

Oracle Integrates Virtual Tape Storage with Public Cloud Economics

By John Webster

June, 2016



Evaluator Group

Enabling you to make the best technology decisions



Executive Summary

Enterprise data center administrators continue to exploit tape for its economic information storage value. However, automated tape storage systems are traditionally bound to the enterprise data center. If tape-based information leaves the data center for deep archive purposes, it is usually by tape cartridges packed in a box and shipped to a secure off-site facility. If any of this information needs to be retrieved, it has to be found manually and shipped back to the data center. With the emergence of Big Data analytics applications that need historical data for additional context and accuracy, this practice becomes increasingly untenable.

Oracle recently announced Oracle's Virtual Storage Manager (VSM) 7 System with the ability to transfer data natively to Oracle's Archive Cloud. As a result, the enterprise storage administrators can use the Oracle Public Cloud as an unlimited near-line capacity extension to VSM 7 at the price of tape-based storage. Archived data can be retrieved from the cloud for emerging Big Data and IoT applications. Indeed, VSM 7 with an Oracle cloud "back end" can be seen as an immediately available repository for the huge amounts of data that an IoT application could generate, whereas disk may not be economically viable for this application.

This announcement also addresses mainframe storage administrators. Tape operations are integral to their operational practices. They can now leverage the Oracle Archive Cloud for mainframe storage to greatly improve economic efficiency while maintaining current management practices.

In this Evaluator Group Technical Insight, we explore the potential value in creating a new tape storage model for the integration of cloud-based archival storage resources with data center-based virtual tape systems. We review Oracle's StorageTek Virtual Storage Manager 7 (VSM 7) System announcement, including new features and performance increases, value gained from its native extension to cloud-based archival storage and the expected use cases.

A Rational Approach to the Coming Enterprise Storage Environment

The default information retention policy in most enterprise storage environments is "save everything forever." Many storage administrators in these environments believe that doing so with disk as the only storage option is not economically sustainable—even with aggressive data compression and deduplication. Therefore, tape persists and as a result is enjoying something of a renaissance. It has displayed remarkable resilience as an enterprise data center storage media largely because of its economies of scale. And, as tape persists, so do advancements in tape technology.

At the same time, public cloud storage has emerged as a viable save-everything-forever repository. Cloud Services Providers (CSPs) are addressing the security and regulatory compliance issues that gave

enterprise storage administrators pause. From publically available survey data, we see increasing use of the cloud for archival storage, while at the same time we also see growth in the amount of data stored on tape. Clearly, enterprises see value in both options and this creates an opportunity for a forward thinking storage supplier to create an integrated solution that brings together all of the alternatives—disk, tape, and cloud.

An example of how this new opportunity could be exploited can be seen in the new applications that generate huge amounts of data and can benefit from—and can perhaps be enabled by the use of cost-reduced data archives. Consider the Internet of Things (IoT) applications now coming on-line. A system of sensors monitoring a manufacturing plant generates data continuously. But, a monitoring application may only need to know when the data shows an anomaly that requires immediate corrective action. For the rest of the time, the data produced in this context is useless and need not be stored.

However, plant administrators may want to go back in time and analyze the whole data set to see if it shows what caused the anomaly and if the same patterns are occurring elsewhere. For this kind of analysis, saving everything on disk could well be cost-prohibitive. Tape is an alternative in these situations where preserving large volumes of infrequently accessed but potentially valuable data is required, and tape can do so cost efficiently at scale. Tape solutions also consume an infinitesimally small amount power on a per TB basis when data is off-line. When integrated with disk for immediate performance and cloud storage for long-term near line archive, tape can be an enabler of IoT applications where users want to store and preserve huge amounts of data for both present and future analysis.

A more rational approach is needed for the new enterprise data storage environment that includes support for both the day-to-day critical business functions and envisions the new data intensive applications. The announcement and general availability of Oracle's StorageTek Virtual Storage Manager Generation 7 (VSM 7) System provides that approach.

