



Controlling the TCO of Long-Term Storage — IBM Introduces the TS4500 Tape Library

Analyst: David Reine

Management Summary

Today's family does not have to be hoarders to experience the demands that a lack of adequate storage space places upon our apartments, condos, and homes. We are all keeping more "stuff" than ever before, both physically and digitally. Whether it is family photograph albums to see where we have been, tax records to enable us to defend ourselves against tax audits, or grandma's china that we are hoping to hand down to our children, we are holding on to more and more. Our attics and basements are overflowing with the memorabilia of our past and possibly those from earlier generations. For some this means converting the garage into a storeroom and leaving the car in the driveway. However, this too has its limits as we still may need to rent space from the local storage facility. The more space you need, the more costly it is, and sometimes, you have no idea if the added space that you commit to is going to be sufficient.

This scenario is not new to the management of any enterprise data center. Recently, data growth has gone through the roof. The quest for more knowledge about enterprise customers has led to the deployment of any number of new servers to run new analytical software. The proliferation of these, and other new enterprise applications, across a virtualized server network, has led to the installation of more and more disk arrays to house this information. Between servers and storage arrays, the typical data center is reaching the limits of available energy and floor space, as well as the IT budget. *Something has to be done!* The data center cannot deploy any additional business- and mission-critical applications without creating a new supplementary data center, or replacing the existing facility, at a significant cost.

There is not much that the data center staff can do about the primary tiers of data. The urgency of providing management with the desired information in a timely fashion, i.e. *right now*, requires the continued deployment of solid state discs (SSDs) and spinning media (HDDs), with the costs of acquisition and replacement, floor space, and energy weighing heavily on the IT budget. However, there is something that can be done with the information that is not very active, but must be protected and retained for a very long time, perhaps forever. This long-term data can, and must, be moved off-line to conserve as many resources as possible. The most efficient way to accomplish this is by archiving this data, perhaps petabytes (PBs) of it, to tape.

With the recent, continued, development of new tape media technology, the data center can store in excess of 6TBs of data on a single open systems cartridge. With continued development of new technology for tape libraries, the data center can now store in excess of 2PBs in 10 square feet of data center floor space. One company at the forefront of this continuing evolution is IBM. With the recent announcement of the *IBM TS4500 Tape Library* and the enterprise edition of *LTFs*, supporting both LTO and enterprise media, IBM continues to lead the way in library automation. To learn more about the IBM tape solution, please read on.

IN THIS ISSUE

➤ Long-Term Storage Growth in the Data Center.....	2
➤ IBM Tape Media Technology	3
➤ IBM TS4500 Tape Library	4
➤ IBM LTFs Enterprise Edition	5
➤ Conclusion.....	6

Long-Term Storage Growth in the Data Center

The enterprise data center has been faced with many problems since the inception of the first mainframe. In the past, these problems have been attacked and cleared like so many hurdles on a race track. Today, however, the issue of long-term storage of data is more like a high jump than a hurdle, and the data center staff needs to find new ways to clear it. In 1965, Dick Fosbury faced a height in the high jump that he knew he could not clear using conventional methods, so he created a new way to jump – backwards – which has become the standard for high jumpers today. Before we attempt to clear the bar in this high jump, however, we need to determine how high the bar is and how high it might be tomorrow.

The widespread growth of data storage is being fed by any number of causes, not the least of which is the voracious appetite for knowledge demanded by the business units of every mid-sized and larger data center, and the need to access it *right now*. This can be seen in a wide range of industries, including: entertainment, video surveillance, medical imaging, and analytics. Some of this data must be available online, some can be stored offline. They both have one thing in common, the need to be preserved and protected, perhaps forever.

We all agree that urgent or “hot” data – that which is being used for mission- and business-critical purposes on a daily basis – must be retained on the fastest media that the enterprise can afford. **In most instances, cost must be factored in.** The most urgent data needs to be stored on SSDs, or at least on HDDs. That with little or no urgency, or “cold” data, can be stored offline, on less expensive, and slower, media. *You may ask: How much long-term storage do I need?* The answer probably is: *I don't know.*

Well, join the club, because very few CIOs could have predicted what has happened over the past five years and, most certainly, few would dare to predict what might happen over the next five with any sense of surety. All we can say is that the long-term storage solution that you do deploy must be flexible. It needs the transparent scalability and density to preserve current investments and allow for the capacity required to meet the expansion needs of tomorrow, without excessive manual intervention.

