Virtual Desktops at the Speed of Memory

A New Storage Approach to Achieve a Responsive User Experience with Scalability to 1,000s of Desktops at a Low TCO

July 2013

Executive Summary

Virtual Desktop Infrastructure (VDI) continues to gain traction in enterprises due to the benefits of centralizing desktops. A successful VDI project will provide a user experience comparable to physical desktops, scale from 100s to 1,000s of desktops and offer a TCO comparable to physical desktops. VDI is very I/O-intensive. The critical challenge is to provide storage with sufficient IOPS performance, scalability and cost-effectiveness to ensure success. Legacy disk-based SAN arrays simply cannot deliver sufficient performance without resorting to expensive over-provisioning or other configurations. However, Violin Memory's revolutionary all-silicon flash Memory Arrays improve boot-up and application access times by 4x, scale to 1,000s of virtual desktops and deliver the lowest TCO per virtual desktop – enabling enterprises to fulfill the promise of a successful VDI implementation.

The Challenge of VDI: Heavy Demand for Storage I/O

Virtual Desktop Infrastructure (BDI) Benefits and Success Criteria

VDI continues to gain traction in enterprises where about 10 to 20% of desktops are now virtualized. By consolidating user desktops into virtual machines (VMs) running on centralized servers, VDI eliminates the need for traditional PCs and laptops, simplifies desktop administration and updates, and lowers IT management costs. Low-cost thin or zero clients can be used instead of PCs, or users can bring their own devices (BYOD), such as smart phones and tablets, and run virtual desktops there without creating a support headache for IT. Security for VDI is more robust because data is contained within the walls of a data center. It supports local and remote access, which gives users a single, consistent desktop environment no matter where they are located.

Successful VDI projects will achieve three objectives:

1) **User desktop experience comparable to a physical desktop, or better.** Users will judge virtual desktops based on their experience with physical desktops. If virtual desktops are perceived as comparatively slow or unresponsive, even intermittently such as during boot storms, users may not accept them. The key to a great user experience is high storage I/O throughput (i.e. 100+ IOPS) per desktop with low latency for fast processing and responsive applications.
2) **Scalable from 100s to 1,000s of desktops.** VDI rollouts are typically done in phases, starting with a small pilot project followed by adding incremental groups of users in production. Storage capacity and performance must be able to scale linearly to support all current and future users.
3) **TCO comparable to physical desktops.** While hardware costs for physical desktops are generally lower, the centralized administration of virtual desktops reduces management overhead and enables a competitive TCO in the long run, with one important caveat: *The server and storage infrastructure must deliver exceptional random IOPS performance in a cost-effective manner.* An infrastructure with high cost/IOPS will not have a competitive TCO, and an infrastructure with insufficient IOPS will not provide a comparable user experience.

VDI Performance and Cost Challenges

VDI is a famously IOPS-hungry application. Each desktop runs in a VM. With hundreds or thousands of VMs running on physical servers and concurrently accessing shared storage, the storage system is flooded with heavily randomized I/Os. During intensive, concurrent processes such as boot storms and anti-virus scans, the IOPS demand surges to multiple times steady state conditions (*see Figure 1 below*). Peak demand for a large number of desktops can reach into the millions of IOPS. In particular, Windows 7 is optimized for maximum performance, not minimizing I/O, and tuning Windows parameters to reduce performance demands (e.g. SuperFetch feature, disabling indexing) can be counter-productive.

Figure 1: Windows Desktop IOPS Demands

Required IOPS	Boot-up	Logon	Stable	Changes (Virus scan/ updates)	Logoff
1 desktop	150-200	100	50	200+	75
1,000 desktops	200,000	100,000	50,000	200,000+	75,000
10,000 desktops	2,000,000	1,000,000	50,000	2,000,000+	750,000

If performance requirements are not met, desktops become slow or unresponsive, especially during periodic performance-intensive processes like boot ups. Uses complain, productivity declines, and the success of the VDI project is jeopardized. Workers may even demand their PCs back!

Local storage for physical desktops has a distinct performance advantage over shared storage arrays for virtual desktops – even hybrid arrays with SSDs and HDDs (*see Figure 2*). While a hybrid array provides around 150,000 IOPS, this performance is shared concurrently among 1,000s of desktops at perhaps 20 or 50 IOPS each. This is compared to physical desktops with local, dedicated HDDs at 75 to 250 IOPS or SSDs at 10,000 IOPS. To provide a similar user experience, IOPS must match or beat physical desktops.

