Violin Memory Operating System

Storage at the Speed of Memory

June 2013

Executive summary

Violin Memory Operating System (vMOS) delivers hardware-accelerated storage management optimized for Flash Memory. This white paper provides an overview of the storage management, data protection and storage optimization capabilities powered by vMOS in Violin Memory Arrays.

1 Introduction

Violin Memory Operating System (vMOS) provides an intelligent management platform for enterprises to benefit from the memory-like performance of Violin flash Memory Arrays, delivering sustained application latencies measured in microseconds. vMOS powers Violin Memory 6000 series arrays with a comprehensive suite of hardware-accelerated software capability, providing various features including:

- Flash optimization
- High Availability
- Ease of management
- Data protection
- Storage optimization

2 Flash Optimization

Violin Memory has been at the forefront of driving innovation in maintaining sustained high throughput over extended array lifetimes. With patented flash optimization technology, in-place fault handling and multi-layer wide striping, vMOS drives Violin Memory Arrays to unprecedented levels of sustained performance, endurance, and resilience. Unlike legacy storage operating systems designed for spinning disk and later retrofitted to handle SSD, vMOS is architected from the ground up to expose the full capabilities of flash memory.

All data written to Violin Memory Arrays is automatically protected by vRAID – a patented RAID algorithm, specifically designed and architected to maximize flash endurance while minimizing application I/O latency. If any flash die/block fails, the data can be seamlessly reconstructed by using the parity information earlier calculated by vRAID. Traditional RAID5 and RAID6 algorithms use a read-modify-write algorithm – while they suit the HDD storage that they were written for, these legacy RAID algorithms are highly inefficient for flash. The read-modify-write operations of legacy RAID algorithms negatively impact flash endurance and application I/O performance. In contrast, Violin's vRAID algorithm is architected to optimize the important attributes of a flash storage system, such as latency under load, bandwidth, storage efficiency and reliability.

vMOS implements wide-striping at each layer of the architecture to actively distribute incoming I/O across the entire Memory Array and ensure that no specific flash device is worn out due to repeated I/O. Besides providing effective wear leveling, wide striping removes the administrative overhead of creating LUNs of different sizes and stripes to ensure that each LUN receives the full bandwidth of the array. Patented techniques such as Erase-hiding and Write-hiding allow the

system to perform all necessary flash management operations in the background without blocking user reads or writes, providing latencies in the microseconds for any workload.

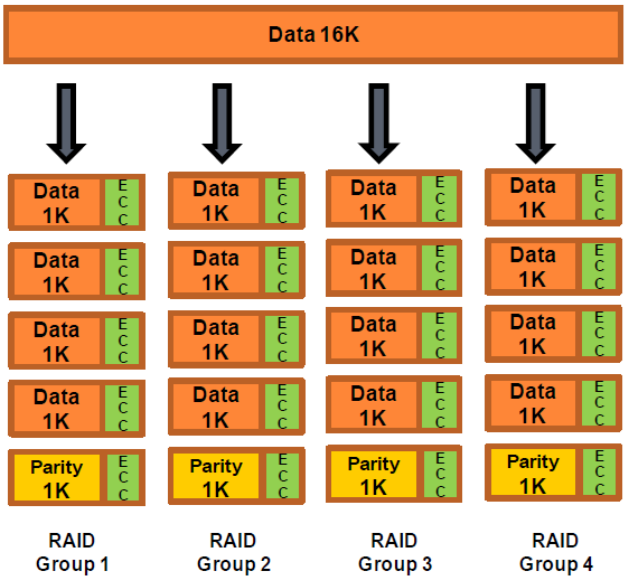

Figure 1: vRAID calculates 1K parity for each incoming 4K write

3 High Availability

Violin 6000 Series Memory Arrays are built to ensure that there is no Single Point of Failure in the architecture. In addition, vMOS adds multiple levels of redundancy to ensure 24x7 High Availability for business critical applications. All active hardware components – fans, memory gateways, power supplies, array controllers, vRAID control modules and flash memory modules - have built-in hardware-controlled redundancy and are hot-swappable. Each 6000 Series Memory Array has four spare VIMMs, which are automatically brought into action by vMOS in the event of any failure. The highly available, clustered pair of memory gateways can be used in symmetric Active/Active as well as asymmetric Active/Active ALUA modes, to meet the dual demands of performance and availability. Upon failure of one gateway, LUNs are transparently made accessible through the alternate gateway, ensuring no disruption in I/O access for the clients and applications.

