ORACLE
STORAGETEK

# Improving Mainframe Disaster Recovery with StorageTek VSM 7 and Brocade

ORACLE WHITE PAPER | DECEMBER 2016

ORACLE

## Table of Contents

## Introduction

Complexity and challenges are growing for enterprise data centers trying to meet business continuity, disaster recovery, batch, primary data, and archival storage requirements. Without careful planning, each of these requirements could entail using its own unique storage solution, creating more complexity. As a result, versatility is quickly becoming the new mandate for meeting many of these enterprise class (mainframe) data center storage requirements. A combination of storage virtualization, disk, tape, robotic libraries, and integration within the network have been integrated into an easy-to-use, cost-effective solution to address these issues. This paper discusses Oracle's StorageTek Virtual Storage Manager System 7 (VSM 7) and the synergies/benefits of using it in conjunction with Brocade's latest extension products to create a high performance, highly secure, highly versatile storage solution.

## Business Continuity and Disaster Recovery Define Higher Levels of Protection and Availability

Business continuity and disaster recovery are, in most cases, the two most pressing IT issues. Fail-proof data backup and recovery is more critical to an organization's survival than ever before, because few businesses can survive for any period of time without their IT function. 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 security breaches. With a primary focus on backup architectures and protecting data, it's important to understand the various roles of backup. "Backup" can be defined as simply a data protection process for executing two primary data center functions:

- *Business continuity (BC):* Provides a local copy of data to be used if an application or infrastructure component fails or data becomes corrupted. For business continuity, fast initial access time is critical. Therefore, disk is the preferred choice. An organization needs to consider how it handles disk hardware failures, application failures, and network failures. Tape could be used as well in conjunction with disk in order to ensure business continuity, enabling fast backup and restore of small chunks of data such as system and user files, emails, and incremental backups.
- *Disaster recovery (DR):* Provides a copy of data that can be maintained off-site and restored from another location, if the primary data center facility no longer is available. For DR, fast data transfer time and high availability are critical. Therefore, tape has become the optimal DR choice, enabling fast backup and restore of large amounts of data (large databases, servers, or even entire data centers).

It's important to distinguish between a business continuity event and a disaster recovery event because they are different. 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. These activities can include file and database backups, project management, change control, and help desk services. 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 (QoS), consistency, and recoverability, that is, activities that keep the business running. Disaster recovery consists of the process, policies, and procedures that relate to preparing for a recovery after a natural or human-induced disaster occurs that may impact the entire data center functionality rather than impacting 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 critical business functions.

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

Oracle's StorageTek team working with Brocade helps organizations consider, select, and deploy networked storage solutions for current and planned needs. Providing validated designs which are repeatable reference architectures that have been engineered and tested ensures organizations have the highest levels of service.

## IT Security is a CIO's #1 Concern

Organizations and the general public have experienced large numbers of online security breaches in the last few years, and such attacks continue to increase in size and severity. The high-profile data breaches of companies prove that no one is safe from cyberattacks as organized activities to access data become more sophisticated.

These violations not only jeopardize personal customer data, severely damaging consumer trust, but they also result in hundreds of millions of dollars in lost revenue for the companies.

Managing security risks requires an effective business resilience and data protection plan covering continuity of operations, backup and recovery of critical information, disaster recovery, and routine testing.

## Would Your Business Recover if Data Records Were Breached, Corrupted, or Destroyed?

Data loss costs companies billions of dollars per year worldwide, while unplanned downtime costs the average company nearly US$1 million per year. It is easy to see that data protection and DR are more important today than they have ever been. Yet dissatisfaction with data protection options is at an all-time high. Today's data protection architectures are overly complex, expensive, and outdated, causing businesses to spend too much time and money on backup and recovery. Many architectures focus solely on short term business continuity, typically using replication, snapshots, and deduplication to scale-out their backup environments. This leaves businesses inadequately prepared for online attacks and major disasters. Additionally, these systems treat all data the same, creating multiple backups of the same data rather than identifying and moving original data to an archival storage system. Storing many copies of data in backup files has major implications when it comes to storage costs and efficiency, especially considering much of the stored data is seldom, if ever, accessed after 90 days. Sprawling backup and recovery systems also lack modern data protection features, such as seamless cloud integration. Leveraging the cloud as part of a backup solution can improve data protection while lowering costs.