Every data center has been undergoing this period of unprecedented data growth, causing an expansion of the storage facilities required to

house it. It doesn't matter whether your data center needs to store ten terabytes or ten petabytes; you need a flexible, scalable and reliable system that will enable each tier of your storage to scale to meet the needs of your business. Many data centers have been doubling their storage requirements every 12-to-18 months, while some are experiencing even greater growth, adding to the management complexity and issues involved in accessing data.

Many applications require a backup solution that may require a disk-to-disk (D2D) architecture – to expedite the availability of the data when a recovery is required. In addition, D2D can be used to better manage data center operations, as the amount of data being backed-up often stretches the bounds of the backup window and improves workload performance when a recovery is required. Unfortunately, the data center staff also has to cope with the challenges of TCO and data protection in a D2D environment: acquisition, maintenance, and licensing costs. In addition, there are ancillary costs, such as the power required to drive and cool the expanded architecture, the space needed to house the disks, and the additional technology needed to protect the data from unauthorized access. This may create an unacceptable burden on the IT budget for the long-term preservation of data, whether backups or archives. Preserving this information may demand a lower cost storage target, such as tape¹, which can be used with a disk-to-tape (D2T) or disk-to-disk-to-tape (D2D2T) storage architecture.

Many IT organizations have deployed or are looking to deploy an automated tape library with the most current technology in order to reduce the TCO and energy consumption of long-term storage as well as to enable easy-to-use data management strategies. In fact, with the expanded capacity and improved throughput of the newest tape architectures, some data centers are even looking to adopt, or return to, a more direct D2T solution. See Exhibit 1, on the next page, for the requirements required to satisfy the needs of storing data over the long term.

With the need to acquire additional primary storage, as well as to provide for the long-term preservation of business- and mission-critical

¹ For a discussion on lowering TCO by using tape, see the issue of *The Clipper Group Calculator* dated May 13, 2013, entitled *Revisiting the Search for Long-Term Storage – A TCO Analysis of Tape and Disk*, and available at <http://www.clipper.com/research/TCG2013009.pdf>.

Exhibit 1 — Long-Term Storage Requirements

- **Non-Interference** – Be able to provide sufficiently high performance to meet enterprise workload scheduling requirements.
- **High Capacity** – Be able to meet the storage growth demands of the data center.
- **Affordability** – Be able to live within budget limitations.
- **Reliability** – Be able to ensure uninterrupted data access.
- **Scalability** – Allow the data center to protect its investment with seamless upgrades as capacity grows.
- **Security** – Be able to ensure and protect the integrity of the data and prevent unauthorized access.
- **Energy Efficiency** – Be able to lower the demand on electrical resources and extend the life of the data center, even as backed up and archived data capacity grows.
- **Portability** – Allow the transportation of media offsite to facilitate disaster recovery.
- **Ease-of-Use** – Do this all while minimizing additional demands being placed upon the IT staff.

Source: The Clipper Group

information and other archival data, some enterprises are looking for a single vendor who can satisfy all of their needs, across all tiers of storage, minimizing the pains caused by multiple management strategies and multiple ordering paths. IBM is one company that is addressing all of the enterprise storage needs. Let's take a look at their latest tape innovations designed for long-term data protection.

IBM Tape Media Technology

IBM has been a leader in tape innovation for the past 60 years. They have been a driving force behind the creation of the LTO Program and support all six generations (and counting) of the *Linear Tape Open* format. The IBM *TS1060 Ultrium 6 Tape Drive*² (LTO-6), with an 8Gb, dual-ported FC interface, has a native capacity of up to 2.5TB, with a compressed capacity of 6.25TB on a single cartridge. The IBM *TS4500*

² See [The Clipper Group Navigator](#) entitled *Big Data Requires Big Storage – IBM Increases TS3500 Capacity - Again* dated October 24, 2012, and available online at <http://www.clipper.com/research/TCG2012024.pdf>.

Tape Library, with LTO-5 (TS1050) and LTO-6, offers capacity and performance for enterprise and mid-range open systems environments. These drives also feature encryption and WORM³ functionality to preserve and protect the vital corporate assets written to it. In addition, with the development of multiple track media, LTO-5 and LTO-6 tapes support media partitioning and self-describing tape with the *LTFS Library Edition (LE)*, *LTFS Storage Manager*, and *LTFS Enterprise Edition (EE)* software. This enables files stored on tape to be accessed like files stored on disk media, improving access time to critical business data over traditional tape methods. In addition, LTFS EE enables the data center to create an LTFS tape tier (as discussed in “LTFS Enterprise Edition” on page 5).

