



Enterprise Tape for Archival Storage? — Why This Just Might Make Sense

Analyst: David Reine

Enterprise Archive Storage Requirements

We need to put an end to at least one data center myth: **Tape is not dead!** Tape, in fact, is thriving in a myriad of environments, including as an active platform for NAS on tape, long-term archiving, and even backup. While some vendors continue to propagate the myth, a number of companies are availing themselves of a billion dollar market opportunity, including HP, IBM, Oracle, Quantum, Spectra Logic, and a litany of others, with innovation arriving yearly in terms of both hardware and software.

This innovation is necessary because over the past decade, storage requirements have exploded throughout the enterprise, not only in the data center, but also in departments and remote locations, such as regional branches. However, this is most evident in the enterprise data center. Moreover, within the data center, requirements for Tier-3 storage in support of long-term archiving have placed a tremendous strain on enterprise resources, both human and financial. In December 2010, The Clipper Group published an exhaustive study comparing the effects of both disk and tape on the enterprise budget, on a total cost of ownership (TCO) basis¹. For disk, we used 2TB SATA drives, the most commonly used high-capacity disk storage unit, *at that time*. For tape we used LTO-5 technology², the most commonly used tape architecture with the highest capacity open tape media, *at that time*. **Tape TCO came out the winner by well over an order of magnitude!** Well, LTO has changed over the past year with the specification announcement (in October, 2012) of LTO-6³ technology by The LTO Program, directed by a consortium of HP, IBM, and Quantum, followed more recently by product announcements. LTO-6 technology (when compared to LTO-5 technology) has increased the native capacity of a single cartridge from 1.5TB to 2.5TB, increasing the compressed capacity from up to 3.0TB to up to 6.25TB, with a 2.5:1 compression ratio; a significant increase in storage capacity. This increased capacity raised an interesting question: *Are there any economic or technical advantages to using an open tape architecture (i.e., LTO-6 tape) over enterprise tape, as offered by IBM and Oracle, in an open systems environment⁴? Or, perhaps, might enterprise tape offer a better value than LTO tape for open systems data?*

In 2011, we at The Clipper Group followed up our TCO study by publishing a bulletin potentially justifying the use of enterprise tape over LTO-5 tape, in any environment, open systems or

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* dated January 29, 2010, entitled *LTO Program Announces Next Gen Tape – LTO-5 Raises the Bar for Tier-3 Storage*, available at <http://www.clipper.com/research/TCG2010002.pdf>.

³ See *The Clipper Group Navigator* dated July 12, 2012, entitled *Magnetic Tape Turns 60 – The IT Industry Receives Another Gift*, available at <http://www.clipper.com/research/TCG2012015.pdf>.

⁴ The origins of enterprise tape lead back to the IBM mainframe and, later, to the *System i* operating environment (which now runs on *Power Systems*).

mainframe.⁵ That *Captain's Log* compared IBM's *TS1140*⁶ enterprise tape drive and media, with a native capacity of 4TB, and Oracle's *StorageTek T10000C* enterprise tape drive and media⁷, with a native capacity of up to 5TB⁸, to the industry-standard *LTO-5* architecture. Based upon capacity alone, the older enterprise drives would still appear to be superior when compared to the *LTO-6* technology. However, **capacity is definitely not the only factor that determines an enterprise decision on this tier of storage.** There are many factors that go into this decision, including acquisition cost, performance, availability, reliability, the TCO, and what already is on the floor of your data center, any of which may be the most important factor in your next tape procurement decision. We have now renewed that study to ascertain what advantages, if any, enterprise tape has over *LTO-6* tape. To learn more, please read on.

Total Cost of Ownership Factors

First, we must define the data center environment: To be clear, this is not about small businesses with a standalone drive or autostacker or a very small library with up to 10 or 12 slots and a single drive. However, there are some small enterprises with large collections of historical data that may benefit from the performance, reliability, and capacity available from enterprise drives. Strictly speaking, this is about the largest data centers; the ones with mission- and business-critical tape libraries, with more than a petabyte of data on hundreds or even thousands of cartridges.⁹ These data centers are using tape libraries as the primary long-term storage target for most, if not all, of their enterprise archiving requirements¹⁰. Some, but not all, of these data centers will be equipped with both mainframes and open systems servers, the latter running *Windows*, *Linux*, or some variant of *UNIX*, but others *only* will have open systems storage architectures. **If you already have both enterprise tape and LTO in your data center, it may be easier to consider using enterprise tape to store open systems data, as you already are committed to the enterprise format. However, if you only have LTO libraries, then adding enterprise tape will require further justification.**

