Josh Krischer & Associates
Enterprise Servers, Storage and Business Continuity

White Paper

Oracle's StorageTek Virtual Storage Manager 7 (VSM 7)–the most affordable mainframe storage with native cloud integration

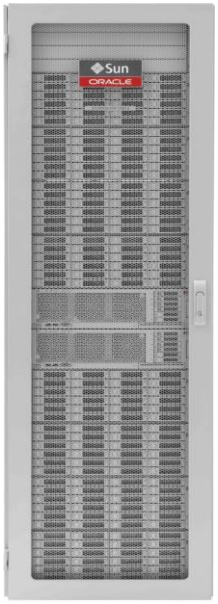

Josh Krischer
November 2016

Long Term Storage Issues 3
Mainframe Virtual Tape Systems (general) 4
Oracle's StorageTek Virtual Storage Manager 7 (VSM 7) 5
VSM 7 Hardware 6
Oracle's StorageTek Virtual Library Extension (VLE) System 7
Physical Tape Support 8
Availability 8
Disaster Recovery Infrastructures 8
VSM 7 Cloud Connectivity 10
Economies 10
Mainframe Market 11
Competitive Landscape 11
Conclusions and recommendations 12

Long Term Storage Issues

The average IT organization's data storage requirement is growing at approximately 40 percent annually. The majority of this is unstructured data. One of the main concerns of IT executives is how to cope with the annual storage growth in economically challenging times. Regulatory Compliance causes organizations to keep more data for longer periods of time; however, in most cases storage budgets have not grown (and in many cases remained unchanged or even shrank). For quite some time now, a routine part of litigation has been e-discovery, the search for, disclosure of, and dispute over email and other electronic records. Inability to present such evidence may lead to defeat in court or heavy penalties[1]. Compliance regulations such as the revised (2006) U.S. Federal Rules of Civil Procedure (FRCP)[2], provide new reasons to build a solid and intelligent archiving infrastructure. Similar regulations exist in other jurisdictions as well, e.g., the German code of commerce (HGB, §146.5.2,3) requires that business-related documents be available and quickly accessible throughout the entire legally-mandated period. Keeping information for longer periods of time may be difficult and in the mainframe market this may have a heavy financial impact.

The only sound solution to close the gap between the rising requirements and the cost of mainframe data storage is to build more efficient storage infrastructures such as tiered storage with a tape tier and thinking outside the box to leverage cloud economics.

No other IT products have been declared dead as vehemently or as often as mainframes and magnetic tape (mainly by suppliers that do not sell mainframes or tape storage). However, both of them refuse to vanish, and, in fact, current requirements are the best advocates for their usage. Mainframes are holding the world economy together. Storing data on tape has several unique benefits: It is more economical than storing on disk (even in comparison to capacity 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. In fact tape storage is experiencing resurgence in the last few years due to much higher capacities, reliability and performance in addition to the characteristics mentioned above.

---

[1] In June 2005 the 15th Judicial Circuit Court in Palm Beach County awarded Ronald Perelman $604.3 million in compensatory damages and $850 million in punitive damages from Morgan Stanley. The judge ruled that Morgan Stanley failed to produce e-mail documents, claiming they had been destroyed, only to find them later. Morgan Stanley had rejected a $20 million settlement offer from Mr. Perelman during the early stages of the case. The Florida State Appeals Court revised this verdict in March 2007. The June 2005 verdict created enormous world-wide media exposure, whereas the final verdict has been ignored mostly; the litigation costs and costs resulting from loss of reputation are unknown.

[2] In Dec. 2006 the Federal Rules of Civil Procedure were revised to require organizations to preserve and discover any type of electronically stored information (ESI) on any media and clarify the obligations of parties to turn over emails and other electronically stored records in the event of litigation.