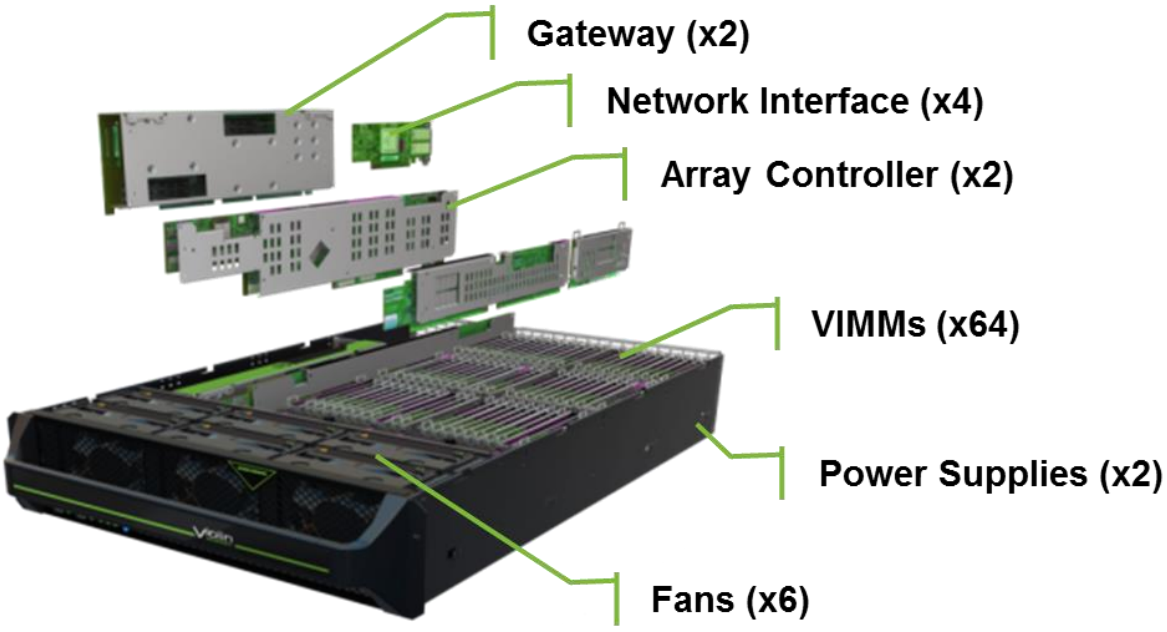

Figure 2: Violin Memory 6000 series arrays are architected with built-in hardware redundancy

vMOS implements robust ECC and CRC algorithms to detect and automatically correct bit errors, caused by flash media idiosyncracies, in the system. All data in the array is routinely read and scanned for errors. If errors are found, they are automatically repaired using ECC or vRAID without any noticeable impact on performance. The automatic self-healing increases data endurance and overall system availability. All read, write and erase statistics are continuously captured and reported – pro-active alerts are sent if any unusual behavior is detected. When the health degrades below permissible levels, a spare VIMM will be automatically brought into service to ensure that there is no I/O disruption and therefore, no compromise to high availability.

The redundancy of Violin Memory 6000 Series architecture also enables non-disruptive upgrades of vMOS, allowing the application I/O to continue without any disruption during the upgrade process. Upgrades can be triggered either through the vMOS CLI or through the WebUI as well.