## Implementing a Modern Data Protection Strategy

What are the best data protection options available? In its most basic form, "backup" can be defined as simply a tool or method for executing business continuity processes (with disk) or disaster recovery processes (with tape). 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 improvements are now available. These offer a variety of choices depending on what operating system is used, as well as the type of storage technologies and network connectivity that can be used. This includes when and how the data is protected; whether compression, encryption, or write once, read many (WORM) are used; and whether any 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 and steady amount of digital data growth. Flash and spinning disk have become the preferred business continuity backup targets for smaller data files demanding the fastest Recovery Time Objective (RTO), while tape is the optimal backup choice for large files and disaster recovery processes. An all-disk data protection and archiving solution is an increasingly expensive option. Remember, data that is not being used should not consume energy. To best address business continuity, disaster recovery, and archival data, the optimal solution is available and deploys a tiered storage approach using disk, tape, and even cloud storage.

## 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 and recovery purposes. A VTL presents a storage component (using hard disk arrays) as tape libraries or tape drives for easy use with existing backup software. A

VTL enables the storage hardware to be switched from tapes to disks while continuing to use the existing tape backup software and processes. The benefits of this virtualization include storage consolidation, faster data restore times, and better RTOs. Note that for large data sets and DR operations requiring large amounts of data to be moved, the streaming data transfer speed of tape makes it faster than disk.

## StorageTek Virtual Storage Manager System 7 Has Arrived

Oracle's StorageTek Virtual Storage Manager System 7 (VSM 7) was announced in March 2016, and is the latest generation of the mainframe virtual tape system. VSM 7 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 IBM z Systems environments. It provides a unique multitiered storage system that includes both physical disk, optional tape storage, and auto-tiering to Oracle Storage Cloud.

VSM 7 was designed for business continuity, disaster recovery, long term archival and big data requirements for z/OS applications. Each VSM 7 emulates up to 512 virtual tape transports but actually moves data to and from disk storage, back-end real tape transports attached to automated tape libraries, as well as auto-tiering to the Oracle Storage Cloud. As data on VSM 7 disks ages, it can be migrated from disk to physical tape libraries such as Oracle's StorageTek SL8500 and SL3000 Modular Library Systems, as well as being auto-tiered to the cloud in order to meet long-term data retention, archive, and compliance requirements.

The VSM 7 enables nondisruptive, on-demand capacity scaling and can grow to 825 TB of customer-usable (without compression) disk storage in a single rack. VSM 7 provides 2x the performance and more than 2x the disk capacity than the previous generation VSM 6 system and scales disk capacity on demand to 211 petabytes of native capacity (scalability up to 256 VSM systems). A simple and seamless migration to physical tape extends system capabilities to provide nearly unlimited capacity to meet an enterprise's long-term data retention and compliance needs. Protecting data throughout its lifecycle has become a reality—with a single architecture.

## StorageTek VSM System 7 Platform

StorageTek VSM System 7 is a follow-on to the previous generation VSM 6 system, replacing the VSM 6 hardware stack with new Oracle servers, storage disk enclosures, and I/O cards. This new hardware stack provides increased performance and capacity. For example, storage capacity is more than doubled from VSM 6; replication performance is improved with the move to all 10 Gb IP networks; and higher processor clock rates, faster internal memory, and faster I/O bus speeds provide noticeable improvements in system performance. As a replacement for the VSM 6, the customer view and functionality of the VSM 7 are generally equivalent, except where performance, connectivity, and serviceability have been improved. The predominate VSM 7 platform differences from the VSM 6 are as follows:

- Increased performance and storage capacity with upgraded servers and storage disk enclosures
- 16 Gb/sec FICON connectivity with eight ports
- 10 Gb/sec IP connectivity throughout the virtual tape storage subsystem (VTSS) and into the customer's network environment
- Two network switches that aggregate or fan out network connections from the servers to the customer's network environment

The VSM 7 is packaged as a standard rack mount system built on existing Oracle server, storage, and service platforms. The servers, storage disk enclosures, and standard rack mount enclosure are delivered as a packaged system. The Oracle Solaris 11 operating system is the foundation of the VSM 7 software environment, which also includes Oracle Solaris infrastructure components and VSM specific software. The VSM 7 software environment is

preinstalled and preconfigured for VSM functionality so that limited site-level configuration is required to integrate the product into the customer's managed tape environment.

