



IBM Delivers Advanced Enterprise Tape for Open Systems and Mainframes

Analyst: David Reine

Management Summary

We Americans certainly do like our automobiles. We spend literally hours a day in our cars. We drive to work; we drive our children to school. We drive to the supermarket; **we even drive to the gym!** We spend so much time in our cars that we have turned our highways, our transportation arteries, into candidates for bypass surgery.

Here in Boston, we are quite familiar with the problem of putting more vehicles on a single stretch of highway than it was designed to support. Built in 1959, The Central Artery, running north to south through the middle of the city, consists of three lanes, and multiple potholes, in either direction. The city designed the road to handle (support) 75,000 vehicles a day – no problem in 1959. Today, the elevated highway supports 190,000 vehicles a day. A highway which opened with minor traffic jams – normal morning and evening rush hour– evolved into 8 – 10 hours a day of transportation chaos: not enough lanes, a road in disrepair, **a crash waiting to happen!** If the city made no changes, by 2010 there would be 245,000 vehicles a day, creating 15-16 hours of traffic jams every weekday. **That is every waking hour!**

The answer for this automotive constipation in Boston was a 10 lane depressed highway designed to support the expected capacity, and more. The city designed the new system to manage the required throughput by speeding the vehicles onto and off the highway, via a network of 14 traffic ramps. The cost for the Big Dig to restore mobility to Boston: **\$14B and counting.**

If you are a CIO or IT Director, this scenario may sound very familiar. How do you manage the ever-increasing traffic in your I/O subsystem? Do you know how you are going to complete a 10-hour backup in 4 hours? Are you having issues with cartridge capacity? Is the throughput to your existing library meeting your current requirements? What will you do about tomorrow?

To find out how IBM plans to solve both your capacity and throughput problems with their new Enterprise Tape Drive, the 3592, a single, low cost (<\$14B) solution that is available for both mainframe and open systems markets, please read on.

IN THIS ISSUE

➤ Traffic Jams on the Info Bi-Ways	2
➤ An Evolving Technology	3
➤ A New Tape Drive – The 3592	3
➤ Conclusion	4

Traffic Jams on the Info Bi-Ways

If we look at the Internet as *The Information Highway*, then the various I/O paths within a single enterprise make up an *Information Bi-Way Infrastructure*. This combination of avenues, streets, and paths interconnects the CPUs within the most critical enterprise servers to memory, cache and I/O controllers, for a variety of storage devices. Like our nation's highways, **the Information Infrastructure needs to have the throughput and capacity to handle a perpetually increasing workload, both in the amount of data traveling along its paths and the number of servers trying to access it.**

Corporate CIOs and IT directors are constantly trying to cure the pain caused by a variety of demands on their capabilities:

- The need to get the information required by management into their hands in a timely fashion;
- The requirement to protect critical corporate assets (data) within a designated time window;
- The need to improve the system infrastructure within the approved corporate business needs, guidelines and processes.

As the amount of data managed by IT explodes from gigabytes to terabytes and beyond, the cost to store it and to access it continues to grow. **IT is increasingly under pressure to provide access to the historical data in the library not only to the mainframes in the data center, but also to the large open systems servers that are propagating throughout the enterprise, including the data center.**

IT departments implement new multi-tiered storage architectures using policy-based hierarchical storage management (HSM) rules in order to try to control storage costs. They continue to adjust the definition of on-line, nearline, and off-line, concerning where this expanding data needs to be stored. (See box at right.) This is especially true in environments where historical data needs to be accessible within a *reasonable* amount of time, where reasonable can be seconds as opposed to milliseconds. New advances in high performance computing (HPC) have enabled new applications that demand high volumes of

Tape Storage Status

- **Online** – Resident on hard disk, SCSI, ATA; access time in milliseconds, for critical business needs
- **Nearline** – Resident on tape in a cartridge loader or library; access time in seconds.
- **Offline** – Resident on tape on a shelf available by hand to the data center; access time in minutes.
- **Remote** – Resident on tape, in a vault at a remote location; access time in hours, perhaps a day.

data, for example, in the area of the study of weather patterns or in seismological research. Some of this data is historical, some dynamic. In either case, the volume of data is so high that not all corporations can justify the business value to keep all of the data on high-speed disk devices; nor do they need the access time associated to those devices. IT establishes rules to keep some data, say five-days worth, on high-speed disk. These rules also specify a period for less urgent data, perhaps 30-days worth, retained on slower disk devices, like ATA-based arrays, for example. Beyond that, all data may be archived to tape, in a library perhaps, for retrieval as required.

In the public sector arena, two prime examples demonstrate the need for expanding storage. In the aftermath of September 11th, intelligence gathering organizations are collecting every scrap of information, recording every call, documenting the travels of every suspected terrorist. In the short term, this data must be available for instant access. As time passes, however, the need to have immediate access to the data becomes less urgent, but the data must be retained. It can be archived to tape, should the need for it return.