The most effective way to cope with the imbalance between storage requirement growth and available budgets is to step out of the box and to consider cloud deployment as part of the total solution. The capacity, performance, and native tiering to cloud storage benefits of the VSM 7 should draw the attention of every IBM z Systems mainframe user. Existing VSM 5 and VSM 6 users may consider upgrading to VSM 7 (while retaining investment protection in their existing physical tape assets). New customers who have no existing investment in virtual tape systems (VTLs) and are looking to lower costs and decrease their environmental footprint should evaluate the cloud economics benefits of such solution.

Mainframe Virtual Tape Systems (general)

Mainframe virtual tape systems, also known as virtual tape libraries (VTLs) or Virtual Tape Storage Subsystems (VTSS) have been popular in the mainframe domain for 20 years. Simply put, a virtual tape system is a disk/tape storage subsystem that emulates tape storage in a mainframe environment. Usually it 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 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. The virtual tape system presents disk storage to the host as if it were tape drives and/or tape libraries. This allows high performance disk resources to work with existing mainframe tape operations software including backup and other tape-oriented batch processes. There is no need to make changes to the mainframe tape operations software or procedures.

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 its leading tape libraries and high-end storage system. The advantages of virtual tape systems led to their adoption in most large mainframe sites. What are the reasons for this popularity? The answer is the remarkable user benefits derived from virtual tape systems, for example:

- Fewer physical tape drive requirements.
- Fewer slots in the physical tape library (smaller library).
- Lower energy requirements and CO^2 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 storage systems.

"Oracle's new Virtual Storage Manager 7, with native support for cloud storage, brings virtual tape systems into the 21st century, seamlessly managing mainframe data between on-premises and the cloud."

The 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 virtual tape system's position as an effective tiered-storage, optimization solution. 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 in comparison to VTLs. Oracle's new Virtual Storage Manager 7 with native support for cloud storage, brings virtual tape systems into the 21st century, seamlessly managing mainframe data between on-premises and the cloud. VSM 7 provides the best of both worlds, allowing the building of an effective tiered storage infrastructure at an affordable price.

Oracle's StorageTek Virtual Storage Manager 7 (VSM 7)

VSM 7 is the latest model of the field-proven Oracle virtual tape system for mainframes. It ensures transparent, seamless optimization of storage resources using automated policies and supports multiple Oracle tape libraries and tape drives on the backend. It is a high availability system utilizing state-of-the-art SPARC 7 processors, advanced fault tolerance, and disaster recovery techniques.
In addition to enhanced performance and larger capacity, the new VSM 7 supports native connectivity to the Oracle Storage Cloud Service (Object and Archive Storage), a new capability that significantly enhances flexibility and reduces costs (see Figure 1).

Oracle's StorageTek VSM supports emulated tape connectivity to IBM z/OS hosts, attachment to Real Tape Drives (RTDs), and attachment to other VSMs and StorageTek Virtual Library Extension systems (VLEs). Up to 256 VSM 7 systems can be combined and managed together (using automated load balancing), satisfying organizations of every size, including the world's largest, most demanding, mainframe environments. To ensure investment protection the additional VSMs can be older models. If only more capacity is required, a VLE(s) can be added. VLEs can be shared by more than one VSM. The VSM 7 supports connectivity over FICON to IBM z/OS hosts and also FICON attachment to Real Tape Drives (RTDs) and IP attachment to other VSMs and VLEs.

"Two major enhancements of the VSM 7 from the previous VSM generation are open systems connectivity and native cloud connectivity."

Two major enhancements of the VSM 7 from the previous VSM generation are open systems connectivity and native cloud connectivity. Cloud connectivity enables mainframe users to leverage Oracle's enterprise class public cloud infrastructure for near-infinite access to storage capacity.