In addition to LTO drives, IBM also has a line of enterprise tape drives, with the latest version being the *TS1140 Tape Drive*⁴. All of the functionality of LTO was available via IBM enterprise tape before it became available in LTO. The TS1140 media has more capacity, throughput, and reliability than LTO-6 with a native capacity of 4TB (12TB using a compression ratio of 3:1) and a native tape drive data rate of 250MB/second. This compares quite favorably with an LTO-6 data rate of 160MB/second. In fact, enterprise tape may be the more economical solution for your data center⁵, even when it is deployed in an open system environment.⁶

Furthermore, this is just the beginning. In 2010, IBM and Fujifilm demonstrated a capability to record 35TBs of data on a single cartridge⁷, following that with an announcement in May

³ WORM=Write Once Ready Many.

⁴ See [The Clipper Group Navigator](#) entitled *IBM's New Enterprise Tape Extends Data Retention Capabilities and Lowers the Cost of Data Protection* dated June 6, 2011, and available online at <http://www.clipper.com/research/TCG2011021.pdf>.

⁵ See [The Clipper Group Captain's Log](#) entitled *Enterprise Tape for Archival Storage? – Why This Just Makes Sense* dated March 31, 2013, and available online at <http://www.clipper.com/research/TCG2013005.pdf>.

⁶ See [The Clipper Group Captain's Log](#) entitled *Ten Reasons Why You Should Consider Enterprise-Class Tape for Open System Storage* dated July 12, 2011, and available online at <http://www.clipper.com/research/TCG2011025.pdf>, and also see the latest issue of [Clipper Notes](#) entitled *Is Tape the Best Low-Cost Technology for the Preservation of Data?* dated July 5, 2014 and available online at <http://www.clipper.com/research/TCG2014015.pdf>.

⁷ See [The Clipper Group Navigator](#) entitled *IBM and Fujifilm Increase Tape Density – Raising the Bar on Tape Capacity* dated February 5, 2010, and available online at <http://www.clipper.com/research/TCG2010005.pdf>.

2014 of the ability to pack up to 85.9 billion bits per square inch on a standard length tape cartridge, up to 154TBs of uncompressed data, a giant leap forward in improving the density of the data center. This followed an announcement by Sony Corporation in April 2014 of their capability, with the assistance of IBM, to record approximately 185TBs of data on a single cartridge, 74 times the capacity of LTO-6 technology. **Tape media is alive and well, thriving in enterprise data centers – maybe even yours!**

IBM TS4500 Tape Library

The IBM TS4500 Tape Library was designed as a next-generation tape solution with increased, transparent scalability to satisfy the requirement of high data volumes and the growth of the data center, while reducing the TCO of the IT infrastructure. See the photos in Exhibit 1 at the right.

It is the successor to the IBM *TS3500 Tape Library*, originally introduced in 2007. Over the years, the TS3500 has been enhanced to provide support for the latest innovations in magnetic tape, in the open systems arena and the enterprise tape market, as well. In October, 2012, IBM updated the TS3500 to add support for the TS1060 LTO-6 Tape Drive, as well as the TS1140 Enterprise Drive.

Now, IBM has improved upon the basic architecture of the library with the availability of the TS4500, supporting TS1050, TS1060, and the enterprise TS1140 drives.⁸ The base frame is identified as an “L” frame, an *L25* for the TS1140, and an *L55* for either the TS1050 or TS1060. Up to three “D” expansion frames can be added to expand the number of tape drives and/or storage capacity – *D25* for the TS1140 and *D55* for the TS1050 and TS1060. Additionally, using IBM’s high-density, capacity-on-demand feature, the total capacity of the library can be increased significantly with zero downtime – there is no library outage or impact to the application. Like its predecessors, the TS4500 includes dual grippers for improved performance and availability.

The TS4500 can support up to 2.2 petabytes (PBs) of long-term storage in a single frame library⁹, consuming only ten square feet of data center floor space, with less than one hour of downtime per year for maximum reliability. It

⁸ There is an upgrade available to migrate some frames to the TS4500 models.

⁹ This is 3.4 times the capacity of previous IBM libraries.

Exhibit 2 —
The IBM TS4500 Tape Library



Source: IBM

Exhibit 3 — TS4500 Management Options

- **TS4500 command line interface (CLI)**—providing access to TS4500 library functions such as moving cartridges to and from the slots and drives;
- **IBM Tape System Reporter** – enabling administrators to monitor/ report on storage devices in an enterprise environment;
- **IBM Tape System Library Manager** – allowing administrators to manage multiple libraries as a single system; and
- **IBM Security Key Lifecycle Manager** – to simplify encryption key management with an intuitive user interface.

