THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2012024 October 25, 2012

Big Data Requires Big Storage — IBM Increases TS3500 Capacity – Again!

Analyst: David Reine

Management Summary

Have you gone shopping recently? Looking for a bargain? Why not head to the Big Box store. With prices on the rise, we are all doing what we can to control expenditures and stay within our budget. One way to do that is to shop at one of the many warehouse stores that are spreading like wildfire across the country to enable the average consumer to get more for less. You can go to Sam's Club or Costco or BJ's Wholesale Club and save money by buying a larger quantity of whatever you are shopping for: a 12-pack of paper towels, a 24-count supply of AAA batteries, or even a 15-pound package of steaks. The more you buy, usually the cheaper it is, at least on a per unit basis.

A similar scenario exists within the data center of every enterprise trying to cope with the ever-increasing demands of Big Data. With data coming in from every corner of the enterprise, in fact, every corner of the globe, the IT staff is constantly fighting the battle of the budget to keep up with this inexorable growth with enough primary storage to enable immediate access to it, and to enable the long-term preservation of this data for archive, statistical analysis, and disaster recovery. The data center has seen the evolution of storage from megabytes to gigabytes to terabytes to petabytes, with exabytes and zettabytes looming larger and larger on the horizon. The data center has also seen the deployment of 1TB, 2TB, and now, 3TB disk technology (with 4TB on the way) and tape drives with even larger capacities. It is critical that the enterprise deploy a big data solution that will not only provide immediate access to petabytes of data, but also long-term data protection that will protect the investment that the enterprise makes in that infrastructure. Far from being dead, tape has proven to be a most resilient media, one that continues to lower the total cost of ownership (TCO) to any larger data center. One company that is trying to stay one step ahead in the long-term storage of big data is "Big Blue".

IBM has been delivering a scalable tape library solution for both mainframe and open system environments for decades. Their latest enterprise model is the *TS3500*, which can utilize both of IBM's tape architectures: IBM's *TS11x0* Enterprise Drive and IBM's open systems *Linear Tape Open (LTO)* family of tape drives. The TS3500 has been available with IBM's TS1140 enterprise drive and the industry-standard *LTO-5* tape drive. Now, IBM has increased the total capacity of the TS3500 with the implementation of the industry-standard *LTO-6* drive. With 16 frames per library, 12 tape drives per frame, and up to 20,000 cartridges in each TS3500 Library, **IBM's TS3500 can provide long-preservation for up to 125PB of LTO-6 compressed data per library**. To learn more about the new, improved version of the TS3500 with LTO-6 drives and media, please read on.

Growth of Big Data in the Data Center

Every enterprise data center has been experiencing a period of unprecedented data storage growth. Most are doubling their storage capacity every 12 to 18 months, while others are experiencing even greater growth, compounding the access issue. This growth places a tremendous burden on the IT budget, as the business attempts to control the TCO of the

IN THIS ISSUE

entire data center. The IT staff not only has to acquire additional primary storage, but they also must provide for the long-term preservation of both business- and mission-critical information. Many enterprises have adopted a disk-to-disk (D2D) storage solution, to better manage their operations, as the amount of data being backed-up often stretches the bounds of the backup window and, impacts workload performance when a recovery is required. **The TCO and data protection challenges of a D2D environment, however, can create an unacceptable burden on the IT budget for the long-term preservation of data.**[1] The IT staff must cope with D2D acquisition and licensing costs plus the ancillary costs, such as the power required to drive and cool the expanded architecture, the space need to house the disks, and the additional technology needed to protect the data from unauthorized access.

For these and other reasons, the IT staff must look for an affordable, high-capacity backup and archive capability, such as a D2D2T environment, or expand their existing tape solution. Many are looking to deploy an automated tape library with the latest technology in order to reduce the TCO and energy consumption of long-term storage as well as to enable manageable data protection strategies. In fact, with the expanded capacity and improved throughput of the newest tape architectures, many data centers are even looking to adopt, or return to, a more direct D2T architecture. There are many features that are important to the enterprise data center beside simply improved capacity and performance in order to lower the TCO of the storage infrastructure. Please see Exhibit 1 (at the top of the next column) for a partial list.

The data center must be able to cope with all of the components of the entire storage environment. This requires the acquisition of an integrated infrastructure – in order to enable the data center to retain and secure adequately the IT infrastructure with high-end disk arrays (including SSD), Tier-2 disk for a D2D target, and tape libraries, along with data deduplication and encryption appliances. For the tape library, this spans not only the physical confines of the data center, but also needs to facilitate off-site access to a redundant, remote location.

For the library, the data center needs a reliable solution supporting the latest tape technology with the highest capacity and throughput, possibly including both proprietary technology, such as the IBM *TS1140*[2] Enterprise Drive, or open systems, with the newest *LTO*[3] technology, in order to protect their investment in existing media deploying a new *LTO-6*[4] drive can complement or replace older, less performant LTO drives and facilitate data preservation within the existing IT budget, even enabling the data center to protect previous investments with the capability to write to *LTO-5*[5] cartridges and read from LTO-5 and *LTO-4* cartridges, while LTO-5 drives can write to LTO-4 cartridges and read from LTO-4 and *LTO-3* cartridges. With more than twice the compressed capacity of LTO-5 tape, LTO-6 tape enables the data center to reduce the number of cartridges required, reducing the floor space, and prolonging the investment in the tape library. In addition, with a roadmap that includes *LTO-7* and *LTO-8* technology, the IT staff can plan for even more capacity and performance in less space in the years to come.

