’ random access method of read/write; with a quantity of data, tape’s serial streaming is the best pick. For handling a lot of data, tape drives are fast. At streaming rates of 140 MB/s (with compression 280 MB/s), tape is the king of the speedway for writing and recovering large quantities of data. SATA storage disk that adheres to the current generation specifications: “SATA 3Gb/s [can

² Panzer-Steindel, Bernd, “Data integrity,” CERN/IT Draft 1.3 8. April 2007.

³ *Focus on Preserving Our Digital Data Workshop*, National Institute of Standards and Technology (NIST) Gaithersburg, Md., 3/29-31.

support] sustained throughput rates of 250-260 MB/s.”⁴ A single LTO-5 tape drive, running with compression, exceeds this speed with 280 MB/s.

Even the next generation SATA 6 Gb/s technology supports only “faster transfers of short bursts of data [at] 580 MB/s.”⁵ Tape is still faster, given the Fibre Channel LTO-5 drives burst transfer rate of 800 MB/s.

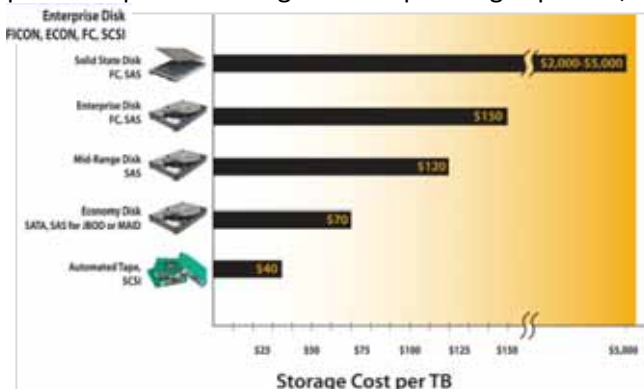
Capacity and Affordability

LTO tape, along with the IBM and Oracle/STK proprietary tape formats, continues to provide high capacity in a very small area. When compared with disk, LTO is significantly denser: “LTO-5 [...] yield{s} a total density by volume of about 107GB per cubic inch. The largest standard internal 3.5-inch SATA disk [has] a density of about 86GB per cubic inch.”⁶

Right now, LTO-5 can store 1.5 TB native /3 TB compressed, all in a half-inch cartridge. All that data requires no power whatsoever after it’s written to tape, reducing long-term power costs when compared to spinning disk, which always demands power. The savings in power alone is enough to remember that tape plays a key role in the data center –particularly when looking at storing large quantities of data that need to be accessible and searchable online for many years into the future.

Tape⁷ costs remain the lowest both in terms of purchase price and longer term operating expenses, independent of statements made by some disk vendors.

Experts in the storage industry predict that this will continue as new generations of tape and disk are released. It is worth noting that, in the storage cost per TB chart, the disk *and* library costs per unit of data assume working systems, to take into account otherwise hidden infrastructure costs. Typically, disk



vendors include only the cost of a single gigabyte of disk without including the cost of controllers and other components required to actually use a gigabyte of disk. When using this method to determine tape cost per gigabyte, the value drops to about 4 cents per gigabyte for an LTO-4 cartridge.⁸

⁴ Serial ATA International Organization (SATA-IO) “Fast Just Got Faster: SATA 6Gb/s,” *White Paper*, May 27, 2009. www.sata-io.org

⁵ Ibid.

⁶ Prigge, Matt. “Tape is dead! Long live tape!” *Infoworld*, January 25, 2010. [Hhttp://www.infoworld.com/d/data-explosion/tape-dead-long-live-tape-090](http://www.infoworld.com/d/data-explosion/tape-dead-long-live-tape-090)

⁷ Moore, Fred. *Storage Navigator*, Horison Information Strategies, 2008.

⁸ Beech, Debbie. “The Evolving Role of Disk and Tape in the Data Center,” *Sylvatica White Paper*, 2009.

Active Archive Implementation

Tape provides some definite advantages, in light of:

- Data growth
- Tape's high speed and reliability
- Tape's low cost and energy use

This leads to the natural conclusion that tape can, and should, be used for storing a primary copy of data. This requires the right environment and the right software.



For years, specialized markets have used a file system on tape. Notably, the high-performance computing (HPC) market—a.k.a. supercomputing—has always had to balance cost and access time given huge quantities of data that always need to be accessible. One of the key applications used by HPC sites is the High Performance Storage System application, HPSS. The application has been jointly developed over a decade through collaboration among five Department of Energy laboratories and IBM, along with universities and other laboratories worldwide. It supports a “Virtual File System Switch (VFS) to present a standard POSIX file system interface.”⁹ HPSS has long utilized both disk and tape as the storage for the VFS.