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

VSM 7 represents the next iteration of Oracle's converged mainframe and open systems virtual tape platform. As such, it advances two primary Oracle initiatives going forward in 2016:

1. Expanding the reach of Oracle's public cloud services

VSM 7 integrates Oracle public cloud storage services as an unlimited capacity extension to its integrated SSD, disk, and tape storage systems. Two tiers are available in the Oracle public cloud: an object based performance tier for data that is likely to be recalled, and a deep archive tier for very infrequently accessed data and data that is only saved for compliance or historical analysis but is rarely if ever accessed again. Oracle cloud storage is fully ISO and IEC compliant. Cloud integration is discussed in more detail below.

2. Inclusion of the VSM virtual tape storage platform into the family of Oracle engineered systems

As with Oracle's Zero Data Loss Recovery Appliance¹, VSM 7 has undergone optimization by Oracle engineers for better backup and recovery performance to meet stringent RPO and RTO SLAs, reduced data loss exposure, increased data integrity, and greater application user satisfaction.

However, VSM 7 does not represent a data island with respect to previous versions. It can co-exist with all previous generations of VSM by supporting full data interchange.

VSM 7 Features and Functions:

Significant features of VSM 7 include:

- Integrated storage continuum—disk, tape, storage cloud, and archive cloud—with automated data copy or migration across all four tiers
- Unified management of the entire storage continuum across all four tiers, including the cloud tiers. Management was enhanced with a powerful new GUI built using the Oracle software stack. Additionally, up to 256 VSM 7s can be combined and managed together.
- Up to 825 TB native, usable disk storage in a single rack;

VSM 7 Announcement Highlights

- Native integration with Oracle Public Cloud storage
- Four storage tiers with disk; tape; storage cloud; archive cloud
- Automated storage policy and simplified management for the entire environment (onsite and cloud)
- Performance and capacity enhancements
- Diverse set of enterprise disaster recovery and business continuance options

¹ See Evaluator Group Technical Insight "Oracle's Engineered Systems Approach to Maximizing Data Protection."

scalable to 211 PB in 256 frames managed and operated as a single system

- Increased performance based on implementation of Oracle's new SPARC M7 processors
- Redundant, dual node clustered servers
- Triple parity RAID for the SAS-3 disk storage tier
- Fully integrated flash SSD as non-volatile cache
- Data deduplication
- Active/active, redundant processing nodes with high-speed interconnect
- Multiple levels of data integrity checking
- Hot swappable components
- Extended support for heterogeneous environments that include open systems in addition to mainframe

Capacity

VSM 7 offers twice the disk storage capacity over the previous generation VSM 6 with up to 0.8PBs of native, usable capacity for each on premise frame. It is scalable to up to 256 frames for over 200 PB native disk capacity, managed and operated as a single system, which can be added in 2.5TB increments.

Connectivity

VSM 7 supports both 16Gb/s FC and FICON ports, as well as 10GbE IP ports.

Performance

Each VSM 7 frame is capable of transferring data at a rate of 4000MB+/Sec, and up to 1TB/Sec at maximum scale.

