

White Paper

StorageTek Virtual Library Extension (VLE)– a Single Common-Management Tiered- storage Infrastructure



Josh Krischer
December 2010

2010 © Josh Krischer & Associates GmbH. All rights reserved. Reproduction of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Josh Krischer & Associates GmbH disclaims all warranties as to the accuracy, completeness or adequacy of such information. Josh Krischer & Associates GmbH shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice. All product names used and mentioned herein are the trademarks of their respective owners.

StorageTek Virtual Library Extension (VLE)–a Single Common-Management Tiered-storage Infrastructure

Executive Summary

An average IT organization’s storage growth is currently around 60 percent p.a., mainly in unstructured storage. In the last few years, average disk storage prices fell at between 30 to 35 percent per year. Many recent surveys show the same result: The main concern of IT executives is how to cope with the annual storage growth in economically challenging times. The slowdown of the global economy has had minimal impact on storage growth, and market research from the last three quarters shows that, despite budget constraints, organizations cannot continue to delay storage procurements any longer. Storage budgets have not grown (and in many cases remained unchanged or even shrank) to allow to close the gap between the rising requirements and the storage price erosion; therefore, the only sound solution is to build more effective storage infrastructures such as tiered storage, for example.

No other product in IT has been declared dead as vehemently as the magnetic tape (mainly by vendors that do not sell tapes). However, tape storage refuses to die, and, in fact, current requirements are the best advocates for tape usage. Storing data on tape has several unique benefits: It is more economical than storing on disk (even in comparison to SATA disks with deduplication), is portable (which allows for export/import), can be stored offline for a relatively long time, and has the smallest floor space and energy requirements per stored capacity, which improves sustainability and complies with corporate social responsibility.

The only way to cope with the imbalance between storage requirement growth and available budgets is to step out of the box and think differently. Existing VSM users may consider upgrading to the VSM5 (retaining investment protection for their physical tape library) and adding Tier-2 disk storage with the VLE. New customers who have no VTLs and are looking to lower costs and decrease their environmental footprint should evaluate the benefits of a VSM/VLE configuration.

Contents

Executive Summary	2
The Virtual Tape Library	3
Oracle Virtual Library Extension (VLE) Subsystem	4
Oracle’s StorageTek Virtual Storage Manager 5 (VSM5).....	5
Mainframe Market	6
Conclusions and recommendations	7

The Virtual Tape Library

For fifteen years, Virtual Tape Libraries (VTLs) have been popular in the mainframe domain. Simply put, a VTL subsystem contains a tape library, a server with external (or internal) storage that act as a cache buffer for the library, and host management software. At retrieval, files are

“Oracle’s new Virtual Library Extension (VLE) subsystem provides the best of both words, allowing the building of an effective tiered storage infrastructure.”

loaded from the tapes to the disk storage to allow for fast data transfer between the host and the VTL. The files on the disk are written in tape format. After modification, the files (compressed and compacted) are staged to new physical tape. To the host, the VTL looks like a usual tape library.

The first virtual tape system, the VTS, was launched by IBM in 1996 as part of the *Seascape* architecture, which, at that time, was based on standard IBM components such as the RS6000. It was followed two years later by StorageTek’s Virtual Storage Manager (VSM), which was designed around their leading tape libraries and the *Iceberg* high-end storage subsystem. The advantages of VTLs led to

their adoption in most large mainframe sites. What are the reasons to this popularity? The answer is the remarkable user benefits derived from VTLs, for example:

- Fewer physical tape drive requirements.
- Fewer slots in the library (smaller library).
- Lower energy requirements and CO² emissions.
- Better usage of tape media.
- Lower storage management costs.
- Improved performance (disk-to-disk data transfer).
- No waits for a “free” drive due to large pool of virtual tapes.
- Ease of migration to new tape technology.
- Better security.
- Ability for applications to support legacy tape formats on state-of-the art tapes.

The VTL benefits listed above translate into significant reductions in CapEx and OpEx and better environmental friendliness, which are even more important now than in the past, and entrench the VTL’s position as an effective tiered-storage optimization solution.

