



## Making Use of Virtual Tape in a D2D2T Environment

Analyst: David Reine

### Management Summary

Technology is constantly evolving some of the most basic elements in our everyday life. Take, for example, the Radio. Invented over 100 years ago by Guglielmo Marconi, the radio was one of our principle means of communications, and entertainment, for decades. It provided its listeners with the immediate distribution of information that was hours, perhaps days for some, faster than printed forms of communication, such as newspapers and magazines. When television became commercially available in the 1940's, over 60 years ago, many people in the entertainment and news business could have predicted the eventual demise of radio. Radio, however, did not disappear, but over the last six decades, it has changed its format and content, and thrived. It has evolved into a means of entertainment for a more mobile population, with car and portable radios providing us with news and entertainment wherever we travel, when television is unavailable or impractical. Starting with a few stations broadcasting a variety of programs such as news, soap operas, variety shows, and serialized adventures, we now see more specialized sources, with channels on Internet radio and satellite radio supplying a wide variety of homogeneous entertainment, such as all sports, all news, or all "soft rock" music. One hundred years later, radio continues to fill a void in the communications world.

Technology continues to evolve in the data center of every enterprise, large or small, as well. For example, the data center first used magnetic tape to record computer data in 1951, shortly after the commercial introduction of television. It remained the principle means of collecting and storing data for the next two decades until the introduction of the disk drive provided a more rapid, and more expensive, vehicle for the retrieval of random data. Tape, however, remained the primary means of protecting this valuable enterprise asset for more than 50 years. Recently, some data centers have evolved to a Disk-to-Disk (D2D) strategy to protect these enterprise assets, due to a reduction in the acquisition cost of disk drives. However, many external influences are now causing all data centers to rethink the entire process that they use to collect and provide long-term protection for enterprise data.

The rapid increase in the cost of energy has forced every CIO to rethink the total cost of ownership (TCO) for their IT infrastructure, where the protection of enterprise data means more than just saving it. The data center is now responsible not only for ensuring the ability to retrieve enterprise information, but it also must ensure compliance with industry and government regulations that the information has not been altered. Furthermore, they have to protect executives from fines and possible litigation, as well as the enterprise's reputation, if any personal data for employees or clients is lost or stolen. In order to reduce data center TCO and protect enterprise assets, the data center is turning to a blended virtual tape environment, where enterprise SLAs require a fast response time, yet also require low-cost, long-term storage. To see if now is the time for your data center to turn to a VTL environment, please read on.

### IN THIS ISSUE

- Solving the Data Protection Enigma in Your Data Center ..... 2
- The Role of a Virtual Tape Library in a D2D2T Environment ..... 4
- Conclusion ..... 5

## Solving the Data Protection Enigma in Your Data Center

The uncontrolled proliferation of servers and storage devices in the data center, and throughout the enterprise, has led to a rise in the total cost of ownership (TCO) for IT infrastructure for under-utilized servers and storage. With data center information growing rapidly, in some cases doubling annually, along with a meteoric rise in energy costs, CIOs of all enterprises are rethinking how they spend their limited budget. You are consolidating legacy applications onto multi-core, multi-processor, open systems servers, with each system virtualized to make better use of valuable IT resources, and replacing direct-attached storage devices with storage area networks (SANs), based upon either Fibre Channel (F.C.) or iSCSI, in support of the virtualized environment. If this sounds like your situation, the increased complexity of your data protection infrastructure may now be inadequate to support the increased demands for performance and capacity. **Often, the problem is that legacy tape drives simply do not have the capacity or throughput to perform nightly incremental and weekly full backups within their assigned window of opportunity.**

**In fact, due to the very nature of a 7x24x365 Internet economy, the backup window is shrinking or non-existent!** You need to implement a data protection environment that will preserve the enterprise's most valuable assets, its data, with an architecture that will not only provide sufficient capacity and throughput for today, but will also provide adequate scalability for years to come. To accomplish this, you must address a set of issues affecting data protection and total cost of ownership, including, but not limited to the following.

- **System Performance** – to meet backup window limitations;
- **Scalability** – to ensure that the enterprise is protected for years to come;
- **Compatibility** – to enable the continued use of data center infrastructure and to enable the replacement of obsolete components;
- **Reliability** – to eliminate downtime;
- **Energy** – to reduce the amount of electricity consumed by the data center;
- **Security/Compliance** – to protect clients and employees, as well as the business itself, from data theft or loss; and
- **Disaster Recovery** – to enable the creation of duplicate data to be stored off-site.