Shared storage is more expensive than local storage, adding to the challenge from a TCO perspective. This is especially a factor for persistent desktops, which require 10x more capacity than stateless desktops and on par with physical desktops.

Figure 2: Shared Storage Amplifies Performance and Cost Challenges

Physical desktops use local storage

- IOPS: HDD 75-250 SSD (256GB) 10,000
- Boot time : HDD 4-5 min SSD 1-2 min
- Capacity: HDD or SSD 40-150 GB

Virtual desktops need shared storage

- IOPS for typical hybrid SAN array (HDD/SSD): 150,000
- Capacity: Stateless 2-4 GB Persistent 40+ GB

In short, VDI storage infrastructure must be sized for peak demand instead of average or steady-state workloads. It must deliver IOPS comparable to or better than physical desktops for a responsive user experience. Moreover, it must be cost-effective on a per-IOPS basis to meet TCO objectives.

Traditional Disk Arrays Fall Short

Traditional disk-based SAN arrays are unable to deliver sufficient IOPS performance for a demanding application like VDI. While the rest of a computer system is solid-state silicon, disk drives still operate like vinyl record players. They work well for sequential I/O where the head can run along the same track, but random I/O and latency are limited by rotational speeds (i.e. latency) and head placement (i.e. seek time). Moore's Law, the doubling of silicon chip transistor density every 18 months or so, has driven the rapid historical development of CPU and memory performance. However, it does not apply to disk drives. Latency and seek time determine IOPS performance for disk drives, and areal density determines sequential throughput (MB/s). Annual growth in disk latency has flat-lined and seek time improves only about 5% per year, so IOPS improves less than 4% per year. Disk areal density grows 25-40% annually, which translates to modest 10-15% growth in MB/s. The net effect is that, compared to the exponential CPU and memory growth afforded by Moore's Law, disk drive performance has lagged and created a very large IOPS performance gap (*see Figure 3 below*). Disk drives have become the performance bottleneck in the computer architecture.

Figure 3: Disk Drive Performance vs. Moore's Law

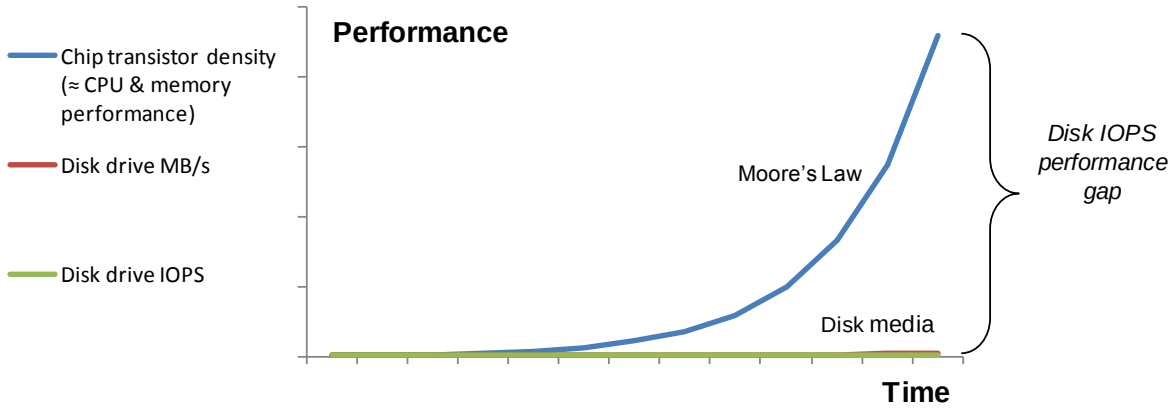

IT departments use several approaches to circumvent the disk bottleneck and improve random I/O performance, but all have disadvantages.

Striping data across a large number of drives and short-stroking (placing data only on outer sectors) aggregates and maximizes their performance. However, it requires over-provisioning capacity, which leads to wasted capacity, high operating costs for power, cooling and data center rack space as well as regular system tuning.

DRAM caches in storage arrays provide fast turnaround times for I/O while the array slack time to write I/Os to disk. However, when the write cache fills up or when a read operation does not involve select data staged in cache, the result is a "cache miss" that incurs the additional latency of disk access. Performance becomes unpredictable under sustained workloads. DRAM is also expensive.