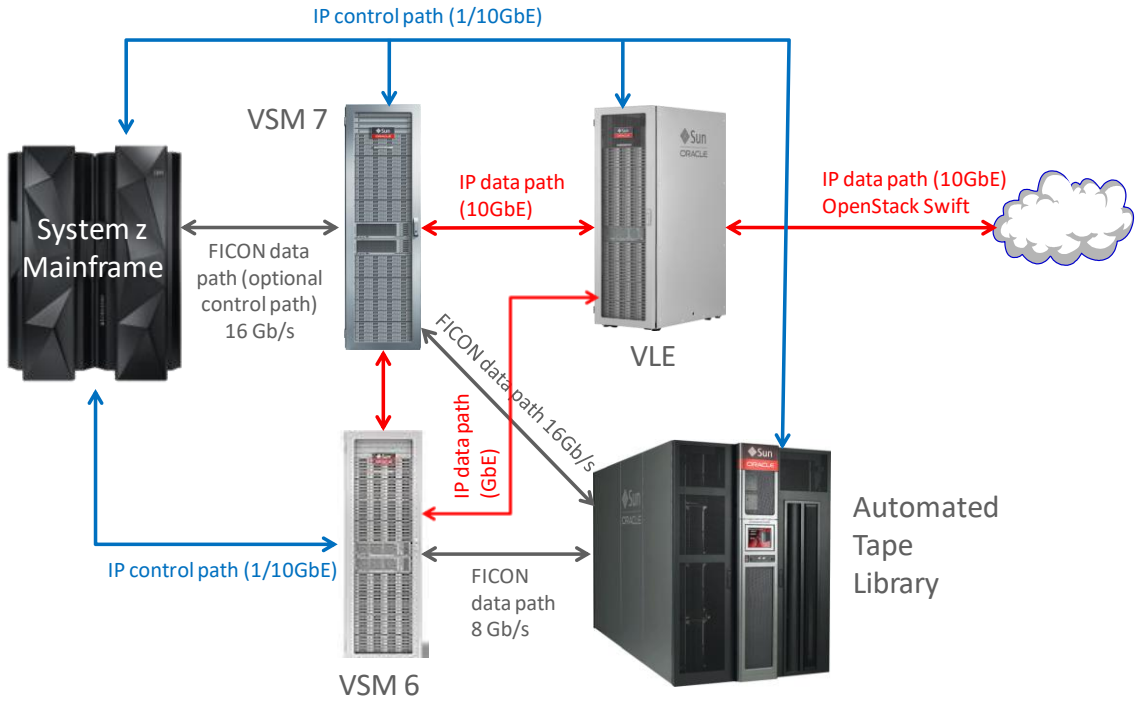

Figure 1: VSM connectivity options.

VSM 7 Hardware

VSM 7 is based on active-active dual node clusters of SPARC M7[3] processors. Oracle's SPARC M7 chip is a 4.1 GHz 32-core/256-thread processor that addresses the most demanding workloads with a balanced, high-performance design across memory, IO, and scalability. The processor modules include Flash disks, which act as non-volatile cache to increase performance and to protect "writes" which were not yet stored on magnetic media. A small file can be fully located on the Flash disk with Tier 0 performance.

In addition to the two control nodes, up to eight storage trays are available. These are populated with SAS-3, 8 TB HDDs protected by triple parity RAID. The base configuration includes two HDD enclosures with a total of 150 TB of useable capacity. Additional drive trays are added in pairs. Each pair adds an additional 225 TB of usable capacity. A fully loaded VSM 7 includes the base configuration plus six additional trays (added in pairs) for a total useable capacity of 825 TB. It is important to note that users can activate capacity in 2.5 TB steps using the Capacity-on-Demand (CoD) concept. Also, up to 256 VSM 7 systems can be logically grouped together for a total capacity of 211 PB. When the Oracle Cloud is enabled as a tier, capacity is nearly infinite.

In addition to capacity and performance enhancements, VSM 7 improves front and back-end connectivity. 16 Gb/s FICON and FC channels double the bandwidth capabilities from the previous 8 Gb/s connections and 10GbE IP enhance the control and disaster recovery performance. Please see table 1 for the major differences between VSM 6 and VSM 7.

[3] SPARC M7 microprocessor, announced October 2015.