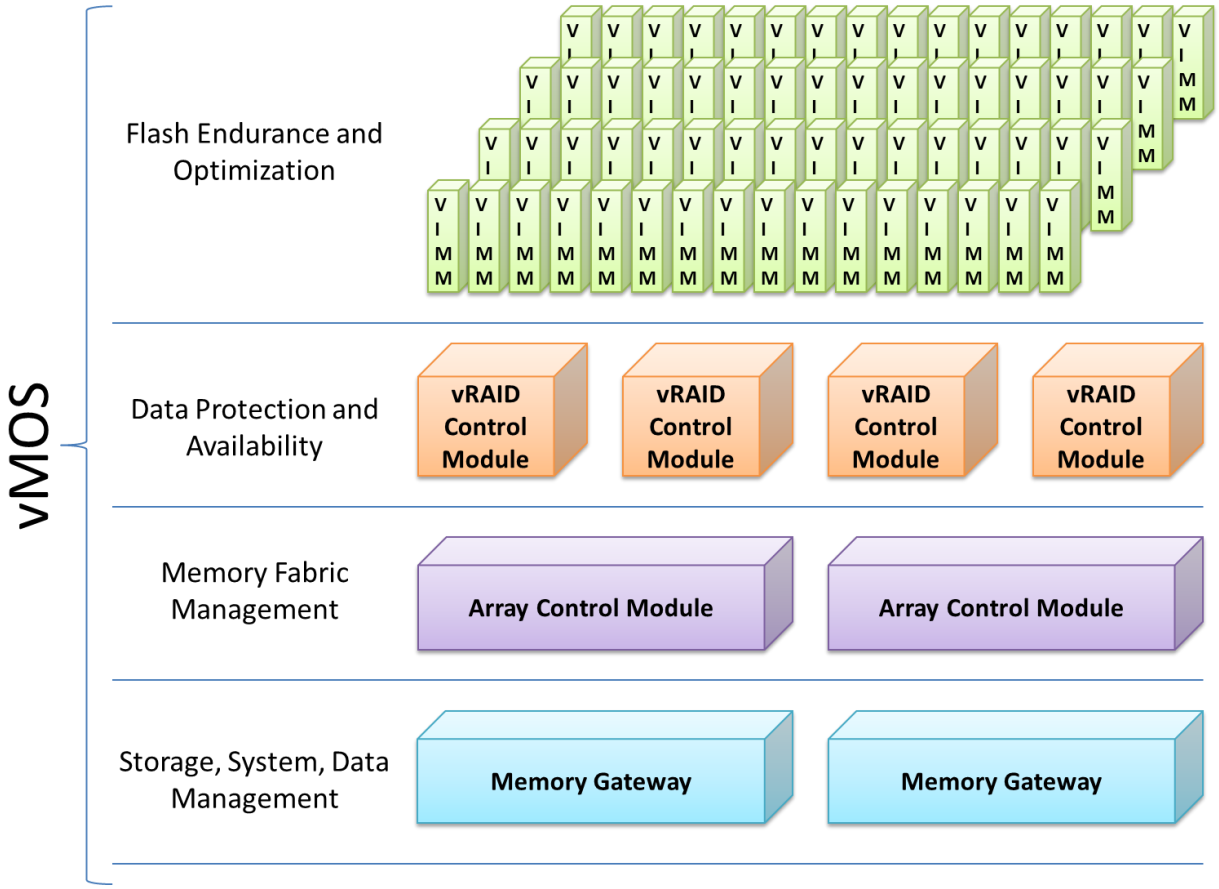

Figure 3: vMOS modular architecture distributes software functionality across all active components

4 Ease of Management

Violin flash Memory Arrays can be deployed as regular SAN-attached primary storage, without the need to re-architect or modify existing storage architecture. vMOS supports all leading forms of connectivity including Fiber Channel, 10 GbE iSCSI, 40 Gbps InfiniBand and PCIe direct attach. vMOS also enables the inter-operability of Violin Memory arrays with all major Unix, Linux and Windows operating systems.

vMOS offers several ways to perform regular storage operations on the Violin flash Memory Arrays:

