



## Magnetic Tape Turns 60 — The IT Industry Receives Another Gift

Analyst: David Reine

### Management Summary

On a typical sixtieth birthday for your father or a favorite uncle, you might have a cake and gifts for the old man. These gifts might include a tie or something else appropriate for a senior citizen. However, this is 2012 and the old rules no longer apply. Sixty is no longer a way-station on the journey to retirement; it is the new “fifty”, more like a springboard to the second half of your life, with new adventures ahead. No rocking chair for this generation’s sixty.

The same thing is true in the enterprise data center, where, in most cases, a technology loses its significance after a few years. There are any number of brands and technologies that have been introduced over the last sixty years that long ago disappeared from the enterprise landscape. Magnetic tape used for data storage is, most certainly, not one of them! Evolved, yes, disappeared, no. Today, tape continues to evolve, providing the data center with the storage capacity and performance required to support the many “Big Data” requirements of this era.

The data center has seen vast improvements in tape over the past sixty years. The IBM 701 Tape Drive, introduced in the early 1950s, used one-half inch tape, in a reel-to-reel configuration. Over the years, tape has evolved from reels to cartridges and the capacity of the cartridge has grown significantly. In fact, the basic tape design remained the same until ½” reel-to-reel tapes came to end-of-life in the early 1990s. The QIC cartridge was followed by DLT, DDS, and DAT, settling on *Linear Tape Open (LTO) Ultrium* as the open system standard that all vendors could utilize. The original LTO cartridge released in 2000 could hold 100 GBs of data. LTO-5 tape, released in 2010, has a native capacity of **1.5 TB** in a cartridge, with a throughput of 140 MB/second. With more than four million LTO drives deployed and nearly 200 million LTO tape cartridges representing more than 80,000PB of storage capacity<sup>1</sup>, LTO technology has become a widely accepted standard. Now, as tape celebrates 60 years of providing data storage to the enterprise, the LTO Program has decided to bestow a birthday gift on the IT community, with the pre-announcement of plans for the LTO-6 specifications availability.

Last month, the LTO Program announced their official plan to make licensing specifications for next-generation LTO-6 technology available in August. These specifications are planned to define the LTO-6 format with a compressed tape cartridge capacity of up to 6.25TB, assuming a compression ratio of 2.5:1, as a result of a larger compression engine history buffer, more than twice the compressed capacity of LTO-5 tape. LTO-6 drives also will support a data transfer rate of up to 400 MB/second, also compressed. This compares favorably to LTO-5 drives, with a compressed throughput of 280 MB/second.

How do these compare to other storage media? How does it affect the total cost of ownership (TCO) of the enterprise data center? How does it affect the data center mission? To learn how LTO-6 technology will address the challenge of storing more data while reducing the TCO of the data center, please read on.

### IN THIS ISSUE

➤ Data Center Challenges .....	2
➤ LTO-6 Functionality and Comparison...	2
➤ LTO Roadmap .....	4
➤ Conclusion .....	5

<sup>1</sup> According to information published by the LTO Program. See <http://www.lto.org/pdf/LTO%20PR%20Maintains%20Top%20Momentum.pdf>.

## Data Center Challenges

With the increasing demands being placed upon the data center by “Big Data” (both analytics and large blobs of data, like images), industry standards, and government regulatory compliance, the challenge to store all mission- and business-critical data within a stagnating or reduced IT budget has become overwhelming. Combine that with the requirements for rapid recovery in case of a data failure or the restoration of an entire system in case of a disaster, and you have a potential catastrophe in the making. Disk remains the option of choice for short-term backup and recovery where the need for speed is paramount.<sup>2</sup> However, low-cost, high-capacity tape may be the better choice to handle long-term storage requirements in the data center where budget constraints are critical and time required for access is a little less demanding. When it comes to tape with an open interface, Linear Tape Open (LTO) technology is the format of choice across a wide range of open systems requirements. LTO-5 tape is the latest available version of the open format, with a native storage capacity of 1.5TB and a native throughput rate of 140 MB/second.<sup>3</sup>

In December, 2010, The Clipper Group published a detailed Note on the advantages of tape vs. disk<sup>4</sup>. This study was done using LTO-5 technology as the tape solution and 2TB SATA drives for the disk solution. In this analysis, we made certain assumptions at that time. One of those assumptions was that both tape and disk solutions would be upgraded on the basis of a repeating three-year cycle. We projected that LTO-5 tape would be the format of choice for at least three years. It appears that we were wrong. The LTO Program has already announced the availability of specifications for LTO-6 tape in August<sup>5</sup>. We estimate a projected availability of

---

<sup>2</sup> There are two kinds of speed. First is the speed to retrieve (read) or store (write) specific data. This speed is measured in, at most, seconds. Second is the speed to restore after a failure of some sort. This speed is usually measured in minutes or hours (or even days). Notice that both are measurements of elapsed time.