In order to examine these issues in more detail, this bulletin looks at the impact of both disk and tape in a backup and recovery environment. Particular attention is devoted to the impact that tape has on the TCO due to the long-term archiving of data that must be preserved in order to comply with industry and government regulations, but is rarely, if ever, read.

Replacing legacy backup and recovery infrastructure with bigger and faster models with the same format is no longer always possible. Many of the formats used in the past are no longer available, and those that are do not satisfy today's system requirements for capacity and performance, let alone what will be required in the years ahead. Replacing the legacy backup/recovery and archive systems has become mandatory. The question remains, however: *With what?* Do you deploy an up-to-date version of your existing architecture, replacing your legacy tapes with LTO-4<sup>1</sup>, the reigning champion of open systems tape formats, or do you hop onto the current wave and transition your data protection environment to a D2D architecture?

One fact is certain – the cost of downtime has become prohibitive, ranging from hundreds of dollars per hour to thousands, depending upon the data center workload. Rapid recovery from any outage in the Internet Age is essential. Your competition is only one-click away, and loyalty is as fleeting as your response time. Does that mean that you *need* to implement a D2D architecture for your enterprise? We know that recovery from disk is significantly faster than from tape, but what about the costs? Acquisition costs for a scalable D2D platform are higher than that for a Disk-to-Tape (D2T) system and with data growing by a factor of two annually, or at worst every 18 months, how much disk will be required to satisfy enterprise demands two, three, or even five years down the road? In addition to the acquisition costs, will the CFO accept the energy costs required to keep a multi-terabyte, or perhaps petabyte, storage subsystem spinning? **It may take longer to recover from tape, but at least you can afford it!**<sup>2</sup> It is clear, however, that a system crash or data loss requires the immediate avail-

<sup>1</sup> 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*, and available at <http://www.clipper.com/research/TCG2007073.pdf>.

<sup>2</sup> See the issue of *Clipper Notes* dated February 13, 2008, entitled *Disk and Tape Square Off Again – Tape Remains King of the Hill with LTO-4*, and available at <http://www.clipper.com/research/TCG2008009.pdf>.

ability of a backup copy, as well as duplicate copies removed from the data center and transported to a data vault or remote site in the event of a total disaster.

Many data centers, both enterprise and SMB, are dealing with the problems intrinsic to a heterogeneous legacy environment, with *UNIX* systems running mission-critical, transaction-oriented applications, while a variety of *Windows* and *Linux* systems run business-critical infrastructure applications. Many of these systems run different backup applications, with an eclectic mix of tape drives and libraries providing backup and recovery support. The complexities incurred in supporting and maintaining such a variety of hardware and software adds significantly to the total cost of ownership of running a modern data center. Every data center, however, requires a comprehensive backup/recovery/archive solution that will work with all legacy platforms, including some, such as *DLT*, which are no longer offered.

That comprehensive solution includes the requirement to meet all Service Level Agreements (SLAs) to have instantaneous access to the most recent backup data for recovery. These SLAs may require you to use high-speed disk, rather than tape, for immediate access. This disk cache, however, cannot be a JBOD<sup>3</sup>; it must be protected from data loss via RAID<sup>4</sup>, typically RAID-1, a mirrored set, or RAID-5, a striped set with parity, not to mention the possible requirement for a high-availability solution. With the advent of 1TB disks in the data center, the length of time to recreate 1TB of data becomes so long, that many data centers have begun to deploy RAID-6 implementations, with striping and dual parity to ensure against the very long rebuild time required when a single parity drive fails (and a possible second failure in the RAID set during that recovery period). **No matter which RAID architecture you deploy, it adds more cost to the TCO for the IT infrastructure in an era of open-ended growth.**

Any recovery system must be easy to deploy and easy to manage. It must also be compact to conserve valuable data center floor space. It has to be scalable in order to protect the investment that the enterprise makes in its storage infrastructure. It must be fast in order to meet backup window and recovery requirements, and it needs to have a distributive component in order to be able to send recovery and archive data off-site in

the event of a disaster. Please keep in mind that remote locations also require data protection, but may not have adequately trained IT personnel on-hand to manage the process.

