

VIRTUAL TAPE AND PHYSICAL TAPE: POWER IN COMBINATION

White Paper
February 2007

Abstract

Over the past few years, the approach IT organizations have taken to backup, restore, and disaster recovery technologies has evolved significantly. This paper describes the need for virtualization, and the current landscape of tape virtualization solutions, as well as why using both virtual and physical tape solutions are ideal for most open systems environments.

Table of Contents

Executive Summary	1
Introduction	2
The Evolution of Virtual Tape Libraries	4
Virtual and Physical Tape—A Powerful Combination	5
Sun StorageTek Virtual Tape Library Solutions	7
Summary	9

Executive Summary

The growth of data in today's enterprise datacenters remains unabated. Driven by email, messaging, fixed content, database and compliance considerations, data growth continues to put immense pressure on backup windows and disaster recovery plans.

Most enterprises are looking for ways to shorten their backup windows, improve their restore operations, and ensure their company's long-term viability with effective disaster recovery solutions. In the early part of this decade, the development of low-cost, high-capacity disk products was an economic enabler to incorporating disk into data protection strategies. Early disk-to-disk and disk-to-disk-to-tape solutions brought with them a new set of problems, which at times were worse than the problems they were designed to solve.

Storage virtualization has emerged as a powerful and practical way to get more from existing storage systems, consolidate resources, and simplify environments. The concept of storage virtualization is the pooling of multiple physical storage resources into a single "virtual" storage resource that can be centrally managed. While this may sound like storage Utopia, finding the right solution isn't always as easy as it seems. Every environment is different, with a unique set of needs that varies from enterprise to enterprise. To identify the right type of solution for an enterprise, it helps to understand the history of storage virtualization solutions as well as what new capabilities and products now make the concept of virtualization much simpler.

Since 2003, many providers have offered an approach that provides an easily implemented alternative—the Virtual Tape Library (VTL). First-generation VTLs were primarily standalone devices that required dual writes or media server involvement to create physical tapes. Second-generation devices, called integrated virtual tape libraries, began to provide the ability to write from the VTL to a physical tape library. Many of the early "integrated" implementations caused some exposure to unmanageable physical tape and single points of failure in backup, restore, and disaster recovery processes.

The Sun StorageTek™ Virtual Tape Library (VTL) is an integrated solution that addresses the problems of early second-generation integrated VTL systems. Backup, restore, and disaster recovery operations are provided in a totally integrated management solution. Not only are backup windows shortened and handled more efficiently, some organizations now perform backups more frequently due to the system's ease of use. Disaster recovery operations are automated via a policy-based management system that enables management with less human intervention.

The StorageTek VTL solution, unlike first- and second-generation VTL systems, leverages the value of both virtual and physical tape. This powerful combination helps organizations reduce operating costs through virtualization and consolidation activities as well as labor costs via automation and policy management, increase the reliability and efficiency of backups, and enhance the security of their data.

Chapter 1

Introduction

For backup, restore, and disaster recovery operations:

- The right question is not “disk or tape?”
- The right question is “how much business value can your company achieve from the powerful combination of disk and tape?”

The debate over the relative cost of disk and tape has been going on for several years. Some studies indicate that high-end tape automation solutions have a lower Total Cost of Ownership (TCO) per capacity versus low-end disk devices when broadly considering all operational aspects. There are other studies that indicate the opposite. This debate may be about individual differences in operating environments.

Unfortunately, the debate misses the mark. Both disk and tape have advantages and limitations. And both play a key role in today's datacenters.

The pros and cons of tape

- Tape automation has the advantages of providing a low-cost, data-streaming storage medium with removable media/portability, a physical break in data flow, and extremely stable media. With some technologies, tape may have a shelf life of thirty years.
- On the other hand, tape can be perceived as slow, difficult to manage, hard to secure, and complex to administer.

The pros and cons of disk

- Disk has the advantages of effective multithreading, random access, and frequently, raw speed.
- The disadvantages of disk (from a backup/disaster recovery perspective) are that it is expensive, impossible to “make portable” without high bandwidth, and too easy to erase or accidentally delete.