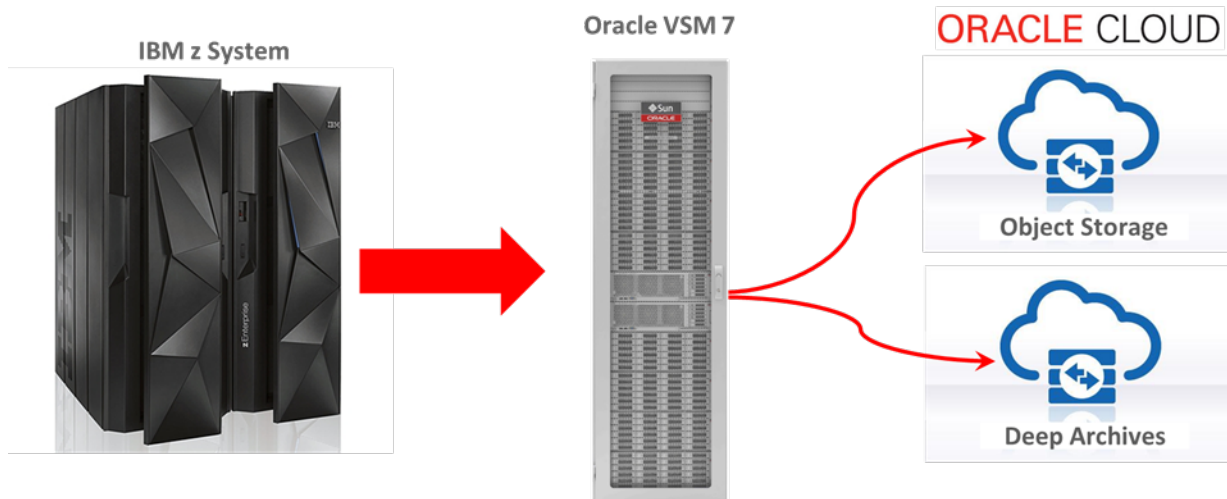
Data Replication Between and Among Systems

Each VSM 7 system has its own synchronous and asynchronous data replication capabilities across separate systems via 10GbE connectivity. Three-way mirroring is supported.

VSM 7-Oracle Cloud Integration

While Oracle's StorageTek VSM 7 represents the implementation of continuing advancements in virtualization technology, the major new capability now available is full storage cloud integration. VSM 7 enables enterprise storage administrators to natively integrate both mainframe and open systems virtual tape storage resources with Oracle's storage cloud. In fact, every aspect of an on-premises VSM 7 can have an analog in the cloud. In contrast, competing tape virtualization systems continue to provide only on-premises hardware upgrades without a choice for cloud economics.

For VSM 7, data stored in the Oracle storage cloud resides in either object-based or deep archive storage as shown below.



(source: Oracle)

Figure 1. VSM 7 allows IBM z Systems mainframes and open systems data to be stored in the Oracle Storage Cloud.

Data placement in the Oracle Storage Cloud (Object and Deep Archive) is automated based on access patterns and/or data set sizes. In keeping with Oracle optimization, VSM 7 enables Oracle database administrators to customize their own cloud storage solutions.

Oracle Cloud Storage integration provides the flexibility for an enterprise storage administrator to immediately add an essentially unlimited amount of capacity as needs require. Data transfers are accomplished directly between VSM 7 systems at the customer site and the Oracle cloud site. For zOS users, no mainframe is required at the Oracle Cloud site. Use cases include:

Cloud Bursting

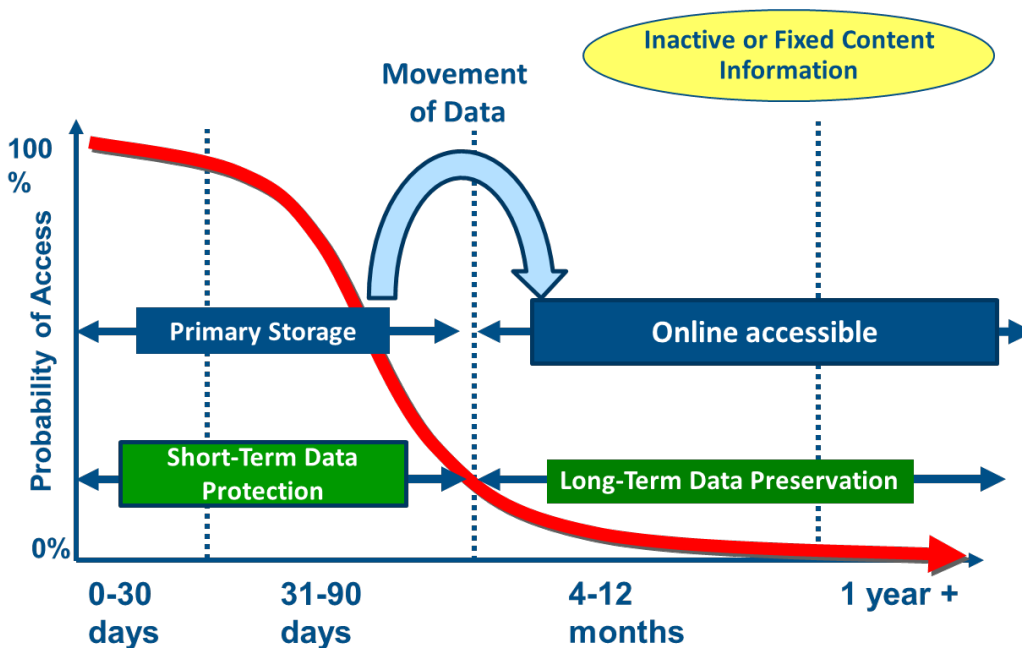
Adding cloud-based capacity on demand can be particularly useful during variable and peak load conditions—a process often referred to as “cloud bursting.”

