

SUN STORAGE TEK™ VIRTUAL STORAGE MANAGER SUPERIOR ARCHITECTURE DRIVES VALUE

White Paper
December 2006

Abstract

Sun StorageTek™ Virtual Storage Manager is the market leader in mainframe tape virtualization. This paper describes the architectural superiority that has driven Sun's market leadership. World-class devices, control software, and integration make StorageTek Virtual Storage Manager a virtual tape system that truly delivers business value.

Table of Contents

| | |
|--|-----------|
| Overview | 1 |
| Introduction | 3 |
| Sun's Virtual Storage Manager | 4 |
| System Architecture | 4 |
| Tape Libraries | 5 |
| Options | 5 |
| SL8500 Availability | 6 |
| Load Balancing | 6 |
| Innovative Shape | 7 |
| Summary | 8 |
| Tape Drives | 8 |
| Data Life Cycle | 9 |
| Capacity Centric Drives | 11 |
| Access Centric Drives | 11 |
| Summary | 13 |
| VSM Tape Buffer | 14 |
| VSM Software | 15 |
| Management Classes | 17 |
| Storage Sub-Pools | 18 |
| Summary | 19 |

Overview

Sun StorageTek™ Virtual Storage Manager (VSM) delivers outstanding value through its superior architecture. Central control of StorageTek VSM enables efficient resource utilization and delivers advanced data protection and disaster recovery. Total systems throughput is delivered via the unique approach that Sun has used to integrate world-class devices and software into a truly unique system. The systems architecture of StorageTek VSM is what drives its business value, enabling flexibility in operations and removing management complexity through policy management.

Organizations do not need to make trade-offs between their level of data protection and their business operations. Data protection applications run faster, reading and writing data to disk, while the StorageTek VSM system manages the downstream copying to tape based on policies. The combination of features available in the StorageTek VSM control software facilitates Disaster Recovery (DR) and Business Continuity (BC) applications. A broad set of utilities enables a DR/BC solution that is tailored to business needs and automated via policy management. StorageTek VSM allows users to test their disaster recovery processes without interrupting systems operations and while continuing normal data protection operations.

StorageTek VSM enables organizations to lower their operational costs while meeting service-level requirements. Improving the utilization of storage assets contains costs. Service levels are met by accessing tape data at disk speeds and utilizing virtual tape drives to handle peak-processing workloads. Batch times can be reduced, freeing up mainframe resources. StorageTek VSM helps organizations adapt to business changes by providing flexibility in how IT resources are used to meet shifting IT priorities.

What you need to know about Sun's StorageTek VSM

A virtual tape solution is a fairly complex system involving several key components. Virtual tape solutions integrate tape libraries, tape drives, disk buffers, and control software to create value that is greater than the sum of the individual components. Understanding how a solution will fit business needs and deliver value requires an understanding of the components and how they are integrated.

Sun's StorageTek VSM focuses on delivering superior value through a strategy that involves central control, resource balancing and sharing, industry-leading reliability, and unsurpassed total system throughput. When considered all together, these attributes create an industry "best value."

In today's economy, investment decisions are scrutinized over and over to find the "best value." From 2001 through 2005, Sun's StorageTek VSM sales outshipped IBM VTS by approximately 2 to 1, as reported by Gartner.

This document will demonstrate how Sun's StorageTek VSM delivers the best value in tape libraries, tape drives, virtual tape disk buffer, and centralized controlling feature rich software when combined together as a solution.

Chapter 1

Introduction

Virtual tape solutions, like Sun's StorageTek Virtual Storage Manager, place disk devices into the backup/restore/disaster recovery path to facilitate tape operations while still maintaining the cost advantage, portability, and removable media advantages of traditional tape. The StorageTek VSM system is the market-leading virtual tape solution in mainframe environments.

The approach used by leading virtual tape vendors in the mainframe space is based on an architectural design that off-loads data movement from the host, including all migration (writing physical tape from the disk buffer) and recall (reading files from physical back-end tape storage to the virtual tape disk buffer). Figure 1 depicts this approach. Data directed to a tape by a mainframe application is written to a virtual tape disk buffer. The application is then able to continue with other processing while the virtual tape system manages and, based on user-established policy, controls the creation of physical tapes. When an application requests data from a tape, the virtual tape system serves up the virtual tape (if it is still in the disk buffer) or a physical tape is mounted—transparent to the application. At a high level, only two vendors, Sun Microsystems and IBM, embrace this architectural design.




Figure 1.

The key to the value that StorageTek VSM drives is in the way that the various features of this design are integrated. The integration architecture delivers flexibility and advanced functions that are not available from any other supplier of mainframe virtual tape.