As mentioned above, VTLs are popular for the mainframe platform, but never achieved big success in UNIX or Windows infrastructures. For these platforms, disk subsystems emulating tape (virtual tape), with or without compression and deduplication, are more popular. Both solutions have their place in IT, with VTLs ensuring scalability and better economics, while the tape-emulation disks provide better performance, but with limited capacity. Oracle’s new Virtual Library Extension (VLE) subsystem provides the best of both words, allowing the building of an effective tiered storage infrastructure.

Oracle Virtual Library Extension (VLE) Subsystem

The Oracle Virtual Library Extension subsystem is a SAS-2-disk based Tier-2 virtual tape extension to the proven VSM5. The effective storage pool capacity (assuming 4:1 compression ratio) ranges from 220 TB to 3.52 PB (four VLEs), and can be shared among multiple VSM5 subsystems using IP connections (see Figure 1).

The host Virtual Tape Control Software (VTCS) sees the VLE as a single tape library storing volumes in tape format on virtual tape emulated on a disk storage subsystem. There are two options of where to store the Virtual Tape Volumes (VTVs): either on disk and tape, or disk-only. A VSM5 can stage to and recall volumes from a VLE, in a similar manner as is done with a real tape library.




Figure 1: Example of a single host with one VSM5 and two VLEs

Scalability and reliability: The VLE's minimal physical capacity is 55 TB, which can be upgraded in 55 TB steps up to a capacity of 220 TB in a single cabinet. The effective capacities should be much higher assuming an average 4:1 compression ratio. The data is protected using ZFS RAID Z3 (triple-parity RAID embedded in ZFS). This technique uses RAID groups in an 8D+3P configuration. The 3 parity drives ensure higher reliability for the 2TB HDDs RAID groups. To ensure data integrity on data transfers, Cyclic Redundancy Check (CRC) information is added to each page of data transferred to or from the VLE. For better availability and disaster protection, the VSM5 can store the same volumes on two VLEs simultaneously.

Connectivity: Multiple VSM5 can share VLEs; a single VLE can connect up to eight VSMs, and each VSM5 can connect to up to 4 VLEs to increase effective buffer space up to 4.52 PB. The VLE supports 16 IP connections, but having at least 4 connections per VSM5 is recommended. The native performance of each port is rated at 40MB/s¹; however, due to compression, the effective data rate can reach 160 MB/s per port. Multiple ports can be aggregated to achieve even higher data transfer rates.

VLE Management: VLE management is completely transparent to an existing VSM policy management. Initially, the VLE is “loaded” with virtual volumes as with a normal library. The existing archive policies can be exploited to transfer the data first to a VLE, and then, after a pre-set period, the data can be migrated to real tape; a tape-less configuration without any physical tape library is also supported.

Operation: VSM/VLE’s flexible policy management can move data to the VLE only, or to the VLE and a physical tape as well. It can also de-stage data from VLE-based to tape-based storage on preset time parameters. User-defined policies enable certain applications’ files to be directed to the VLE (e.g., backups whose recovery times are critical), while less performance-sensitive data such as archiving data can be directed to the VSM and tape library only. As a specific data set stored on a VLE becomes less active over time, its re-use is less likely, and it is migrated to tape. Temporary volumes can be simply deleted from the VLE. A VSM configuration without VLE (VSM and physical tape library only) requires more mechanical operation of the library as data is being more frequently migrated and recalled from tape. Using a VLE reduces wear and tear on tape resources and increases the reliability of the VSM5 subsystem.

Oracle’s StorageTek Virtual Storage Manager 5 (VSM5)

VSM5 is the latest model of the field-proven StorageTek virtual tape library for mainframes. It ensures transparent, seamless optimization of storage resources using automated policies and supports multiple StorageTek tape libraries and drives. Although it is true that the current VSM has its origins in the Iceberg, it is now a high availability subsystem utilizing state of the art RAID technology and advanced fault tolerance. The VSM architecture allows for up to 256 VSM subsystems to connect to a single system, which represents enormous scalability.