Clearly, you must be able to encrypt any privileged data whenever it travels, either electronically or by transport, in order to protect the C-level staff from potential litigation and embarrassment if the information is lost or stolen. In addition, if you are going to encrypt the data, then you must also be able to manage the encryption keys. Any data protection system must include an effective key management solution. Protecting enterprise executives also includes being able to prove that enterprise records remain untampered. Compliance to regulations includes the requirement to prove that all data is in its original state. The data center has that capability with WORM<sup>5</sup> technology.

The data protection architecture also needs to be flexible, to enable you to continue to use legacy backup components until you can identify and budget for replacements. Last, and by no means least, the data protection infrastructure must be “green”: it has to be energy-efficient.

**One way to help control the TCO is to implement a recovery strategy based upon a SAN using iSCSI rather than a traditional F.C. SAN.** Most SMB data centers do not have the budget for F.C. or personnel trained to support a F.C. SAN. Enterprises may see iSCSI as a way to contain costs. Regardless, these sites prefer an iSCSI solution based on the lower costs associated with the acquisition, administration, support, and maintenance of an iSCSI SAN, even if a F.C. SAN is viewed as more reliable. iSCSI provides the affordability, scalability, and interoperability that you need to maintain a growing data store, especially when you realize that, since iSCSI is supported in all operating environments, it is operating system independent. A simple F.C. infrastructure needed to connect 10 hosts costs at least \$15,000, as opposed to \$2,500 to connect those same 10 hosts, through an Ethernet switch, to an iSCSI array.

Another method of reducing the TCO of the storage network is by compressing all data as it is stored. Many tape architectures, such as LTO-4, have an automatic tape compression capability, reducing the amount of tape required by a factor of 2:1, on average. Compression improves scalability to the point that you can store up to 1.6TB

<sup>3</sup> Just a Bunch of Disks.

<sup>4</sup> Redundant Array of Inexpensive Disks.

<sup>5</sup> Write Once, Read Many.

of information on a single LTO-4 cartridge.

There are significant time concerns, however, in considering backing up directly to tape. Even if the rated throughput of the tape device indicates that the application has enough time to complete the process in the available window, the host must be able to maximize the throughput stream to tape in order to avoid frequent stop/start conditions. Not only would frequent stops destroy performance, it would also create additional wear and tear on the tape infrastructure, creating out-of-service problems while awaiting repair. In some cases, your legacy hosts may not even support the streaming rates required to keep current media technology spinning at rated levels.

The CIO cannot resolve all of these issues in deploying a single technology. It requires a blended architecture of both disk and tape to provide the recovery performance required by enterprise SLAs while curbing the insatiable appetite for energy demanded to keep the disks spinning in a high-capacity environment. A blended environment where data is written to disk as it would be to tape, referred to as a Virtual Tape Library (VTL), and then is saved to tape for long-term storage, is a possible solution. When you add a physical tape drive, or physical tape library, to the VTL, it becomes a D2D2T, or Disk-to-Disk-to-Tape, environment. In some cases, the enterprise can procure a VTL solution with a physical tape component already integrated.

### The Role of a Virtual Tape Library in a D2D2T Environment

A Virtual Tape Library enables a centralized, shared backup and recovery facility for the mid-sized data center. Combined with tape, in an integrated D2D2T environment, a VTL also becomes an integral component in an archival architecture. By connecting a VTL to a SAN, a single VTL appliance can serve as a backup/archival target for multiple heterogeneous servers throughout the data center or across a campus. You have a significant amount of scalability to work with, as you can configure the VTL with as few as 4 entry-level SATA drives, or as many as 16 high-capacity 1TB drives. The data center has the ability to add new drives to the configuration as they are needed, or upgrade existing drives, as higher capacity models become available. **By deploying the VTL in an iSCSI SAN environment, you can implement a low-cost heterogeneous storage solution.** With backup windows shrinking and data expanding geometrically, you have the

plot line for a disaster epic, or an epic disaster! A VTL, however, averts this disaster by enabling the enterprise to recover from an outage, via an incremental or full backup, in significantly less time, and with less difficulty, than from tape. **Keeping mission-critical applications on the air may be the most mission-critical application in the data center.**

The virtual tape library appears to the iSCSI network, and to the storage application, as a tape device, except, through virtualization, it improves backup efficiency, because the data is written to disk in sequential tape format with no file system overhead. A VTL also simplifies operations in a data center that has multiple tape drives deployed, as it enables you to configure multiple virtual drives, and multiple virtual libraries, within a single VTL. You do not have to worry about re-configuring existing backup processes or maintaining multiple legacy tape drives, as the VTL supports multiple backup applications. There is no need to search for tape cartridges, load and unload tapes, or even wait to rewind a cartridge. A VTL eliminates all time-consuming physical tape activities from the on-line backup process. In addition, using a D2D2T environment enables the data center to create a dedicated stream of data to tape *offline*, improving the throughput and the reliability for the archival process. The VTL to tape process reduces the number of service calls and extends the life of the drive, improving the quality of service measurement for the data center staff. IT personnel can still perform backups and recoveries at RAID disk speed, where the only limitation is the streaming rate of the host.