Virtual tape has provided the benefits of storage consolidation in the mainframe space for several years. The quite simple yet effective ability to direct small flat files to a disk subsystem, that emulates tape to the host, and subsequently migrate those small files to achieve full cartridge utilization has given users an exhaustive list of benefits, such as media savings, floor space reduction, environmental savings, labor savings, faster batch processing, total automation through rack elimination, CPU and disk savings, lower total cost of storage, maintenance savings, and growth avoidance.

StorageTek Virtual Storage Manager

StorageTek VSM is composed of four key components: tape libraries, tape drives, virtual tape disk buffers, and software. Its superiority is based on the individual qualities of these components as well as the way they are integrated.

System Architecture

Highlights

- Single image of up to 256 StorageTek VSM systems enables workload balancing and flexibility in designing data protection and disaster recovery scenarios
- High availability and nondisruptive growth
- Simple migration between releases, while continuously protecting data

The systems architecture of StorageTek VSM is focused on delivering a single image of tape to multiple applications and processors. The control software is host-based, enabling the virtualization of up to 256 StorageTek VSM systems as a single image. Intelligence in the control software enables multiple StorageTek VSM systems to balance workloads, move data freely between other StorageTek VSM systems (both local and remote), and support a wide variety of data protection and disaster recovery scenarios.

Another key focus of the StorageTek VSM systems architecture is to enable high availability. Growth within the subsystems does not generally involve disruption or any loss of access to data during implementation of additional resources. High-availability functions, such as multiple pathing and failover, enable highly reliability and continuous operations in the event of a component failure.

The StorageTek VSM system is designed to provide a simple migration path between releases. Releases can be run concurrently to enable a sound and protected migration path. In fact, Sun has continued to facilitate migration in StorageTek VSM through four major releases over the last seven years.

Tape Libraries

Highlights:

- StorageTek VSM tape libraries scale performance with capacity
- Library options range from 326 to 70,000 slots using Sun's pass-thru capabilities to create single-image libraries from multiple physical libraries
- Redundant hot-swap components minimize scheduled downtime — add more slots, drives, or robots while the StorageTek StreamLine™ SL8500 modular library system continues to operate
- The StorageTek StreamLine SL8500 system's redundant and layered robotic design drives unparalleled availability and access to cartridges
- Pass-thru capabilities enable automatic load balancing between libraries
- The StorageTek StreamLine SL8500 “tight oval” architecture is designed to provide high performance, save floor space, and provide consistent performance to all slots in the library while minimizing robotic movement and wear

In the mainframe environment, the ratio of cost/capacity for disk is many times higher than tape storage. A high-performance tape library enables a system to achieve its lowest cost point by minimizing the disk buffer size. The key characteristics of the tape library in a virtual tape system should support this by ensuring high availability, efficiency, and systems throughput.

Options

With the integration of StorageTek, Sun has a history of industry-leading libraries. Sun provides a broad range of tape library and drive options for back-end tape processing in the StorageTek VSM system. The StorageTek PowderHorn™ library led the industry for several years. Today, Sun's flagship library system is the StorageTek StreamLine SL8500 (“SL8500”). In addition, Sun also offers the smaller Timberwolf 9740 automated tape library (“9740”) for mainframe StorageTek VSM environments. Both the 9740 and the SL8500 libraries have pass-thru capabilities that allow the physical linkage of multiple physical libraries into a single image. This provides an extremely broad range of library options for StorageTek VSM systems, ranging from 326 slots to 70,000.

Sun offers more options in performance and capacity of its automated cartridge systems than IBM. The small footprint Sun's 9740 offers a base of 326 slots as well as an incremental boost within a single library up to 494 slots. A strong feature of this robotic system is that just like the larger PowderHorn or SL8500, the 9740 offers pass-thru capabilities that enable a user to grow the library configuration up to six pass-thru connected libraries. The pass-thru library architecture of the 9740 can grow to 2964 slots and has six robotic arms providing service data. In smaller configurations, the IBM TS3500 offers acceptable mount times, but as it grows robotic performance diminishes quickly due to its linear design.

The IBM TS3500 library ranges from 160 cartridges to a maximum of 6681 cartridges. 6681 slots are achieved with 16 frames, but can only configure four tape drives in that situation. In a more realistic scenario, for a ratio of 6 drives to 250 slots per frame and one control unit, the maximum capacity would be approximately 3900 slots. Comparing an IBM TS3500 to Sun's SL8500 clearly demonstrates that IBM has trade-offs in drives to slots, whereas the SL8500 slot count is independent of the number of drives. Moreover, since the SL8500 drive bays are built in, drives can be added while the system is operational.