The discussion of tape versus disk as a backup/disaster recovery media is ideally a matter of the balance of tape and disk as opposed to an “either/or” question. The problem remains that the management cost of maintaining a multi-tier backup/disaster recovery model is significant. Simply put, data management tools have not yet effectively addressed the needs of moving and managing data from media servers to disk to tape and back again.

A well-designed virtual and physical tape library solution can exploit the benefits and reduce the limitations of both disk and tape. We call this “power in combination.”

Partial solutions to the issues associated with backup, restore, and disaster recovery processes include:

- Disk to tape
- Disk to disk
- Disk to disk to tape

Why is the backup/disaster recovery model changing in most organizations?

Traditional media server to physical tape library implementations can suffer from long backup windows, high labor requirements, large space requirements, and media reliability issues. Tape continues to improve in speed, capacity, and in its management capabilities. At the same time, pressures from 24x7 operations and the changing requirements from users for self-managed and near-instantaneous data recovery have driven the restore requirements and backup window pressure to the point where most organizations have added disk to their backup/recovery operations in order to exploit the power in combination of disk and tape.

Disk to tape, disk to disk, disk to disk to tape—Finding the right data protection solution

The concept of disk-to-disk backup emerged in the early 2000s. The need was driven by the commercialization of low-cost, high-capacity disk technologies that make disk-to-disk (D2D) backup affordable.

The D2D approach has its limitations. It involves changing backup/disaster recovery processes; it does not have an effective approach to removable or transportable media; and it does not meet some of the compliance requirements for offsite data storage. Asynchronous mirroring and remote copy technologies are being used by some D2D operations to address these issues. Communications and capacity-related expenses still limit broad adoption of a pure disk backup/disaster recovery model.

Most large organizations, unwilling to accept the risks of a full D2D environment, adopted a disk-to-disk-to-tape (D2D2T) backup model. This was intended to mitigate the exposures of a pure D2D model and gain the benefits of adding a multithreaded, random, faster device to the backup/restore process, while also providing for effective offsite storage. In most implementations, the D2D2T approach merely added buffering in the backup restore stream, which may have improved recovery speed but added significant complexities.

Tape remains an important part of the D2D2T strategy: it is still the most cost-effective, long-term storage option, plus it offers additional redundancy and portability.

So how do organizations use the power of disk and tape in one easy to implement, manage, and support solution?

Enter virtual tape library solutions.

Chapter 2

The Evolution of Virtual Tape Libraries

Integrated virtual tape libraries can help organizations:

- Reduce technology costs
- Shorten the backup window
- Reduce labor cost
- Increase backup, restore, and disaster recovery reliability
- Increase operational efficiency
- Enhance data security

In late 2003, a broad set of vendors began to offer virtual tape library solutions. VTL solutions have the clear benefit of minimizing the disruption of adding disk to the backup/disaster recovery process.

A VTL system is viewed by the media server application as another library. Disk management (e.g., provisioning, volume management, etc.) is handled by the VTL system. The use of disk is actually more efficient and generally has higher performance in a virtual library than it has in other D2D2T applications. This is due to the way the backup applications write to tape (or an emulated tape cartridge in a VTL environment) because fragmentation is eliminated. Unlike the processing approach used by file systems in other D2D2T applications, the StorageTek VTL system contains software that controls RAID volumes and ensures that a given RAID group is only written to by a single virtual tape. This explains why the fastest file systems are considerably slower than the fastest virtual tape libraries.

Early VTL solutions did not integrate with existing tape libraries, but operated in what was called “stand-alone” mode. Creating a physical tape required either a dual-write operation or using the media server application to generate backup copies after the backup window was “closed.” This practice was not a good use of time or money. Companies wanted to leverage the tape assets they already had, not alienate them.

Chapter 3

Virtual and Physical Tape — A Powerful Combination

Integrated virtual tape libraries manage the movement of data to downstream physical tape libraries based on policies set by storage managers. Media server backup software is updated with volume information so that a “bare metal” restore (directly from physical tape) capability is able to occur.

In addition, advanced integrated VTL solutions reduce management costs and complexity by managing both virtual tape and physical tape. They can reduce the backup window via increased performance and provide automated management to reduce labor costs.