<sup>3</sup> The reader needs to keep in mind that enterprise tape from either IBM or Oracle provides greater capacity, throughput, and reliability, though usually at a greater cost.

<sup>4</sup> 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>.

<sup>5</sup> See the LTO-6 Announcement Letter at <http://www.lto.org/pdf/LTO%20PR%20Maintains%20Top%20Momentum.pdf>.

new devices as early as December 2012, more than a year earlier than we had assumed. Unfortunately, we missed on the projected capacity, also. Using the roadmap that existed at that time, we expected a figure of 3.2TB for the uncompressed LTO-6 capacity. This now appears to be coming in at 2.5TB.

On the other hand, our projections for disk may have been a little optimistic. Our assumption of having 4TB disk drives in place in the data center by the beginning of 2013 seems to be premature, at best. This year, 3TB drives began to appear in arrays and now only one disk vendor discussing the availability of 4TB enterprise disks. Last year’s floods in Southeast Asia contributed to, and aggravated, the progress of near-line disk devices for high-capacity, mission-critical applications. As a result, you might say that growth in disk drive capacity is progressing more slowly than for tape.

In addition, many system manufacturers now are fabricating disk arrays using 2.5” disk drives, with a current maximum capacity of 1TB. This is a significant trend that, at the present, seems to reduce the total capacity of an array (i.e., the maximum TBs that can fit into a drawer (chassis) or a rack).

The purpose of this bulletin is to look at the projected benefits of LTO-6 tape in light of the expected announcements from the vendors of spinning disks with regard to 3TB and 4TB devices. A detailed revised TCO study is planned, based upon the criteria established in our 2010 study but using LTO-6 tape and available disk arrays once pricing and the “speeds and feeds” for the newest devices are established.<sup>6</sup>

## LTO-6 Functionality and Comparison

Like LTO-5 tape, LTO-6 tape is expected to come with hardware-based encryption and WORM<sup>7</sup> functionality to protect your data from unauthorized access and ensure compliance with government regulations and industry standards. In addition, LTO-6 tape continues the capability introduced in LTO-5 tape for linear partitioning, enabling the data center client to utilize a tape-based file system, the *Linear Tape File System (LTFS)*, enabling the data center to use tape as NAS, because of its ability to access files in a manner as if they were on disk. LTO-6 technol-

---

<sup>6</sup> Thus, you should look for this report in 1Q2013, based upon present information and intentions.

<sup>7</sup> Write Once, Read Many.

ogy will deliver data storage that is accessible, safe, reliable, secure, and portable for both backup/recovery and archive applications, enabling the enterprise to avail itself of active archiving functionality<sup>8</sup>, helping to reduce the TCO of the data center<sup>9</sup>.

In addition to functionality, the data center is obviously most concerned about both the TCO and total cost of acquisition (TCA) of the storage backup/recovery and archiving environments. The IT staff needs to factor in the capacity, performance, and reliability of their long-term storage infrastructure in making any acquisition decision.

What follows is an analysis of what LTO-6 tape (based on the newly-stated capacities and capabilities) might mean to you. Because pricing won't be available or necessarily comparable before the end of the year, the analysis herein is not based on TCO. Check back in about 6 months for that detailed study.

### Capacity

LTO-6 tape has increased the native (also called "raw") capacity of a single cartridge to 2.5TB (6.25TB compressed). This represents an increase of 67% when compared to LTO-5 tape in native mode and an increase of over 100% in compressed mode.<sup>10</sup> LTO-6 tape is the largest open systems storage device on the horizon. This capacity increase is critical to the data center in view of the rapid growth of data in the enterprise data center, where data at many enterprises is doubling every 12-to-18 months, but often without additional funding in the data center's budget. The majority of this growth is occurring with the growth of unstructured data in the data center, which, in many cases, is ideal for long-term archiving<sup>11</sup> (utilizing the partitioning of LTO-6 tape under LTFS).

The IT staff also must consider the storage capacity of "enterprise tape" as offered by IBM,

with its *TS1140*, and Oracle, with its *T10000C*. Currently, Oracle's T10000C<sup>12</sup> with a native capacity of 5TB holds the lead in cartridge capacity, followed closely by IBM's TS1140<sup>13</sup> at 4TB. This gives the T10000C a 66% capacity edge over LTO-6 capacity, while the TS1140 has a 60% advantage. Unfortunately, there is a kicker – both enterprise drives carry a steeper acquisition cost than the projected estimates for LTO-6 technology. We will discuss the TCA and TCO comparison a little later in this bulletin.

