Dragon Slayer Consulting

Marc Staimer

W H I T E P A P E R

Why IBM z System Mainframe Cloud Storage Is No Longer An Oxymoron

Why IBM z System Mainframe Cloud Storage Is No Longer An Oxymoron

Marc Staimer, President & CDS of Dragon Slayer Consulting

Table of Contents

Introduction

IBM mainframe computing has been with us now for more than five decades. Mainframes (now called IBM z System) started the IT revolution. Many IT personnel, especially those from the millennial generation, today view IBM z System mainframes as dinosaurs, big ungainly obsolete technology. Like most misinformed prejudices, they're wrong. For unlike the dinosaurs, the mainframe has continuously evolved staying relevant through industry transitions from minicomputers, to PCs to the client/server architecture and now to the era of cloud computing, mobile applications, and hyper-convergence. In fact, the IBM z System mainframe was the original virtualized hyper-converged infrastructure. Today it supports the latest enterprise dev ops environments including Java, C++, JSON, AJAX, etc. While the IBM z System mainframe's continued resilience in the marketplace may be surprising to some, it is definitely not to those organizations that rely on it. Here are just a few examples:

- 92 of the world's largest 100 banks;
- Top 10 insurance companies;
- 18 of the 25 largest retailers.
- Most of the world's major airlines.

These IT organizations and thousands more like them, rely on IBM z System mainframes for sound, cogent reasons. Transactional performance is one. The customer information control system (CICS) handles more than 1.1 million transactions per second worldwide. That's more than 95 billion transactions per day. To put that in perspective, Google searches average approximately 60,000 per second. Facebook likes average approximately 30,000 per second. Consider that a single IBM z System mainframe CICS can handle roughly as many transactions—up to 2.5 billion/day—as all of the Facebook servers combined. Other reasons include production proven, end-to-end security, reliability, availability, and relatively low operational costs.

But like all technologies, the IBM z System mainframes are not perfect. A fundamental IBM z System mainframe problem has been, and continues to be, the high cost and complexity of its storage. That storage cost consistently averages more than 2 to 3 times that of Unix/Linux/Windows storage costs. Some of the reasons behind this high cost are the IBM z System mainframe unique protocols such as extended count key data (ECKD) instead of the more common SCSI (small computer system interface) protocol; and FICON channel protocol. ECKD and FICON require significantly more vendor expertise and because the market size is considerably smaller there are fewer experts. These all contribute to make costs higher. It also complicates IT resilience planning—business continuity and disaster recovery—making it far more expensive.

The IBM z System mainframe has no file storage (NAS) or object storage interface at the time of this writing. That has made it incompatible with fast-growing, low-cost cloud storage. It also cannot take advantage of very low-cost, cloud-based warm, cool or cold storage for data that needs to be saved but is rarely accessed. This means the IBM z System mainframe users must pay a huge premium for primary as well as secondary storage and cannot leverage cloud storage economics.

This white paper takes a deeper look and examines how Oracle is uniquely solving these difficult problems.

The IBM z System Mainframe Storage Dilemma

Having already established the rationale behind why many users still utilize the IBM z System mainframe, why then do so many pundits, analysts, and users believe it is too costly? The fundamental reason is the cost of mainframe storage.

❍ Economic Reality

Storage for the IBM z System mainframe utilizes proprietary ECKD storage protocol and FICON channel protocol. No other server vendor uses these protocols. That limits the market and demand for storage systems that support these protocols. And supporting them is non-trivial, requiring significant development and support costs. There are very few vendors willing to take this on. This limited supply, coupled with limited demand, plus additional cost to develop, support, and maintain equals higher storage costs. In fact the storage that supports IBM z System mainframes is very expensive with costs that typically exceed that other types of enterprise-class storage by as much as 5x.

This cost issue has historically been a minor or mid-level deterrent to IBM z System mainframe users. It was an issue but not an overwhelming one. Technology inevitably advances and changes. Storage has arguably changed the most radically with advances in storage media—flash SSDs, high density flash SSDs optimized for reads, high density drawers of storage media—highly-scalable object and file storage, and of course cloud storage. These newer technologies are rapidly reducing the cost of data storage consumption with a lower cost per TB. This is particularly true for passive warm, cool and cold data. Cloud storage is especially changing the economics of secondary storage like never before. Unfortunately, the IBM z System mainframe does not participate with these far more cost effective technologies. It has no object or file storage interface. There is no IBM z System mainframe RESTful API support. This has resulted in no non-proprietary cloud storage options all of which are quite costly. The lack of cloud storage options has left IBM z System mainframe users at a serious cost and competitive disadvantage.

