



TAPE IT

HP OEM NEWSLETTER - February 2011

Contents

- Digital Archiving provides a solution to the data growth challenge
- The benefits of Digital Archiving
- LTO Tape – the Digital Curator of the information age
- Target applications and vertical markets for Tape Digital Archiving

Tape: well positioned to meet the needs of the exploding digital archiving market

IDC's 2011 Storage Trends Report predicts a resurgence of interest in tape storage; "large volumes of data to be stored for long-term archiving will drive users to leverage tape as a less expensive alternative to disk for infrequently accessed data".

"The pace of data creation is not expected to slow down, creating new challenges for organizations large and small. The data needs to be stored, which implies more disk drives, more data center space, and more energy. Once stored, it needs to be protected, and as the data sets get bigger and bigger, using traditional methods for data protection become cumbersome.

Data created and stored represents some intrinsic business value, and there will need to be a way to extract that value in a consistent and timely manner. Organizations will buckle under the weight of the data and costs associated with storing and protecting it unless there is a cost-effective way to store digital assets. Storing such data on disk quickly becomes inefficient and complex.

The market (including many cloud-based storage service providers) has recognized that tape must play an important role in providing a way to store data efficiently and effectively."

IDC Report, Top 10 Storage Predictions for 2011

Digital archiving provides a solution to the challenge of data growth


A recent Gartner¹ survey shows data growth as the largest data centre infrastructure challenge. More than half of the 1,000 survey respondents are planning to expand capacity at their existing data centre sites by the end of 2011. The Gartner survey findings further suggest that data archiving and retirement projects are the most popular response to this data growth challenge with 62 percent of the global respondents reporting that they will begin investing in such projects or continue their current investment in them through year-end 2011.

The capacity of archived electronic information in the commercial and government sectors worldwide is already huge, and is forecasted to grow significantly. ESG's research² forecasts that

the total worldwide cumulative digital archive capacity will grow to 300,000 Petabytes by 2015, driven predominantly by the increase in file based unstructured archive data.

ESG's study shows that tape will remain a significant storage media for these expanding digital archives. According to ESG, tape currently commands a 38% share of the overall digital archive volumes and is set to experience a six fold increase in digital archive petabytes stored from 2010 through to 2015³.

Forecast of Total Worldwide Digital Archive Capacity on Tape 2010-2015 (in Petabytes)



Source: Enterprise Strategy Group 2010

¹ <http://www.gartner.com/it/page.jsp?id=1460213> ² ESG Research Report, Digital Archive Market Forecast 2010-2015, July 2010
³ <http://www.enterprisestrategygroup.com/2010/12/nesc-proving-tape-as-cost-effective-and-reliable-primary-data-storage>

Defining Digital Archiving

So what is digital archiving and how does it differ from backup? ESG⁴ defines digital archiving as:

“The long-term retention and management of electronic information that has been purposefully retained to satisfy records management, data management, regulatory compliance, or litigation support requirements. Archived information differs from backup data in that backup data is typically a temporary copy of a data set that is ultimately overwritten, while archived information is moved—not copied—from one system to another and is often a permanent copy of a record or data set that is stored without alteration or deletion for a specified period of time.”

In other words, digital archive capacity represents historical digital assets that have been moved and/or retained to satisfy records management, data management (e.g., removing inactive or infrequently-accessed data from production systems, improving end-users’ access to information, etc.), regulatory compliance, or litigation support requirements.

The following table⁵ compares the key characteristics of backup and archive:

Backup	Archive
Copies data for protection	Moves data from the disk to another media for longer term storage
Supports operations and data recovery	Supports business and compliance
Supports availability	Supports operational efficiencies
Shorter-term storage	Long-term storage
Data typically overwritten	Data typically secure, not overwritten
Poor solution for regulatory compliance	Ideal solution for regulatory compliance
No historic relevance	For historic information purposes
Difficult to search	Easier to search



⁴ ESG Research Report, Digital Archive Market Forecast 2010-2015, July 2010 ⁵ http://www.activearchive.com/common/pdf/TimeValue_WP_US_4229.pdf

The Benefits of Archiving

Digital archiving serves to reduce overall 'data footprint', and in particular the amount of data that is being stored on online primary storage.