## StorageTek VSM System 7 Components

Oracle's StorageTek Virtual Storage Manager System 7 is a collection of hardware and software products that comprise a disk-based virtual tape system to provide enterprise-class storage management capabilities for the IBM mainframe environment. VSM 7 optimizes streaming workloads and backup and recovery functions, reduces management overhead, and maximizes tape capacity utilization to reduce data protection costs in a wide range of storage environments. VSM 7 stores virtual tape volumes (VTVs) on a disk buffer on the VTSS and can optionally migrate them to Oracle's StorageTek Virtual Library Extension (VLE), real tape drives (RTDs), or both. VTVs can be up to 32 GB. When needed by the host, if the migrated VTVs are not VTSS-resident, they are then automatically recalled to the VTSS. VSM 7 includes the following subsystems:

- Virtual storage management hardware and software. VSM 7 supports emulated tape connectivity to IBM MVS hosts over fiber connection (FICON) interfaces as well as Fibre Channel connection to select backup applications. It allows attachment to RTDs, IP attachment to other virtual storage managers (VSM 7, VSM 6, or VSM 5) and to VLE, and remote host connectivity using Extended Control And Monitoring (ECAM) over IP and replication of one virtual storage manager to another.
- StorageTek Enterprise Library Software (ELS).
- Security for IP Extension flows.
- Virtual Library Extension (VLE) hardware and software. VLE is IP-attached to VSM 7 and functions as a migration and recall target for virtual tape volumes (VTVs). VLE supports migration and recall to and from Oracle Cloud Storage. A VSM 7 that is connected to a properly configured VLE can use the VLE to migrate and recall VTVs to and from Oracle Cloud Storage instead of local disk. Refer to VLE documentation for details about this feature.
- Real Tape Drives (RTDs) connected to physical tape libraries. RTDs serve as migration and recall targets for VTSS VTVs. RTDs are FICO and or IP attached to VSM 7.

## Enhanced Replication Feature

StorageTek VSM System 7 supports an enhanced replication link (RLINK) feature that extends the replication capabilities of VSM 7. With this feature, synchronous replication begins replicating data to the target VTSS upon first host write to the VTV and provides host acknowledgment to the rewind unload operation once all data has been successfully replicated to the target VTSS. An RLINK is composed of all IP paths defined to the target VTSS. There is only one RLINK between the primary and target VTSS. With RLINKs, the number of replications is limited only to the number of virtual tape devices (VTDs) supported within the VTSS.

**Note**: RLINK functionality cannot be used concurrently with synchronous cluster link (CLINK) replication. RLINK functionality is available for use between two StorageTek VSM System 7 or System 6 VTSSs, where each VTSS can be both a primary and a target for bidirectional synchronous VTV-level replication.

Oracle VSM 7, working in conjunction with the Brocade 7840 Extension Switch and/or Brocade SX6 Extension Blade, 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 multisite support for disaster recovery. In addition, physical tape FICON/Fibre Channel over IP (FCIP) extension can be used to extend the VSM 7 storage to span onsite and off-site repositories that are geographically dispersed. Up to 256 VSM VTSSs can be clustered together into a tapeplex, which is then managed under one pane of glass as a single large data repository using StorageTek Enterprise Library Software (ELS).

**Note**: A tapeplex is a collection of StorageTek tape hardware represented by one StorageTek ELS. Mainframe systems place the highest demands on tape, using it for a variety of applications such as primary data, batch, data backup, Oracle Hierarchical Storage Manager (HSM), archive and big data applications. VSM 7 uses automated user-defined policy-based management to simplify day-to-day operations by automating tasks that were previously handled manually and to enable IT resources to be deployed on higher value, more-strategic assignments.

## Brocade 7840 Extension Switch

Brocade 7840 Extension Switch leverages the core technology of Brocade Gen 5 Fibre Channel platforms, consistently delivering 99.999 percent uptime in the world's most-demanding data centers. It combines enterprise-class availability with innovative features and the industry's only WAN-side, nondisruptive firmware upgrades to achieve always-on business operations and maximize application uptime. These capabilities enable a high-performance and highly reliable network infrastructure for disaster recovery and data protection.

Brocade 7840 Extension Switch is an ideal platform for building a high-performance data center extension infrastructure for replication and backup solutions with the Oracle VSM 7. It leverages any type of interdata center WAN transport to extend open systems and mainframe storage applications over any distance. Without the use of extension, those distances are often impossible or impractical.

In addition, this extension switch addresses the most demanding disaster recovery requirements. Twenty-four 16 Gb/sec Fibre Channel/FICON ports, sixteen 1/10 Gigabit Ethernet (GbE) ports, and two 40 GbE ports provide the bandwidth, port density, and throughput required for maximum application performance over WAN connections. Designed for maximum flexibility, this enterprise-class extension switch offers "pay-as-you-grow" scalability with capacity-on-demand upgrades. To meet current and future requirements, organizations can quickly and cost effectively scale their WAN rate from 5 Gb/sec to 40 Gb/sec per platform via software licenses. With compression enabled, organizations can scale up to 80 Gb/sec application throughput, depending on the type of data and the characteristics of the WAN connection.