SL8500 Availability

Sun's SL8500 is designed for high availability. Redundant hot-swap components, designed to minimize scheduled downtime, provide the ability to add more slots, drives, or robots while the SL8500 continues to operate. The redundant and layered robotic design of the SL8500 drives unparalleled availability and access to cartridges.

SL8500 capacity upgrades take place with minimal to no disruption to ongoing operations through capacity-on-demand library code, pass-thru ports and additional SL8500 libraries that can be attached while the system is operational, and near continuous operations software functionality.

The IBM TS3500 will scale to 6681 slots, but the design is linear, with more robotic movement required to service tapes and therefore a higher probability of wear and breakage. In a dual-robotic system, two robots traverse the rail to service tapes, reducing the length of travel from about 40 ft. or so to 20 ft. or so per robot. This is a less desirable design than the tight oval of Sun's SL8500 RaceTrack architecture, which requires less physical movement of a robot. Through its feature evolution, the SL8500 library provides eight robots/hands, for redundancy and to minimize queuing, because the work is spread out over multiple robots working in parallel. This creates the need for less required robotic movement to get to a tape, which increases library reliability over time.

Load Balancing

Sun's pass-thru port also provides system load balancing. During normal day-to-day operations, it is possible that activity will hit one library more frequently than another, creating "hot spots" and degrading service times. The pass-thru port eliminates these "hot spots" by automated system load balancing. Automatic load balancing allows the tape library system to use all tape transports and robots in the system, minimizing contention throughout the system and over time evening out the "hot spots."

The SL8500 library adds even more load balancing capability by providing up to four levels of robotic performance within a single library, allowing cartridges to be passed from one level to the next. This feature accomplishes load balancing within a single library in addition to cross-library balancing using pass-thru.

While Sun's pass-thru port connects the libraries in an Automated Cartridge System and permits all units to share all transports and access to all cartridges, its significant benefit lies in its system load balancing. This attribute can best be understood by looking at how IBM tape libraries would operate.

Assume that a customer has two IBM TS3500 tape libraries. Since they have no pass-thru ports, they are independent islands of automation. When each of the IBM libraries is loaded with tape data, it is difficult to predict how future end-user tape activity will spread out over both libraries. It is entirely possible that at certain periods, end-user activity will hit one library more frequently than the other, creating "hot spots" and degrading service times. The only way to eliminate these "hot spots" would be to eject tapes from one library and load them into the other. This manual task defeats the purpose of tape automation, which is to eliminate operator intervention.

Innovative Shape

The innovative shape of Sun's SL8500 allows for growth of capacity, performance, or both, all in a "tight oval" design that provides multiple robots working in parallel to minimize queuing. This design provides consistent performance across all slots and saves floor space—about 25-50% less than a comparable PowderHorn configuration and in excess of 60% relative to IBM's linear design. The combination of eight robots/hands, for redundancy and to minimize queuing, and the tight oval design lowers the requirements for robotic movement to get to a tape, which increases library reliability over time.

Each SL8500 can grow to 10,000, and can be connected via pass-thru to grow beyond that (up to 70,000). This provides for virtually unlimited scalability in one system, which means there are no worries that the customer will outgrow the solution and have to manage a second, third, or fourth library system. The performance, redundancy, footprint efficiency, and availability of this design far exceeds anything that is possible with linear-designed robotics from IBM, or any other vendor for that matter. As linear -designed robotics are upgraded to accommodate larger numbers of cartridges, access times to cartridges diminishes and cartridges that end up at the "end of the line" cannot be mounted as quickly. Figure 2 demonstrates the footprint and cartridge access differences between an IBM TS3500 library with 2 virtual storage subsystems versus Sun's SL8500 library and 2 StorageTek VSM subsystems.




Figure 2.

Summary

Sun recognizes that one size does not fit all when it comes to its tape libraries, and offers a broad portfolio of solutions. This enables the strategy of placing data on devices that meet service requirements while minimizing costs. The superior architecture of Sun's libraries enables high availability, minimizes planned downtime, automates workload balancing, and ensures that performance scales with capacity. Most importantly, Sun's architectural approach drives business value.

Tape Drives

Highlights

- Sun supports a hierarchy of tape drives in the “back end” of the StorageTek VSM system, minimizing the cost of storing data while maintaining service-level requirements
- The fastest access-centric tape available helps reduce disk buffer requirements and drive lower system cost
- Automated management of the hierarchy results in almost no management cost to operate the hierarchy
- Sun's capacity-centric drives are, in general, faster and higher capacity than comparable drives

StorageTek VSM system throughput is a function of all of the components of the systems architecture, and tape drives are an integral part of that equation. Sun's libraries deliver superior performance, but without a similar strength in tape drives that performance would not result in system throughput.

Sun's StorageTek tape drives, and the way that they are integrated with libraries, drive three key business values.