❍ Most Common Workarounds

The way most IBM z System mainframe users have gotten around some of the high storage cost limitations has been to utilize mainframe tape/automated tape libraries (ATL). The tape/ATL workaround has further evolved to include mainframe virtual tape libraries.

Mainframe Tape/Automated Tape Libraries

Mainframe tape drives, tape cartridges, and ATLs are much lower cost than mainframe HDD or SSD storage systems. The characteristics of tape make it well suited for sequential writes and reads. Tape is not very good at the random dynamic variety. This is why non-mainframe servers utilize tape strictly for backup/recovery and archive. IBM mainframes past and present are a different beast. They too have used tape in the traditional manner. However, since mainframes were developed in an era when disk drives were massive as well as massively expensive with relatively low capacities, they were architected to uniquely utilize tape as a form of primary albeit slower storage. They still do. The fact that magnetic tape cartridges are a removable media with increasing reliability and availability makes tape storage a relatively cost effective methodology for getting data offsite. The usual methodology for moving large amounts of data offsite is via a courier service. That methodology tends to be more cost effective than attempting to move large amounts of data electronically over wide area networks.

Keep in mind that mainframe-class tape drives are generally installed in robotic ATLs that can hold thousands of tape cartridges. This can provide storage densities well into the Petabytes (PB) or even Exabytes (EB) of data. Take the example of the Oracle StorageTek SL8500 Modular Library System or ATL in the more common vernacular. The Oracle SL8500 scales from a relatively pedestrian 1,450 tape cartridges for 12.3 PB of uncompressed (30.75 PB compressed) storage capacity to 100,880 cartridges for 857.5 PB of uncompressed (2.144 EB compressed) storage capacity in a relatively small footprint.

As previously stated, mainframe tape/ATLs are only fast when it comes to sequentially writing or reading large amounts of data. Conducting backup and restore operations,

archiving data, providing disaster recovery services and processing other tape-based batch operations, are relatively slow on tape drives or ATLs, as compared to primary disk storage. That performance gap led to the emergence of mainframe virtual tape libraries.

Mainframe Virtual Tape Libraries

Mainframe virtual tape libraries (MVTL) optimize physical tape operations solving several of these issues. They do so by behaving exactly like actual mainframe compatible automated tape libraries while performing at much higher speeds. The reason MVTLs are so much faster than tape or tape libraries is that they are in reality writing and reading to/from HDDs instead of tape cartridges. This is a boon to IBM z System mainframe users. MVTLs are providing HDD performance at much lower cost than IBM z System mainframe storage systems direct access storage devices—DASD or mainframe disk. Thus eliminating some of their high storage burden.

MVTLs also have costly limitations. Most do not scale beyond a few PBs per system reducing their value as secondary storage. Scaling beyond those limitations requires more systems, more supporting infrastructure, more data center real estate consumed, more management, more power, more cooling, more maintenance, more troubleshooting, more capital expenditures, more operating expenditures, and inevitably a lot more headaches.

Mainframe Virtual Tape

Noun | main-frame \ vir-tu-al \ tape | ˈmān-ˈfrām \ ˈvər-chə-wəl \ ˈtāp

1. A disk storage subsystem that emulates tape storage in a mainframe environment
   – Presents disk storage to the host as if it were tape drives and/or libraries
   – Allows high performance disk resources to work with existing mainframe tape backup software, processes and policies

When it comes to security, those MVTLs that actually offer encryption do so either in software or via encrypted HDDs. Encrypting in software noticeably reduces MVTL performance. Encryption is processor intensive. Encrypting HDDs encrypts the data only when it's at-rest. There is no encryption in-flight.

Cloud storage integration is a non-starter to most MVTLs. Only one non-Oracle MVTL has capability of utilizing standard cloud storage in any manner. The one that does integrate with cloud storage does so in a halfhearted or minimalist manner. It does not automatically move data to and from the cloud storage. It unfortunately severely constrains cloud storage capacity to no more than 2x the capacity of its active data tier (less than 2PB). And, it's not a standard feature costing several thousands of dollars more. Worse yet, retrieving that data from the cloud is not automated requiring a manually labor-intensive process.