Brocade 7840 Extension Switch is a robust platform for large-scale, multisite data center environments implementing block, file, and tape data protection solutions. It is ideal for the following:

- Data protection for open systems and mainframes
- Multisite synchronous and asynchronous storage replication
- Centralized tape backup, recovery, and archiving
- Consolidation of Fibre Channel, FICON, and IP storage data flows from heterogeneous arrays

Brocade 7840 Extension Switch maximizes replication and backup throughput over long distances using data compression, disk and tape protocol acceleration, WAN-optimized TCP, and other extension networking technologies. Advanced features and technologies include the following:

- IPsec, which ensures secure transport over WAN links by encrypting data-in-flight with a hardware-implemented version of the standard 256-bit AES encryption algorithm—without a performance penalty or excessive added latency.
- An unparalleled, extremely efficient architecture, which uniquely permits the high-speed, low-latency processing of IP datagrams and Fibre Channel/FICON frames, making extension of synchronous applications possible.
- WAN-optimized TCP. The switch provides an aggressive TCP stack, optimizing TCP window size and flow control, and accelerating TCP transport for high-throughput storage applications.
- Streams. The switch supports IP Extension, a feature of WAN-optimized TCP, to prevent Head-of-Line Blocking (HoLB) across the WAN.

- Per-Priority TCP Quality of Service (PTQ), which provides high-, medium-, and low-priority handling of Fibre Channel and Brocade IP Extension flows within the same tunnel for transmission over the WAN using autonomous individual TCP sessions per QoS priority.
- Extension Trunking, an advanced feature of the switch that combines multiple WAN connections into a single, logical, high-bandwidth trunk, providing active load balancing and network resilience to protect against WAN link failures.
- Lossless Link Loss (LLL), which is a part of Extension Trunking that provides recovery of data lost in-flight when a link goes offline. From the perspective of the storage applications, nothing ever happened because all data is delivered—and delivered in order.
- Failover/failback with failover groups. Circuits are assigned metrics and put in a failover group. If all circuits of the lower metric within the failover group go offline, the higher metric circuits take over. This process uses LLL, and all data is delivered and delivered in order. The storage application will not know that a failover/failback has occurred.
- Adaptive Rate Limiting, which dynamically adjusts bandwidth sharing between minimum and maximum rate limits to optimize bandwidth utilization and maintain maximum WAN performance during disruptions.
- An advanced compression architecture, which provides multiple modes to optimize compression ratios for various throughput requirements.
- An advanced accelerator for FICON. The switch uses advanced networking technologies, data management techniques, and protocol intelligence to accelerate IBM zGM, mainframe tape read and write operations, and z/OS host connection to Teradata warehousing systems over long distances.

## Ensuring Continuous Availability Between Data Centers

Today's organizations depend on fast, reliable access to data whenever needed, regardless of location. The ramifications and potential business impacts of an unreliable DR and data protection infrastructure are greater than ever. Brocade 7840 Extension Switch provides a suite of features from predeployment validation to advanced network failure recovery technologies to ensure a continuously available storage extension infrastructure. The switch has built-in tools to validate the condition of the WAN links and network paths, and to ensure proper configurations prior to deployment. Administrators can validate and troubleshoot the physical infrastructure with the built-in Flow Generator and a WAN performance test tool called Wtool to ease deployment and avoid potential issues.

Extension Trunking, a Brocade technology originally developed for mainframes, is now broadly used in open systems environments. Extension Trunking shields end devices from IP network disruptions, making network path failures transparent to replication traffic. Multiple circuits (two or more) from Brocade 7840 Extension Switch are applied to various paths across the IP network. The most common example of this is redundant Data Center Local-Area Network (DC-LAN) switches. For instance, one circuit goes to DC-LAN Switch A, and the other circuit goes to DC-LAN Switch B. This is a simple and effective architecture for redundancy and increased availability. Of course, as needs dictate, the application of circuits over various paths and service providers (up to eight circuits per IP extension tunnel) can establish a highly available infrastructure for IP storage.

Extension Trunking supports the aggregation of multiple WAN connections with different latency or throughput characteristics (up to a 4:1 ratio), allowing organizations to use WAN circuits from multiple service providers with different physical routes, to ensure maximum availability. If all WAN circuits are from the same service provider, then chances are high that a single failure event (for example, equipment failure, power loss, or cable cut) can take down all WAN circuits at one time. With Extension Trunking, organizations can protect their replication traffic from these kinds of outage events.