It goes without saying that enterprise tape has a higher per-cartridge cost than *LTO* tape, even when compared to *LTO-6* tape with its increased per-cartridge cost over *LTO-5* tape. However, calculating the contribution of tape media towards the TCO requires us to consider the quantity of cartridges and library slots needed to achieve the desired storage capacity. Deployed in enterprise-class tape libraries, Fibre Channel (FC) enterprise drives¹¹ also have somewhat higher per-unit costs than FC *LTO* drives. But, using the same logic as is used to determine media costs, the tape drive contribution towards the TCO must consider the quantity of tape drives needed to achieve a desired overall throughput. **As you will see below, when comparing the TCO of enterprise tape to LTO, cost tends to fall away as a key factor for deciding between the two technologies.** Enterprise tape simply delivers more in quantifiable values, such as capacity per cartridge and throughput, than *LTO-6* technology. This, in turn, could mean fewer drives and cartridges to manage, fewer library slots to hold the cartridges, and potentially fewer frames to acquire than with *LTO-6* tape. With fewer drives, cartridges, slots, and frames, and less floor space, the cost per TB for archiving might be reduced significantly.¹² Then there is reliability. If you have a mainframe, you already know about the reliability attributes of enterprise tape drives. They are worth every penny that you paid for them. Moreover, if you have deployed spare

⁵ See the issue of *The Clipper Group Captain's Log* dated July 12, 2011, entitled *Ten Reasons Why You should Consider Enterprise-Class Tape for Open Systems Storage*, and available at <http://www.clipper.com/research/TCG2011025.pdf>.

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

⁷ See **The Clipper Group Navigator** dated January 31, 2011, entitled *Oracle Fulfills Commitment – StorageTek T10000C Takes Leap Ahead*, available at <http://www.clipper.com/research/TCG2011003.pdf>.

⁸ Up to 5.5TB, with the Oracle T10000C StorageTek Maximum Capacity Feature.

⁹ As a general rule, SMBs do not have the requirements to justify the deployment of enterprise-scale tape.

¹⁰ To be clear, this is not a reference to storing backup data for a long time. This is about archiving business data, usually files or objects, for a long time.

¹¹ In an open systems environment, FICON, a more expensive interface, is not required.

¹² This is situation dependent. While we think that these conclusions are generally true, there may be situations where the savings might be less. For example, if you need a certain number of drives for simultaneous access (writing and/or reading), using drives that can write or read faster might not result in reducing the number of drives that are required, potentially negating the TCO savings.

drives (LTO or enterprise drives), the drive availability concern is lessened significantly in terms of downtime.¹³

The roadmap for enterprise tape is also a significant factor in the IT solution. The current generation of enterprise drives and cartridges (released between January 2011 (Oracle StorageTek T10000C) and June 2011 (IBM *System Storage TS1140*)) has significantly higher native capacity and delivers significantly higher throughput than LTO-6 technology. LTO-6 technology is at 2.5 TBs per cartridge (uncompressed) and has a native transfer rate of 160 MB/second. Oracle T10000C cartridges have a 5TB capacity (native) with a native transfer rate of up to 240 MB/second, while the IBM TS1140 cartridge has a native capacity of 4TB, with a native transfer rate of up to 250TB/second.

In addition, we can expect that the next generation of enterprise drives, likely out by 2014, will have even greater densities than they do today¹⁴, somewhere in the range of 50% to 100% more. That will put the next generation of enterprise cartridges at about 6-10 TBs (uncompressed)¹⁵. That will be about three times the uncompressed capacity of LTO-6. Furthermore, LTO-6 is offered with an average compression ratio of up to 2.5:1; today's enterprise drives have a compression ratio of up to 3:1¹⁶. For large data collections, greater compressed capacity equates to fewer library slots, less floor space consumption, and, probably, fewer drives. In addition to these quantitative capacity calculations, there also is the benefit of using what likely will be enterprise drives that are capable of delivering higher reliability and availability. In addition, future enterprise drives are expected to use the existing 4 – 5.5TB media at higher capacities, while LTO media is expected to remain at currently stated limits for several years. **Nonetheless, the economic question almost always comes down to TCO per TB. Put more specifically, if enterprise media costs about the same per TB as LTO media, and if the performance per dollar invested in libraries and tape drives is about the same, which should you buy for large-scale archiving?**

Let's take a deeper look at the TCO considerations that go into any tape storage acquisition and the components that comprise it.