- They enable an automated hierarchy of storage within the system, reducing cost while maintaining service-level requirements. The layers of StorageTek VSM hierarchy are virtual (disk buffer), high speed (or access centric) and high capacity (or capacity centric).
- They have superior raw speed in the access-centric layer
- They have superior raw capacity and speed in the capacity-centric layer

Data Life Cycle

One of the powerful capabilities of Sun's StorageTek VSM system is the ability to drive business value via a hierarchy of storage devices. The strategy involves placing data on the device that will minimize the cost of storing it while meeting service-level requirements for retrieval. In StorageTek VSM systems, this is done automatically via policies.

Figure 3 depicts this approach. Initially, all data is stored in virtual tape volumes in the disk buffer. Based on storage policies, the data is migrated to either access-centric drives (the StorageTek 9840), capacity-centric drives (the StorageTek T10000 or StorageTek 9940), or it is moved to the access-centric drives and later migrated to the capacity-centric drives.

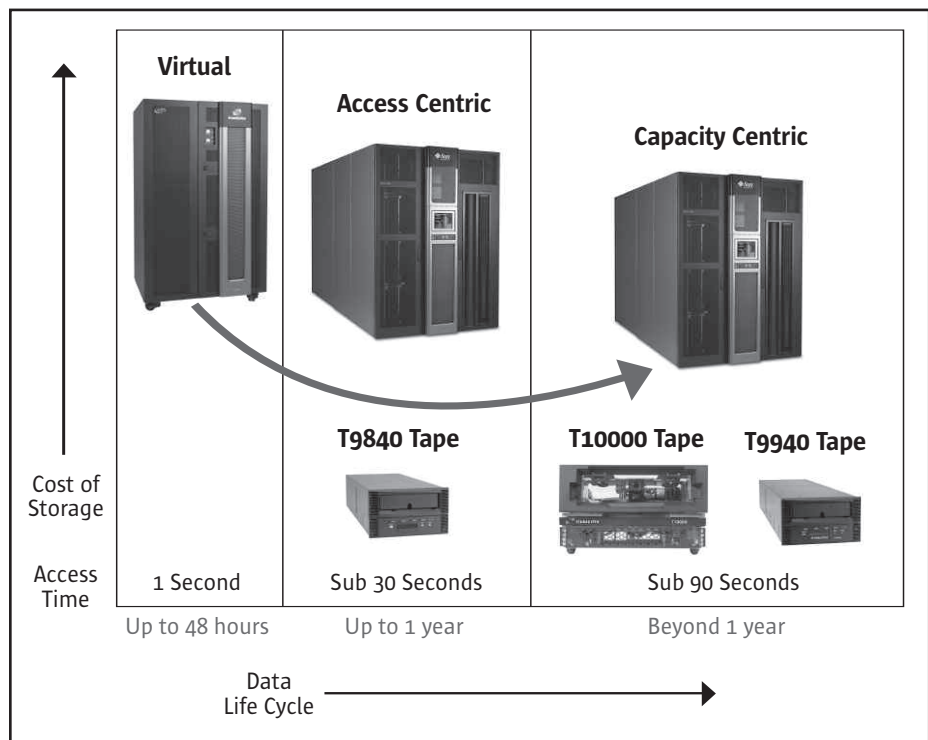


Figure 3.

With the StorageTek VSM system, service requirements can be matched to device characteristics, while minimizing cost. Thus delivering business value.

StorageTek VSM provides a feature that allows a customer to dictate where a virtual volume that has been migrated to tape will reside. This feature is known as “media preferencing.” This achieves lowest cost of storage with maximum performance and is only available from Sun on its StorageTek VSM. Although IBM offers 3592 and 3590 technology behind the VTS engine, both are single-hub architectures, and therefore cannot offer the same granularity as Sun or achieve this cost-balancing feature as readily. The IBM 3590 technology is a capacity-centric drive with relatively slow access times. The IBM 3592 is also a capacity-centric design with a short tape feature, neither of which approach the 12 second access times of the StorageTek 9840 product line. As it relates to the virtual subsystems, the new IBM TS1120 tape drive is limited to IBM 3592 emulation and the capacity and performance of the IBM TS1120 drive is unavailable to the IBM VTS Tape Hierarchy*.