Extension Trunking offers more than the ability to load balance and failover/failback data across circuits. Extension Trunking is always a lossless function, providing in-order delivery within an extension trunk (defined by a Virtual Expansion_Port, which is also called a VE_Port). Even when data in-flight is lost due to a path failure, data is

retransmitted over remaining circuits via TCP and placed back in order before it is sent to the Upper Layer Protocol (ULP). IP storage applications are never subjected to lost data or out-of-order data across the WAN. Extension Trunking protects against WAN link failures with tunnel redundancy for lossless path failover and guaranteed in-order data delivery using LLL. Extension Trunking allows multiple network paths to be used simultaneously, and when there is a failure for a network path, Extension Trunking retransmits the lost packets to maintain overall data integrity. The storage application is protected and there is no disruption.

With Adaptive Rate Limiting (ARL), organizations can optimize bandwidth utilization and maintain the full WAN performance of the link during periods when a path is offline due to the outage of an extension platform, IP network device, or array controller. ARL is used with Extension Trunking to maintain available bandwidth to storage applications. For example, if DC-LAN Switch A goes down, then as long as DC-LAN Switch B remains online and has ample bandwidth connectivity, it should be able to maintain all the original bandwidth to the application. In this case, it is necessary for rate limiting to readjust upward and compensate for the lost pathway. ARL uses dynamic bandwidth sharing between minimum (floor) and maximum (ceiling) rate limits to achieve maximum available performance during failure situations. In addition, with unprecedented amounts of storage data crossing extension connections and consuming larger, faster links, Brocade has enhanced ARL to react 10 times faster to varying traffic patterns that compete for WAN bandwidth or use shared interfaces. ARL adjusts from a normal condition that is not oversubscribed to an outage condition that maintains the same bandwidth. Clearly, this is essential to continuous availability.

## Brocade IP Extension: IP Replication with Local Performance over Long Distances

IP storage arrays with native replication applications are not built to handle latency and packet loss. Brocade 7840 Extension Switch provides a robust IP extension solution called Brocade IP Extension (IPEX), which delivers local performance over long distances along with strong encryption for comprehensive DR. It leverages Brocade TCP Acceleration to help achieve the fastest replication speeds possible from IP storage devices, and Brocade WAN-optimized TCP to ensure in-order lossless transmission of IP extension data. Brocade IPEX helps to significantly increase the performance of IP storage applications across the WAN even with encryption turned on. The more latency and packet loss between the data centers, the greater the gain. Brocade 7840 Extension Switch can move 50 times more data than native TCP/IP stacks to meet rigorous recovery objectives. Such performance gains enable use cases that at one time were deemed unfeasible.

Brocade IPEX also offers other, more-far-reaching benefits. The switch supports and manages Fibre Channel/FICON- and IP-based data flows, enabling storage administrators to consolidate I/O flows from heterogeneous devices and multiple protocols. The consolidation of these applications into a single, managed tunnel between data centers across the WAN has real operational, availability, security, and performance value. Consolidating IP storage flows, or both IP storage and Fibre Channel/FICON flows, into a single tunnel contributes significantly to operational excellence. Operational advantages are gained with Fabric Vision, MAPS (Monitoring Alerting Policy Suite), Wtool, and Brocade Network Advisor. Using custom, browser-accessible dashboards for IP storage or combined Fibre Channel and IP storage, storage administrators have a centralized management tool to monitor the health and performance of their networks.

## Acceleration of IP Extension Flows

Acceleration of flows across the WAN improves IP storage performance dramatically. Long distances increase latency and are prone to packet loss. Tested applications have demonstrated improvements of up to 50 times, due to the ability to handle latency and packet loss without performance degradation. This performance has nothing to

do with compression; any compression achievable is in addition to flow acceleration. Flow acceleration is purely a function of enhanced protocol efficiency across the network. Brocade IPEX on Brocade 7840 Extension Switch terminates IP storage TCP flows locally and transports the data across the WAN using a specialized TCP transport called WAN-Optimized TCP (WO-TCP). The primary benefit here is the local acknowledgement (ACK). By limiting ACKs to the local data center, TCP flows that originate from an end IP storage device have to be capable of merely high-speed transport within the data center. Most native IP storage TCP stacks are capable of high speeds only over short distances. Beyond the limits of the data center, "droop" becomes a significant factor. *Droop* refers to the inability of TCP to maintain the line rate across long distances. Droop worsens progressively as distance increases. WO-TCP is an aggressive TCP stack designed for the movement of big data, operating on the purpose-built hardware of Brocade 7840 Extension Switch. The switch offers 64 processors and 128 GB of RAM to support WO-TCP. In comparison, WO-TCP has no droop across two 10 Gb/sec connections, up to 160 ms round-trip time (RTT) per data processor. This is equivalent to two fully utilized 10 Gb/sec WAN connections (OC-192) between Los Angeles and Hong Kong.