	VSM 6	VSM 7
Peak Write @4:1	2000 MB/s	4000 MB/s+
FC/FICON Ports	8 Gb/s (FICON Only)	16 Gb/s (FC or FICON)
Replication Between VSMs	IP	IP
Redundancy	Dual Node	Dual Node
Data Integrity	Triple Parity	Triple Parity
IP Connectivity	1 GbE IP ports	10 GbE IP ports
Capacity – Effective, Per Frame	2.5 TB – .4 PB*	2.5 TB – .8 PB*
Maximum Capacity	102 PB*	211 PB*
Native Cloud Tiering	Yes	Yes

Table 1 Comparison of VSM 7 and VSM 6. *native useable capacity

The enhancements in comparison to VSM 5 are enormous; hundreds of times more effective capacity and more than four times faster data rate.

Oracle's StorageTek Virtual Library Extension (VLE) System

The VLE System can add capacity to installed VSMs. Each VLE appears as a single tape library to the VSM, so the VSM can migrate/recall virtual tape volumes just as with a physical tape library. Multiple VSMs can share VLEs; a single VLE can connect up to eight VSMs, and each VSM can connect to up to four VLEs to increase effective capacity. All generations of VSM systems can co-exist and share a single system. For better availability and disaster protection, the VSM can store the same volumes on two VLEs simultaneously. In fact, up to four copies of any one volume can be stored across multiple tiers of storage (VSM, VLE, physical tape, cloud storage).

Flexible policy management can move data from VSM to the VLE only, or to the VLE and physical tape. 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 VSM and VLE (e.g., backups whose recovery times are critical), while less performance-sensitive data such as archived data can be directed to the VSM and physical tape library or to the Oracle Storage Cloud Service—Object Storage. As a specific data set stored on a VLE becomes less active over time, its re-use is less likely, and it can be migrated to tape or to the Oracle Storage Cloud Service—Archive Storage.

Physical Tape Support

VSM 7 supports Oracle's StorageTek tape libraries; the high-end Oracle's StorageTek SL8500 Modular Library System and the mid-range SL3000. The SL8500 supports from1,450 cartridges (total 12.3 PB) up to 100,880 cartridges (total 857.5 PB). The SL3000 has from 200 to 5,925 customer-usable slots.

Availability

VSM 7 provides full redundancy across the architecture, with various data correction and data integrity techniques included. The active-active dual node processor cluster protects against isolated controller outages. In addition, the cluster allows non-disruptive microcode updates, a critical high availability feature for when firmware upgrades are required. The storage disks are protected by triple parity RAID, which allows recovery from one, two or even three simultaneously failing HDDs. Cyclic Redundancy Checking (CRC) added to the data ensures data integrity to and from the host(s). And built-in HDD encryption provides an additional layer of data security and protection.

Disaster Recovery Infrastructures

One of VSM 7's strengths is the myriad of disaster recovery options available. VSM 7 offers multiple DR configurations including one-to-one VSM *Cross* T*apePlex* replication and one-to-many multisite configurations. DR options can include multiple tiers of secondary storage including the Virtual Library Extension (VLE) and/or StorageTek tape storage, and now with the introduction of VSM 7, cloud-based DR leveraging the Oracle Storage Cloud Service. Different utilities such as *Physical Vaulting, Concurrent DR Test Utility* and *Recovery Utility* ease the management and testing of the different DR solutions. See Figure 2 for a simple example of Cross TapePlex replication. More sophisticated DR scenarios might include many-to-many site DR configurations with synchronous or asynchronous replication. For example, three VSMs can be connected in a triangle configuration with bi-directional data transfers among all three[4]. Another option is the cascading configuration, in which each of the three is connected to an additional VSM leg.

[4] Recently a three-site solution has been gaining popularity, particularly within US financial institutions and telcos. 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.

