

The Oracle StorageTek VSM6

Providing Unprecedented Storage Versatility



Fred Moore, President
Horison, Inc.
www.horison.com

Introduction

Have you considered the growing complexity and challenge for enterprise data centers trying to meet their data protection, batch, primary data, and archival storage requirements? Without careful planning, each of these tasks could entail pursuing its own unique and costly storage solution creating even more complexity. As a result, versatility is quickly becoming *the* new mandate for solutions to meet many of these enterprise class (mainframe) data center storage requirements. Fortunately it's the 21st century and the combination of storage virtualization, disk, tape, and robotic libraries have successfully been integrated into an easy to use cost-effective solution to address these issues. This report will discuss the Oracle StorageTek Virtual Storage Manager 6 (VSM6) subsystem and how it addresses these challenges with particular focus on data protection and advanced backup architectures. With Oracle StorageTek Virtual Storage Manager 6 (VSM6) you can provide data protection and archive services for all of your data in one integrated architecture.

Defining Higher Levels of Data Protection and Availability

The two primary components of data protection are business continuity and disaster recovery. In most cases these are (still) the two most pressing IT issues. It's important to distinguish between a business continuity event and a disaster recovery event as they are different. With increasing reliance on what is in most cases a company's most valuable asset – digital data - an organization can quickly lose millions of dollars associated with lost access to data along with much of its competitive advantage and credibility as has been the case in many highly visible and publicized security breaches. Fail-proof data backup and recovery is more critical to an organization's survival than ever before as few businesses can survive for any period of time without their IT function. With our primary focus on backup architectures and protecting data, it's important to understand the various roles of backup. We can define "backup" as simply a data protection process for executing two primary data center functions:

Business continuity - Provides a local backup copy of data to be used should an application or infrastructure component fail or data become corrupted. Business continuity is the activity performed by an organization to ensure that critical business functions will be available to customers, suppliers, regulators, and other entities that must have access to those functions. Backup makes copies of data that can be used to recover or restore the original data in a required timeframe after a data loss event. Business continuity is not something implemented at the time of a disaster; rather business continuity refers to those activities performed *daily* to maintain quality of service, consistency, and recoverability. For business continuity, fast initial access time to back up data is critical. Therefore, a disk array is often utilized to enable fast backup and restore of smaller chunks of data such as system and user files, emails, and incremental backups. When using disk for backup, a business needs to consider how to handle disk

hardware failures and or application failures when they occur meaning that tape is optimally used in conjunction with disk in order to provide the highest availability levels.

Disaster recovery - Provides a copy of data which can be maintained at a different off-site or physical location which can be restored to another location, should the primary data center facility no longer be available. Disaster recovery (DR) is the process, policies and procedures that are related to preparing for a recovery after a natural or human-induced disaster occurs that may impact the entire data center functionality rather than specific files. While business continuity involves planning for keeping all aspects of a business functioning in the midst of disruptive events, disaster recovery focuses on the availability of the IT infrastructure systems that support the critical business functions. For disaster recovery (DR), fast data transfer time and high availability are the critical components. Therefore, tape has become the optimal DR choice, enabling faster backup and restore of large amounts of data (large databases, servers, or even the restoration of entire data centers).

Many mainframe enterprises have adopted a three site strategy and have chosen to configure their DR strategies this way. In a recent survey, forty percent of enterprises locate two of their three data centers in the same metropolitan region but in two separate data centers that are within a 40 km radius of each other while 60 percent usually have two discrete data centers within the same building but on different floors or on the same floor with a firewall separating the data centers.

[Source: Enterprise Tech Journal March April 2013, pp 18-19].

Current Problems Related to Data Protection are Cited

An independent survey of 500 CIOs conducted in November 2012 across USA and Europe found that enterprises (large-scale IT organizations with more than 1000 employees) are still not experiencing the full benefits that storage virtualization offers or the advantages that modern data protection tools and architectures can provide. Key findings from the study shows that capability, complexity and cost issues are hampering IT departments:

- 68% of CIOs feel that their backup and recovery tools will become less effective as the amount of data and servers in their organization steadily grows
- 88% of CIOs experience capability-related challenges with backup and recovery, 87% with cost and 84% with complexity showing that data protection is still not a simple task
- 58% of CIOs are planning to change their backup tool for virtual environments by 2014
- Financially, CIO respondents stated the cost per hour of downtime - for their business critical servers that are not been protected by replication - as \$324, 793. Coupled with a recovery time of 5 hours or more this means that, on average, each outage is costing organizations at least \$1.6 million. Unless data protection evolves, these costs will remain high and continue to increase. This is even more significant in mainframe environments where many of the world's most mission-critical applications are hosted.