The need for the retrieval of archival data has never been as current as with the investigation into the shuttle disaster by NASA. The need to examine maintenance records, for example, was crucial to the investigation. However, the need to see them in milliseconds was not a requirement. It was perfectly acceptable to expect five-year-old

records to be available in a few minutes.

By using tape in a data processing environment such as this, the enterprise, or the government, historically has saved money not only in media costs, but also in terms of power requirements, floor space, air conditioning, etc. Tape also provides IT with a more secure historical environment than disk, especially when you are required to preserve 10, 20 or even 30 years of data. What is the cost of trying to keep that much data for that much time on disk? Can a disk even be kept spinning for that length of time? **This results in one unmistakable conclusion: there is still a requirement for tape.**

However, as the volume of data increases, the costs to store that data, both nearline and offline, can rise, as the collection of data outstrips the capabilities of existing technology. IT needs to find a better solution in terms of capacity and throughput.

Another example of data outgrowing the resources needed to manage it is the mundane, scheduled backup procedure. As the amount of data expands, so does the time needed to back it up. Unfortunately, there is nothing that your IT director, CIO or CEO can do about time. **There are only 24 hours in a day, 8 hours in a shift, and, most significantly, perhaps only 4 hours slotted to complete a weekly or monthly backup procedure that currently requires 10 hours to complete. Moreover, what about recovery? If you lose data or have delays in recovering it, your enterprise will lose money!** Staging backups to disks may be part of the solution, but tape still has a significant role in backup retention, especially when backup copies are being *snapped* to disks with ever-greater frequencies and size.

As the size of the backup or recovery expands, the number of cartridges needed to complete the backup increases. Time and volume equate to throughput and capacity – and increased costs. What can be done about it? The typical enterprise has defined their business practices, for example:

- Three iterations of data will be maintained nearline for recovery purposes;

- Historical data, including email and the like, will be preserved according to corporate, industry, and government regulations;
- Some of those cartridges may transition to a shelf, but a sufficient quantity will stay in the library to satisfy the corporate policies.

The important question here is: *How many cartridges?*

If guidelines and procedures govern the process, you must stay within those limits. The policy cannot be re-engineered dynamically. You must rely upon technology to solve the problems, to cure the pain. You must:

- Improve the drive throughput to your storage media;
- Expand the capacity of that media;
- Extend the access of that data to the Open Systems servers in your network;
- Stay within the budget in completing that task; and, most importantly,
- **Preserve the investment that the enterprise has made in its storage architecture.**

The implementation may have to evolve, but the enterprise cannot afford radical change. IT must increase the horsepower of the systems overall and must increase the capacity of its media tank. *But how?*

An Evolving Technology

Today, many enterprises are using the IBM 3494 *TotalStorage Enterprise Automated Tape Library*, IBM 3584 *UltraScalable Library*, or an *STK 9310 PowderHorn Silo* as a part of their architecture. They have a variety of drives from which to choose. The primary need for capacity or for throughput dictates the choice between:

- IBM 3590 B, E, or H
- STK 9840C
- STK 9940B
- IBM LTO Gen 2
- Quantum SDLT 320

<u>Drive</u>	<u>Capacity</u>	<u>Data Rate</u>
<i>IBM 3590H</i>	60 GB	14 MB/s
<i>STK 9840C</i>	40 GB	30 MB/s
<i>STK 9940B</i>	200 GB	30 MB/s
<i>IBM LTO-2</i>	200 GB	35 MB/s
<i>Qtm SDLT</i>	160 GB	16 MB/s

**Table 1 –
Enterprise Tape Speeds & Feeds**

Please see the latest speeds and feed for these drives in Table 1, above.

StorageTek (STK) had refreshed their technology by migrating from the 9840B to the 9840C, increasing cartridge capacity from 20 to 40 GB, and introducing the 9940B at 200GB. The throughput for the 9840C remained the same as the 9940B, at 30 MB/sec. IBM has significantly improved cartridge capacity in the commodity LTO arena using fibre channel, matching STK at 200 GB per cartridge, but with a slightly faster throughput at 35 MB/sec. However, there has been no progress in the IBM 3590 arena until now. With about 100,000 3590s, and earlier tape drives installed¹, and aging, there are many enterprises waiting for relief.

A New Tape Drive – The 3592

With their recent announcement of the *Enterprise Tape Drive 3592*, IBM has reclaimed the enterprise-tape leadership position, not only in cartridge capacity and drive throughput, as shown in Table 2 above, but in drive access and density as well.