Hybrid storage arrays add SSDs alongside HDDs and layer on complex software workarounds like automated storage tiering to leverage the faster performance of flash memory, but these systems are based on legacy disk-based architectures and only able to achieve "disk-plus" performance with millisecond latencies.

A New Storage Approach

It is clear that a new approach is needed for VDI. Imagine if storage could run at the speed of memory, if storage I/O constraints were lifted in a revolutionary all-flash storage system that could:

- Increase virtual desktop boot-up and application access times by 4x
- Scale to 1,000s of virtual desktops
- Decrease power and cooling by 9x and rack space by 36x
- Cut TCO by ½ or more compared to traditional storage arrays

Violin Memory Runs VDI at the Speed of Memory

Violin flash Memory Arrays deliver storage at the speed of memory, shattering the disk I/O bottleneck. Designed for sustained performance with high reliability, expanding to hundreds of terabytes and millions of IOPS with low, spike-free latency, Violin flash Memory Arrays are the ultimate storage solution for VDI.

Next-Generation Flash Architecture

Violin flash Memory Arrays are all-silicon shared storage systems built from the ground up – chip to chassis – to harness the full power of flash memory and deliver memory-like performance with data access latencies measured in microseconds, a full order of magnitude lower than legacy Tier 1 storage solutions. These systems are plug-compatible with existing infrastructure, supporting Fibre Channel, iSCSI, Infiniband and PCIe connectivity. At the system's core lies a resilient, highly available deep mesh of thousands of flash dies that work in concert to continuously optimize performance, latency and longevity:

- Violin Intelligent Memory Modules (VIMM) organize individual dies into intelligent flash management units with garbage collection, wear leveling and error/fault management. With greater storage density than SSDs, VIMMs enable a highly scalable solution in a compact profile ideal for VDI deployments.
- VIMMs are integrated into the patented Violin Switched Memory (vXM) architecture, designed from the ground-up for power efficiency and exceptionally high and sustained IOPS performance.
- The vXM layer and VIMMs work in conjunction with vRAID, Violin's patented hardware-based RAID algorithm specifically designed to increase reliability and reduce latency for responsive and dependable virtual desktops.

Collectively, these deeply integrated layers make up the Violin Flash Memory Fabric, which delivers industry leading performance, 99.9999% availability and space-efficient capacity.

Violin Flash Memory Array Portfolio

The Violin Memory 6000 series flash Memory Arrays deliver up to 1 million IOPS, latency as low as 100 µsec and up to 32TB of raw flash capacity in only 3 Rack Units (3U). They include 6 different models with a variety of capacity and performance profiles to fit your particular workload. The 6600 Series uses single-level cell (SLC) flash memory to maximize performance, and the 6200 Series uses multi-level cell (MLC) to optimize density and capacity.

Deduplication Reduces Capacity for Persistent Desktops

Violin flash Memory Arrays now offers inline data deduplication for greater storage efficiency and reduced storage capacity costs – with virtually no impact on performance. Like a hybrid car with better mileage or a house with thicker insulation, deduplication is feature that makes more efficient use of scarce resources. It scans written data for redundancy and stores only original blocks of data, while redundant data is substituted with a pointer referencing the original data.

For VDI, this means reduced capacity for persistent desktops. Persistent desktops allocate full capacity for each virtual desktop. The benefit is desktops can be customized, which users like, but persistent desktops consume more storage capacity and raise the infrastructure costs of VDI. In contrast, stateless desktops use a single “golden” image and create linked clones for each user desktops that only store changed data. In effect, stateless desktops are deduplicated by design. The benefits of stateless desktops are standardization, easier IT administration for IT and minimized capacity, but there is no customization.

With Violin deduplication, enterprises can have their VDI cake and eat it too. In-house testing has found deduplication ratios of 5:1 are typical for VDI environments using persistent desktops. This means users can be given customizable persistent desktops with only one-fifth the capacity normally required, reducing flash capacity and operating costs by 80%.

Violin Benefits for VDI

The exceptional IOPS performance, cost-effectiveness, scalability and power- and space-efficiency of Violin flash Memory Arrays enables you realize your virtual desktop goals:

- *5x – 10x improvement in user desktop experience*
- *2x – 5x greater desktop consolidation*
- *Scalable to 1,000s of virtual desktops* – Easily handle I/O-intensive processes for massive desktop consolidation and consistent, responsive desktop experience at scale.
- *Lowest TCO per virtual desktop* – Less storage is needed to meet your performance requirement, so less power, cooling and floor space are consumed and operating costs are lower. Ultra-high storage IOPS enables greater server utilization and consolidation, so fewer servers and software licenses are needed.
- *Ease of management* – Violin Symphony for centralized management. Seamless integration with VMware APIs.

