



IBM's New Enterprise Tape Extends Data Retention Capabilities and Lowers the Cost of Data Protection

Analyst: David Reine

Management Summary

Traveling around the United States is relatively easy. We have an Interstate Highway System that was initiated by President Dwight Eisenhower in the mid-1950's to facilitate the movement of people and goods from state to state, and coast to coast, supplementing the local roads that previously bore the brunt of the load. It can be a simple, straight shot from east to west on I90, or an equally easy ride from north to south on I95. We move people along these freeways in cars; vans and buses can be used to make more efficient use of fuel and improve the person miles per gallon ratio. We consolidate freight into larger tandem trailers in order to get our food to market as quickly and economically as possible. Unfortunately, there are times when so many of us decide to use these roads that traffic jams are created, leaving our cars and trucks stuck in a hopeless quagmire, not to mention what happens when snow makes one of these highways impassable or an accident ties up traffic for hours. Fortunately, there are always alternative routes via parallel highways. US1, for example, parallels I95 along much of the East Coast. Many connector roads extend from the local highway, with all of its strip malls and traffic lights, to the on-ramps of the Interstate. The local roads may be slower when normal conditions prevail on the Interstate, but when accidents occur, the local road can be a superior alternative. Sometimes, slower is the better!

The same scenario is true in our enterprise data centers. With the glut of data proliferation in the IT infrastructure, moving it, storing it, and protecting it has become a serious issue for the CIO of every enterprise, and a huge consumer of data center budget, putting the development of new, revenue-producing, mission-critical applications at risk. Whether the application calls for a simple backup or the creation of a data archive, with more data to store, the IT staff must resort to higher capacity vehicles (or cartridges) to consolidate big data requirements into the smallest possible footprint, while ensuring the fastest possible creation of, and access to, that data, when the requirement arises.

Last quarter, Oracle introduced a new, high-capacity tape technology with improved performance. However, IBM has attacked the problem of protecting and retaining more and more data, perhaps for decades, with even faster access from the tape drive and the library flexibility that the data center needs to support many of their enterprise applications. With this latest announcement, IBM provides a new, high-capacity tape drive/cartridge offering and new library features in support of a highly-efficient backup and recovery system, continuous data availability, and smart archiving. To learn more about the business value of the new IBM *System Storage TS1140 Tape Drive* and enhanced *System Storage TS3500 Library*, please read on.

IN THIS ISSUE

- **Data Protection in the Data Center** 2
- **New System Storage TS1140**..... 3
- **System Storage TS3500 Functionality** 4
- **Conclusion** 6

Data Protection in the Data Center

The proliferation of application servers and storage devices throughout the IT infrastructure has had a profound effect upon the budget for enterprise data center managers. The total cost of ownership (TCO) of these infrastructure components has risen precipitously due to inefficient servers and the growth of data storage, creating a complex infrastructure that is difficult to manage. The issue of server inefficiency has been addressed with a continual improvement in the capabilities of x86 microprocessors from Intel and AMD. Extended processing power through a multi-core design has been the key component in the ability of servers to consolidate and virtualize applications, reducing the extent of server sprawl, simplifying the infrastructure, and lowering the TCO of the environment.

In the past few years, many enterprises have seen a growth rate of 40-50% annually in the amount of storage being consumed by mission-critical and business-critical applications.¹ This is reflected in the IT budget by the acquisition and support of high-performance Fibre Channel (FC) and SAS arrays for Tier-1 storage, and SATA disk and high-capacity tape for Tier-2 or Tier-3 storage capacity. The growth rate for secondary storage has been even greater than the overall rate, often reaching 50%. This is reflected in the acquisition of SATA disks and enterprise tape products being used for data protection and data retention (the backup of critical data and the archiving of historical data) of these mission- and business-critical applications. Storage capacity for this class of data has grown from gigabytes to petabytes; exabytes of storage are predicted for 2012 and the specter of zettabytes is on the horizon by 2015 (Can you even comprehend a billion trillion bytes?). The challenge of addressing this growth in data protection and retention, and simplifying the storage infrastructure, is at hand – *today*.