Data protection and high availability are yet other MVTL limitations. Most only protect data with RAID-6 dual parity. With the notable exception of Oracle and one other MVTL vendor, all other MVTLs have single points of failure such as a single controller. The prevailing vendor perception is that the data is secondary data making data protection a low priority. Regrettably, that perception fails to take into account that once again, IBM z System mainframes utilize tape for more than backups and archive. Even if that were the case, recoveries are never a low priority. The worst time to discover that data cannot be recovered is during a recovery. When data protection is a priority, MVTL data protection is typically quite expensive. With few exceptions, it requires duplicating the equipment onsite and/or at another location, or making extensive use of ATLs. The tapes can be shipped offsite as required.

Neither ATLs nor MVTLs have done enough to enable the IBM z System mainframe to leverage the convenience, geographic distribution, security and low cost of cloud storage. Oracle changes all of that with their StorageTek Virtual Storage Manager System 7 (VSM 7).

How Oracle StorageTek VSM 7 Completely Changes IBM z System Mainframe Cloud Paradigm

The Oracle VSM 7 is architected specifically to address and solve each and every one of these IBM z Systems mainframe issues. It is engineered from the ground up to deliver the best possible IBM z System

mainframe to MVTL performance utilizing up to 4 different storage tiers which includes native hybrid cloud storage integration. The Oracle VSM 7 delivers unmatched security, simplicity/automation, availability/reliability, and flexibility with that native hybrid cloud storage integration. It does all this while concurrently speeding up time-to-value when compared to any other MVTL or ATL.

Oracle VSM 7 Engineered For IBM z System Mainframes

Performance

It starts with performance and for any mainframe operator, security. The Oracle VSM 7 is based on the latest high performance Oracle SPARC M7 processor. The SPARC M7 is up to 4 times faster than other processors. Each Oracle VSM 7 comes with dual node active/active processing nodes, which allows for failover/failback as well as non-disruptive code load capability. They are more than capable of handling the most demanding IBM z System mainframe workloads for tape transactions and operations.

Security

More importantly, each Oracle VSM 7 comes with native encryption. That encryption delivers wide-key encryption for data-at-rest in the VSM 7 as well as for removable tape media. The VSM 7 uniquely provides this encryption without degrading system performance.

Simplicity/Automation

Intuitive, policy-driven automation from a single control point in each VSM 7 simplifies deployment, operations, and management. That simplification extends across tiered implementations that might include VSM 7 backended by, the Virtual Library Extension (VLE), physical tape or cloud storage or any combination of thereof.

Automation is essential to today's under-staffed flat budget data centers. VSM 7 reduces IT professional system expertise and training requirements and accelerates time to value from deployed systems. That's why VSM 7 is highly automated. It automates data deduplication and compression, as well as the copying of critical business information and transparently migrating it to tape or cloud storage tiers.

Availability/Reliability

Every StorageTek VSM 7 system has been deliberately designed with the same levels of reliability, availability, and serviceability that users expect from IBM z System mainframes. The main concern of IBM z System mainframe administrators and users is continuous data access. To meet and exceed those expectations on availability and serviceability Oracle provided VSM 7 with active/active redundant processing nodes with very high speed interconnect; fully redundant disk communication paths; automatic failover of critical system components; hot-swap componentry; triple parity RAID (protects against data loss with up to 3 concurrent disk drive failures); and unique "on-the-fly" dynamic job recovery upon hardware failure.

Availability is critically important to every IT administrator; however, data reliability is even more so. It doesn't do anyone much good if the data is available but is error prone. Oracle ensures the data stored on VSM 7 is error-free by employing multiple levels of data integrity checking. Every virtual tape block has a read/write cyclical redundancy check (CRC) both before and after every compress/decompress cycle. A minimum of one CRC is always protecting the virtual tape blocks making sure the data reliability matches its availability.

Flexibility

Oracle recognizes that workload performance and data protection requirements vary. Few organizations have the resources, time, patience, or budgets to implement unique storage infrastructure for every workload. VSM 7 is a highly flexible storage system. It can be used with both IBM z System mainframes as well as non-mainframe servers such as Linux, OpenStack, Windows, VMware, etc. From an efficiency perspective, it can deliver disk only, disk plus tape, or disk plus storage cloud performance and any combination thereof. From a DR perspective, VSM 7 offers one-to-one and one-to-many multi-site configurations as well as combinations with Oracle StorageTek modular tape libraries or Oracle Storage Cloud Services. And from a capacity perspective, it scales from hundreds of TBs to PBs of on-premise capacity and nearly unlimited capacity when connected to the Oracle Storage Cloud Service or Oracle Storage Archive Service.