- Graphic, web-based interface (WebUI)
- Web-based multi-array management system (Symphony)
- Simple, intuitive CLI
- Script-ready REST API

vMOS eliminates complicated storage operations such as setting up RAID and striping since these features are natively provided by Violin Memory arrays. vMOS supports up to 1024 LUNs

on a single Memory Array – with wide striping, vMOS ensures that every one of these LUNs gets the full bandwidth and performance of the array. Each LUN can be shared by many servers in the cluster, with full support for SCSI3 Persistent Reservations and compatibility with all leading cluster file systems.

vMOS empowers administrators to monitor, manage, and configure Violin flash Memory Arrays anywhere, any time through a simple, elegant GUI that offers a comprehensive view of the entire system, from provisioning and reporting to on-going management. vMOS provides administrators with access to all elements of the array, enabling full visibility into the system. With support for multiple browsers and an iPad application, as well as integration into 3rd party management environments like VMware vCenter, vMOS provides all the flexibility and simplicity expected of an enterprise-grade primary storage system, including several options to ease administration with live monitoring and dynamic administration capabilities.

Figure 4: vMOS WebUI is also available through smartphones and Tablets for ease of administration

With a fully customizable dashboard and a large library of gadgets to choose from, users can customize the WebUI to report the information that's most relevant to them, in the format that they are most comfortable with. WebUI provides real-time performance monitoring trends for throughput, latency and bandwidth across the entire array as well as for specific LUNs. vMOS continuously monitors the health of all active components, space utilization in the array and various performance metrics. In addition to reporting all the information through easy-to-consume graphs in the WebUI, users can also configure SNMP triggers and email alerts to be sent if any of the health criteria are not met.

Figure 5: Customizable dashboard in vMOS WebUI gives complete flexibility to administrators

The RESTful API included in vMOS provides all of the same functions as provided by the WebUI and CLI. It is a powerful scripting tool which can be used to integrate with orchestration tools for automated cloud and SAN environments such as VMware vSphere, Microsoft SCOM and OpenStack.

For large deployments, Violin Symphony provides centralized management of hundreds of arrays in a next generation web interface delivering automation of common tasks and advanced rule based analytics and monitoring. More information on Violin Symphony is provided in a later section.

5 Data Protection

Enterprises depend on Violin Memory Arrays to store and manage the most business critical data. vMOS offers enterprises a wide range of data protection capabilities to reliably protect their data, without affecting the impressive performance of Violin flash Memory Arrays:

- Encryption
- Snapshots and Clones

5.1 Encryption*

Compliance regulations, such as HIPAA, require businesses to protect all personally identifiable information, such as customer data and healthcare information, from any unauthorized access. vMOS Encryption extends the data protection capabilities of Violin Memory Arrays to provide high performance data-at-rest encryption across an entire system. vMOS Encryption works seamlessly in the data path to encrypt all writes before the data is written to flash and to decrypt the data that is being read off flash, providing:

- Built-in data-at-rest encryption for privacy protection
- Passphrase protection for reliable access authorization
- Latencies measured in microseconds for all access to encrypted data
- Always-on flash optimization for high performance and endurance

vMOS Encryption uses AES-XTS 256-bit algorithm, as outlined in the IEEE 1619 encryption standard and as required by most leading compliance regulations such as HIPAA and FIPS. vMOS uses a combination of two encryption keys to encrypt every write before it is written to Violin Intelligent Memory Modules (VIMM). The data on the VIMM cannot be decrypted without the encryption keys – this protects the data from any unintended access in the event of a VIMM reuse or theft. Administrators have the flexibility to enable passphrase protection to prevent any unauthorized access to the array.

The 6000 Series encryption-ready arrays are powered by high-performance internal Memory Gateways. The higher computing power of these gateways, coupled with all the performance capabilities of the Violin flash Memory Arrays, delivers up to 1 Million encrypted 4 KB IOPS and provides sub-millisecond latencies to applications – even with encryption enabled. Once encryption is enabled on the array, all the LUNs are enabled for encryption. When a LUN needs to be re-used or moved to a different business unit, discarding the key effectively destroys access to all the data stored in this LUN.