There are two models: the entry-level VSM5e, and the VSM5. The scalable VSM5 supports up to 90 TByte of effective cache-buffer for the attached library, up to 32 physical drives are supported² (seen as 256 virtual tape drives), ESCON- and FICON-type channels, and the ability to perform up to 613 MB/s sustained writes. The cache-buffer uses 450GByte FC-disks in a fault-tolerant RAID-6+ configuration. In addition to the mainframe channels, each VSM5 supports 4 IP connections for remote mirroring between VSMs and for data transfer to/from VLEs. Each VLE

¹ The VLE IP ports can run at full 1 GigE speed however the VSM5 is limited to 40 MB/s/

² Both the 9840 and T10K are supported, which ensures investment protection and creates two tiers of tape storage.

appears as a single tape library to the VSM5, so the VSM can migrate/recall virtual tape volumes just like with a physical tape library. All generations of VSM subsystems can co-exist within a single system.

Disaster Recovery infrastructures: VSM offers several options to deploy DR configurations. The two simplest one are simple physical export/import, or remote electronic vaulting at a safe location such as a shelter or vault. More sophisticated options include adding VSM (with or without VLE) at a remote site, clustering between two sites, data duplication (*Cross Tapelex Replication*) or virtual volume multiplexing. Different utilities such as *Physical Vaulting*, *Concurrent DR Test Utility and Recovery Utility* ease the management of the different DR solutions. The VSM also supports many-to-many site DR solutions. For example: Three VSMs can be connected in a triangle configuration with bi-directional data transfers among all three³. Another option is the cascading configuration in which each of the three is connected to an additional VSM leg.

Mainframe Market

Despite hard economic times, the mainframe market has enjoyed a constant growth over the last six years. In the third quarter of 2010, IBM's sales represented a 54 percent MIPS growth, the highest since 2003. This growth is attributed to equipment renewals, as well as to winning new customers, mainly in the BRIC countries (Brazil, Russia, India & China). India and China, in particular, have experienced a strong expansion of IT infrastructure in their respective banking, telecommunications, transportation, and energy sectors.

³ Recently a three-site solution has been gaining popularity, particularly within US financial institutions. In this scheme, there are two replications: a synchronous disk-storage replication at a near distance, which protects against a local disaster with minimum loss of data, and a remote replication at a safe distance (usually more than 200km) as protection against large-scale disasters. In case of a total loss or long outage of the primary site, the secondary (near) site becomes the primary, and the (far) remote site takes up the role of the main recovery site. If the disaster is wider-spread and impacts both production and near sites (flood, hurricane, power grid, etc.), then production recovers at the far site.

Conclusions and recommendations

Through its acquisition of Sun Microsystems, Oracle inherited Sun's StorageTek division, which can look back at 40 years of experience in tape, disk, and automated tape libraries technologies. The VSM5 is its fifth generation of virtual tape libraries with industry-leading tape and disk scalability. The VSM with migration to tape is best suited for volumes whose activity diminishes over time but whose data may still be required later; the new VLE allows users the flexibility to keep more active data sets on disk for longer periods of time.

Some additional benefits of deploying VSM/VLE in the data center:

- Single common-management tiered-storage infrastructure.⁴
- Scalable performance/capacity within a single system image.
- The only Tier-2 subsystem using triple-parity RAID.
- Advanced and flexible disaster recovery/business continuity options.
- Policy-driven management.
- Investment protection: all Oracle enterprise tape drives & libraries are supported.
- Lower energy requirements than disk-only tiered storage infrastructures, thus, reducing CapEx and OpEx costs, and decreasing CO² emissions.

The only way to cope with the imbalance between storage requirement growth and available budgets is to step out of the box and think differently. Existing VSM users may consider upgrading to the VSM5 (retaining investment protection for their physical tape library) and adding Tier-2 disk storage with the VLE. New customers who have no VTLs and are looking to lower costs and decrease their environmental footprint should evaluate the benefits of a VSM/VLE configuration.

Josh Krischer is an expert IT advisor with over 41 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Josh Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering mainframes, enterprise servers and storage from 1998 until 2007. During his career at Gartner he was responsible for high-end storage-subsystems and disaster recovery techniques. He spoke on these topics and others at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, as well as major vendor events.

Find more on: www.joshkrischer.com

⁴ The four tiers of storage—cache buffer (Fiber Channel disks as tier 1), VLE (SATA drives as tier 2), and tape (9840 and T10000 tapes as tier 3 and 4)—are managed by the VSM's common storage management. Once the mainframe writes data to the VSM, it is moved between the disk tiers and the tape library without mainframe involvement.