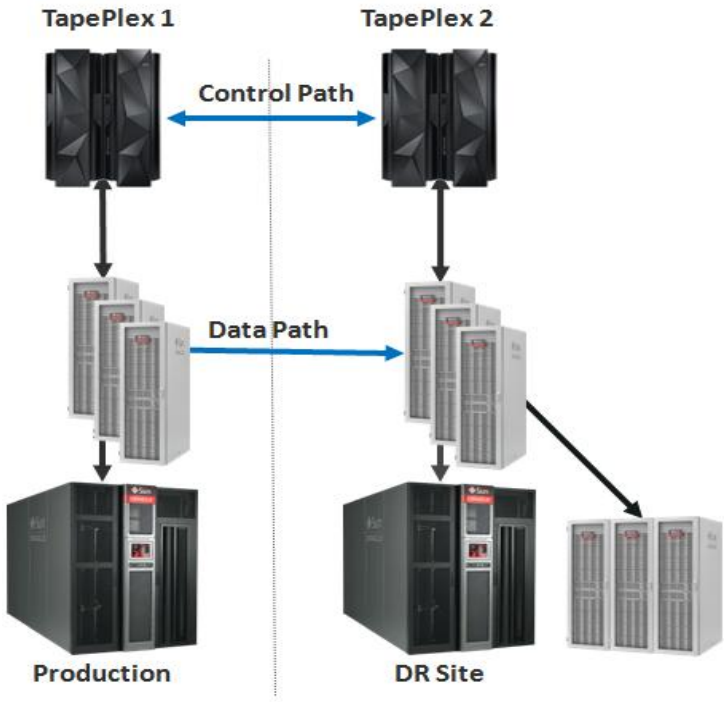

Figure 2: Cross TapePlex replication

In this configuration, the host writes once and the virtual tape solution makes all the copies. Multiple VSMs per site operate as a single environment. Such configuration simplifies DR testing.

A new option for Oracle VSM 7 users is ability to use the Oracle Public Cloud service as an external vault. It is implemented via the Virtual Library Extension (VLE) and utilizes OpenStack Swift (object storage) API to talk to storage cloud. See Figure 3.

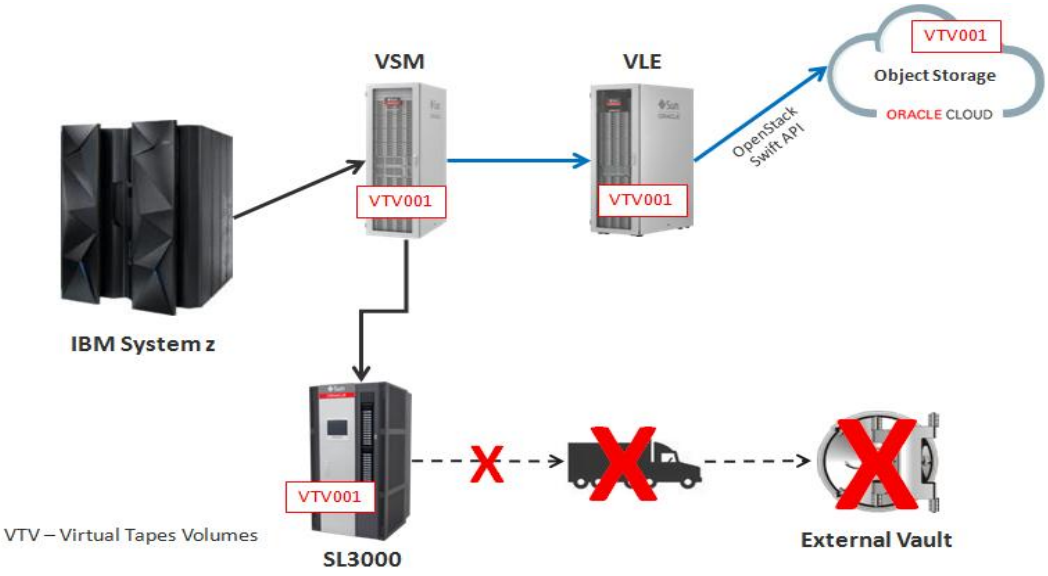

Figure 3: Oracle Public Cloud instead of external vaulting

VSM 7 Cloud Connectivity

According to International Data Corporation's *Worldwide Quarterly Cloud IT Infrastructure Tracker*, total spending on IT infrastructure products for deployment in cloud environments will increase by 16.2% in 2016 to $37.4 billion. In comparison, spending on enterprise IT infrastructure deployed in traditional, non-cloud, environments will decline 1.8% in 2016, but it will still account for the largest share (63.1%) of end user spending. For the long-term forecast, IDC expects that spending on IT infrastructure for cloud environments will grow at a five-year compound annual growth rate (CAGR) of 13.6% to $60.8 billion in 2020, or 49.7% of the total spending on enterprise IT infrastructure, while spending on non-cloud IT will decline at 1.8% CAGR during the same period[5].