[Survey Source: The third edition of Veeam Annual Virtualization Data Protection Report 2013].

Implementing a Modern Data Protection Strategy

What are the best data protection options available? In its most basic form, we have defined "backup" as a tool or method for executing business continuity processes or disaster recovery processes. The

challenges of effective backup are numerous given the range of application availability requirements, but so are the options. Effective data protection plans address both business continuity and disaster recovery by defining operational procedures, implementing hardware redundancy, and practicing/testing the recovery processes. With mounting pressure to reduce the amount of time required and the amount of storage consumed by backup/recovery methods, many new and improved solutions are now available. These offer a variety of choices depending on what operating system is used, the type of storage technologies used, when and how the data is protected, if compression, encryption or WORM is used, and if additional geographic locations are involved for data redundancy.

Backing up and later restoring potentially huge amounts of data in the least disruptive manner is becoming increasingly difficult given the tremendous amount of digital data growth. Disk has become the preferred business continuity backup target for smaller data files demanding the fastest RTO (Recovery Time Objective), while tape is the optimal backup choice for larger files, disaster recovery and archive processes. An all-disk data protection *and* archiving solution is an increasingly expensive option to acquire and to operate. Recently published studies indicate that the TCO (Total Cost of Ownership) for disk ranges up to 26x higher than tape according to the new Clipper Group Calculator dated May 13th, 2013 than tape for backup and archiving. See [Analyst Report \(Clipper\): Revisiting the Search for Long-Term Storage - TCO Analysis of Tape/Disk](#) (June 2013). The initial acquisition cost or purchase price per GB is also much lower for automated tape libraries. Remember - data that is not being used should not consume energy. To best address business continuity, disaster recovery, and archival data requirements the optimal solution is available and deploys a highly scalable tiered storage approach using *both* disk and tape.

Meeting RTO and RPO Requirements When a Disruption Occurs

A data recovery event involves restoring the most current image copy from disk or tape, then applying any applicable log data to get to a valid recovery point. With tape, when a backup image expires, the tape cartridge is returned to the backup scratch pool for reuse. With disk, when a backup image expires, the image is erased from disk and the disk space is freed up for more backups. Unplanned downtime, by its nature, is a surprise and can have numerous causes including hardware failures, software errors, intrusion, user errors, poor maintenance, or a system upgrade as shown in the chart below.

Type of Storage Disruption or Outage	% of all Incidents	Possible Protection Solutions
Hardware, network, or system failure	44%	Mirroring, RAID, backup copies, transaction logs, clusters, failover, virtualization
Human error, equipment theft	32%	Snapshots, CDP, surveillance, facilities security
Software and coding errors	14%	Snapshots, CDP, replication
Intrusion – security, virus, web scams	7%	Firewalls, authentication, anti-virus, encryption
Natural disasters, power outage, flood, building damages, fire....	3%	Off-site compute facilities, emergency power, UPS

Source: Data from a variety of industry sources

RTO (Recovery Time Objective) – How long can businesses successfully operate without this data? Once the problem is discovered, RTO is the time required to recover from a data loss event and return to service. In other words, this requires classifying data or an application by its criticality or value to the business and determining how long the business can survive without having this data available. RTO is a key availability metric as there is constant business pressure to reduce the length of time it takes to recover a critical file or an application. The more critical data, the smaller the RTO becomes.

RPO (Recovery Point Objective) – What is the desired amount of time between data protection events? Keep in mind that RTO and RPO are user defined policies based on the criticality or business value of the data being protected. Defining the RTO and RPO are key components of a successful data classification plan. The user doesn't have much control over how long it takes for the problem to be discovered. The user does have control over what means (technology solution) from which to recover and the level of redundancy required based on the value of data.

Establishing RTO and RPOs for applications is an essential process to attain the highest availability levels and meet SLAs (Service Level Agreements). With optimal data protection architectures, there's always a tape safety net in case a recovery from disk fails (regardless of the data's value). In addition, HSM (Hierarchical Storage Management) software allows you to set pre-defined policies to take advantage *both* of the performance of disk and the economics and superior reliability of tape where it makes the most sense from a RTO and RPO perspective. Taking it to the next level, why not develop *one* subsystem to do it all? This builds a strong case for a multi-tiered and integrated storage system which includes both physical disk and tape storage (optional) for business continuity, disaster recovery, long term archival and big data requirements for today's z/OS applications. That is precisely what the Oracle StorageTek VSM6 subsystem is all about – combining disk and tape with storage virtualization to provide a seamless backup architecture for all RTO and RPO requirements for any data classification category. The following table provides traditional suggestions for which types of recoveries should optimally be supported by disk and which ones should be optimally supported by automated tape.

