



Tape Virtualization in the Enterprise — Reducing Data Center TCO

Analyst: David Reine

Management Summary

For sailors, there is nothing better than a sunny sky, a fresh breeze, and a clear sea. Unfortunately, the ocean has many surprises, thus a sailor must plan ahead, plotting his course carefully, checking navigation charts to be sure to avoid underwater obstacles. There is always the possibility to unexpectedly meet a sandbar or other hidden disaster. This can cause damage to the hull, leaving him stuck, with the last three words any sailor wants to hear: “high and dry”. When that happens, a sailor can do nothing except wait for help to set him back on course or a rising tide to lift him clear. But what happens to an enterprise data center when the ocean unexpectedly comes to you? What happens when the calm sea turns into a flood and threatens the very existence of your business? With the recent spate of disasters which have struck the Gulf Coast, being left “high and dry” might not be such a bad thing.

If you run an enterprise data center in the path of an oncoming hurricane, an effective disaster recovery plan may be the only factor that enables your business to survive. Recent technological advances, including the introduction of high-speed, low-cost disk arrays, have led to a rise in the implementation of disk-to-disk architectures for backup and recovery. Unfortunately, these technologies only protect you from errors due to file corruption or deletion, or a head crash which physically destroys a drive. Many of these disk-to-disk installations occur in a single room, or different rooms within the same facility or campus. Disk drives are not removable media; they can not be stored on a shelf. If, as seen in September, floods hit a large area, your data center *and* backup site stand to lose the most valuable asset of the enterprise – information. Data centers that employ tape as an integral component of their information technology (IT) strategy would be spared this tragedy. Tape provides the portability unavailable from disk, plus the security provided by WORM and encryption technologies. Tape now also has improved capacity and throughput not previously available for that medium.

IBM, with 50 years of tape experience, has now announced new enterprise tape drives capable of 500GB of raw data (uncompressed), with a throughput of 100MB/sec, and a new tape virtualization environment capable of supporting up to 8,192 virtual cartridges in a disk-to-disk-to-tape (D2D2T) environment. These products complement the 3592 enterprise drive and *Virtual Tape System (VTS)* in use in thousands of data centers. In order to determine whether IBM’s new family of tape subsystems (TS) can provide your enterprise with data security and improved tape efficiency, please read on.

IN THIS ISSUE

➤ Dual Dilemma of the Data Center	2
➤ New Tape Storage Family	3
➤ Conclusion	5

Dual Dilemma of the Data Center

The typical data center today is faced with two, equally important objectives. One, driven by the enterprise CFO, is to **reduce the total cost of ownership (TCO) of the IT infrastructure**. The other, driven by the CIO, and equally important to the very existence of the enterprise, is the need to **protect and preserve the information that is the very lifeblood of corporate health**. Are these objectives mutually exclusive, or are they complimentary to each other in their joint benefit to the enterprise?

With an ever-increasing demand for storage in order to meet the information requirements of mission-critical applications and to adhere to industry standards and regulatory compliance, an emphasis has been placed upon the need for data center simplification and fiscally responsible information lifecycle management (ILM). *Simplification* can be achieved through removing complexity in the infrastructure and minimizing the number of vendors integrated within the IT environment while an *effective ILM policy* requires the data center to ensure that Tier 1 information be kept on the most capable and thus most expensive, storage resources while Tier 2 data can be accessed from less performant devices. High-speed disk arrays situated on a storage area network (SAN) have become the linchpin for access to Tier 1 information for many enterprises in a two-tiered storage environment. Tape traditionally has been the inexpensive Tier 2 alternative. With the advent of low-cost, less expensive disk storage, such as SATA, some enterprises have entertained the idea of using lower performing SATA drives for archiving of rarely-accessed, secondary-tier data and implementing backup and recovery solutions online, replacing tape as the Tier 2 alternative.

Unfortunately, the TCO of the IT environment extends beyond just the acquisition costs. Controlling data center staffing costs through minimizing the backup window is preeminent among them. The data center needs to reduce operator interaction and increase data throughput in order to reduce scheduling demands. Maintenance of spinning media can also become a drain on any IT budget, not to mention the energy required to

keep the devices spinning, and the air conditioning needed to dissipate the heat created. Tape, on the other hand, only needs to be in motion when it is actually being written or read. It does not require the data center to waste valuable resources. Further, with the latest developments in drive and media technology, each tape cartridge can store more data than ever before, significantly more than disk drives, and transmit that information at rates far in excess of the previous generation. Tape thus remains a vital component in backup and archiving strategies within a new Tier 3. Using a shared D2D2T virtual architecture, data centers can take advantage of the best of both worlds: the instantaneous recovery of lost data available from a D2D strategy and the security and economy of tape from a D2D2T implementation.

However, what if you lose an entire data center? How long will it take to recover the lost data? Can you recover it? How much will it cost while you await retrieval? What is an immediate recovery worth? If a disaster of the proportion of hurricane Katrina struck your data center, and the surrounding community for miles around, what would you need to do in order to keep your business afloat? How long would it take to recreate your IT environment? Could you? Would you sink or continue to surf in the electronic seas?