VSM 7 with native cloud tiering enables users to participate in this industry trend. VSM 7 supports automatic and native connectivity to the Oracle Storage Cloud Service, as well as the Oracle Storage Cloud Service—Archive Storage which is comparable to Amazon's Glacier storage service, but at a 7Xlower price point.

Economies

VSM 7 with attached physical tape and/or connected to the Oracle Storage Cloud Service introduces new opportunities for mainframe storage at extremely attractive price points. Additional efficiencies can be realized via:

- The uses of capacity-optimized HDDs in VSM and VLE which are better priced than performance-optimized HDD solutions often found in other high-end storage subsystems for mainframes.
- The system includes embedded compression with ratios of up to 4:1
- Optional data deduplication can be enabled, but the experience shows low deduplication factors for mainframe data.
- Highly economical physical tapes are supported as a secondary tier of storage. Physical tape adds cost savings in terms of at a much lower price per GByte and significantly less energy consumption than with disk-only solutions
- Support for the Oracle Storage Cloud Service and Oracle Storage Cloud Service—Archive Storage with prices as low as $0.001/GB/month—the lowest in the industry

[5] IDC Press Release: Expected Recovery in Hyperscale Service Provider Spending Will Lead to Another Year of Strong Growth for Cloud IT Infrastructure, According to IDC; October 5, 2016

Mainframe Market

Many industry analysts have declared the mainframe as "dead." However, half a century after the mainframe's birth, as much as 80 percent of the world's corporate data resides on mainframes and 71 percent of the Fortune 100 use it. This is incredible if you consider that mainframes absorb only 6% of the total worldwide IT spend. IBM's zEnterprise is the most robust, most secure, and, at the same time, the most versatile platform with the most advanced virtualization features. The platform is highly automated, requiring little management, which translates to lower staff costs. These characteristics are reflected in sales; in 2015 – 35 percent more revenue, 36 percent more MIPS and 50 net new mainframe customers.

The February 16th, 2016 IBM announcement introduced new z13s with more computing power and boost for cryptographic coprocessors among other security enhancements. IBM targeted it for mid-size organizations as the new entry point for the company's z System.

Consequently, mainframes are still incredibly relevant for IT environments that require support for multiple, high volume, real-time transactional applications, especially where extreme resiliency and security are required.

Competitive Landscape

As opposed to the "crowded" non-mainframe storage market, only three companies deliver products that compete with VSM: EMC, IBM and Fujitsu.

Dell/EMC

The **Dell/EMC D**isk **L**ibrary for **m**ainframe (DLm) is a virtual tape storage system available in two models; the DLm8100 for large enterprise data centers and the DLm2100 for smaller data centers. DLm connects directly to IBM mainframes via Virtual Tape Engines (VTE) using FICON channels, and it appears to the mainframe operating system as standard IBM tape drives. The basic configuration includes two VTEs and is upgradeable to eight. It should be noted that the VTE controllers are not clustered and configurations with multiple VTEs provide only limited high availability.

The DLm controller is stateless and does not retain any information about the status of the tape volumes. DLm supports FICON only (no FC or native cloud connectivity). As opposed to the triple RAID protection of the VSM 7, the DLm supports RAID 6 only and will tolerate two HDD failures per RAID group; whereas VSM will tolerate up three HDD failures in a RAID group. DLm controllers support Dell/EMC VNX, VMAX and Data Domain disk storage systems, but provide only limited (one tape drive) support for physical tape on the backend and no integrated connectivity to public cloud storage. Replication for data protection techniques depends on the type of storage being used. Therefore, SRDF if VMAX is employed, VNX Replicator for VNX or DD Replicator for Data Domain. Licensing for the replication software can add significant cost to the overall solution. High availability and disaster recovery are impeded due to lack of clustering, lack of support for physical tape on the backend, and lack of support for public cloud storage. It is

also interesting to note that Dell/EMC doesn't include the DLm in its strategic products list which raises concern about its future as Dell/EMC integrates the two companies and their product portfolios.