More Efficient Operation

- Backup can be completed within the allotted time, avoiding impact to application availability
- Removing inactive data from primary storage enables faster backup and quicker restore
- Removing inactive data from primary storage helps improve application performance

Reduce Business Risk

- Manage the legal, regulatory and business risks associated with information
- Retain, protect and secure business critical information
- Adhere to internal records management policies
- Reduce the costs of e-discovery and litigation
- Index, search and retrieve business records with accuracy and efficiency
- Ensure the records are unaltered and undeleted

Cut Costs

- Free up disk space, and delay the need to purchase more primary storage
- Move inactive data off-line, saving on power and cooling
- Move inactive data to media that has a lower cost per Gbyte

As the digital universe grows by an order of magnitude, organisations will have to fundamentally adjust the way that they deal with information to maximise efficiency and reduce costs. Jon Toigo, CEO and Managing Principal of research firm Toigo Partners International, suggests that one part of the solution is to actively archive data from online 'production disk', to a lower cost solution such as tape:



“40 percent of the planet's data should be archived according to our research, and the best archiving practices encompass solutions that offer high capacity and high data migration speed. Data archiving is no longer an option but has grown into a strategic imperative at the forefront of data center planning.”⁶

⁶ http://www.activearchive.com/analyst_view

Active Archiving: putting cold data on ice

A consistent message from industry experts is the need for active data management as the first step in controlling storage costs. This begins by considering the value of the different types of data that an organization must store, and how that value changes over time. The 'Time Value'⁷ of data means that when data is fresh, it is typically more valuable. The value of some data decays over time (e.g. old receipts, outdated reports, etc). The value of other data increases over time (e.g. NASA footage of the first men walking on the moon). In addition to

understanding the time value of data, it is also important to understand whether that data needs to be active for instantaneous access, or whether a delay of a few minutes, or hours is sufficient. A key discovery often made when analysing stored data is that a great deal of the data is much older and less frequently accessed than originally assumed. A 2008 study by the University of California, Santa Cruz reinforces these findings.

MEASUREMENT AND ANALYSIS OF LARGE-SCALE NETWORK FILE SYSTEM WORKLOADS OF CALIFORNIA, SANTA CRUZ⁸

The study was based on an analysis of two large-scale network file system workloads. The team measured CIFS traffic for two enterprise-class file servers deployed in the NetApp data center for a three month period. One file server was used by marketing, sales, and finance departments and the other by the engineering department. Together these systems represent over 22TB of storage used by over 1500 employees.

KEY FINDINGS:

- Files are rarely reopened once created and closed
 - 66% are reopened only once
 - 95% are reopened fewer than 5 times
- File reopens are temporally related – over 60% of reopens occur within a minute of the first
- A small fraction of clients account for a large fraction of file activity – 1% for 50% of file requests

CONCLUSION:

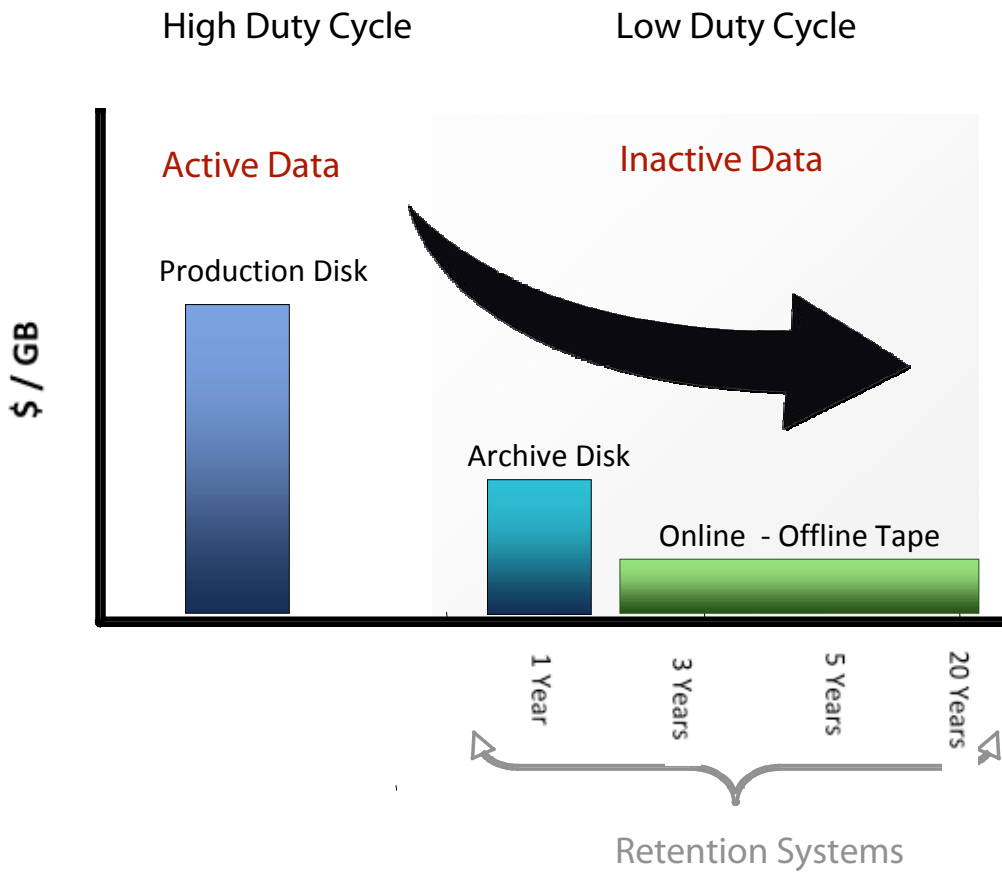
Over 90% of the active storage was untouched during the study – the data is cold