Source: IBM

can support up to 42.6 PBs in a multi-frame, enterprise-drive configuration with a 3:1 compression factor.¹⁰ In addition, they have increased the density and capacity, of an expansion frame, providing better utilization of data center floor space, with support for up to 12 drives in the leftmost frame and up to 16 tape drives in the others, with a current maximum of 60 drives in any library string. This can alleviate some of the complexity of managing an increasing data store with a decreasing management staff. In addition, with the TS4500, the data center gains the flexibility of being able to extend the library with additional frames on either the right or left. The TS4500 also provides flexibility in terms of configuration. The TS4500 rack reduces cabling and footprint with the availability of an additional 10U of rack space available on top of the library for FC switches, tape data movers, and/or LTFS nodes.

IBM has also improved integrated management of the library for both drives and media within a single-pane-of-glass management console, improving ease-of-use. The *Advanced Library Management System*, featuring dynamic storage management, enables users to dynamically virtualize tape to create and change logical libraries and configure any drive into any logical library. The TS4500 offers business continuity and disaster recovery with automatic control path and data path failover, carried over from previous IBM tape solutions. See Exhibit 3 above for a list of other available management capabilities.

¹⁰ The TS4500 has a maximum capacity of up to 29.3 PBs using LTO-6 technology with a 2.5:1 ratio.

At the same time, the IBM TS4500 also is the answer to the security dilemma facing every enterprise: How can the enterprise ensure that their archives and long-term backups are protected and preserved in conformance to government regulations and industry standards? With integrated encryption and WORM media, also continued from previous generations, the enterprise knows that it can trust its archives to be secure from prying eyes. Furthermore, it does all of this within the framework of the same graphical user interface (GUI) used by other IBM storage solutions¹¹.

LTFS Enterprise Edition

The LTFS Enterprise Edition can play a significant role in reducing the TCO of storage for file data. It provides a seamless integration of LTFS and IBM's *General Parallel File System (GPFS)* for both the TS3500 and TS4500 tape libraries, creating an LTFS tape tier for "cold" data that does not need the access performance of primary disk. LTFS EE enables the use of LTFS for the policy management of this tape storage tier in a GPFS environment. It has support for the open system TS1050 and TS1060 drives, as well as the TS1140 enterprise drive.

LTFS EE enables the data center to replace disks with tape for Tier 2 and Tier 3 storage, improving data access over other storage solutions, improving efficiency, and simplifying file management on tape. LTFS EE makes tape transparent to the user and simplifies the administration under a single infrastructure. This is especially important when LTFS LE serves as a migration target for GPFS to transparently archive data to tape based upon a specific set of predetermined policies.

LTFS EE uses an enhanced version of LTFS LE. Its scale-out architecture enables the addition of nodes and tape devices as needed to satisfy the growing bandwidth requirements between GPFS and LTFS. The benefits of LTFS are defined as:

- Speed time to value;
- Reduce TCO;
- Improve data management;
- Improve data access without extended management; and
- Easily enable data sharing in an open format.

¹¹ Based upon IBM's *XIV* technology.

Exhibit 4 — LTFS EE Functionality

- Create and define tape storage pools for file migrations;
- Migrate files in the GPFS namespace to the LTFS tape tier and recall files that have been migrated back into GPFS;
- Reconcile file inconsistencies between files in GPFS and their equivalents in LTFS;
- Reclaim tape space that is occupied by non-referenced files and non-referenced content that is present on tape;
- Export cartridges to remove them from your LTFS EE system;
- Import cartridges to add them to your LTFS EE system with no disruption; and
- Obtain inventory, job, and scan status.

Source: IBM

The significant functions of LTFS EE are listed in Exhibit 4, above.

Conclusion

Today's CIO and data center management staff are experiencing a level of data growth never seen before. They are also faced with a limited budget that is being severely challenged by the need for more floor space and more energy to run the expanding IT infrastructure. They require a flexible solution with the scalability to preserve existing investments.

The IBM TS4500 Tape Library is a next-generation storage solution to the storage challenges facing the 21st-century data center. It enables the mid-sized and larger data center to deploy the long-term storage density that the enterprise requires in order to help control the TCO of the IT infrastructure, while improving the cost efficiency and manageability desired. The TS4500 provides the density and scalability that every enterprise needs to meet the challenges of an expanding storage requirement. The TS4500 also provides for extensibility of the tape library for the next decade, and beyond, protecting the investments that the enterprise has already made, and will make, in IBM tape storage hardware.

If these challenges sound familiar, The TS4500 may be the answer you are seeking.



About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-second year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- **The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.**

About the Author

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- **Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)**

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, The Clipper Group Calculator, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.