Data Classification by Value	Mission-critical	Vital	Sensitive	Archive
Avg. Data Distribution by class	15%	20%	25%	40%
Availability index	99.999%	99.99%	99.9%	99%
Downtime minutes/year	5.256	52.56	525.6	5256

You can use VSM to meet *any* RTO and RPO for all classes of data

Any RTO Any RPO	VSM	VSM	VSM	VSM
--------------------	-----	-----	-----	-----

Or do it the traditional way...by using disk and automated tape separately to meet specific RTO and RPO requirements

30 min 0 min	Disk	Disk	Disk	Disk
2 hours 15 min	Disk	Disk	Disk	Disk
12-24 hours 2-6 hours	Disk	Disk	Tape	Tape
>1 day 12-24 hours	Disk/Tape	Disk/Tape	Tape	Tape
> 1 week > 1 day	Tape	Tape	Tape	Tape

Typical RTO
Typical RPO

The Advantages of Virtual Tape in Achieving High Availability

The Virtual Tape Library (VTL) concept first appeared in the late 1990s for mainframes and has become a successful storage virtualization technology used primarily for backup, recovery and archive purposes. A System z mainframe VTL presents a storage component (disk arrays) as multiple virtual tape drives integrated with a physical tape library for easy use with existing backup and HSM software. The benefits of this virtualization includes storage consolidation, faster data restore times and successful RTOs. Note: For open systems, no physical tape is used with the VTL and these VTLs are sometimes referred to as a “tapeless” environment. A VTL enables the storage hardware to be switched from tapes to disks while continuing to use the existing and familiar tape backup software and processes but this approach doesn’t efficiently address tier 3 and long-term archive data requirements.