While keeping all data close at hand on high speed disks might seem ideal for access purposes, in reality to do so could be prohibitively expensive in terms of both hardware purchases and the cost of power, cooling and physical space, especially when compared with tape storage.



⁷ http://www.activearchive.com/common/pdf/TimeValue_WP_US_4229.pdf
⁸ <http://www.ssrc.ucsc.edu/Papers/leung-usenix08.pdf>

Migrating inactive data to lower cost storage frees up primary storage, and storing inactive data on lower cost storage saves money.



LTO tape - the digital curator of the information age

LTO tape technology is expanding its role from a pure backup solution to that of a premier long term storage technology and archive. To quote one storage analyst **“The 21st century data explosion is here – and tape is well positioned to become The Digital Curator of the Information Age.”**⁹ Tape is strengthening its position as the ideal solution for data archiving as a result of the following features and benefits:



⁹Tape: The Digital Curator of the Information Age - Fred Moore, President Horison, Inc.

Lowest TCO

Tape acquisition and on-going operating costs remain very low. At approximately 6 cents per Gbyte (native) for LTO-5 tape media¹⁰, LTO Tape offers one of the lowest costs per Gbyte for long-term storage, particularly when factoring in energy costs. Tape storage is cool – literally - and has been shown to decrease storage power requirements by 99% when compared to disk-based storage. Tape helps meet the goal of many data centers that inactive data should not consume energy.

In a new report, published in December 2010, The Clipper Group¹¹ Inc studied the Total-Cost-Of-Ownership (TCO) of using disk or tape to archive large binary files with a 45% annual growth rate over a 12 year period. The disk comparison point began with 2TB disks assembled into fully-configured arrays. The tape comparison point began with LTO-5 tape

cartridges and drives housed in a fully-configured automated tape library. In conclusion, The Clipper Group discovered that:

- Disk is more than fifteen times more expensive than tape for archival.
- Disk uses 238 times more energy - costing more than the total cost of the tape solution.

“In every dimension, the TCO of the tape solution was found to be less expensive than the TCO of the disk solution for long-term data retention, especially for energy consumption, where disk consumes 238 times as much energy as tape under assumptions that lean toward favoring disk.

For most uses, we believe that the vast majority of archived data should reside on tape.”

High capacity, small footprint

With a compressed capacity of 3TBs of storage on a single LTO-5 Ultrium cartridge, organizations can scale and keep pace with data growth, simply by adding more LTO tapes to a library. LTO Ultrium palm sized tapes can be safely stacked and stored in a vault at regular ambient conditions; they can reduce the cost and amount of premium physical floor space required to house data in archive. Data archived on LTO tape does not require the data center floor space, power or cooling that’s required for data stored on disk.

Portable and durable

Tape media is completely removable and the lightweight yet rugged cartridge designs are easily transported without risking data loss. The promise that bandwidth would replace trucks for moving large amounts of data has not come true. The growth in data that needs to be backed up exceeds the growth rate of bandwidth, and will continue to for the foreseeable future.

Removable media also has the advantage that data is held ‘offline’ which means that archive data on tape has the additional level of protection from the threats to on-line data from viruses, hackers and cyber-attacks.



¹⁰ Based on prices obtained from www.cdw.com in January 2010

¹¹ <http://www.clipper.com/research/TCG2010054.pdf>

WORM and Data Encryption for added data security

Both LTO-5 and LTO-4 Ultrium tape drives include hardware based data encryption to prevent unauthorized access to data at rest. They use the Advanced Encryption Standard (AES) with the longest and most secure keys, 256 bits, designed to be compliant with the emerging standard for tape drive security, IEEE 1619.1.