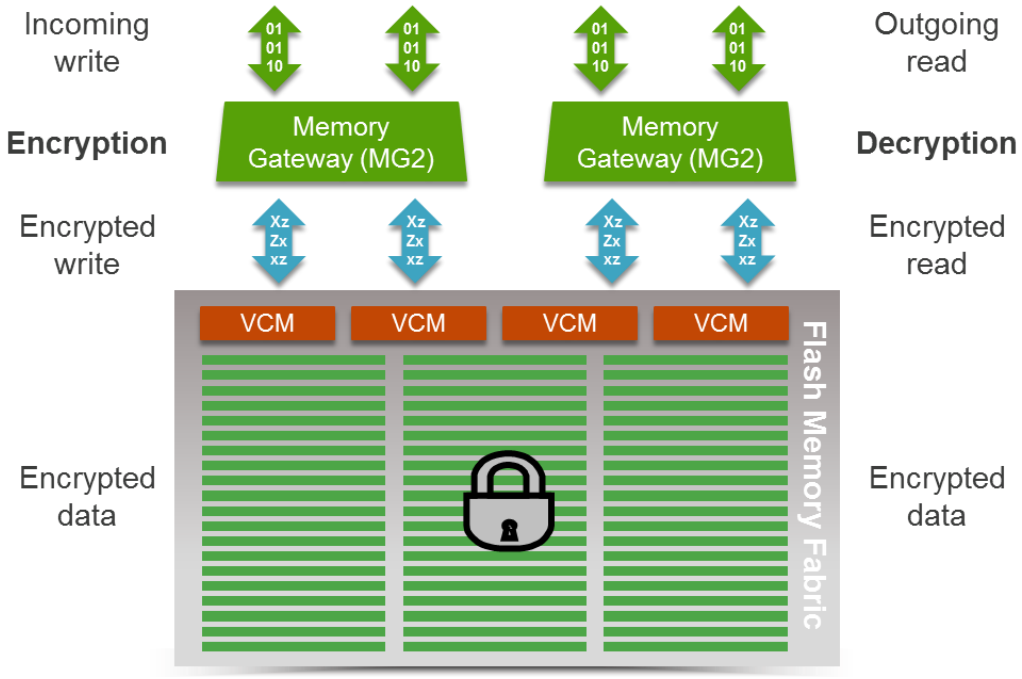

Figure 6: vMOS Encryption ensures complete protection of all data written to Violin Memory Flash Arrays

5.2 Snapshots

vMOS Snapshots empower enterprises to benefit from sustained low application latencies, without compromising their needs for advanced enterprise data protection. vMOS Snapshots are point-in-time, space-optimized, instant copies of the active data. vMOS Snapshot captures the state of a LUN at the point of time when the snapshot was created – any changes to the active LUN after snapshot creation leaves the snapshot unchanged. vMOS internally manages storage allocation through a set of pointers – when data in a LUN is updated, vMOS simply needs to manipulate the pointers to ensure that the snapshot retains the earlier data. This ensures that any number of vMOS snapshots can be created instantly. vMOS Snapshots are space optimized by default – they only use as much space as the changes in the LUN after the snapshot was created.

vMOS (Version 6.0 and later) supports up to 1024 snapshots for each LUN and up to 10,000 snapshots for the entire array. When a rollback is needed, any of the snapshots can be exported for selective recovery or with a single command, the entire LUN can be recovered from the selected snapshot. Snapshots can be managed through the vMOS CLI, WebUI and REST API. vMOS provides soft limits to alert the administrator when the used capacity of the array, including the space used by snapshots, exceeds the specified limit. vMOS also provides hard limits – when the used capacity of the array exceeds the hard limits, vMOS automatic space

reclamation kicks in and deletes the oldest unprotected snapshots to reclaim space. vMOS also provides the ability to schedule automatic snapshot creation for specified LUNs.