**A blended D2D2T architecture utilizes the inherent strengths of a RAIDed disk environment for fault tolerance and redundancy, along with the convenience provided by tape to eliminate any one potential single point of failure.** First, and foremost, tape is portable. You can remove it from the drive and store it in a remote location for future use in disaster recovery. Recovery then becomes transparent as the backup/recovery application retrieves the data through the VTL<sup>6</sup>. Further, it is inexpensive, capable of storing up to 1.6TB of compressed data on a single LTO-4 cartridge at a cost of about \$100. It is also highly performant, with a compressed throughput of 240MB/sec. **As an alternative, you could implement a remote replication strategy, although this would require an invest-**

<sup>6</sup> In some cases, the physical tape can be accessed directly from the recovery server.

### ment in additional spinning media, and energy consumption at the remote site.

Typically, there is little need to read archived data, so it can be stored on a shelf in the data center, or remotely, with no additional cost for energy. If there is a need to have the archived information in a *near-line* state, in a typical auto-loader or tape library, for example, an entry-level solution typically can access up to 24 physical cartridges, giving the data center near-line access to 38.4TB of compressed data. Obviously, the data center could implement a D2D2T environment with a much larger library, with multiple drives and hundreds, or thousands, of cartridge positions, giving the host network near-line access to petabytes of data.

In addition to the advantages that tape provides for high capacity and portability, LTO-4 also possesses several unique characteristics that make it extremely attractive to data center personnel, and senior management. LTO-4 cartridges can be written in WORM mode, providing management with data credibility in the event of litigation. You can also write encrypted LTO-4 cartridges, protecting the enterprise from embarrassment, and expense in notifying any individual whose personal data may have been exposed. Both the WORM process and encryption are done in-line at the drive, with no server processing degradation.

### Conclusion

You can derive many benefits from the implementation of a virtual tape library in the data center.

- **Continuous application availability** – Because all short-term backups are kept on disk, the data center staff can meet and exceed all established SLA norms for the enterprise;
- **High scalability** – With up to 16 disk drives in an entry-level system, the data center can upgrade or expand the VTL on-demand to configure as many virtual volumes as they need to support the enterprise server network;
- **Reduced backup/recovery media costs and management** – Only those files that need to be archived have to be written to tape, enabling you to control tape media expense;
- **Compatibility with existing backup applications** – Because a VTL appears to the host as a tape device, no changes are required to the backup process or policies, simplifying the deployment of the VTL and further reducing the TCO of the infrastructure.

- **High Redundancy** – With a RAID architecture and redundant components such as controllers, power supplies, and fans, the VTL provides a highly-available recovery platform;
- **Reduced management costs** – With a centralized backup/recovery engine, the data center can eliminate the maintenance and support of multiple media formats;

**As we have seen, however, the VTL can also have its drawbacks in terms of long-term storage capacity, portability, and energy consumption. In these areas, the VTL needs the high-capacity complement of physical tape to keep the total cost of ownership down.**

The advantages of tape are:

- **Total Cost of Ownership** – Tape has a lower cost per gigabyte for long-term storage;
- **Portability** – Tape cartridges provide a durable and secure mode of transport for valuable enterprise assets, with WORM and encryption to protect the data; and
- **Energy conservation** – Tape cartridges on a shelf or in a library use no energy.

A D2D2T infrastructure allows the archival and non-critical data to flow out to a lower cost and lower energy medium (tape), while delivering backup and recovery performance to meet enterprise SLAs for the VTL. A D2D2T environment also enables the data center to retain older backups for a longer period without a significant impact on the TCO.

Clearly, a VTL alone cannot satisfy all of the demands of a modern data center, small or large. A D2D2T solution, however, can provide your data center with all of the latest standards in data protection as it delivers high performance, high capacity, and reliability, while controlling the TCO of the IT infrastructure.



### ***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](http://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](mailto: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.