LTO-5, LTO-4 and LTO-3 Tape cartridges are also available with write-once-read-many (WORM) capability to prevent accidental overwriting of data archived on the tape. Both WORM and hardware data encryption features help organizations to comply with increased data security regulations.

30-year shelf-life

Organizations need to be confident that data in archive is still there when it's needed. With media specified to a 30-year shelf-life, LTO Tape provides an ideal long-term data retention solution that has proven to be reliable. With over 150 million LTO Ultrium cartridges shipped since the technology was first launched in 2000, customers are happy to trust their data to LTO Tape.

The amount of data at the back end of the data life cycle is now growing, not shrinking, as was the case a decade ago. For some data, the storage preservation requirement will become infinite. For the majority of digital data, the generalized axiom of "90 days on disk and 90 years on tape" refers to lifetime data management and to tape as the curator.

Tape drives have longer useful life than disk drives and reduce the need to move data to new media as frequently.

HP LTO-5 with LTFS—innovation that heralds a new era for tape archiving

HP Linear Tape File System (LTFS) delivers a new class of portable storage media based on open standards. It combines the economy, robustness, high density and low power of tape, with much of the functionality and usability of a hard drive. LTFS has captured the imagination of many customers, vendors and analysts in the storage industry and it's not hard to understand why. HP LTFS provides a self-describing file system on a HP LTO-5 cartridge, which enables the combined archiving benefits of application independence, transportability and protection from obsolescence.



Key Digital Archive Applications

While archiving data to tape has a positive benefit for the majority of organizations, there are a few key applications where the ability to archive large volumes of data is particularly critical.

“New approaches to incorporating tape into traditional storage architectures are under development (e.g., LTFS) and will gain greater visibility and traction, especially in verticals where data is the product such as media and entertainment, healthcare, and life sciences.” IDC

Media and Entertainment

From classic film archive, to new digital films with complex special effects, to video gaming, social networking and user generated content including images and video – data is growing exponentially in the communications, media and entertainment industry. The benefits of HP LTFS are particularly relevant for media and entertainment companies that need storage solutions that simplify operations, improve manageability and meet their long-term data retention requirements.

There is already significant industry momentum behind developing the full potential of HP LTFS in the media and entertainment vertical market segment. In January 2011, Cache-A Corporation, a leading supplier of network-attached archive appliances for the digital film, broadcast and video professionals, announced collaboration with HP to develop an easy-to use implementation of LTFS (Linear Tape File System) for the professional media and entertainment industries.

When Cache-A and HP complete the integration of LTFS into the Cache-A archive appliance, customers will benefit from a seamless file management that is critical to managing a secure archive where quick accessibility is key. The LTFS solution will enable clients to more effectively safeguard data, increase data mobility and share content organization-wide. For more information visit: www.cache-a.com or www.hp.com/go/lfts

Healthcare

The average US hospital has roughly 150 different applications, all generating fixed and dynamic data, each with their own requirements for data retention. In addition, the average hospital requires approximately 60TBs of storage for their image archives. With hospitals around the world transitioning to electronic record management, hospital data volumes are doubling every two years. Similar issues are being faced by research centers, healthcare trusts and larger local medical centers.

Implementing a tiered storage approach that uses tape for archive and compliance data (images and retired patient records), can free up valuable storage space on more expensive online systems. Moving inactive data to less costly tape storage with hardware-based data encryption can provide cost savings while preserving data security.

Video Surveillance

Many factors are contributing to the fast growing storage requirements in the video surveillance market, including higher resolution images from IP cameras, the transition from analogue to digital, and increased industry regulation requiring longer retention periods for digital video evidence.

Once the primary video data is captured, it may be centralized for storage and post-acquisition processing. The video is then archived for long-term retention. Tape provides a 'cheap and deep' solution that offers a cost-effective way to archive and scale to meet growing data volumes without disrupting existing infrastructure.



“Instead of being killed off by the disk drive, tape is being kept around for both traditional uses as well as taking on new roles where it is best suited, such as long-term or bulk off-line storage of data in ultra dense and energy efficient, not to mention economical manners.”

Greg Schulz, StorageIO¹²

In Summary

The exploding digital archiving market is now the largest growth market for tape. LTO-5 delivers a compelling solution for markets and applications where cost-effective, long-term and secure data retention is a requirement. HP LTFS adds to this solution by bringing prompt access to data in archival for even greater ease-of-use.



¹² <http://storageio.com/blog/?p=1060>