Together, virtual tape and physical tape provide a powerful combination. They can help organizations:

- Reduce cost
 - By consolidating tape media, drives, and libraries in a Storage Area Network (SAN)-based VTL solution, organizations can reduce equipment, floor space requirements, power consumption, and maintenance costs.
- Close the backup window
 - A SAN-based VTL solution can support multiple media servers simultaneously. Virtualized resources providing performance in parallel can “close” the backup window and migrate the tape to physical tape outside the back up window.
- Reduce labor costs
 - Storage management productivity is frequently measured by the volume of data managed by a single administrator. With a virtualized, SAN-based VTL solution, the backup/disaster recovery process can be managed by a single group of specialists who can usually increase storage manager productivity. In combination, virtual tape and physical tape provides a policy based, automated total solution for backup, restore, disaster recovery, and compliance.
- Increase the reliability of the backup/disaster recovery process
 - VTL solutions with disk buffer redundancy and the inherent high availability of today's disk devices, help organizations experience higher levels of availability and better performance of service-level objectives.
- Increase efficiency
 - Virtual tape library solutions increase media and tape drive utilization by buffering and stacking processes. Tapes written to the physical tape drives are full and not interleaved. Restore operations are enhanced through this process.

- Enhanced data security
 - By providing multiple copies of data driven by policy, organizations can address their needs for protecting data with reduced management costs. Accidental data deletions are minimized by removing or lessening the human element from the backup/disaster recovery processes. Bad media problems are mitigated and compliance needs are served by automatically creating encrypted offsite copies.

Chapter 4

Sun StorageTek Virtual Tape Library Solutions

Sun StorageTek VTL solutions provide:

- Automated, policy-based management
- Integration of a total solution for backup, restore, and disaster recovery
- Industry-leading capabilities, including data protection expertise, scalability, non-disruptive operations, and automated disaster recovery

Sun's virtual tape library solutions are ideal for organizations that need the high performance of both physical and virtual tape for their backup and restore operations as well as an integrated solution that manages multiple copies of that data across local and remote sites. Sun's enterprise-class reliability, robust policy-based management, and seamless integration of physical and virtual tape as a single solution creates significant benefits over other alternatives.

StorageTek VTL systems were developed with real customer inputs, to operate in open systems environments. Our research determined that a VTL solution needs to focus on improving data protection by leveraging the powerful combination of both virtual and physical tape. It must be simple to implement and manage, without any slow down to applications or access of data.

Sun's VTL solutions offer:

- Automated, policy-based management. With the StorageTek Virtual Tape Library appliance, organizations can manage and track backup data for disaster recovery, including local and remote copies.
- Integration of all components of the backup, restore, and disaster recovery process. The Sun solution looks at the data protection processes of the organization with a total view to help manage all of the technology supporting the process.
- Sun's commitment to open system solutions. Sun's strong relationships with other providers of the process, including application and infrastructure providers, helps ensure that Sun's StorageTek VTL solutions can integrate with the solutions that organizations already have.
- A history of great products coupled with great support. StorageTek VTL solutions are backed with a long history of leadership in data protection and archive solutions. Sun's expertise and focus on quality service and support help organizations get the most out of their backup, restore and disaster recovery investments.

Sun understands that most enterprises want to leverage both tape and disk for their data protection strategy. This understanding of tape as a complement to disk helps leverage the combination of virtual and physical tape and helps ensure that the complementary nature of these technologies leverages value for customers.

Sun possesses a set of industry-leading virtual tape library capabilities that further deliver customer benefits:

- Grow as you need, not as your vendor dictates. Sun's StorageTek VTL solutions offer industry leading scalability. Scalability means upgrades rather than system replacements, saving planning costs and potentially reducing subsequent hardware purchases.
- No more costly downtime. StorageTek VTL solutions support continual operations, thus removing scheduled downtime for upgrades and maintenance. This helps organizations deliver better service to their users and customers.
- Automated disaster recovery means shorter backup windows with less failed backups. With Sun's StorageTek VTL solutions, organizations can spend less time planning, have better protected data, and are exposed to less opportunity for human error.

Summary

Sun's StorageTek Virtual Tape Library systems do more than integrate virtual tape with physical tape; they make the combination powerful. Sun's solutions offer a complete, end-to-end storage virtualization solution that can help reduce cost by integrating virtual and physical tape into your backup, restore, and disaster recovery processes.