Perhaps you can count on your business partners to support you in an emergency with replacement infrastructure. But whom can you turn to for help to recreate your data? It is your responsibility! No one else! There can be no blame game or finger pointing. The enterprise data center must take whatever steps are necessary to preserve safely not only mission-critical enterprise data, but also the historical archived information that is necessary to resume normal operations. The IT staff must ensure the availability of the most reliable, and economical, storage solutions.

IBM has been providing innovations in tape solutions for over 50 years, and they haven't stopped. With the recent announcement of the *TS7510¹ Virtualization Engine for*

¹ Following the lead of their brothers on the disk side of the storage house, IBM has created a new nomenclature for all new tape products, designed to simplify the configuration

Tape Systems, IBM has introduced new capabilities to the data center in terms of a centralized storage medium for backup and recovery, as well as the sharing of valuable IT resources between multiple servers and applications. The *TS3310 Tape Library* provides the open systems community with a new modular LTO Gen 3 library with improved design and scalability. IBM also has introduced follow-on products for the *Enterprise Tape Drive 3592* and the *3995 Optical Library*. The *TS1120 Tape Drive* provides the data center with faster throughput capability and higher capacity, enabling the data center to reduce the number of drives required for IT operations as well as the number of cartridges required to be maintained in the enterprise library. The *3996 Optical Library* returns optical storage to center stage for iSeries customers who have been calling for improved functionality, performance and reliability for optical technology.

New Tape Storage Family

While the 3996 Optical Library is being introduced for follow-on business for the iSeries, the remainder of IBM's recent announcements addresses cross-platform solutions covering all market sectors from entry, to midrange, to enterprise. IBM is providing a set of scalable platforms to enable their customers to protect the investment being made in all applications requiring tape media.

TS7510 Virtualization Engine

In order to reduce the TCO of the IT infrastructure, data centers have been adopting virtualization strategies in a number of areas, most notably in servers, with virtualization solutions such as *VMWare*, and in disk storage, with solutions such as IBM's *SAN Volume Controller*. Now, starting from a base of over 3000 early adopters of the *3494 VTS* subsystem, IBM is extending that technology with the *TS7510 Virtualization Engine*. Designed to provide centralized tape system services, the *TS7510* uses disk cache to reduce the time required to backup and restore, improves the sharing of tape libraries across applications and servers, facilitates data sharing and resource virtualization between

and ordering process. Thus, they have added the letters "TS" in front of a structured four-digit product identifier.

Exhibit 1 –

TS7510 Functional Highlights

- Network replication of a virtual tape cartridge over IP
- Network encryption of the data stream to help ensure end-to-end security and privacy
- Network compression to help maximize bandwidth utilization
- Investment protection for existing 3584 and 3494 libraries
- High-Availability with second engine
- Active-Active mode with second engine

servers, and reduces the TCO of IT resources.

The *TS7510* is a highly scalable, modular solution consisting of the *TS7510 VIR1* software and an integrated set of components mounted in a *Model F05* tape frame², including:

- One/two *3954 CV5* virtualization engines;
- Two *3955 SV5* cache controllers; and
- 14 optional *3955 SX5* disk cache modules;

The *TS7510* virtualization engine can support up to 46TB of user disk cache.

3954 CV5 Virtualization Engine

Users can choose one or two *3954 CV5* engines with the capability for each to support:

- 64 virtual libraries
- 512 virtual tape drives
- 4,096 virtual cartridges and
- 4 Fibre Channel ports

With a second *3954* controller installed, these figures are doubled and the data center gains the capability to operate as an active-active cluster or in high-availability mode for the virtual tape environment with complete redundancy for nodes, disks, power, and I/O connections. (For a list of functionality, see

² A second *3952* tape frame is required for expansion to more than six disk cache modules.

Exhibit 1 on the previous page.) In addition, the data center can attach their 3494 and 3584 libraries to the TS7510, along with 3592 and LTO2/LTO3 tape drives, protecting the investment that the enterprise has made in their tape environment. The TS7510 supports a variety of servers, including pSeries, and xSeries, and operating systems, including Sun's *Solaris*, *Windows 2000* and *2003*, *Linux* (Red Hat and SUSE), and *HP-UX*, plus Linux on the mainframe. IBM's tape virtualization engine is intended to interoperate with a number of standard storage management applications including *Tivoli Storage Manager (TSM)*.

TS1120 Tape Drive

Since the summer of 1993, the enterprise tape market has been dominated with tens of thousands of shipments of IBM's *Enterprise Tape Drive 3592*³. With a native capacity of 300GB and a native throughput of 40MBps, the 3592 could outperform its competition, with the *STK 9940B* the closest with a capacity of 200GB and a throughput rated at 30 MBps. In May 2004, IBM responded to enterprise requirements for security and regulatory compliance with the addition of WORM functionality to the 3592⁴. Now, IBM has updated this data center fixture with a new version, renamed to the TS1120. Designed to provide even higher levels of performance, reliability, and capacity than the 3592, the TS1120 addresses the needs of data centers across a broad spectrum of IT environments. With a native cartridge capacity of 500 GB (with higher capacities on the horizon) and a native throughput of 100 MBps, the TS1120 reduces infrastructure costs and the complexity of any enterprise data center by minimizing the number of drives required, and the number of person-hours required, to complete a data center backup. In fact, the TS1120 will enable more frequent backups, minimizing the damage, and cost, to the enterprise when a recovery is required, improving the business continuity strategy of any

³ See **The Clipper Group Navigator** dated August 31, 2003, entitled *IBM Delivers Advanced Enterprise Tape for Open Systems and Mainframes* at <http://www.clipper.com/research/TCG2003039.pdf>.