Figure 1 illustrates the performance gains realized with VSM 7 and Brocade 7840 IPEX and WO-TCP. This testing was done using one GbE CLINK with 4 GB VTV size data being transferred from a VSM 6 to a VSM 7. For a series of distances/latency measurements, a "pristine" network with no packet loss was used for a base comparison, followed by two more-likely network scenarios: .1 percent packet loss and .5 percent packet loss. These results illustrate the enhanced replication performance that can be realized by VSM 7 customers when using the Brocade 7840 Extension Switch's IPEX technology for connectivity between data centers and VSM 7 devices.

| Latency ms RTT | No packet loss | | 0.1% packet loss | | 0.5% packet loss | |
|---|---|---|---|---|---|---|
| 1ms=~100 miles | VSM + IPEX | VSM Native | VSM + IPEX | VSM Native | VSM + IPEX | VSM Native |
| 0 | 110 | 112 | 110 | 83 | 106 | 82 |
| 5 | 110 | 110 | 103 | 14 | 102 | 14 |
| 10 | 104 | 109 | 101 | 7.4 | 89 | 7.5 |
| 25 | 110 | 90 | 103 | 3.1 | 82 | 3.1 |
| 50 | 99 | 51 | 92 | 1.6 | 67 | 1.6 |
| 100 | 86 | 35 | 80 | 0.8 | 60 | 0.8 |
| 200 | 75 | 23 | 73 | 0.4 | 45 | 0.4 |
| 250 | 55 | 10 | 47 | 0.3 | 38 | 0.3 |

Figure 1. Performance gains realized from StorageTek VSM System 7 and Brocade 7840 Extension Switch

## Security of IP Extension Flows

IPsec secures data end to end, to ensure that flows leaving the confines of the local or remote data center and connecting to a service provider are secure from eavesdropping and attack. The WAN infrastructure for a service provider in itself is not secure. The most common mistake is assuming that a "private" WAN connection is secure and cannot be eavesdropped on or attacked. Data encryption services such as the IPsec provision on Brocade 7840 Extension Switch prevent eavesdropping, altering, and outsider attacks of all kinds. IPsec also prevents the need for costly and complex firewalling. Because firewalls are software-based, they tend to provide poor performance.

Brocade 7840 Extension Switch with IPsec encourages the use of encrypted data protection and does so with no performance penalty. The Brocade IPsec capability is hardware-implemented and operates at line rate (20 Gb/sec) per data processor with 5 microseconds (µs) of added latency. IPsec is included in the Brocade 7840 Extension Switch base unit, with no additional licenses or fees. The IPsec capabilities are ideal for storage replication solutions that have performance degradation when turning on their in flight encryption capabilities, or for storage replication solutions that do not offer native in flight encryption.

## Summary

StorageTek VSM System 7 offers a number of compelling replication solutions to provide the highest levels of protection and availability. For business resumption and disaster recovery purposes, the versatility of VSM 7 provides the capability to make additional backup copies of primary data, keep data in multiple sites in synch, and copy data offsite to multiple geographic locations. Synchronous or asynchronous node clustering across FICON/Fibre Channel or GbE/10 GbE, cross-tapeplex replication (CTR), and RTD channel extension for remote site physical tape support all provide exceptional levels of enterprise data availability and redundancy.

The optimum solution for the future will address the needs for business continuity and disaster recovery as well as archive needs by taking advantage of disk, tape, and cloud storage technologies. Versatility in a storage solution has become the key ingredient for attaining the highest availability levels for meeting looming storage requirements. Tape, disk, and cloud storage are being used to meet most of the archive demand, and today these technologies are optimally integrated into a single architecture with VSM 7 utilizing Brocade networking capabilities.

Data protection has become the most-critical IT discipline because most businesses can no longer survive without their IT function. As a result, the next generation of data protection solutions are beginning to appear as legacy processes become increasingly burdensome, expensive, and unreliable. Today's optimum data protection solution ensures business continuity (using disk) and disaster recovery (using tape), and it provides the most cost-effective means of protecting tier 3 data (using tape and or the cloud) while providing foolproof network capabilities from Brocade.