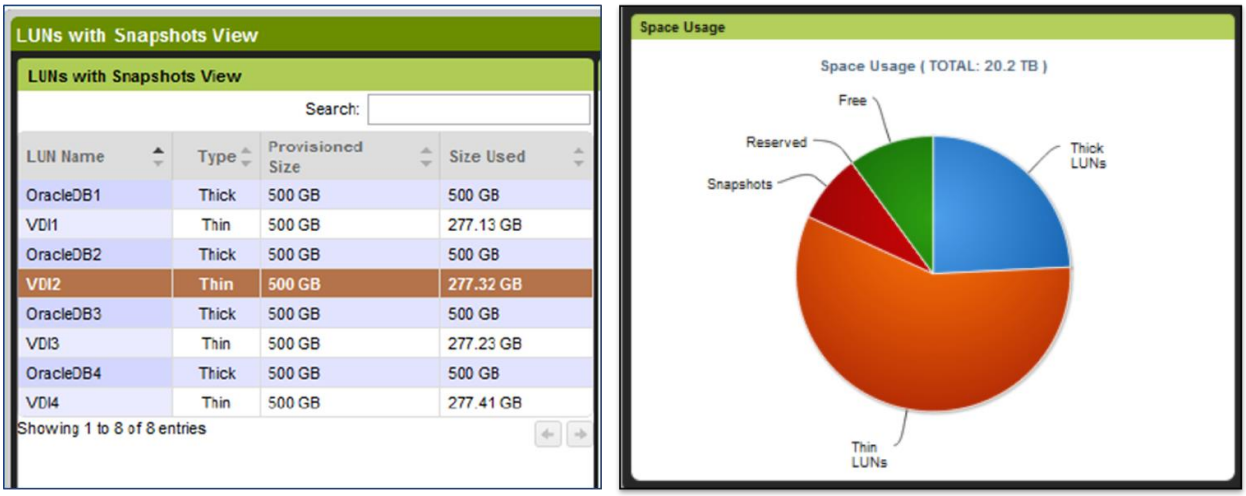

Figure 7: vMOS Snapshots are always space optimized; WebUI provides detailed insight into space usage

vMOS Snapshots are read-only, by default, to protect against any unintended updates to the snapshot data. However, vMOS also provides the flexibility of making any snapshot writeable. Writeable snapshots can be used as clone copies for testing production data, testing patch upgrades or running test simulations, without impacting production servers. Unlike other technologies that often require a full copy to mask performance issues, vMOS writeable snapshots are also fully space optimized. The underlying performance capabilities of Violin Memory Arrays ensure consistent high performance for applications using the snapped LUN as well as the writeable snapshot, while vMOS minimizes space utilized for the copy through shared space optimization.

vMOS provides crash consistency to all the snapshots – i.e., all writes to the snapshot will follow the same sequence as the writes to the production LUN and application recovery using the snapshot will only require the standard steps that follow recovery from a system crash. vMOS also provides the ability through the REST API to drive further integration with specific applications to capture application-consistent snapshots. These snapshots provide for faster application recovery by ensuring that the snapshot was taken at a point in time that is deemed consistent by the application.

vMOS also supports LUN Groups, to manage all LUNs that are used by a single application as a single management entity. In addition to allowing a LUN Group to be exported through a single click to an initiator host, vMOS also enables more advanced operations such as creating a

snapshot of an entire LUN group and providing aggregated performance monitoring across all LUNs that comprise the LUN Group.

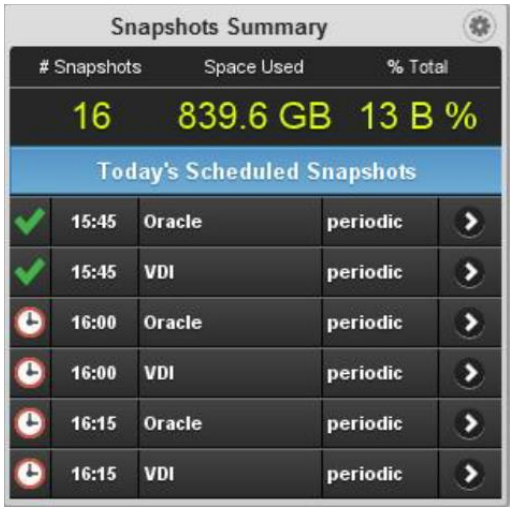

Figure 8: vMOS Snapshots can be directly managed through the automatic scheduler gadget in the WebUI