Enterprises of all kinds are keeping information on hand for longer periods of time, as the value of that data continues to grow, for the purposes of data protection, continuous availability, and data retention. Enterprises now are once again turning to tape for its value as a medium for data protection and are planning the strategic use of tape libraries to consolidate historical data into

*active archives*² to meet the growing requirement for the searching and analyzing of historical business data, with the intention of increasing its business value. For example, oil and gas companies are now required to analyze up to twice the amount of information before committing to expensive bids for leases. The U.S. healthcare industry is under great pressure to store patient health records accurately, to protect and preserve them, and to share them safely, yet promptly.

Even though there may be additional costs associated with deploying an adequate secondary storage capability, it is far less than the cost of preparing for, and possibly losing, for example, lawsuits for failure to provide an audit trail of deletion of embarrassing emails or financial records. The best practices for every enterprise now include the requirements for a more highly-efficient backup and recovery system to help to reduce the TCO for data protection and retention. These practices also now account for the continuous data availability of data that may be too large to provide a rapid recovery and provide the real-time protection required by Internet commerce. Furthermore, smart (or active) archiving is needed to lower the TCO of writing, searching, and reading, all of this historical data.

In addition, the requirement for *immediate* recovery of data to replace corrupted or deleted data will never go away. We will need to retain some level of Disk-to-Disk (D2D) data protection, perhaps to a VTL³, in order to maintain a 24x7x365 availability. However, in most cases, back up to tape in a disk-to-tape (D2T) or disk-to-disk-to-tape (D2D2T) hybrid environment will be sufficient to control the spiraling costs. In order to lower costs in this category, the enterprise needs to deploy higher-capacity cartridges, with higher data transfer rates, and engage in better management of those cartridges in high-density frames. Reducing the costs of the enterprise data protection and retention environment will mean more budget availability for the development of additional mission-critical applications.

Until recently, the industry has been fairly stagnant in the improvement of enterprise-class tape devices⁴, which have traditionally been

¹ For some enterprises, the growth rate approaches 100% or more.

² 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>.

³ Virtual Tape Library.

⁴ Enterprise-class tape products also are called *proprietary*, because only one vendor provides each of the product lines, in

desired for their faster access time and higher performance in reading and writing, as well as higher level of reliability. We have seen a steady improvement in the LTO format, with capacities rising from the original version, which was released in 2000 and could hold 100 GB of uncompressed data in a single cartridge, to the most recent version, released in 2010 and capable of holding 1.5TB in the same size cartridge. We have also seen LTO add new features such as WORM, encryption, and partitioning over the years. Enterprise tape capacities had been holding at the 1TB level, uncompressed, since the introduction of IBM's *System Storage TS1130 Tape Drive*⁵ and Oracle's (née Sun, née STK) *StorageTek T10000B* in 2008, until the announcement of Oracle's new *StorageTek T10000C* at 5TB earlier this year.

IBM has now released the new *System Storage TS1140 Tape Drive* optimized with a 4TB native cartridge capacity, along with industry-leading improved performance characteristics and low energy consumption to reduce the TCO of the storage infrastructure, while still protecting the investment of mission-critical data centers in existing technology. If history is a correct indicator, this new 4TB tape should be reusable in future generations of drives at a significantly higher density, extending the return on the investment being made today.

New System Storage TS1140

IBM's *TS1140* is the fourth generation of IBM's *3592* tape format. With four times the capacity of the *TS1130*, from 1TB to 4TB, the data center can reduce the number of library frames required, conserving precious data center floor space, or storing four times more data in the same floor space. The *TS1140* also has 2.67 times the capacity of the *LTO-5*.

With the industry's highest native tape drive data rate of 250MB/second and a maximum compressed data rate of 650MB/second, the *TS1140* is 56% faster than the *TS1130*, which has a maximum data rate of 160MB/second. It also has almost 80% higher performance than *LTO-5*, which has a native data rate of 140MB/ second.