The following chart lists characteristics of the StorageTek 9840, 9940, and T10000 from Sun and the 3592, 3590, and TS1120 from IBM.

| | StorageTek T10000 | StorageTek T9940B | StorageTek T9840C | IBM TS1120 | IBM 3592 | IBM 3590 |
|--|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------------|-------------------------------|-----------------------------|
| Cartridge capacity | 500, 120 GB | 200 GB | 40 GB | 500*, 100 GB | 300, 60 GB | 10/20, 20/40, 30/60 GB |
| Data throughput/sec | 120, 50 MB/s | 30 MB/s | 30 MB/s | 100 MB/s* | 40, 30, 20 MB/s | 14, 9 MB/s |
| Compressed data throughput/second | 360 MB/s | 75 MB/s | 75 MB/s | 220 MB/s | 120 MB/s | 40 MB/s |
| I/O ports | 2 ports | 2 ports | 2 ports | 2 ports | 2 ports | 2 ports |
| I/O port speed | 4 Gb | 2 Gb | 2 Gb | 4 Gb | 2 Gb | 2 Gb |
| Direct attach FICON | Yes | Yes | Yes | No, J70 required | No, J70 required | No, J70 required |
| Library support | StorageTek 9310, SL8500, L700, L180 | StorageTek 9310, SL8500, L700, L180 | StorageTek 9310, SL8500, L700, L180 | IBM 3584, IBM TS3500 | IBM 3584, IBM TS3500 | IBM TS3500 |
| Tape emulations | 3490, 3592E | 3490E, 3590B | 3490E, 3590B and F | 3592 | 3592, 3490E, 3590E | 3490E, 3590E |
| Virtual back hitch | No | No | No | Yes | Yes | No |
| Data access time | 500 GB–62 sec 120 GB–29 sec | 59 seconds | 12 seconds | 500 GB–49 sec 100 GB–27 sec | 300 GB–61 sec 60 GB–33 sec | Short–61 sec Long–91 sec |
| Multiple media suppliers | No | Yes | Yes | Yes | Yes | Yes |
| Multiple media manufacturers | Yes | No | No | Yes | Yes | Yes |
| Single brand media | Yes, Sun only | No (StorageTek, Imation) | No (StorageTek, Imation) | No | No (IBM, Fuji, Imation) | No (IBM, Fuji, Imation) |
| Physical cartridge size | 3490 cartridge | 3490 cartridge | 3490 cartridge | 3490 cartridge | 3490 cartridge | 3490 cartridge |
| Media reuse | Next generation only | Yes | Yes | Yes | Yes | Yes |
| Product line | T Series | T Series (T9840, T9940) | T Series (T9840, T9940) | Magstar (3590, 3592) | Magstar (3590, 3592) | Magstar (3590, 3592) |

Figure 4.

Capacity Centric Drives

For archiving or long-term retention of data purposes, in a Multiple Virtual Storage (MVS) virtual tape environment, the StorageTek 9940 and T10000 also deliver superior device characteristics. In general, the StorageTek T10000 delivers faster access and higher capacity than comparable drives.

In an mainframe virtual tape environment, the relevant competitive comparison is the StorageTek T10000 versus the IBM TS1120. The StorageTek T10000 offers a 20% faster native data rate and equal native capacity, and also has the advantage that the media capacity on the current StorageTek T10000 cartridge will expand to 1 TB with the next-generation drive, doubling capacity by retaining near-term media investment. It is also important to note that the new IBM TS1120 will operate in 3592 mode (300 GB and 40 MB/s) when placed behind VTS, which is nearly half the capacity and 1/3 the performance when compared with the T10000 behind VSM.

Access Centric Drives

In order to minimize the size, and in turn the cost of the disk buffer, the access-centric drives in a StorageTek VSM system are key. As the next layer of storage in the product hierarchy, the characteristics of the access-centric drives, more than any other factor, determines the disk buffer requirements (for a given workload). This is one more area where Sun excels at delivering business value.

The StorageTek 9840C tape drive offers the fastest access to mainframe tape data available. This raw speed enables the implementation of a smaller disk buffer, driving down storage cost and delivering business value. Sun's policy software automatically manages the hierarchy, so there is almost no management cost associated with achieving this value.

In the access-centric comparison, Sun far exceeds the tape drive offerings of IBM 3592 and TS1120 products with the StorageTek 9840C tape drive. Only Sun offers and can mix access-centric (StorageTek 9840) and capacity-centric (StorageTek 9940 and T10000) technologies behind a virtual tape buffer at the same time. In effect, IBM only offers capacity-centric drives.

In a similar comparison, access time to data is key, especially with data that is between 48 hours and 33 (or so) days in age (month-end processing). The StorageTek 9840C drive has an average access to data of 12 seconds with a 40 GB uncompressed tape cartridge, about twice as fast as the nearest competitor, whereas an IBM TS1120 drive with a 500 GB uncompressed tape or a 100 GB uncompressed tape cartridge has a 49 and 27 second average access to data, respectively. This overwhelming difference (when comparing access times of the StorageTek 9840C 40 GB tape and IBM TS1120 100 GB tape) of 180% is significant to overall system processing and throughput. It also points out that IBM would

need to offer a much larger VTS disk buffer than Sun to provide the same type of overall system throughput, while losing the ability to offer an equivalent cost of data storage to the user.