In a data center with a deployed network of 2TB disks, and with 3TB now generally available, LTO-6 capacity compares very well. Even with future 4TB hard drives, the existing LTO roadmap for tape capacity (as discussed on the next page) far outstrips the expected future growth of disk capacity. The IT staff needs to keep in mind that the compressibility of data varies with the type of data being compressed. Therefore, it is always best to compare native capacity between disk and tape.

### Performance

In terms of performance today, tape is very much in the middle of the mix, with the potential for a greater advantage tomorrow. LTO-6 drives are planned to have increased the native throughput to 160 MB/second, with a compressed throughput of up to 400 MB/second. This represents a 14% improvement, in native mode, as compared to LTO-5 drives, and a projected increase of over 40% in compressed mode. Meanwhile, when compared to the available enterprise tape capabilities, Oracle's T10000C has a native throughput of 240 MB/second and the TS1140 is rated by IBM at 250 MB/second, both substantially higher than LTO-6 drives. The LTO-6 tape drive *should* come in at a substantially lower acquisition price than either the T10000C or the TS1140, but we'll have to wait until the end of the year (at the earliest) to draw this conclusion. We say "should" because the LTO-6 Program partners have not yet revealed final pricing for either the LTO-6 drive or media.

<sup>8</sup> See [The Clipper Group Navigator](#) entitled *Simplified, Online Access to Archived Data – Turning History into an Active Archive* dated June 11, 2010, and available at <http://www.clipper.com/research/TCG2010029.pdf>.

<sup>9</sup> When compared to storing only on disk drives.

<sup>10</sup> This is possible because LTO-6 has improved compression capabilities over LTO-5. LTO-6 is capable of delivering up to 2.5 times the raw capacity through compression while LTO-5 is specified as having a maximum compression ratio of 2:1.

<sup>11</sup> See the issue of [Clipper Notes](#) dated February 29, 2012, entitled *Regaining Control Over Hoards of Enterprise Data – Why You Need an Archive*, and available at <http://www.clipper.com/research/TCG2012005.pdf>.

<sup>12</sup> See [The Clipper Group Navigator](#) entitled *Oracle Fulfills Commitment – StorageTek T10000C Takes Leap Ahead* dated January 31, 2011, and available at <http://www.clipper.com/research/TCG2011003.pdf>.

<sup>13</sup> 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 at <http://www.clipper.com/research/TCG2011021.pdf>.

**LTO-6 drive's increase in throughput will enable the data center already committed to LTO technology to shrink backup windows along with recovery time. It also will expedite the creation of archives, enabling mission- and business-critical systems to return to their primary functions, improving enterprise profitability.**

How does this compare to enterprise disk? Seagate's *Constellation ES.2* is offered with both 2TB and 3TB models, and with both SAS and SATA interfaces, running at 6 Gb/second. All four versions have a maximum sustained transfer rate of 155 MB/second. This is about 10% faster than that of LTO-5 drives, but a few points slower than the projected performance of LTO-6 drives. Hitachi, another respected disk provider, has recently announced the availability, in limited quantities, of a 4TB disk (also available in a 3TB or 2TB model). Hitachi's *UltraStar 7K4000* hard drive specifies a typical sustained transfer rate of 171 MB/second. This is about 6% faster than LTO-6 drives, but only 10% faster than the previous generation of disk drives.

### **Reliability**

Is there anything more popular these days than the bashing of tape reliability? When tape opponents are not crowing "Tape is dead", the proponents of disk-for-all-purposes claim that tape is unreliable. Unfortunately, in many cases, the statistics they use are old, in fact, very old. What may (or may not) have been true before the turn of the century, is certainly a figment of the imagination of the disk mavens today. Let's take a look at the reliability figures quoted in a variety of vendors own specifications.

The reliability of high-capacity SATA drives is quoted at 1 error (non-recoverable, bits read) in  $10^{15}$  bits. The reliability of high-performance FC/SAS drives comes in at 1 error in  $10^{16}$ . What about the error rate for "unreliable tape"? LTO-5 is offered at a bit error rate of 1 error in  $10^{17}$  bits, while enterprise tape comes in at an even more reliable figure of up to 1 error in  $10^{20}$  bits.<sup>14</sup> This represents a huge factorial difference in favor of tape. *No wonder the enterprise IT staff is concerned about selecting the RAID technology to deploy for their (unreliable) disk farms!*

<sup>14</sup> The Oracle T10000C drive is rated at 1 error in  $10^{19}$  bits, while the IBM TS1140 is quoted at 1 error in  $10^{20}$  bits.