Disaster Recovery and Business Continuance

Bi-directional data movement is supported between customer on-premises VSM 7s and the Oracle storage cloud for an additional level of risk mitigation as well as three-way mirroring. The cloud can be used as an additional layer of protection against instances of data corruption or loss as well as partial or complete data center loss or destruction. Regular disaster recovery testing is also supported.

Storage Asset Optimization

Using policy-based automation to move mainframe and opens systems data bi-directionally between the production data center and an off-site cloud can deliver significant capacity efficiency benefits. This equates to a quicker return on the investment in mainframe and open systems storage resources. As shown in Figure 2 below, short-term protection data not required for immediate work—backups for example—can be moved to a more economical repository such as the cloud.



(source: Evaluator Group)

Figure 2. Movement of data from primary to secondary data tiers based on probability of access to maximize the efficiency of primary and backup storage platforms. Movement of both primary (blue) and backup data (green) is shown.

Moving both primary and short term backup data (green box above) from a higher to a lower cost tier on the basis of probability of access saves yearly costs related to storage hardware, software licensing fees, maintenance and support, machine room environments, and administration. These data sets can also be preserved to meet long-term archival requirements. In the case of VSM 7, the archival storage tier lives in the Oracle Storage Cloud at a price that can be as low as \$12K per year per PB.

Replace Off-site Physical Tape Storage for Compliance

VSM 7's native archive cloud tiering eliminates the need for off-site data retention for regulatory compliance that requires physical tape transportation, auditing and other associated administrative costs. Enterprise storage administrators can finally stop putting tapes in a box and boxes on a truck—a process that assures retention but creates what is essentially a data dumpster. Off-siting tapes is also nearly useless for today's Big Data analytics and IoT applications.

Evaluator Group Assessment:

Enterprise executives level are increasingly showing an understanding of the intrinsic value of data ownership. What we have yet to see is an equal commitment to protection and preservation of this valuable resource. Risk of data loss is ever present and can't be outsourced to anyone. The good news is that, a strategy that minimizes the risk by combining on-and off-site data protection and preservation options is now economically viable at large scale. In addition, integration of these two options allows data to flow bi-directionally so that, not only is it protected in multiple ways, it is continually available to applications and users.

Oracle's StorageTek VSM 7 System not only integrates data center storage with Oracle cloud storage, it does so for both open systems and mainframe data. As we have shown, the use cases for this solution include cloud bursting to immediately overcome capacity shortages; data protection and availability for disaster recovery and business continuance requirements; and the elimination of off-site physical tape storage for long-term data retention. Beyond those, we also believe that VSM 7 can enable enterprise storage architects to create a cost efficient mainframe and open systems storage environment that encompasses data protection and data preservation requirements.

The need for a rational approach to an enterprise storage architecture that recognizes both traditional production applications and new data intensive applications is real. In contrast to competitive offerings that are confined to the data center, Oracle's VSM 7 offers an integrated multi-tiered solution for both mainframe and open systems data that provides near-infinite capacity expansion in the Oracle Public Cloud, enabling customers to take advantage of cloud economics.

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. www.evaluatorgroup.com Follow us on Twitter @evaluator_group*

Copyright 2016 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.