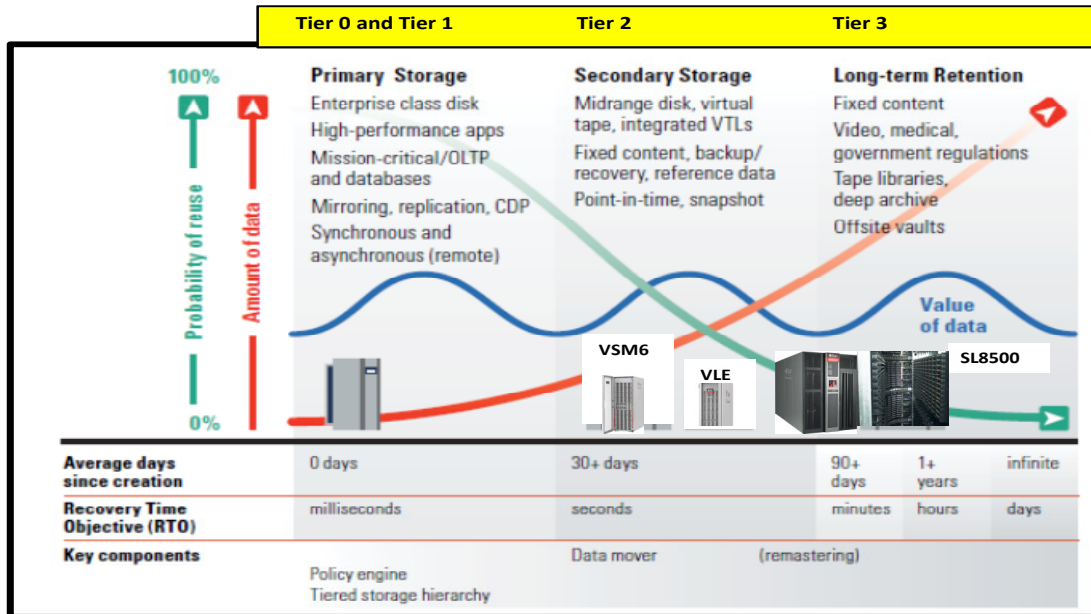
StorageTek VSM 6 Arrives

Oracle's StorageTek Virtual Storage Manager 6 was announced in Nov. 2012, and is the latest generation of its successful mainframe virtual tape system. StorageTek Virtual Storage Manager 6 (VSM 6) is the first and only mainframe virtual tape storage system to provide a single point of management for the entire system that leverages the highest existing levels of security provided by System z environments. It provides a unique multi-tiered storage system which includes both physical disk and tape storage (optional) for business continuity, disaster recovery, long term archival and big data requirements for z/OS applications. Each StorageTek VSM 6 emulates up to 256 virtual tape transports but actually moves data to and from VSM disk storage and backend real tape transports attached to automated tape libraries. As data on StorageTek VSM disk ages, it can be migrated from VSM’s disk buffer to physical tape libraries such as a StorageTek SL8500/SL3000 Modular Library Systems that are attached to the StorageTek VSM 6 to meet long-term data retention, archive and compliance needs.

The StorageTek VSM 6 system is dual-node architecture which enables non-disruptive, on-demand capacity scaling and can grow to 1,200 TBs (1.2PB) of disk storage in a single rack. StorageTek VSM 6 provides 2x the performance and 13x more disk capacity than the previous generation StorageTek VSM 5 system and scales disk capacity on-demand to 307 petabytes (with scalability up to 256 StorageTek VSM systems) providing more than one exabyte or “exascale” capacity . A simple and seamless migration of data from disk to physical tape extends system capabilities to provide nearly unlimited capacity to meet your enterprise’s long-term data retention and compliance needs. With StorageTek VSM 6, protecting data throughout its lifecycle has become a reality - with a single architecture. The extensive versatility of StorageTek VSM 6 is evident.

Lifecycle Data Protection with VSM6

Integrates storage virtualization, disk, and tape



Source: Horison Information Strategies

StorageTek VSM 6 Delivers Data Protection by Providing Business Continuity, Disaster Recovery, and Archive Capabilities

StorageTek VSM 6 offers the flexibility to address multiple enterprise configuration requirements including high performance disk-only or massively scalable disk and physical tape, as well as single or multi-site support for disaster recovery. In addition, physical tape channel extension can be used to extend the StorageTek VSM 6 storage to span onsite and offsite repositories that are geographically dispersed. Up to 256 StorageTek VSMs can be clustered together into a *Tapeplex* which is then managed under one single point of management as a single large data repository using Oracle's StorageTek Enterprise Library Software suite (ELS). Note: Tapeplex is a collection of StorageTek tape hardware represented by one StorageTek ELS. Mainframe systems place the highest demands on tape storage using it for a variety of applications such as primary data, batch, data backup, HSM, archive and big data applications. StorageTek VSM 6 uses user-defined policy-based management and simplifies day-to-day operations by automating tasks that were previously handled manually thus enabling IT staff resources to be deployed on higher value, more strategic assignments.

With StorageTek VSM 6's optional StorageTek Virtual Library Extension (VLE) from Oracle, a second-tier SAS disk array storage feature, StorageTek VSM 6 is the *only architecture* supporting multiple tiers of storage under a single point of management. Many data types typically need to be stored on the disk subsystem longer because of their usage patterns, which could be from 45 to 90 days or more for example, before the probability of re-use diminishes. StorageTek VLE is a shared disk storage resource for StorageTek VSM 5 and StorageTek VSM 6, but is also a replication engine that migrates data between StorageTek VLE's, which takes the backup and replication workload outboard from the StorageTek VSM itself therefore freeing up more StorageTek VSM resources for host related activities. StorageTek VLE adds a new disk-based higher performance dimension to the StorageTek VSM 6 architecture. Without StorageTek VLE there is more reliance upon tape storage, which can create situations where data is

being constantly migrated and recalled back and forth from tape which is not an effective use of tape resources. The StorageTek VSM disk buffer and StorageTek VLE disk array both use triple bit parity for an even higher availability level.

StorageTek VLE also allows customers to not utilize tape or not have directly attached tape, if they desire. This is often called “going tapeless”. Going tapeless is not always recommended since tape is the by far the most cost-effective means to store data for longer periods of time. The larger the storage environment, the more benefits tape provides. Remember data that is not being used should not consume energy!

StorageTek VSM 6 offers several compelling replication solutions to provide the highest levels protection and availability. For business resumption and disaster recovery purposes, the versatility of StorageTek VSM 6 provides the capability to make additional backup copies of primary data, keep data in multiple sites in synch, and can copy data to offsite and multiple geographic locations. StorageTek VLE to StorageTek VLE disk replication, synchronous or asynchronous node clustering across FICON or GbE (Gigabit Ethernet), Cross Tapeplex replication (CTR), and real tape drive channel extension for remote site physical tape support all provide exceptional levels of enterprise data availability and redundancy.