Third-Party Performance Validation Results

The industry analyst and consulting firm Enterprise Strategy Group tested and validated the performance of Violin flash Memory Arrays in a VDI environment (see Figure 5 below). The results show Violin delivered an unmatched user desktop experience and continued to perform exceptionally well at scale and during intensive operations like boot storms:

- Sustained low application access times (< 5 seconds) even at 10,000 desktops

- Faster boots, logins, virus scans and patching
- Boot up in only 36 seconds at 10,000 desktops
- Ability to execute a significantly large number of concurrent user actions

Figure 5: VDI Performance Results for Violin flash Memory Array (from ESG, July 2012)

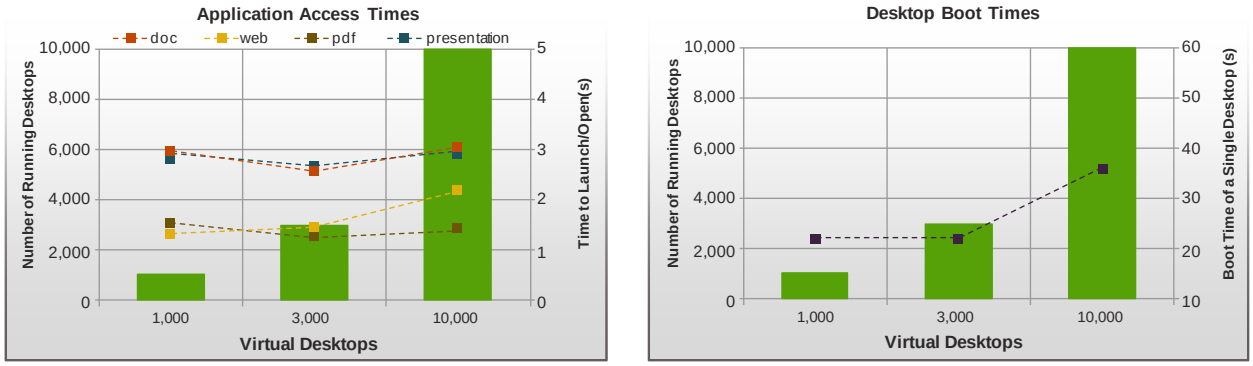

The All-Flash Impact: Benchmark Performance Records

When leading technology vendors like HP, IBM, Oracle and VMware want to set performance records for industry benchmarks like TPC-C, TPC-E and VMmark, they build their test configurations with Violin flash Memory Arrays. For the simple reason that nothing else is faster.