Exhibit 1 — Long-Term Storage Requirements

- **High-Performance** – Adhere to enterprise requirements to meet workload schedule;
- **High Capacity** – Meet the storage growth demands of the data center;
- **Affordability** – Adhere to budget limitations;
- **Reliability** – Ensure continuous data access;
- **Scalability** – Enable the data center to protect its investment with seamless upgrades as capacity grows;
- **Security** – Ensure and protect the integrity of the data and prevent unauthorized access;
- **Energy Efficiency** – Lower the demand on electrical resources and extend the life of the data center as capacity grows;
- **Portability** – Enable the transportation of media offsite to facilitate disaster recovery; and
- **Ease-of-Use** – Minimize any additional demands being placed upon the IT staff.

***Source**: The Clipper Group*

[1] See the issue of Clipper Notes dated December 20, 2010, entitled *In Search of the Long-Term Archiving Solution – Tape Delivers Significant TCO Advantages over Disk*, and available at http://www.clipper.com/research/TCG2010054.pdf.

[2] 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; see http://www.clipper.com/research/TCG2011021.pdf.

[3] *Linear Tape Open*, also called *Ultrium*.

[4] See **The Clipper Group Navigator** entitled *Magnetic Tape Turns 60 – The IT Industry Receives Another Gift* dated July 12, 2012, and available online at http://www.clipper.com/research/TCG2011015.pdf.

[5] See **The Clipper Group Navigator** entitled *LTO Program Announces Next Gen Tape – LTO-5 Raises the Bar for Tier-3 Storage* dated January 31, 2010, and available online at http://www.clipper.com/research/TCG2010002.pdf.

In order to take advantage of the higher capacity and higher reliability provided by an enterprise-class (proprietary) drives and the high capacity and lower cost inherent in an open system solution, the enterprise needs a flexible IT infrastructure to support both, in high volume, in a scalable storage environment. Space and power are significant issues for the enterprise due to the limited floor space and energy available. Therefore, the enterprise data center requires the densest configuration possible, along with reduced energy consumption per TB.

Media portability is mandatory for the offsite storage of information for data protection, for both long-term archiving and/or disaster recovery. Portability, however, also implies additional requirements for data protection. The enterprise needs to encrypt any data that will leave the data center to ensure the increased security of data in transit, keeping the enterprise out of the headlines and its executives out of court. The enterprise also needs to prove that retained data has not been altered, requiring support for WORM[6] media to address compliance with industry and government required standards.

Automated identification of tape media is essential and requires the availability of a barcode reader for identification of the correct cartridge and the availability of an I/O station to enable the easy importing and exporting of media while maintaining continuous operation. Management and administration are other important elements in the information infrastructure. Due to the potential complexity inherent in this level of enterprise architecture, the enterprise requires a web-based management capability to enable the remote management and administration of library functionality. IBM has addressed all of these requirements with their latest *TS3500* Library with the addition of their *TS1060 Ultrium 6 Tape Drive*.

The IBM TS3500 Library

The TS3500 Library, along with IBM storage applications, delivers a functionally rich tape storage solution that includes both the TS1060 and the TS1140 technology. It was designed to provide a highly-scalable, automated tape library for both mainframe and open systems environments, in midsize and large enterprises. Since being introduced in August, 2000 as the IBM 3584, IBM has continued to extend the capacity and functionality of this library, most recently, by adding a *Shuttle Complex* to the library, with the result being configurability to up to 300,000 cartridges[7]. This capability was added at the same time as the introduction of the TS1140 in 2011.

The TS3500 brings a wide variety of unique features to the enterprise data center. In addition to the Shuttle Complex, the TS3500 has a dual-robotic accessor, an advanced library management system, capacity-on-demand, and both control path and data path failover. IBM's *Tape System Library Manager* (*TSLM*) software expands and simplifies the use of TS3500 libraries by providing a consolidated view of multiple libraries, reducing tape pathing maintenance. TSLM now enables enterprises with large Tivoli TSM environments, spanning multiple TS3500 libraries with the shuttle connector, to manage them as a single system. TSLM also simplifies Tivoli TSM tape drive connection definitions, saving operational resources that would be needed to manage pathing over time. In addition, TSLM enables Seismic customers to migrate their IBM *3494 Libraries* to the TS3500, eliminating the need to find and maintain these older libraries. These customers can now utilize both LTO technology and the TS1140.

The TS3500 Library can be configured with up to 16 frames and up to 12 drives per frame, with up to 15 libraries per library complex. Deployed with TS1060 Ultrium 6 tape drives, a TS3500 library can be configured for up to 125PB of compressed data, with a maximum of 1.875EB per complex, with up to 2,800 drives per complex.