### **Pricing**

It is still too early to be able to quote a price for an LTO-6 drive or LTO-6 media. We do know that each new generation has come in at approximately the price of the previous generation. We also know, however, that one company already has offered a program where you can order an LTO-5 drive today with a pre-paid option to upgrade to an LTO-6 drive when it becomes available. In this program, the data center can acquire a standalone SAS drive for \$5,300 with an upgrade cost of \$2,250. The figure for an "enterprise-class" FC drive<sup>15</sup> for a library is \$17,956 for the LTO-5 version, with an upgrade cost of \$4,250 for an LTO-6 drive. Based on these figures, we can assume that the LTO-6 enterprise drive should be coming in at around \$20,000. An enterprise-class drive from Oracle or IBM could have a TCA as much as \$20,000 more. All of this makes a TCO study more relevant. These figures may seem high when contrasted to the cost of a new 2TB disk drive. However, the IT staff needs to keep in mind that the cost of additional media for LTO-6 drives will be much lower than the cost of additional disk drives and that the energy requirements for a tape library with drives is only a small fraction of what is required for spinning disks.

### **LTO Roadmap**

The current, published roadmap from the LTO Program has defined the capacity and performance levels of the next two generations beyond LTO-6 technology: LTO-7 and LTO-8 technology, which both continue to use a compression ratio of 2.5. The LTO-7 generation is being quoted with a compressed capacity of up to 16TB and a throughput of up to 788 MB/second (315 MB/s uncompressed). The LTO-8 generation is being quoted with a compressed capacity of up to 32TB and a throughput of up to 1180 MB/s (472 MB/s uncompressed). That said, we need to keep in mind that the LTO-6 generation carried a roadmap specification calling for 8TB (compressed) with a throughput of 525 MB/s, neither of which are expected in the final specifications. However, we do know that, in 2010, Fujifilm and IBM announced the capability to store 35TBs of data,

<sup>15</sup> These enterprise-class drives are inherently more reliable and usually more performant than the lower-priced "open systems" drives.

uncompressed, on a single cartridge.<sup>16</sup> The potential is there. This reinforces the possibility of fulfilling the goals as estimated on the LTO Program technology roadmap. (Currently there is no published roadmap from Oracle or IBM for their enterprise tape drives or from any disk vendor in terms of setting expectations beyond 4TBs per 3.5" disk).

## Conclusion

The availability of LTO-6 technology should have a major positive effect upon the long-term storage capabilities of the enterprise data center, in terms of TCO, capacity, performance, and reliability. We set this expectation by building upon the results reported in the December 2010 issue of **Clipper Notes**.

In that report, we determined that disk cost 15 times the average tape-based solution, based upon a total cost of ownership (TCO) in a long-term archiving scenario. In fact, **the cost of energy for the disk-based solution exceeded the entire TCO for the average tape-based solution**, using 238 times more energy. This study was undertaken using the latest storage technologies available as stated above, 1.5TB LTO-5 tape and 2TB SATA disks, for the initial 3-year cycle, ending in 2013. With the recent announcement of the availability of final LTO-6 specifications in August 2012 and with shipments expected about the end of this year, LTO-6 tape appears to be well ahead of the schedule assumptions made in our study, capacity assumptions notwithstanding. Even if we discount future roadmap assumptions by 25%, this still represents a vast improvement over what we do know about disk. We still do not have a firm schedule for the availability of 4TB drives in storage vendors' arrays.

LTO-6 technology will provide further impetus for the deployment of tape in the open system data center for long-term data protection and archiving. An enterprise data center could convert their LTO-5 media to LTO-6 media, reducing their library size in half. More likely, this will enable that data center to free up 50% of the slots in their library to handle future growth for the next 18 months, reducing the urgency to expand their library. In addition, with a throughput faster than LTO-5 drives, the data

center can reduce the window to archive that data or proceed with fewer LTO-6 drives. This will free up both the storage servers and storage administrators for more mission-critical applications.

In any case, with capacity and performance comparable to disk, and greater reliability than disk, at a much lower TCO, the enterprise data center cannot stick its head in the sand and ignore the facts. Naysayers aside, tape is here to stay, with LTO technology leading the way for the open systems data center. The architecture of your enterprise storage environment is up to you. As you begin to plan for 2013, 2H2012 is the time to think about what LTO-6 technology can do for your enterprise.



<sup>16</sup> See **The Clipper Group Navigator** entitled *IBM and Fujifilm Increase Tape Density – Raising the Bar on Tape Capacity* dated February 5, 2010, and available at <http://www.clipper.com/research/TCG2010003.pdf>.

### **About The Clipper Group, Inc.**

**The Clipper Group, Inc.**, is an independent 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](http://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](mailto: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.