In terms of access performance, the *TS1140* is

clearly the industry leader, supporting faster data recalls, data searches, and locates to append positions for writing, all very common and time-consuming parts of tape operations. The Oracle cartridge is 31% longer than the *TS1140* cartridge, while the *TS1140* cartridge delivers 80% of the capacity of the longer *T10000C* cartridge by recording at higher density. The high-recording density of the *TS1140* drive, coupled with higher search speed than the *T10000C*, results in significant access time performance advantages. For example, the *TS1140* has a maximum rewind time of 76 seconds while the *T10000C* takes 115 seconds, more than 50% slower. In addition to the 4TB cartridge, the *TS1140* can also use a 500GB economy cartridge to provide an alternative to the much higher-capacity cartridge, providing a lower price-point cartridge with faster search and rewind times. Additionally, the IBM Economy cartridge rewinds 80% faster than the 1TB *T10000C Sport* cartridge.

With media usually being the most significant element in the TCO of a tape library⁶, the *TS1140* may lower the cost per terabyte for storage. With a list price of about \$8500 for 20 cartridges, the *TS1140* cartridge will result in a cost of about \$106/TB. The increase in throughput of the *TS1140* drive will minimize the number of drives required, with each drive carrying a list price of just under \$43K. In order to protect the investment that the enterprise has made in existing media, IBM has enabled the data center staff to continue to re-use *TS1130* media, reformatting the cartridge, and increasing capacity by 60%, to 1.6TB⁷, with an improvement in throughput from 160MB/ second to 200MB/second. In fact, the IT staff can also field upgrade the *TS1130* drive itself with *TS1140* functionality to protect the investment that the enterprise has made in *TS1130* hardware, with a further expectation to upgrade the *TS1140* to the next generation of tape drives when they become available.

A separate 4TB cartridge is also available with WORM functionality. The *TS1140* comes with improved encryption support, enabling *Library Managed Encryption (LME)*, *System Managed Encryption (SME)*, *Application Managed Encryption (AME)*, and *T10 ANSI Standard*

this case, IBM and Oracle. This is in contrast to *open tape* products, like *LTO*, which are sold by many vendors.

⁵ See [The Clipper Group Navigator](#) entitled *Lowering the TCO of the Data Center - IBM Innovates Tape Architecture ... Again!* dated July 16, 2008, and available at <http://www.clipper.com/research/TCG2008037.pdf>.

⁶ 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>.

⁷ The Oracle *T10000C* drive can only read *T10000B* media, and cannot write or reformat the legacy media.

Exhibit 1 – Comparing the Tape Drives – TS1140 vs. T10000C

- T10000C has higher cartridge capacity – 5TB to 4TB, but:
- TS1140 has a faster sustained data rate – 250MB/second versus 240MB/second.
- TS1140 has a faster compressed data rate – 650MB/second versus 360MB/second (80% better).
- TS1140 has a faster I/O interface – 8Gb per second versus 4Gb/second. IBM has dual ports for failover.
- TS1140 has a faster access time – 39 seconds versus 57 seconds (each for their full-length cartridge).
- TS1140 has a higher probability for continuous streaming with 12 speed-matching settings (30-250MB/second), *SkipSync*, and *FastSync* virtual backhitch, while the T10000C uses two speed match settings and a 2GB buffer to improve performance.
- TS1140 consumes less energy – 51W in active mode and 24W in idle mode versus 67W for the T10000C in all modes.

Sources: IBM and Oracle

encryption methods for improved data security. Separate encryption methods can be implemented for each TS3500 logical library.