In any small file, file-stacking application that attempts to maximize media utilization (whether it be a Hierarchical Storage Management (HSM) or virtual tape environment), the fast access/recall times of the StorageTek 9840C significantly outperform any IBM tape drive in the TS1120 family as well as any other tape drive on the market today, when data transfer as well as tape mechanics are considered. The following charts help demonstrate the important point that Sun offers the best tape technology in a virtual tape environment.

Figure 5 covers smaller files with access time comparisons; this is the most common scenario in an OS390-centric processing environment. The information is the outcome of a test performed by Sun using an IBM 3592 60 GB capacity tape. The request was to load the tape, read the first 100 MB using 256 K record sizes, rewind, and unload. Robotics were not considered here nor was the effect of channel or controller (I70) to drive ratios that can impact bandwidth. The StorageTek 9840C was 100% faster than the IBM 3592 drive.

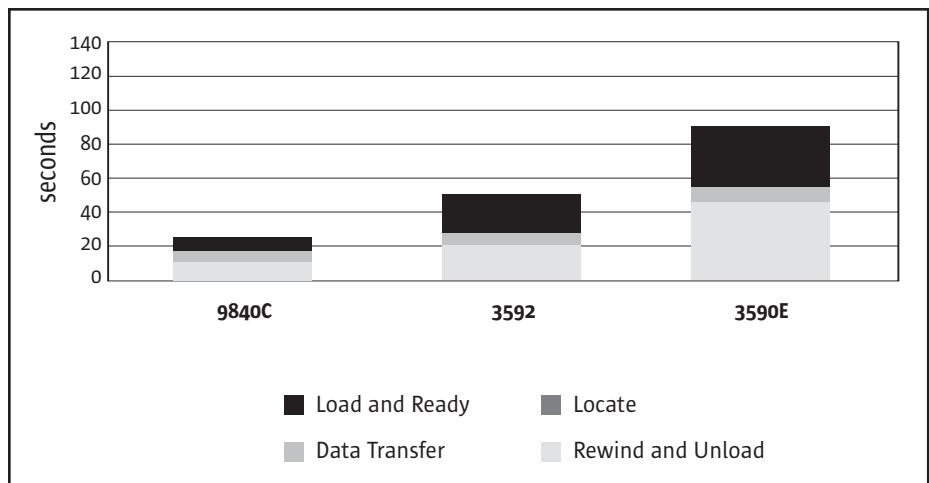


Figure 5.

Source: Sun Internal

Figure 6 is the outcome of a test performed by Sun using an IBM 3592 60 GB capacity tape. The request was to load the tape, search down over 1.25 GB of data, read 100 MB, rewind, and unload. Robotics were not considered here, nor was the effect of channel or controller (J70) to drive ratios that can impact bandwidth. The StorageTek 9840C was 75% faster than the IBM 3592 drive.




Figure 6.

Source: Sun Internal

Summary

Libraries and tape drives are two significant parts of a virtual tape system. With the SL8500, PowderHorn 9310, and 9740 library solutions, and combinations of access-centric StorageTek 9840C and capacity-centric StorageTek 9940B and T10000 drives, Sun’s StorageTek products have clearly demonstrated superiority over the IBM libraries and tape drives.

VSM Tape Buffer

Highlights

- Centralized control information enables the effective integration of up to 256 disk buffers in a single system image
- Central control, coupled with the system's high-performance FICON back-end architecture, enables a multi-site, single-image system

The disk buffer of a Virtual Tape System (VTS) is where applications write and read their virtual tape data. The ability of the system to make this disk buffer appear to applications as a set of tapes and tape drives is the reason the system is described as virtual.

Most of the technology used in a virtual tape system is dedicated to the process of managing this disk buffer: ensuring free space is available, protecting data written to the buffer, retrieving data to the buffer quickly, and sending data to and from the attached application hosts quickly.

Sun's key advantage is in managing the disk buffer, called the Virtual Tape Storage System, as part of a larger VSM system. Up to 256 disk buffers can be managed as a single system.

The distinguishing factor between IBM's VTS disk buffer and Sun's StorageTek VSM disk buffer, the Virtual Tape Storage System (VTSS), is not what they are as much as how they are integrated into the complete virtual tape system and how this integration allows enhanced advantages to be readily available. The largest difference between the two vendors is that IBM's VTS operates as a unique single system, while Sun's VTSS operates as a part of a larger system. For example, in IBM's architecture, once a virtual tape volume is written to a VTS it belongs to that VTS and can only be moved to another VTS by their Export-Import utility. Each VTS has its own dedicated set of Virtual Tape Volumes (VTVs), physical tape drives and physical tape cartridges. (We should note that in IBM's Peer-to-Peer VTS, only two disk caches act as one; VTVs being mirrored between them. The statement remains true as the two caches are managed as one VTS). This restriction is due to the control information for that VTV being held in a database local to the VTS, which has no knowledge of other VTS buffers.