- 1. Drive acquisition cost – The acquisition cost of a tape drive is relatively minor in comparison to the cost of the media and frames required to hold and store the data.** We have noticed that the list price of an LTO-6 tape drive for installation in an enterprise library may be \$5K-\$15K less (per drive) than the price of one of the new enterprise drives. Moreover, because of the throughput advantages of enterprise drives, their larger capacities, and their higher compression ratios, you may need fewer drives to do the same amount of work. In addition, if the data center has a heterogeneous environment with both open systems servers and mainframes, fewer frames may be required as the enterprise libraries from IBM and Oracle can support both LTO-6 and enterprise media.¹⁷ In these cases, **enterprise drives may end up costing less than LTO-6 drives because the drives' cost contribution to TCO per TB may be less when enterprise tape drives are deployed.**¹⁸
- 2. Media acquisition cost – The cost for a drive is not nearly as important as the cost per TB for the media.** An LTO-6 cartridge may only cost about \$110 while some enterprise cartridges might cost \$275-\$325 each. The difference in cost/TB, however, is not nearly as significant. For LTO-6 media, the cost/TB (uncompressed) is about \$44/TB, and about \$17.60/TB compressed at 2.5:1. The cost/TB for an enterprise cartridge (uncompressed) is now about \$55/TB, and about \$18/TB

¹³ A drive can be considered a spare when it is above and beyond the number needed to handle the largest expected peak activity.

¹⁴ At this point in time, LTO tape and enterprise tape generations seem to be on staggered announcement schedules. This is being played out as a game of leapfrog. The last to leap was LTO in 2012 with LTO-6. We expect that the next announcements will come from the enterprise tape providers, IBM and Oracle. LTO-7 might appear in 2015, maybe a little earlier.

¹⁵ The current LTO roadmap call for LTO-7 to have a native capacity of 6.4TB and is expected to be available by 2015.

¹⁶ Actual compression definitely will vary depending upon the data being compressed.

¹⁷ IBM requires that LTO and enterprise tape cartridges be segregated into separate frames. Oracle does not.

¹⁸ The same logic applies to out-of-warranty maintenance. However, because maintenance usually is based on purchase price, the cost of maintaining an enterprise-class drive will be higher per drive than for LTO.

with a compression ratio of 3:1¹⁹. As stated above, some enterprise drives enable the IT staff to reformat the previous-generation media to a higher density. Future enterprise drives are expected to use the existing 4 to 5TB media at higher capacities, while LTO media is generationally bound to its native capacity. This can be a TCO bonus. You get to reuse older media and can store more capacity. LTO drives require new media to achieve the higher capacity of newly released drives.

3. **Number of slots needed – With a capacity of 4 to 5TB²⁰ today, the enterprise drive has a significant advantage over LTO-6 tape**, in terms of the fewer number of cartridges (and slots) needed and, therefore, (potentially) a need for fewer library frames. The number of enterprise cartridges required to hold 20 PBs of archived data today, based upon each drive type's standard compression ratio, is significantly more for LTO-6 tape – 3200 slots compared to between 1200 and 1600 slots for enterprise cartridges. Ignoring compression, today's enterprise drives require only 45-65% of the slots required by LTO-6 tape. It may be assumed that the next generation of enterprise drives expected by 2014 will have 50% to 100% higher capacity and the gap will increase until LTO-7 technology appears (presumably in by 2015).
4. **Less floor space – With significantly fewer enterprise cartridges (and slots), the data center is able to support a digital archive in less floor space with enterprise tape than with LTO tape.** Fewer slots usually means fewer frames need to be deployed, potentially increasing the useful life of the existing floor space, and possibly saving millions of dollars in new construction costs.
5. **Throughput performance – The native sustained data rate of an enterprise drive is approximately 250 MB/second, while an LTO-6 drive has a native sustained data rate of 160 MB/second, potentially giving an enterprise drive more than 50% higher throughput.** With two enterprise drives installed, the data center can deliver up to 500 MB/second of throughput. Achieving a comparable throughput with LTO-6 drives would require three drives (to deliver about 480 MB/second). If the data center needs three enterprise drives (an investment of \$90K or \$120K, depending on the vendor), they would require five LTO-6 drives (at a cost of approximately \$125K). Reducing the number of drives required lowers the TCO, due to lessened acquisition and maintenance costs. Obviously, if you need a lot of drives, this might be a significant reduction in capital costs and future maintenance. This leaves the current generation of enterprise drives with a noticeable edge over the current LTO technology. In addition, the number of exchanges can be significantly less with the higher capacity enterprise cartridges as more data can be written or recalled in a single cartridge exchange. While we can only estimate the throughput for the next generation of enterprise drives, it would not be unreasonable to expect an increase in performance when compared to the previous generation.²¹
6. **Other performance metrics – Comparing performance metrics between LTO and enterprise drives also reveals advantages for enterprise tape.** The average access time for LTO-6 drives is about 50 seconds for a 2.5TB cartridge. The average access time for an enterprise drive is 59 to 70 seconds for a full-length cartridge of up to double the capacity of LTO-6 cartridges. The rewind time is also interesting, with a maximum rewind time of 90-100 seconds for LTO-6 cartridges and a rewind time of 76-115 seconds for standard enterprise cartridges. These may equate to noticeably better operational performance for the enterprise tape solution.
7. **Maintenance – While it comes without saying, fewer drives and possibly fewer frames can mean reduced maintenance costs, after the warranty expires.**
8. **Reliability – Enterprise tape drives and cartridges historically have delivered higher reliability and more features**, such as FICON for mainframe support, than more open designs. Enterprise tape's superiority comes from a mainframe heritage where the required duty cycles are higher and superior reliability is expected. Furthermore, the Uncorrected Bit Error Rate (UBER) for enterprise