The TS1140 also features *SkipSync*, which avoids performing a backhitch and continues to stream data on the same wrap. This results in a significant improvement in performance, but does reduce the overall available capacity on the volume. The *FastSync* command detects that small amounts of data are being transferred between Synchronize operations and changes mode. In the new mode, the drive uses an alternate area of tape as a work area to write working copies of datasets to tape without backhitching. This technique enhances the effective drive data rate while maintaining full non-volatile integrity of the data and full utilization of cartridge capacity. The drive is enabled for *LTFSS*⁸, introducing the file-access functionality into the enterprise-class tape media for the first time. The TS1140 boasts a new 32-channel enhanced ECC recording format and a new third-generation GMR head technology, which help to provide enhanced data integrity.

⁸ Linear Tape File System (LTFS) refers to both the format of data recorded on magnetic tape and the implementation of specific software that uses this data format to provide a file system interface to data stored on magnetic tape.

How does IBM's TS1140 compare overall to the Oracle StorageTek T10000C? While the T10000C does have a 25% edge in raw capacity, the TS1140 leapfrogs over Oracle's T10000C in every other category. As a result, IBM's TS1140 offers an optimized blend of high capacity, high performance, low energy consumption, and low TCO. See Exhibit 1, at the left, for a detailed comparison of many of the specifications.

System Storage TS3500 Enhanced Functionality

With a total consolidated capacity in excess of 2.7 Exabytes⁹, the IBM System Storage TS3500 Tape Library extends the scalability of tape libraries to unprecedented capacities (i.e., number of tape slots). This increased scalability also applies to the number of drives, with configurability for up to 2700 TS1140 drives and the capability to complete 15,000 mounts per hour across 15 library strings¹⁰. These library strings can be interconnected into a shuttle complex via a new library *Shuttle Connection* and a *Shuttle Car* (see Exhibit 2, on the next page), with the shuttle complex capability of holding up to 300,000 LTO cartridges or 225,000 TS11x0 enterprise cartridges. The shuttle complex enables the data center to optimize floor space, consolidate backup, continuous data protection, and archiving functionality, and to share tape drives between library strings. This eliminates the requirement to have duplicate or spare drives installed in each string, where actual usage may be low, or infrequent (as may be the case for drives retained to read and write older formats). Once again, this lowers the TCO of the IT infrastructure.

In a Shuttle Complex, the TS3500 can transfer any cartridge from a high-density frame in one string to another high-density frame in another string. The TS3500 can be accessed as a single library image to facilitate drive sharing, potentially enabling the data center to reduce the number of tape drives required, especially for older drives kept for compatibility reasons and for spare drives.¹¹ This also enables the IT staff to con-

⁹ Configured with 225,000 4TB cartridges and a 3:1 compression ratio.

¹⁰ A "string" is a chain of frames linked in a row together to create a single library.

¹¹ Rather than keeping n-3 generation or older drives (or spares of any age) in several strings, now older drives (and spares) might only be needed in one string, which can then service older cartridges stored in any of the strings. Of course, this depends on reading and writing patterns, but the point remains – you

Exhibit 2 — TS3500 (3 Strings with 3 Library Connections)



Source: IBM

figure their TS3500 to fit the available data center space without worrying about moving a wall to create a contiguous run. Cartridges can be moved between any two high-density frames. The Shuttle Complex supports all versions of enterprise TS11x0 cartridges, all models of LTO cartridges, and all models of high-density frames, without intermediate robot hand-offs, i.e., the cartridge can make a “direct flight” to the desired library string without stops in-between, enabling high-performance operations.

With all of this flexibility and performance, the data center has multiple options in the consolidation of frames to make better use of more data, while lowering the TCO of the data retention infrastructure. In addition to the new, high-capacity tape drive and Shuttle Connection components, IBM is also introducing LTFS functionality for the library for the first time. With the new *LTFS Library Edition (LE)*, the data center can use the library as an active archive, graphically displaying the contents of a portion of the library, or the entire library in GUI format, typically a folder or tree, eliminating the requirement to migrate data to disk for analysis, simplifying the movement of data throughout the storage architecture. The LTFS metadata and its index are cached in server memory, after the cartridges are mounted and inventoried. That metadata remains viewable and searchable, even after ejecting the cartridge. The application can browse the directory and do a filename search without having to remount the tape, improving search and library

performance. Every cartridge is accessible via standard operating system file system commands from the applications. Each cartridge becomes *self-describing* once the cartridge is formatted with LTFS.