The broadcast market uses specialized digital asset management applications that use a file system on tape. Tape's serial nature and speed of data retrieval makes it an ideal match for downloading and uploading large media files rapidly.

Proprietary archival systems are also available—but they are proprietary, typically requiring, or strongly recommending, a single vendor's hardware and software. For example, Oracle/Sun hardware supports SAM-QFS (Storage and Archive Manager-Quick File System). Only in the past year or two have these proprietary solution started to look to open up compatibility with other products that may be used in a customer's environment.

New applications such as File Tek's StorHouse and QStar's Archive Management provide an affordable open-systems platform for active archive, with a single file system that extends across disk and tape. The software runs on standard operating systems, brings much lower costs and much broader compatibility across multiple hardware and software platforms. This provides straightforward and affordable access to archived data while leveraging both disk and tape for the storage of the primary copy of the file data in its native format. No longer is it necessary to go to an offline file that is no longer searchable or accessible in the original file format. Restores from both disk and tape can be a thing of the past when needing to access archived data.

⁹ *HPSS User's Guide: High Performance Storage System*, Release 6.2, Revision 2.0, July 2008.

Active Archive Advantages

Active archiving addresses multiple contemporary concerns, including:


- Rapid access to all data that is increasingly driven by regulatory and organizational mandates
- Compliance and e-discovery through intelligent use of tape storing data that can be quickly accessed
- Organizational efficiencies through the opportunity to access the all of the organization’s data readily
- Savings through appropriate dispersal of data to the proper tier, freeing expensive primary disk for more appropriate data storage and so reducing the need for more expensive disk
- Savings through power reduction given the now straightforward task of moving less-used data to tape while retaining access to it
- Space savings through the use of high-density tape libraries that can be scaled with minimal effect on footprint, while storing accessible data
- Reduced backup windows through the removal of unused data from backed up disk

Tape Libraries’ Role in Active Archive

Currently, a typical storage environment includes primary disk, SATA for secondary or VTL storage, and tape for backup and off-site storage, as illustrated in an adjacent diagram.


By adding Active Archive to a storage environment, you can affordably and rapidly expand data center storage while using less disk and footprint –and add the ability to access all of your data, with a file system that spans both disk and tape. The tape automation now serves dual functions—active data storage and backup copy creation for disaster recovery. Further, with the movement of data to tape, disk volumes and backup times are concurrently reduced.

By making greater use of the tape library, with its highly dense storage and low energy demand, you can expand access and reduce costs. Instead of adding primary or secondary disk to handle data, add tapes to store the older data and make better use of disk for frequently accessed data.



Spectra T-series libraries provide scalability, high density, and the lowest power use compared to competing libraries, making them superior candidates for Active Archive environments.

The T-series rich feature set also serves well in an active archive environment, letting you track drive, media, and library health. The Spectra BlueScale features make it particularly easy to identify problem tapes and drives through the use of icons that display on the graphical interface, which is on the library's front panel and available from anywhere through a web browser.



The Spectra T-series libraries offer a feature that is central to active archive in ensuring data integrity over the life of the data: ScanTape. ScanTape is a set of background processes that provide two levels of verification for LTO data cartridges, PreScan and PostScan. PreScan checks newly imported tapes to make sure they are ready to use—for example, in good health and not write-protected.



PostScan scans each LTO cartridge and confirms that there are no media errors on the tape by reading the entire length of the tape up to the end of the recorded data (EOD). The PostScan process is performed by the library independent of the backup application normally used to read and write data to the tape. The results of the scan are stored in the library's media-health database. Set triggers to check the health of tapes over an interval of time, to ensure that the data is valid. You can also verify a specific tape's data integrity manually/by request.



Over data storage of weeks, months and years, Spectra T-series ScanTape feature is a critical attribute of a successful file archive, ensuring data integrity over time.

Next Steps

The Active Archive Alliance is a vendor-neutral industry group that helps bring together information that customers need to efficiently put together an active archive system that takes advantage of existing data infrastructure and to select the right application/hardware. Founding members of the Active Archive alliance include:

- Spectra Logic: data protection leaders with disk and tape library products, designed and manufactured for active archive

Active Archive

- File Tek: active archive application products
- Compellent: tiered storage and active archive hardware and software products
- QStar: software archiving packages that view media as a single pool, moving static reference data to secondary archive storage technologies such as tape

By promoting best practices, recommended standards, and improving access to information about data archive the alliance can support the implementation of active archive, and its savings, around the world.



www.SpectraLogic.com

Spectra Logic Corporation
6285 Lookout Road
Boulder Colorado 80301 USA
800.833.1132
303.449.6400

For more information, please visit www.SpectraLogic.com