



Optimizing Business Continuity — Virtualizing the Enterprise Tape Environment

Analyst: David Reine

Management Summary

An old proverb says that *a journey of a thousand miles begins with a single step*. In fact, the journey starts *before* that first step. It starts when you decide where you want to go – and put a *plan* in place to get you there. When you go on vacation, for example, you do not decide where to go when you get in the family car. You have already determined that destination before you pack your swimsuit or skis. Your wife even went to AAA and got a multi-page *TripTik* to assist you in mapping out your journey, and she is more than happy to tell you where to go, and the shortest way to get there! Perhaps you have an aversion to paper maps and asking directions (you are a man!), so you mount a GPS so that you do not have to ask for directions. Now, when you get into your car to commence your journey, your trip is automated, with the directions provided dynamically by your fantasy navigator.

A similar journey occurs in the data center of enterprises around the globe, as a rapidly expanding storage environment creates challenges daily for a beleaguered IT staff attempting to find their way to a simpler, automated backup/recovery/archive process. They seek the path to optimized IT operations, with improved performance and reliability, and cost control. With a best practices policy in place for full weekly backups, and daily incrementals, in order to ensure business continuity, it becomes more difficult to meet the specified Service Level Agreement (SLA). The threat of litigation hanging over the corporate staff compounds your task requiring you to retrieve the most arcane transaction or email at a moment's notice. It is becoming increasingly complex to implement an effective ILM plan and to satisfy enterprise policies regarding completing a backup/recovery mission within a specified window while keeping mission-critical applications operating 24 by 7. The IT staff needs to find a way to simplify the architecture and reduce the total cost of ownership (TCO) of the IT infrastructure, optimizing data center operations for backup, recovery, and archiving processes.

Virtualization has proven to be an effective method to improve utilization and reduce the TCO of the server environment. It can also be effective to optimize storage. In order to meet enterprise storage SLAs, many CIO's are supplementing their physical tape infrastructure with a virtual tape library (VTL). A VTL combines traditional tape backup methodology with low-cost disk technology to create an optimized backup and recovery solution.¹ IBM has announced a high-performance, high-capacity VTL that acts like tape but with the performance and flexibility of disk. To learn how IBM's *TS7520 Virtualization Engine* can optimize your data center, please read on.

IN THIS ISSUE

- **Backing Up Today's Data Center** 2
- **IBM's New VTL Solution**..... 2
- **Conclusion** 3

¹ See the issue of *Clipper Notes* dated January 20, 2007, entitled *Virtual Tape Libraries - Current and Emerging Innovations - Seventeen Questions to Guide You*, available at <http://www.clipper.com/research/TCG2007008.pdf>.

Backing Up Today's Data Center

The unrestrained growth of storage in the data center has placed a tremendous burden on an overworked IT staff responsible for maintaining business continuity in a 24-by-7 enterprise environment. Corporate back-up policies define the type of information that must be retained: inventory data, financial records, emails, etc., even defining the length of time that it must be maintained, anywhere from one week to one year, or longer. Many organizations even define *recovery point objectives (RPOs)* and *recovery time objectives (RTOs)* to ensure a speedy return to operational efficiency in the event of an application or operational failure. **These SLAs become increasingly difficult to manage, however, when the window of opportunity to complete a backup is limited – while the data being retained seems limitless.**

Tape has always been a primary destination for both backup and archive processes in the data center. It is an inexpensive and portable medium with high-streaming throughput and high capacity to ensure the continuation of data processing in the event of a disk failure, or, as we have seen in recent years, the need to re-create an entire storage environment for disaster recovery.²

Tape does not always satisfy the RTO goals of a 24-by-7 enterprise. In some instances, the data center has replicated data on a disk-to-disk, or D2D, basis in order to accelerate the recovery process. Many businesses, recognizing the need to improve the utilization of tape resources *and* to archive data for very long periods, employ a D2D2T³ strategy, taking advantage of the economies of tape in a multi-tier ILM environment. Just because there is a need to improve recovery time is no excuse, however, to lose sight of the overall enterprise objectives to simplify the storage infrastructure, reduce management costs, and control the TCO of the storage environment. **The data center staff must balance the requirement to improve overall performance with the need to lower the cost per gigabyte⁴,**

including the costs for electricity and cooling⁵.

One way to improve enterprise storage performance is to automate the backup and recovery processes with a virtual tape library (VTL). A VTL can act as intermediate storage on the way to physical tape to help enable faster recovery. It employs low-cost, high-capacity SATA drives for second-tier storage, complementing your ILM infrastructure. This enables more frequent, automated recovery points and, therefore, a faster recovery time without additional tape mounts.

The data center staff can complement existing physical tape facilities by taking advantage of a multi-tier hierarchy with D2D backup, especially when the existing software cannot take advantage of the high-speed streaming characteristics found in new enterprise drives, such as IBM's *TS1120* or the new open systems LTO-4⁶ drive, offered by IBM as the *TS2340*. With standard backup software such as IBM's *Tivoli Storage Manager*, Veritas' *NetBackup*, or EMC Legato's *NetWorker*, the data center can stage the initial backup to a VTL, **improving the backup and restore processes.** A VTL can **improve the RTO**, eliminating a slower recall from tape, utilizing higher raw disk performance to reduce bottlenecks. A VTL can augment the existing business continuance infrastructure by optimizing IT resources and utilizing physical tape assets more efficiently.