The library comparison between IBM and Oracle is even more important than the drive comparison in terms of TCO, because this is where “the rubber meets the road”. (See Exhibit 3 on the next page.) Even though the T10000C cartridge has a higher capacity than the TS1140 cartridge, the newly-enhanced IBM TS3500 Library has significantly higher capacity than the Oracle StorageTek SL8500 Library. With a maximum slot count of 300,000 (225,000 for TS1140) and a native storage capability with the TS1140 of 900PB (2.7EB compressed), the TS3500 far exceeds the Oracle *SL8500* specification for 100,000 slots and 500PB of uncompressed T10000c capacity. In terms of the number of possible drives (another parameter of scalability, the TS3500 also wins out, 2700 to 640; in other words, the IBM solution can have more than four times as many drives as can Oracle’s solution.

Each individual TS3500 library “string” can support up to 20,000 slots and 192 drives, compared to 10,000 slots and 64 drives for the SL8500. In terms of performance, once again the TS3500 bests the SL8500 with a typical cell (slot)-to-drive time of less than 4.7 seconds versus less than 11 seconds for the SL8500. In terms of cartridge swap time, the TS3500 has a typical exchange time of 4-6 seconds while the SL8500 is

might get by with fewer drives given the new connectivity between the strings.

Exhibit 3 — Comparison of IBM and Oracle Tape Libraries (with maximum interconnections)

Library Complex Comparison			
Key benefits / attributes	Measurement	IBM TS3500	Oracle SL8500
Performance	Max. uncompressed throughput	2430 TB/hour	553 TB/hour
	Max. library-library intermediary hand-offs	0	8
Capacity	Maximum cartridge slots	300,000	100,000
	Max. total uncompressed capacity for largest enterprise cartridge for the maximum number of slots	900 PB	500 PB
	Maximum tape drives	2700	640
	Maximum "strings"	15	10
Single Library Comparison			
Key Benefits / Attributes	Measurement	IBM TS3500	Oracle SL8500
Performance	Typical Cell to drive time	2.4 – 4.7 seconds	< 11 seconds
	Typical Exchange Time	4 – 6 seconds	2-10 seconds
	Audit Times	1 minute per standard frame and 10 minutes per high-density frame	< 40 minutes
Capacity	Minimum slot capacity, base frame	58 slots	1,448 slots
	Maximum slot capacity per "string"	20,000	10,000
	Maximum tape drives per "string"	192	64
	Maximum "strings"	15	10
Complexity	Number of robots	2 for high availability	8 for high availability
	Elevator mechanism	Not required	Required to move cartridge to another rail
	Required Software and server	None	ACSLs on Sun server
Upgradeability	Slot capacity upgrades	60 minutes	4-6 hours

Sources: IBM and Oracle

rated at 25 seconds¹². In terms of ease-of-use, the IT staff can add a frame to the library with less than 60 minutes of downtime, while it might take a shift of downtime to accomplish the same upgrade for the SL8500.

Conclusion

With data exploding all around us, it is more urgent than ever that the data center get control of the data protection and data retention of this valuable business asset. With IBM's TS1140 Tape Drive and TS3500 Library, the data center has the capability to deploy a single system with higher capacity cartridges and a higher capacity library. They can consolidate all of their backup

data, all of their archive data, and all of the images required for continuous data availability, into a single library view, reducing the TCO of the IT infrastructure while also improving the access time and data throughput performance of these processes.

If you need to simplify your data retention requirements *and* ease the burden on your IT infrastructure budget as it accommodates and protects even more data, then you should consider these new high-capacity, high-performance, enterprise tape storage offerings from IBM.



¹² Excluding shuttle time for the Connector, which might add up to 5 seconds.

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.***

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.