Don't Forget About Archive Data, Unprecedented Demand and Moving to the Era of Colossal Content

For the overall storage industry, digital data overall is growing 30-40% annually. However tier 3 data is experiencing a higher CAGR (Compounded Annual Growth Rate) in many cases and includes digital archives, fixed content, digital images, entertainment, multi-media, video, social networks, surveillance, compliance data and are the fastest growing data category. The vast majority of this data reaches tier 3 (tape storage) archival status typically within 30, 60, or 90 days mandating that a more intelligent architecture arrives that can seamlessly migrate data to tape from disk based on user defined policies. VSM6 directly addresses this issue.

It's Important to note that tape technology has surpassed disk in reliability and remains the most secure and cost effective means of storing and protecting long-term archival data. The optimum storage solution for the future will address the needs for business continuity, disaster recovery as well as archival data by taking advantage of the benefits of *both* disk and tape technology. Clearly versatility in a storage solution has become *the key* ingredient for attaining the highest availability levels possible for meeting these looming storage requirements.

Oracle's StorageTek Tape and Libraries

Effective use of modern tape technology is *the* key component that makes the StorageTek VSM 6 tiered data protection solution so cost-effective. Oracle's StorageTek T10000C tape drive provides a native cartridge capacity of up to 5.5 TB, making it the largest capacity tape cartridge ever announced. The StorageTek T10000C has a native data rate of up to 252 MB/sec., offering significant improvement in recovery times and making disk drives the limiting factor in large-scale recovery performance. Uncorrectable bit error rate for the StorageTek T10000C tape drive is three orders of magnitude better than the most reliable Fibre Channel disk drive and The StorageTek T10000C media has a life of 30 years or more. Encryption and WORM are available with the StorageTek T10000C for further protection of data at rest. The StorageTek SL8500 modular library system can scale to a capacity of more than one exabyte (1×10^{18}). A single bit in error can render data that is deduplicated, encrypted or compressed unusable. Relative to disk technology in recent years, tape technology progress has been even more significant on many fronts; tape has a higher cartridge capacity than disk, a faster data rate than disk,

and is more reliable than disk. Oracle’s future StorageTek tape roadmaps indicate that even higher capacities, more extensive enhancements and new functionality lie ahead.

StorageTek VSM 6 – Things to Consider

Architecture	Combines and integrates storage virtualization, disk, and tape
Applications supported	Selected primary data, batch, backup, DR, HSM, archive and big data
Disk only VTL option (tapeless)	Yes – with StorageTek VLE option
Remote site storage	Yes - using Ficon or Gigabit Ethernet
Create non-disruptive backup copies of StorageTek VSM 6 data sets	Yes - without host involvement to multiple locations
Lifecycle data management	Yes - from creation to end-of-life
Migrate aging or less active data to tape	As StorageTek VSM data ages, it can be migrated from disk to physical tape libraries based on user-defined policies
Technology migration	StorageTek VSM and physical tape migrations are made seamless. Different hardware generations can co-exist (i.e; VSM and tape drives and libraries) with the ability to do full data interchange between all generations of VSM’s
Capacity scalability	Can scale capacity to over an exabyte (1x10 ¹⁸)
Cost of ownership	Optimizes extensive TCO advantages of tape with benefits of disk

Summary

Data protection has become the most critical IT discipline as most businesses in the modern world can no longer survive without their IT function. As a result, we are beginning to see the next generation of data protection solutions appear as legacy processes become increasingly burdensome, expensive, and unreliable. Today’s optimum data protection solution ensures business continuity disaster recovery, and provides the most cost-effective means of protecting tier 3 data using tape.

Data protection throughout the lifecycle of data is not an option – it is now a requirement for survival. Engineered with versatility as a fundamental component, StorageTek VSM 6 can be deployed to meet each enterprise’s unique business requirements. From high-performance disk only or massively scalable “disk and tape” configurations to single site or multisite support, the StorageTek Virtual Storage Manager System 6 deployments meet each application’s specific requirements balancing speed, capacity, and availability so that the right data is on the right device at the right cost. Oracle is the only company to offer a fully integrated disk, tape and virtualized data protection solution with its StorageTek VSM 6 architecture. Given this, maybe the StorageTek *Virtual* Storage Manager 6 from Oracle should also be called the StorageTek *Versatile* Storage Manager....

End of report