¹⁹ While the compression algorithms in both LTO and enterprise-class drives are similar, mainframe data typically compresses at a higher rate than open systems data, 3:1 compared to 2.5:1 for LTO-6. Your mileage may vary based on your data and the fullness of your cartridges.

²⁰ The Oracle T10000C is rated at 5.5TB with its *StorageTek Maximum Capacity* option.

²¹ And while we can only look at the past, the price of a drive tends to go up a little between generations, but far less than the proportional increase in capacity, etc.

drives can be up to 100 times more reliable. This may not be a meaningful TCO issue, but it is a noteworthy difference.

9. ***I/O Interface – LTO-6 drives come with both a 6 Gb/second SAS interface and an 8 Gb/ second FC interface for automation. Enterprise drives come with up to an 8 Gb/ second FC interface in addition to FICON for connection to a mainframe.***²² Your existing IT infrastructure may well dictate which direction you need to take. **The question being raised in this paper is whether you should be using enterprise drives for open systems use, which implies Fibre Channel connectivity.** If you have mainframes and open systems, only enterprise drives offer connectivity to both network.²³
10. ***Energy consumption – The total energy consumption of enterprise drives continues to go down given that fewer drives are required with each new generation.*** This factor also is relevant when comparing enterprise drives with LTO-6 drives, based upon the total tape drive throughput required. For an equivalent amount of tape drive throughput, an enterprise solution may require fewer drives, and thus less energy for powering and cooling the tape drives. However, the amount of energy consumed is inconsequential, especially when compared to the energy consumed by disk in an archival environment.
11. ***Other Features – Enterprise tape drives tend to offer a more robust set of features than LTO drives,*** mostly designed to improve efficiency in writing data with a high frequency of syncs and tape marks. While some vendors' LTO drives offer a variety of enhanced data integrity features, enterprise drives offer a more robust set of features for validating data integrity.

Conclusion

From a pure capacity and performance standpoint, enterprise drives from IBM and Oracle have distinct advantages over LTO-6 drives, not only today, but seemingly for the foreseeable future. **From a TCO standpoint, the total cost of ownership of an enterprise tape solution may not be any (or much) higher at all than a midrange LTO solution. In fact, depending upon your configuration and your environment, enterprise tape's TCO could be lower for an equivalent capacity and performance using LTO-6.**

From a functionality standpoint, LTO-6 technology has features and functions quite comparable to those offered with enterprise drives. From a usage and reliability standpoint, however, enterprise drives have been designed to support the unique applications of the mainframe environment and are regarded as superior to midrange drives, proving themselves in the most demanding of mainframe environments.

In the long run, the capacity, throughput, and reliability advantages of an enterprise tape solution could allow you to do more with less – less tape drives, less cartridges, less library frames, and less hassle. Doing the math may prove that you need to spend less, too!



²² This is one area where the approach to connectivity differs between IBM and Oracle. This may add to your TCO, depending on how you plan to handle FICON and Fibre Channel in a mixed environment. For example, do you (either by desire or architectural necessity) have to dedicate certain enterprise-class drives to be used only for mainframes and only for open systems?

²³ There is an exception. If you are only running Linux on IFLs in your mainframe, you don't necessarily need FICON/ESCON connectivity to disk and tape.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-first 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.