Data protection throughout the lifecycle of data is not an option; it is a requirement. Engineered with versatility as a fundamental component, VSM 7 along with Brocade 7840 Extension Switch can be deployed to meet each enterprise's unique business requirements. From high-performance "disk only" or massively scalable "disk + tape + cloud" configurations to single-site or multisite support, VSM 7 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, cloud, and virtualized data protection solution with its VSM 7 architecture.

Brocade 7840 Extension Switch is a purpose-built extension solution that securely moves more data over long distances faster while minimizing the impact of disruptions. With Gen 5 Fibre Channel, IP extension capability, and Brocade Fabric Vision technology, this platform delivers unprecedented performance, strong security, continuous availability, and simplified management to handle the unrelenting growth of data traffic between data centers in Fibre Channel, FICON, and IP storage environments. The switch is a purposely engineered system consisting of hardware, an operating system, and a management platform, which provides capabilities that native storage IP TCP/IP stacks cannot provide.

Oracle's comprehensive approach, paired with regular system testing, allows for protection from instances of data corruption and data center loss or destruction while meeting recovery time objectives (RTO) and recovery point objectives (RPO). When businesses fail to implement comprehensive data protection strategies, disaster strikes. A simplified and modern data protection strategy with off-site and offline backups, archiving, and DR will be the difference between peace of mind and becoming the next front-page story.

# Appendix

This appendix documents Brocade 7840 Extension Switch qualification with Oracle's StorageTek VSM System 6/System 7 and tape drives at the StorageTek Mainframe Customer Emulation Test (MCET) lab in May 2016.

## Overview

Brocade 7840 Extension Switch is the latest 16 Gb channel extension solution from Brocade. The switch was successfully tested for FICON/Fibre Channel and IP interface extension interoperability with the devices and microcode described below.

## Devices and Software Levels Tested

- Host mainframe: IBM Z13 with FICON Express16 channels
- Host OS: z/OS 2.01
- Application software: IEBDG, IEBGENER, DFSMS/DSS Backup/Restore, FDR Backup/Restore, SyncSort, and DFSort
- Library/ StorageTek VSM System control software: StorageTek Enterprise Library Software 7.2
- Brocade switches: **v7.4.1b firmware**, three Brocade 7840 Extension Switches
- StorageTek VSM System 6: 6.3.0.05.000
- StorageTek VSM System 7: 7.0.0.02.000
- StorageTek 9840C: 1.47.502
- StorageTek 9840D: 1.47.702
- StorageTek T10000A: 1.52.106
- StorageTek T10000B: 1.52.206
- StorageTek T10000C: 3.62.111
- StorageTek T10000D: 4.10.106

Network latency was emulated at 10 ms, 30 ms, and 50 ms RTT. For the IP extension testing, network latency was emulated up to 200,ms. All testing was done using the FCIP GbE interfaces connected through a third Brocade 7840 Extension Switch that was used to generate network delay.

"Front-end" refers to the IBM Z13 host channel path identifier (CHPID) to StorageTek VSM System channel interface.

"Back-end" refers to the VSM channel interface card to a Real Tape Drive (RTD).

No performance measurements were taken because that would require a dedicated mainframe host, which was not available.

## Configurations Tested

The following figures show the tested configurations.