⁴ See **The Clipper Group Navigator** dated May 19, 2004, entitled *IBM Adapts 3584 Tape Library for ILM – 3592 Drive Adds WORM Option* at <http://www.clipper.com/research/TCG2004048.pdf>.

Exhibit 2 – TS1120 Features

- Dual power supply connectivity with hot-swap capability for enhanced reliability
- Digital speed matching for enhanced streaming performance
- High resolution tape directory for improved data access
- Streaming lossless data compression algorithm for optimal cartridge utilization
- 512MB data buffer for improved streaming performance (vs. 128MB)

enterprise. (See Exhibit 2, above, for additional features.)

With a capacity of 500GB, the TS1120 can store 66% more data than the older 3592 drive, on the same media, and with a throughput of 100MB/sec, it records data at 2.5 times the speed of the 3592. Combined with a 3:1 compression ratio, the TS1120 is capable of completing a 1TB backup on a single cartridge - within 1 hour. This is an ideal cost-cutting improvement for any harried data center manager. The TS1120 can also be scaled down to a 100GB capacity to reduce data access to only 11 seconds. The TS1120 also has a 4-Gbps FC interface for dual-ported switched fabrics to prevent I/O bottlenecks which may be developing in your server infrastructure. In addition to FC, the TS1120 also supports ESCON and FICON with the *3592 J70* controller or the *3494 VTS Model B10* or *B20*. The TS1120 may be connected to an IBM 3584 or 3494 library or an STK Automated Cartridge System solution.

TS3310 Tape Library

As a founding partner, along with HP and Certance (now owned by Quantum), in the open systems tape consortium which founded the Linear Tape Open (LTO) standard, IBM has always maintained a leadership position in open systems tape libraries. The *IBM Total Storage 3583 Tape Library* was designed to support the new *LTO Ultrium 3 Tape Drive*⁵.

⁵ See **The Clipper Group Navigator** dated February 18, 2004, entitled *IBM Restores Order to Data Center*

With up to 6 drives and 72 cartridges, the 3583 can support 28.8TB of uncompressed data at a sustained data transfer rate of 80MBps. IBM has now moved that model to drydock, replacing it with the TS3310, a modular open systems tape solution supporting LTO Gen 3 SCSI and FC drives.

The base unit of the TS3310 is a 5U, rack-mountable control module which supports 2 LTO3 drives and 36 cartridges. The base unit may be extended with a 9U expansion module, capable of supporting 4 additional drives and 92 additional cartridges. With this expansion module (14U total), the TS3310 supports 6 drives and 128 cartridges, 77% more capacity than the 3583. Furthermore, IBM intends to support additional expansion capability for the TS3310, preserving the investment that you make today. With four expansion modules installed, the TS3310 could support 18 drives and over 400 cartridges, in excess of 160TB of enterprise data in 41U of rack space.

3996 Optical Library

In an attempt to satisfy a growing need from within the *iSeries* community, IBM has also introduced a replacement for the 3995 *Optical Library*. The large-scale, externally-attached 3996 *Optical Library* models are designed for performance and reliability using 5.25" 30GB ultra density optical technology, specifically for the *eServer i5* and *iSeries* families, provided yet another storage tier for those customers looking for a WORM alternative. An entry model of the 3996 has up to 2 drives and 32 slots, supporting up to 960GB of optical storage via 30GB media. This compares quite favorably with the entry model of the 3995, which had 20 slots and supported 104GB of data with 5.2GB media. The 3996 also is available in a mid-range Model 080 with up to 80 slots, or 72 slots with 4 drives, and a high-end Model 174 with a maximum of 174 slots, or 166 with 4 drives, supporting over 5TB. 3996 optical technology provides a low-cost form of media with the following standard features:

- 30GB UDO Drives
- Dual Picker
- Single Import/Export mail Slot

- SCSI LVD Library Interface
- WORM and Read/Write Media Support

Optionally, the 3996 can be configured with a barcode reader.

Conclusion

In an age of open systems and standards, CIOs and data center staffs are still looking for an integrated, comprehensive solution, preferably from a single vendor – finger pointing not allowed. IBM is just that, capable of providing all of the hardware, software, and services required by the enterprise. In a data center with a multi-tiered ILM environment as a key focus to reduce the TCO of the IT infrastructure, a turnkey solution is the answer, especially one in which there is a virtual tape environment for increased business continuity, high-speed, high-capacity tape drives for reduced infrastructure and backup window, and an open systems, modular library for planned expansion. If this describes your IT environment, then you should be looking at IBM, the company that has not forgotten about tape.



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 Director, Enterprise Systems 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 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, 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.