The TS1060 Ultrium 6 Tape Drive

With the TS1060 LTO Ultrium 6 Tape Drive, IBM has retained all of the features and functions of the *IBM System Storage TS2350 Tape Drive Express*[8] *LTO Ultrium 5* drive, and more. IBM has designed the TS1060 for the heavy demands of backup and archive tape storage. They have increased both the capacity and performance of the TS1060 to better facilitate the growth of Big Data[9] within the enterprise data center, while controlling the costs of long-term data storage. The TS1060 combines IBM tape reliability and performance at an open systems price, using an advanced *Barium Ferrite* tape developed to help provide durability and increased capacity. The TS1060 does the following.

[6] WORM=Write Once, Read Many.

[7] See footnote #2 on page 2.

[8] See **The Clipper Group Navigator** entitled *IBM Enables the Data Center with LTO-5 Products- Increasing Capacity, Throughput and File Management* dated April 12, 2010; see http://www.clipper.com/research/TCG2010017.pdf.

[9] Big Data is data that is coming into the data center from everywhere: from sensors used to gather climate changes, posts to social media sites, digital pictures, and videos to name a few.

- Increases the native data transfer rate from 140MB per second (LTO-5) to 160.
- Provides support for up to 14 levels of speed matching to adjust the drive's native data rate as closely as possible to the host data rate. This helps to reduce the number of backhitch repositions and improves throughput performance.
- Supports a cartridge capacity of 6.25TB, when compressed at a compression ratio of 2.5:1 (your mileage may vary) and through increased number of tracks. In addition to the standard cartridge, the TS1060 also supports a WORM cartridge to prevent the alteration or deletion of user data.
- Includes an 8Gbps FC dual-ported interface.
- Uses a Giant Magneto Resistive head to reduce wear on the tape media, with a head actuator designed to provide precision head alignment for higher track density and improved data integrity.
- Provides improvements in power management over the LTO-5 device in idle mode to reduce power consumption.
- Supports media partitioning and self-describing tape with IBM's *Linear Tape File System (LTFS)*, IBM's *LTFS Library Edition*, and IBM's *LTFS Storage Manager* software for media asset management.
- Provides continuing support for encryption and a 1GB internal buffer to help improve job performance.

In terms of costs, LTO-6 has a profound impact on the TCO of the IT infrastructure. First, in terms of acquisition cost, with a list price of $25,855, IBM's TS1060 lowers the cost per terabyte for compressed data to $4,137. This compares quite favorably to the $7,980/TB for LTO-5, and even more so when compared to LTO-4 at $14,250/TB and LTO3 at $24,750/TB.[10] In terms of operational costs, we do not as yet have a cost for the LTO-6 media, but it is expected to be introduced at about the same price at which LTO-5 media was introduced, potentially reducing data center costs for media by a factor of two.

The TS1060 mounts in the TS3500 Tape Library (Model *L53* or *D53*) and the *3584 Tape Library* module (*L52*, *L32*, *D52*, or *D32*). It can read and write to LTO-5 Ultrium cartridges and read LTO-4 Ultrium cartridges. It supports a wide-range of operating environments, including: *IBM System p*, *IBM System i*, *IBM System x*, and *IBM System z*, running *zLinux*. It also supports servers running *HP-UX*, *Linux*, Oracle *Solaris*, and Microsoft *Windows* operating environments that support FC interfaces. In addition, the TS1060 enhances many other LTO features, which are listed in Exhibit 2, above.

Exhibit 2 — Enhanced LTO-6 Features

- **Independent tape loader system** to improve the loading/unloading reliability.
- **Graceful dynamic braking** to help reduce tape breakage and other mishaps.
- **Servo and track layout** increase data tracks from 1,280 to 2,176.
- **Robust drive components** enhance the reliability and prolong drive life/
- **LTO Data Compression** to help provide optimal data compression and increase data throughput.
- **LTO Cartridge Memory doubled** from 8,160 bytes to 16,320 bytes to hold more data about the specific cartridge.

Source: IBM

Conclusion

What are the mission-critical factors that the CIO needs to consider in evaluating a long-term storage solution? As he or she evaluates the enterprise storage requirements, he or she must keep in mind that there is more to an enterprise storage solution than simply higher capacity and more performance. The IT staff needs to ensure the reliability and availability of the library, as well as the integrity of the data and the protection provided by encryption. They need to consider the advantages of the data partitioning, formatting and management provided by LTFS. Cost is always a component of the decision equation; however, the investment protection provided by the continuation of existing tape formats, with additional upgrades on the horizon, is also critical. **Far from "Dead", tape remains a vibrant technology for the enterprise data center.**

These are the specific factors addressed by IBM with their TS3500 Tape Library with LTO Ultrium 6 tape drives. The fact is that the TS3500 excels in addressing these requirements and delivers a functionally rich tape storage solution incorporating both of IBM's tape technologies: the TS1060 and the TS1400 enterprise drive. If your enterprise is concerned about the long-term protection and storage of their data, the IBM TS3500 may be the answer.

[10] IBM provided the pricing and costs per terabyte.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twentieth 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 25 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**, 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.