IBM

The **IBM TS7700** family of mainframe tape storage virtualization systems is offered in three models: IBM TS7720 high-capacity cache using 3 TB SAS disk drives with RAID 6, the IBM TS7740 using 600 GB SAS drives with RAID 6 and the IBM TS7760 using 4 TB SAS disk drives with RAID 6. The TS7760/TS7740 supports backend physical tape including IBM TS1150 and earlier IBM 3592 model tape drives installed in IBM TS3500 and TS4500 tape libraries. Physical tape support is limited and optional on the TS7720. Only FICON channels are supported. It is upgradable from a single node up to a six node (IBM POWER processors) cluster. In contrast VSM 7 offers fully clustered, active/active controllers, up to 256 VSMs grid delivering up to 211PB of storage capacity, and near-infinite capacity in the Oracle Public Cloud.

TS7700 initial configuration performance is 100 MB/s only. However, it can be upgraded in 24 steps up to 2.5 GB/s, with each additional 100 MB/s of performance adding significant additional cost. VSM 7 delivers up to 4.0 GB/s of aggregate performance without additional performance upgrade charges.

Fujitsu Technology Solutions CS8000

The **Fujitsu Technology Solutions CS8000** which was developed in Europe is based on tightly connected clusters of x86 processors, up to ten processors on the front-end and similar numbers on the back-end. It supports most major computing platforms (not just mainframes), and provides support for most backend physical tape libraries. Its market share is limited mostly to EMEA with some presence in Japan, as well. It has almost no measurable market share in the Americas.

Conclusions and recommendations

With its acquisition of Sun Microsystems, Oracle inherited Sun's StorageTek division, which can look back at more than 40 years of experience in tape, disk, and automated tape library technologies. VSM 7 is its seventh generation of virtual tape storage for mainframes. VSM 7 is an extremely robust solution for mainframe data storage and the recent inclusion of native support for the Oracle Storage Cloud Service and Oracle Storage Cloud Service—Archive Storage adds significant new degrees of flexibility and cost savings.

Some additional benefits of deploying VSM 7 in the data center:

- Single, common management for multiple tiers of storage including VSM cache, VLE expansion storage, physical tape and Oracle Storage Cloud Services[6]
- Highest scalability across performance and capacity within a single system image

[6] The four tiers of storage–cache buffer (Flash cache as tier 0), VSM and VLE (HDD drives as tier 2), and tape (T10000 tapes as tier 3) and Cloud as tier 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.

- Native connectivity to the Oracle Storage Cloud Service and Oracle Storage Cloud Service—Archive Storage
- Fibre Channel and FICON connectivity
- Clustered, active/active controllers for application availability
- High performance to meet shorter batch-processing windows requirements
- Advanced and flexible disaster recovery and business continuity options.
- Policy-driven management
- Investment protection: support for connection to previous generations of VSM and VLE.
- Lower costs energy requirements than disk-only tiered storage infrastructures, thus, reducing CapEx and OpEx costs, and decreasing CO^2 emissions
- Ability to export physical tapes for vaulting

The only way to cope with the imbalance between storage requirement growth and available budgets is to think differently investigating different and new technologies. New customers who have no virtual storage systems and are looking to lower costs and decrease their environmental footprint should evaluate the benefits of a VSM 7 vs. data center-bound alternatives. Existing VSM users may consider upgrading from VSM 5 or VSM 6 to VSM 7 and choose among four tiers of storage in one system, including native tiering to the Oracle Public Cloud.

Josh Krischer is an expert IT advisor with over 47 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 mainframes, 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