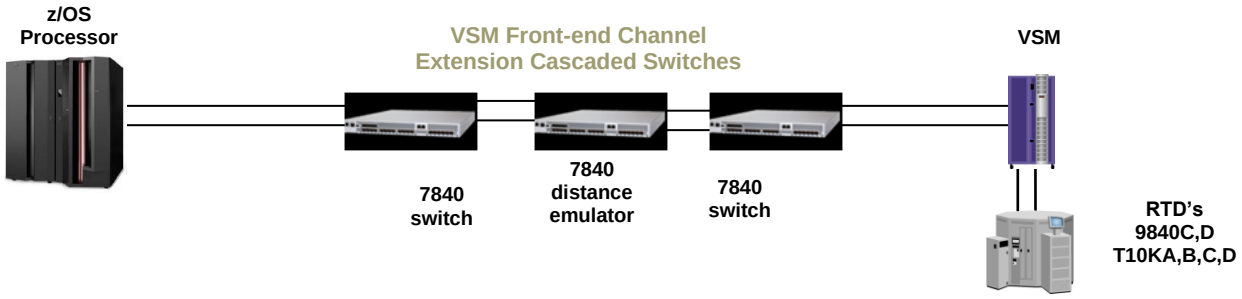


Figure 2. VSM front-end CHPID

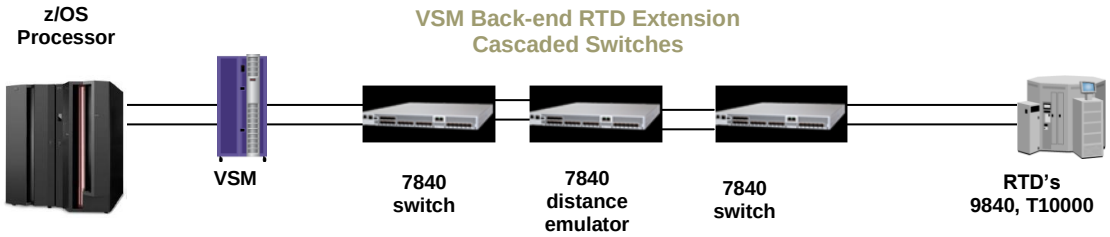


Figure 3. VSM back-end RTD

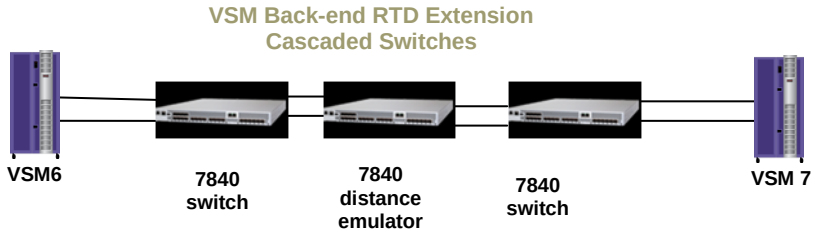


Figure 4. VSM CLINK replication

Brocade 7840 Extension Switches were tested as a method of extending direct-attached VSM systems, as well as "back-end" RTD communication of the VSM system. The GbE ports on the switches were routed through an additional Brocade 7840 Extension Switch to simulate the network latency.

Testing consisted of mainframe jobs that read and wrote to VTDs using IEBDG and IEBGENER, DFSMS/DSS Backup/Restore, FDR Backup/Restore, SyncSort, and DFSort. The total runtime for testing was 72 hours for each configuration tested.

A VSM 7 subsystem was added to the configurations for IP interface extension testing with the Brocade 7840 Extension Switches. Both the VSM 6 and VSM 7 were configured with two standard CLINKs (rep1 from both nodes on each subsystem). Both copper and optical replication ports were used on the VSM 7. Network latency was emulated at 10 ms, 30 ms, 50 ms, and 200 ms round trip time.

## Error Injection Testing

In addition, error injection testing was performed by disconnecting cables, disconnecting GbE links, and disabling switch ports to test the error recovery of the solution. No abnormalities were found. Expected results were seen after each disruption. After re-establishing connectivity, all connections recovered as expected.

Additional testing was performed with Hot Code Load and HAREBOOT on the Brocade 7840 Extension switches five times for each command on both switches in the configuration.

## Limitations of the Testing

Testing was limited to interoperability testing, meaning the testing of whether hardware and software interoperate without error under normal usage and under the errors injected as described.

What was not covered:

- **Performance** testing (whether performance would be acceptable under all circumstances). Acceptable performance is a subjective judgment and also highly dependent on a particular customer's needs.
- **Suitability** testing (whether the tested configuration was suitable for a particular customer). The VSM solution can be configured in many different ways, and because of this, the best-suited solution for a particular customer might not be any of the configurations tested.
- **Network error injection** testing (whether real-life network errors would cause problems). The MCET test lab does not have the resources or the expertise to test injection of errors into the communications network between switches. It is left to the switch vendor to do such testing.

## Issues Opened

Two issues were opened against the VSM subsystems during this test activity. Both issues were subsequently closed as a tape drive hardware/lab infrastructure issue.

- Bug 23140879 - VSM6 - VTSS66A TDX FAILURE Occurred: running recall operation
- Bug 23140829 - VSM 7 - VTSS76A TDX FAILURE Occurred: running recall operation

**Oracle Corporation, World Headquarters**
500 Oracle Parkway
Redwood Shores, CA 94065, USA

**Worldwide Inquiries**
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US

blogs.oracle.com/oracle
facebook.com/oracle
twitter.com/oracle
oracle.com

Integrated Cloud Applications & Platform Services



StorageTek Virtual Storage Manager System 7 and Brocade 7840 Extension Switch
December 2016