Scalability

Oracle VSM 7 completely changes the game on MVTL scalability. For conventional on-premises scalability, it scales HDD capacity from 825TB raw in a single rack to as much as 211PB raw in a single VSM 7 tapeplex. The effective usable capacities are significantly greater with built-in automated deduplication and compression. VSM 7 also seamlessly integrates with Oracle tape drives and automated tape libraries treating it as a storage tier and delivering EBs of on-premises capacity. Where VSM 7 truly changes the MVTL scalability game is with its deep integration with Oracle Cloud Storage Service and Oracle Cloud Storage Service–Archive Storage.

Engineered To Be Hybrid Cloud Storage

Oracle Storage Cloud Services are engineered for enterprise-class workloads including the IBM z System mainframe. Oracle Storage Cloud Services are secure, elastic, and highly durable with aggressive, competitive pricing. This enables the Oracle Storage Cloud Services to be highly cost-effective for off-site data protection and archiving. It considerably reduces large capital expenditures that are necessary for acquiring and maintaining on-premises secondary storage hardware. An example would be a customer that currently has two active/active data centers each backing the other up. The customer has always wanted to have a third copy of their data but could not justify. With the VSM 7 architecture they could have VSM 7s at both active sites and send a third copy cost effectively out to the cloud. This can be accomplished without modifications to mainframe applications or JCL.

The Oracle StorageTek VSM 7 treats Oracle Cloud Storage Services as a nearly unlimited capacity storage tier. A customer will never outgrow it. Data appears to the IBM z System mainframe as local even though it's in the cloud. Visibility is end-to-end across the entire data path with diagnostics from on-premises to the public cloud. Security is also end-to-end.

Accelerated Time-To-Value

Combining VSM 7's unmatched scalability with tight tape integration, plus unique hybrid cloud storage tiering, significantly decreases secondary storage costs of the IBM z System mainframe. Additional cost savings come from reduced storage real estate footprint, supporting infrastructure, power, cooling, management, operations, complexity, and even tech refresh. The real game changer is the co-engineered automated tiering with Oracle Storage Cloud Services. That hybrid cloud architecture converts upfront capital expenditures into more dynamic elastic operating ones. This is a true pay as needed model for mainframe storage.

Summary and Conclusion

The IBM z System mainframe world has been stuck in a late 20th century storage paradigm. It left IBM z System mainframe users on the outside looking in as their x86 server technology peers enjoyed the fruits and benefits of the cloud revolution. They could only watch that revolution in envy—until now. Oracle StorageTek VSM 7 has brought the revolution to IBM z System mainframe users. That's why IBM z System mainframe cloud storage is no longer an Oxymoron.

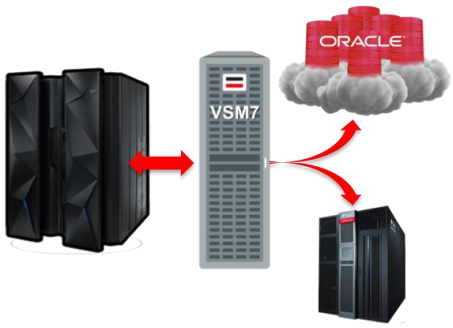

For More Information

- **Oracle StorageTek VSM 7**

Visit www.oracle.com/goto/VSM or contact your local Oracle office.

- **Oracle Storage Cloud Service & Oracle Storage Cloud Service – Archive Storage**

Visit https://cloud.oracle.com/en_US/storage or contact your local Oracle office.

Paper sponsored by Oracle. *About the author:* Marc Staimer, as President of the 18-year-old Dragon Slayer Consulting in Beaverton, OR, is well known for his in-depth and keen understanding of user problems, especially with storage, networking, applications, and virtualization. Marc has published thousands of technology articles and tips from the user perspective for internationally renowned online trades including Searchxxx.com and Network Computing. Marc has additionally delivered hundreds of white papers, webinars, and seminars to many well-known industry giants such as: Brocade, Cisco, DELL, EMC, Emulex (Avago), HDS, HP, LSI (Avago), Mellanox, NEC, NetApp, Oracle, QLogic, SanDisk; as well as smaller, less well-known vendors/startups including: Asigra, Caringo, Cleversafe, Cloudian, Clustrix, Condusiv, DH2i, Diablo, FalconStor, Gridstore, Nexenta, Neuxpower, NetEx, NoviFlow, Permabit, Qumulo, Scality, Tegile, and many more. His speaking engagements are always well attended, often standing room only, because of the pragmatic, immediately useful information provided. Marc can be reached at marcstaimer@me.com, (503)-579-3763, in Beaverton OR, 97007.