Designed with a dual-ported, 2Gb Fibre Channel interface, in addition to ESCON and FICON, and a cartridge capacity of 300GB, **the 3592 has five times the capacity of its predecessor, the 3590H** (also called *Generation 3*). It also has 50% more capacity than its main competitor, the STK *T9940B*. This does not even address older products such as the IBM *3590B* or the STK *T9840B*, each with a capacity of 20GB per cartridge, that many enterprises still deploy. **The 3592 drive has 14 times the capacity of either!**

In terms of connectivity, the Fibre Channel interface allows access to the

<u>Drive</u>	<u>Capacity</u>	<u>Max Data Rate</u>
STK 9940B	200 GB	70 MB/s
STK 9840C	40 GB	70 MB/s
IBM 3592	300 GB	120 MB/s

Table 2 - Enterprise Tape Leaders

automated library from a variety of open systems servers throughout the enterprise. This enables IT to minimize, if not eliminate, the requirement for departmental archive libraries throughout the network. In order to consolidate the library access further, IBM has developed new frames for its 3494 and STK's silos. These increase the density of the drives from six 3590s to 12 3592s in the 3494 frame, and from four 3590s to 20 3592s in the STK frame.

With the new 2Gb/s interface and a unique architecture called *Backhitch Reduction*, implemented to improve start/stop performance, the new IBM drive has almost three times the working performance of the 3590. This could enable IT to complete a 10-hour backup in the much-shrunk 4-hour window. **The 3592 has 33% more throughput than its main competitor, and more than twice the performance of the previous generation**, truly resembling the power and performance reflected by its internal name, *Jaguar*.

In addition, with the implementation of Capacity Scaling, an application can force data to the first 60GB of the cartridge, reducing the average locate time to a random record by 70%. This will significantly improve the read performance of any OLTP application searching for data on tape and improve the response time to the user.

Enterprises can replace multiple drives of an older generation with one 3592, improving their performance while lowering their costs. With a compression ratio of 3:1, the 3592, also known as *Generation 4*, has a cartridge capacity of 900 GB, rapidly approaching the 1TB barrier. The next model, Gen 5, will put that barrier in its rearview mirror with native capacity projected at over 500GB, compressed capacity at 1.5TB. The only question remaining for Jaguar is *how does it hold the track at Daytona*.

In fact, during a beta test at the European

¹ Mostly on mainframes. (Source: IBM.)

Centre for Medium-Range Weather Forecasting (ECMWF), **significant gains in production level performance were reported using five 3592 drives in a StorageTek Silo.** This site consists of a pair of p690 clusters managing 88 library drives of various age, with over 25,000 cartridges. ECMWF generates 1TB per day of new data, with around 800 TB of nearline storage, a perfect example of a very large database.

Even more impressive than the performance was the ease in which the upgrade took place. ECMWF integrated the drives into a Brocade-based SAN connecting to IBM *pSeries 690s* under *AIX 5.1*. There were no hardware problems with this open systems solution using 3592 enterprise drives.

Enterprise CFOs might also be interested in the fact that the 3592 tape drive lists between \$30K – \$35K, significantly less than the 3590H, introduced in 2001, at \$43,500 or the STK T9840C at between \$35K-\$40K. For many with very large data collections, it may be more important that the 3592 cartridges (media) will have a street price of \$125-150/cartridge, reducing consumables cost down to \$.40/GB. IBM plans to reuse the 3592 media in succeeding generations, maintaining read/write compatibility, further protecting user investment. Moreover, your existing robotic library can still be used, from either IBM or STK, protecting a significant investment.

In addition to the mainframe connectivity via *FICON* and *ESCON*, the 3592 works via Fibre Channel with the full range of Open Systems server options: *UNIX*, *Windows*, and *Linux*, enabling their access to the automated library and its terabytes of data.

Conclusion

In an era when the most important business decisions are being made to ensure business continuity, so that you can keep the business in business, the management, in fact the reduction, of IT expenses is critical. **When the economy dictates that every budget dollar be reviewed twice, IT organizations need to make every effort to protect enterprise infrastructure. A new solution that enables an enterprise to maintain its existing architecture while, at**

the same time, increasing both the capacity and the performance of a critical system, deserves attention. This is especially true if you have a mixed environment with both IBM and STK drives installed in an STK library. IBM can deliver a single solution to replace both types of drives, while maintaining the investment that you have made in the silo architecture.

In an era when IT has been forced to choose between high capacity and high performance, IBM has broken the tape paradigm with its new Enterprise Tape Drive 3592. With a compressed cartridge capacity of 900 GB, IBM is rapidly approaching the one-terabyte barrier, providing a space-saving solution managing large volumes of historical data. With native throughput of 40 MB/s (144GB/hr), and impressive random locate capabilities, IBM has reduced the window and reduced the staffing burden for backing up and recalling terabyte-sized databases. The IBM 3592 is simultaneously capacity-centric and performance-centric.

The IBM 3592 presents automated library users with the flexibility to replace, in place, the oldest technology in the data center, enterprise tape drives, with:

- The latest engineering enhancements,
- The flexibility to work with many in-place robotic libraries,
- The capability to service enterprise-wide open systems servers, in addition to mainframes,
- The ability to scale up in performance to as many as 20 drives in a frame, and
- The room to scale out in capacity to over five petabytes of compressed storage, in a very price-performant package.

IBM will definitely have the attention of the IT management with this product.



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. 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 a MBA from Northeastern University.

- *Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 37. (Please dial "1-37" 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, 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.

Disclosure

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.

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.