The StorageTek VSM is controlled by a common shared database, the Host Software Component (HSC) Control Data Set (CDS). All the disk buffers in that system are known to the CDS. This allows virtual volumes to be moved from one disk buffer to another.

This capability enables several potential benefits:

- High-availability failover between multiple disk buffers without loss of access to physical tape volumes, without mirroring disk buffers, reducing the cost of high availability and driving business value
- Flexibility in designing disaster recovery approach with both local and remote disk buffers under the control of one StorageTek VSM system and with both local and remote tape drives and libraries without requiring disk buffer mirroring
- Disk buffer units may dynamically share the tape drives used for migration and recall, which can enable the use of fewer tape drives and reduced system cost

Another key element of the StorageTek VSM architecture is the use of FICON/ESCON connectivity to the back-end tape drives. This enables the connection between disk buffers at distances applicable to disaster recovery applications. The FICON back-end connectivity provides for fast migrations and recalls of virtual volumes, potential channel simplification, and quicker disk buffer-to-buffer data movement for improved disaster recovery protection.

Combining central control and fast back-end connectivity at distance creates the potential for a StorageTek VSM system spanning multiple sites and acting as a single system. This architecture is efficient and cost effective against a single or multiple site disaster.

Sun StorageTek VSM Software

The Virtual Tape Control System (VTCS) is host-based software used to control how the system operates and allocates resources. It is the central point of control that allows the StorageTek VSM to protect data and make efficient use of disk buffer space, tape drives, tape cartridges and libraries.

Highlights

- Host-based control of the StorageTek VSM resources enables integration with mainframe management and monitoring operations
- Central control of the system enables intersystem operations for data movement, replication, and load balancing
- Management Classes (categories of data) enable automated operations and reduce TCO by meeting service-level requirements at the lowest possible cost
- Storage sub-pools (groups of physical tape resources) enables disaster recovery applications and physical separation of stacked data on cartridges by application or user

Some of the important benefits of VTCS software include:

- Enables control of the StorageTek VSM system from the central consoles (or TSO, Time Sharing Option terminals) as with any mainframe software
- Integration of the StorageTek VSM controls and reporting with mainframe-based systems management applications, such as CA/Unicenter or BMC/Patrol, to participate in “single-pane” monitoring with the rest of the infrastructure
- Resource management of multiple StorageTek VSM systems, for example, VTCS can balance workloads across multiple virtual disk buffers

The host VTCS software controls the movement of the data, but plays no role in actual data movement, either from the host to the disk buffer or from the disk buffer to the tape drives performing migration and recall. Thus, no overhead is imposed by VTCS on the host for data movement. Microcode within the StorageTek VSM unit handles data movement between the buffer and physical tape.

VTCS software is the central intelligence that allows the StorageTek VSM to do all the smart things it does to protect data and make efficient use of disk buffer space, tape drives, and libraries. By running this software on the mainframe, the StorageTek VSM establishes a central control point that has knowledge of, and can use and commonly share, all the disk buffers, tape drives, and libraries defined to the system. For example, StorageTek’s VSM system offers buffer load balancing where multiple virtual tape disk buffer systems resources are “workload” balanced as virtual tape processing is required across processing systems. This contrasts to IBM’s VTS where the software runs on separate systems in the disk buffer units and the tape libraries, meaning that each unit is ignorant of others and cannot share data and resources with them. This adopted architecture of IBM creates islands of automation, which become a major limitation to scalability and ease of operation by adding management complexity. Interestingly, in order to implement their Peer-to-Peer VTS, IBM had to add another layer of control, the 3494-AX0/CX1, to have knowledge of and coordinate the activities of two otherwise separate VTS buffers.

Sun’s VTCS software has broad functionality that takes advantage of the simplicity and power of central mainframe-resident software. A subset of the features that deliver high business value includes:

- Controlling use of the disk buffer and physical tape via management classes
- Controlling pools of physical tape resources that enables aggregation of real drives used for migration and recall

Management Classes

Classes of data are established to provide a mass-customized approach to the placement of data within StorageTek VSM. Data objects are categorized and managed based on the service-level requirements of the category. This enables efficient use of the disk buffer within StorageTek VSM.