IBM's New VTL Solution

IBM has recently announced a next-generation VTL, replacing their *TS7510 Virtualization Engine* with the *TS7520*, improving capacity, throughput, and overall scalability while lowering the TCO for the enterprise data center. IBM has announced two new Editions, a *Limited Edition (LE)* for distributed sites within an enterprise, and an *Enterprise Edition (EE)* that offers the full complement of software options and scalability. With a native throughput of 4.8 Gbps, the *TS7520* can connect via 4Gbps FC, for SAN host attach, or 1Gbps Ethernet for iSCSI and LAN attach, to all IBM server families⁷, servers running supported versions of *Windows*

² See the issue of *Clipper Notes* dated February 1, 2007, entitled *The Evolving Role of Tape in the Data Center*, available at <http://www.clipper.com/research/TCG2007013.pdf>.

³ Disk to Disk to Tape.

⁴ See also two issues of *Clipper Notes* dated February 1, 2007, entitled *The Evolution of Backups - Part One - Improving Performance* and *The Evolution of Backups - Part Two - Improving Capacity*, respectively available at <http://www.clipper.com/research/TCG20070015.pdf> and <http://www.clipper.com/research/TCG20070016.pdf>.

⁵ See the issue of *Clipper Notes* dated May 11, 2007, entitled *Watts Up? How Much Electricity Does It Take to Backup Your Data?*, available at <http://www.clipper.com/research/TCG2007013.pdf>.

⁶ See the issue of *Clipper Notes* dated July 12, 2007, entitled *LTO-4 Pounces into the Data Center with New Life, Greater Capacity, and Higher Performance*, available at <http://www.clipper.com/research/TCG2007073.pdf>.

⁷ *System I, System p, System x, and System z (Linux).*

and Linux, and selected HP and Sun servers⁸.

The *Limited Edition* comes with a single virtualization engine and a single cache controller supporting up to 19.5 TBs of RAID5 storage, and can scale to 128 libraries and 64,000 volumes. The Enterprise Edition supports up to 884TBs⁹, 512 virtual tape libraries, 4,096 virtual drives, and 256,000 virtual volumes, enabling each backup server to have the virtual resources needed. Hardware-assisted compression is an option on both models to increase usable cache capacity¹⁰. IBM implemented three major changes: *faster processor*, *bigger disk cache sizes*, and *hardware compression*. These provide the *scalability* lacking in the TS7510 that was limited to capacities of 46TB and 600MBps for a four-node complex. The TS7520 supports all current IBM tape libraries and drives, with support already committed for LTO4. The TS7520 virtually leaves the competition in the dust when comparing functionality and scalability, in terms of drives or throughput. In addition to outstanding scalability and compatibility, protecting the enterprise investment in IBM's virtual tape products, the TS7520 provides many other key benefits for the data center. (See Exhibit 1.)

Taking advantage of IBM's leadership in virtualization, the TS7520 has incorporated market-proven innovations from IBM's *System x* and its *DS4000* storage array into the TS7520, giving IBM a technological edge without having to reinvent the wheel. IBM simplifies processing with a single, scalable solution, from a single node with a pair of dual-core *Xeon* processors, to a four-node cluster with four quad-processors per node. Standard features include local replication for business continuation, tape caching and compression to meet the RTOs, along with optional "in flight" or tape drive data encryption to satisfy regulatory compliance. Additional advanced functionality includes active failover to achieve high availability for both the control path and the data path, ensuring continued communication between the server and the TS7520, along with a call-home feature to report problems before they become a problem.

With up to four *SV6 Cache Controllers* (DS4200 controllers), the TS7520 consists of up to two base frames (52TB each) and ten expansion frames (78TB each). Each frame consists of multiple *SX6 Cache Expansion (EXP420)* drawers, with each drawer supporting 500GB

Exhibit 1 – TS7520 Benefits

- **Improved RPO/RTO** solution reduces the TCO by minimizing recovery time and improving the restore process;
- **Support for real-time compression** reduces disk storage requirements;
- **On demand allocation** of disk storage maximizes storage utilization;
- **Integrated appliance** simplifies storage infrastructure and reduces deployment time and cost minimizing dependency on expensive staffing; and
- **Supports encryption** functionality of TS1120 to secure enterprise data.

Source: IBM

SATA drives, configured in RAID5 for data protection.

Conclusion

With the TS7510, IBM began addressing the requirements inherent in virtualizing the tape environment. With the TS7520, they have delivered on the *requirements* by introducing a scalable platform with the performance and headroom necessary to satisfy the most demanding very-large enterprise. The TS7510 provided IBM with the direction on where they needed to be. The TS7520 is getting them there.

Make no mistake about it; IBM remains a company committed to tape and disk complementary storage solutions. They have been for 55 years; they will continue to be. Innovations such as integrated tape encryption confirm that commitment. However, the requirements of a 24-by-7 enterprise cannot always be satisfied by just one technology alone, and that is why a tiered strategy for data storage is important. VTLs improve the backup/recovery process, reducing the backup window, enabling mission-critical business to proceed as usual. Tape archives can help reduce storage and energy costs while protecting data in the event of a disaster, ensuring business continuity. **The TS7520, along with IBM's physical tape components, provide the total, integrated solution you need to optimize and simplify your mission-critical operations.**



⁸ HP servers running HP-UX 11.0, 11i (64 bit), 11,23i, and 11.23pi and Sun servers running Solaris 8, 9, and 10.

⁹ 1.3PB will be supported with 750GB drives.

¹⁰ Compression will typically double usable capacity.

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, 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.

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.