6 Storage optimization with Thin Provisioning#

vMOS enables storage administrators to power more business-critical applications with Violin's high performance by driving greater storage efficiencies. Seamless integration of functionality across multiple layers of the Violin flash Memory Array enables vMOS to deliver hardware-accelerated and flash optimized thin provisioning. Applications and users typically use less storage than they ask for. As a result, all the allocated storage is not actively utilized, forcing down the effective storage utilization and driving up the cost of storage. vMOS Thin Provisioning breaks this gridlock, by allowing administrators to provision LUNs with the size that the users request while internally only allocating space necessary to support actual application writes to the LUNs. This drives up the effective storage utilization, by keeping the allocated storage in line with what the application is actively using.

Unlike several other storage vendors who also provide Thin Provisioning capabilities, the high performance of Violin Memory Arrays ensures that applications continue to see sub-millisecond latencies even with thin provisioned LUNs. vMOS Thin Provisioning allows administrators to extend the logical size of the array beyond its physical usable capacity based on typical storage utilization patterns. To prevent running out of space, vMOS provides soft limits that can be configured to send alerts when the actual utilization exceeds these limits. vMOS also provides hard limits that can be configured to cap the maximum logical capacity of the array.

Figure 9: vMOS WebUI provides easy administration and management of thin provisioned LUNs

7 Violin Symphony

Violin Symphony, part of the Violin Memory Cluster Management Center management software family, provides a simple, unified experience for managing multiple Violin Memory Arrays. Capable of managing hundreds of Memory Arrays deployed in one or more data centers, Violin Symphony presents a refreshingly innovative approach to storage management, driving high levels of operational excellence and administrative ease. Symphony extends the rich management capability provided by vMOS for individual arrays to facilitate scale-out deployment of large numbers of Memory Arrays in the data center.

Violin Symphony revolutionizes the operational experience of storage administrators with several unique features:

- Dynamic web interface, with next-generation look and feel, as a single point of management for all Violin Memory Arrays in your data center(s)

- Single sign-on and role-based access, offering local authentication as well as full integration with Microsoft Active Directory and LDAP
- Real-time and up to 2 years of granular historic performance trends to map the I/Os per second (IOPS), latency and bandwidth, down to the level of an individual LUN
- Advanced monitoring for Service Levels of Availability (SLA) compliance, including a fully customizable proactive alert system for real-time monitoring of health and performance metrics
- Personalized dashboards for viewing critical information at a single glance – unique Build Your Own Dashboard capability, together with a rich, growing library of dashboard gadgets, gives users full control of their dashboards
- Comprehensive reports library, with more than 20 ready-to-use reports, makes it easy to share information about the performance, health and several aspects of the managed Memory Arrays. All reports can be customized and saved as pdf, xls or sent over email.

Figure 10: Consolidated dashboard of Violin Memory Cluster Management Center

8 Endnotes

* vMOS Encryption is available with vMOS-5 software release

\# Snapshots and Thin Provisioning features are available with vMOS-6 software release

Table of Figures

About Violin Memory

Violin Memory is pioneering a new class of high-performance flash-based storage systems that are designed to bring storage performance in-line with high-speed applications, servers and networks. Violin flash Memory Arrays are specifically designed at each level of the system architecture starting with raw flash memory and optimized through the array to leverage the inherent capabilities of flash memory and meet the sustained high-performance requirements of business critical applications, virtualized environments and Big Data solutions in enterprise data centers. Specifically designed for sustained performance with high reliability, Violin's flash Memory Arrays can scale to hundreds of petabytes and millions of IOPS with low, predictable latency. Founded in 2005, Violin Memory is headquartered in Mountain View, California.

For more information about Violin Memory products, visit www.vmem.com.