VTCS software sets out to do this by establishing management classes that specify how a particular category of data should be handled and by providing the user a place to specify rules, called “policies,” that put files into the appropriate category. IBM’s VTS was introduced treating all data the same. It is only recently that IBM has begun to allow users some control of how their data is treated by VTS.

A typical example is that of data (e.g., a management class) to be archived, which can be migrated to capacity-centric tape (StorageTek 9940 or T10000) at a remote location or to a local capacity-centric sub-pool for exporting to offsite for data protection, and concurrently migrated locally to access-centric tape (StorageTek 9840) for fast recall.

Management classes are used to establish policy for data migration from the disk buffer to physical tape. Policy options available to act on a management class include:

- The movement of data from the disk buffer to physical tape can be controlled at the management-class level. This can be specified as “move immediate,” “hold in buffer,” or “copy to tape and hold in buffer.” This flexibility helps StorageTek VSM deliver business value by ensuring that service levels are met at the lowest possible cost.
- Multiple copies (up to four) can be migrated to different libraries both local and remote. For important data classes this can add a significant level of protection from media failure and enhance disaster recovery. Sun’s StorageTek VSM adds value by automating data protection based on policy without human intervention.
- The ultimate control of a data class is to specify the desired number of hours that it will remain in the disk buffer. This is normally used as an “exception” policy reserved for very critical data. StorageTek VSM adds value by providing flexibility in control based on policy so that data is protected in the way that best fits requirements.

Storage Sub-Pools

When VSM migrates (writes data from the disk buffer to physical tape) it aggregates virtual volumes on physical tapes to improve cartridge utilization. This is commonly referred to as “stacking.” The control of how volumes are stacked is a second area where StorageTek VSM provides control through software.

Storage sub-pools are collections of real tape cartridges that a group of virtual volumes are stacked to during migration. This powerful feature has several uses that can provide real business value:

- The ability to physically separate data from different sources (users, applications, etc.). This can support security, compliance, audit, or user-preference requirements where it may be important for data from different sources not to be intermingled on stacked cartridges.
- Storage pools facilitate the preparation of offsite data for data protection or disaster recovery. The typical scenario is that data is flagged for offsite and then stacked during migration to a second sub-pool of physical cartridges. The second sub-pool of cartridges is exported and ejected to be taken offsite. Since they have already been written (when they were migrated), no copying of data or human intervention is required to select the offsite data. A remote StorageTek VSM allows the cartridges to be imported, making them available to the system in minutes, without copying any data. If the remote site does not have a StorageTek VSM system, Sun has utilities to allow the use of StorageTek VSM’s stacked cartridges with standalone tape automation systems.

Summary

StorageTek Virtual Storage Manager deliver outstanding value through its superior architecture. Central control of StorageTek VSM enables efficient resource balancing and advanced data protection, business continuance and disaster recovery. Total systems throughput is delivered via the unique approach that Sun has used to integrate the world-class devices into a truly unique system.

Enhanced Data Protection

Organizations do not have to make trade-offs between their level of data protection and their business operations. Data protection applications run faster, reading and writing data to disk while the StorageTek VSM system manages the downstream copying to tape based on policies. Recovery is at disk speed. Backup windows are met. StorageTek VSM enables policy-driven copying and migration offsite of critical business information.

Disaster Recovery and Business Continuance

The combination of features available in the StorageTek VSM control software facilitates disaster recovery and business continuance applications. The system's fast data replication capability copies data and migrates it offsite quickly and efficiently, without intervention. A broad set of utilities enables a disaster recovery/business continuance solution that is tailored to business needs and automated via policy management. A key function, Concurrent Disaster Recovery Testing (CDRT), allows users to test their disaster recovery processes without interrupting systems operations, while continuing normal data protection operations.

In addition, StorageTek VSM vaulting utilities allow customers, that do not have a second StorageTek VSM system, to take advantage of the stacking capabilities of virtual technology. The vaulting facilities utilize ExLM and the local tape management system to control the cycling of StorageTek VSM stacked tape offsite, defragmenting the appropriate tapes, and retrieving the proper tapes back onsite through an advanced feature technique. This minimizes tape handling, media requirements, and offsite storage costs as well as improves disaster recovery options. IBM VTS has no facility to achieve this benefit.

Business Productivity

StorageTek VSM enables organizations to lower their operational costs while meeting service level requirements. Costs are contained by improving the utilization of storage assets. Service levels are met by accessing tape data at disk speeds and utilizing virtual tape drives to handle peak processing workloads. Batch times can be reduced, freeing up mainframe resources.

The combination of StorageTek VSM support for multiple drive types, management classes, and storage sub-pools can enable use scenarios that drive significant business value.

StorageTek VSM helps organizations adapt to business changes by providing flexibility in how IT resources are used to meet shifting IT priorities.