Figure 6: Storage of Choice for Performance Records

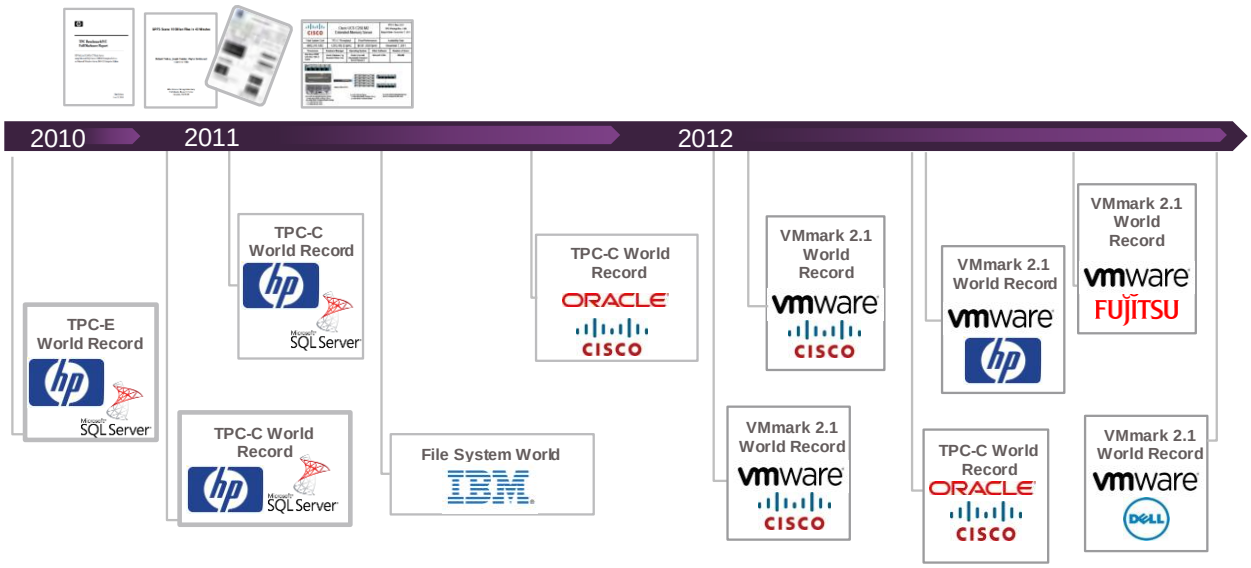

Leading Storage Vendor vs Violin: How Much Storage for 10,000 Virtual Desktops?

To compare a leading storage vendor with Violin Memory, we analyzed how much storage is needed from each vendor to run 10,000 virtual desktops. Assuming stateless desktops and 50 IOPS per desktop, the total estimated performance and capacity estimate to support this workload is 500,000 IOPS and 40 TB. The products compared were a storage system with SSDs & HDDs and automated storage tiering software vs. the Violin 6232 flash Memory Array in standard configuration.

As shown in Figure 7, it takes 5 fully-configured traditional storage arrays consuming five full data center racks to support this workload, but only 2 Violin flash Memory Arrays in 6U, a small fraction of a single rack. The Violin solution also delivers the 500,000 IOPS and 40 TB with 10x lower latency, enabling faster application response times.

The TCO of the Violin solution is dramatically lower than the leading storage vendor:

- *Half the cost* – Save 50% right away on equipment costs.
- *36x less rack space* – Save 97% on data center floor space. If the storage is in a colocation facility, Violin saves 97% in annual rack space fees. If in an internal data center, Violin saves 97% in building overhead and possibly avoids a data center expansion.
- *9x less power and carbon emissions* – Save 88% annually in electricity costs, both to power the systems and to cool the heat generated.

In other words, Violin offers a VDI storage solution with superior performance, lower acquisition and operating costs, smaller profile and lighter environmental footprint.

Figure 7: Leading Storage Vendor vs. Violin flash Memory Array – 10,000 Virtual Desktops

	VNX
Rack Space	214U
Power (kw/yr)	294,187
Carbon (Ton/yr)	127
Latency	8-10 ms

	Violin
Rack Space	6U
Power (kw/yr)	34,264
Carbon (Ton/yr)	14.7
Latency	<1 ms

Violin Memory Has Proven Record for VDI

Violin has a proven record of delivering VDI with superior performance and scalability and low TCO. Multiple customers have 1,000 or more virtual desktops running on Violin flash Memory Arrays, including the US Department of Defense (multiple 1,000s), Great American Insurance Group (2,000), NATS (6,000), Anglia Ruskin University (3,000) and a multi-billion dollar construction company (10,000). NATS and Anglia Ruskin University are profiled below.

NATS: Large-Scale VDI Rollout with £9 million Savings

NATS, the leading provider of air navigation services and solutions in the UK, needed a highly available storage platform to support a mission-critical, cloud-based VDI rollout with 6,000 seats. By selecting Violin Memory 6000 Series flash Memory Arrays that each provide over 1 million IOPS at ultra-low latency, the firm expected to save £9 million over 4 years, shrink its environmental footprint, gain flexibility and meet demanding SLAs for users across multiple geographical regions.

"We needed a high performance solution which would scale to at least 6,000 seats with no appreciable degradation to user experience. Using Violin Memory will also introduce a resilient architecture, avoiding single points of failure."

Gavin Walker, Chief Information Officer, NATS

Anglia Ruskin University: Fast, Energy-Efficient Virtual Desktops at Lower TCO

Anglia Riskin University in the UK needed to provide new desktops capable of responsive video presentation and user experience, while in a building with no additional cooling, so traditional PCs were not an option. The university chose VDI based on VMware View, zero clients and a Violin flash Memory Array capable of 20 times the random write performance of a SAN disk array. As a result, the new virtual desktops delivered 30-40% power savings, boot times reduced by nearly 90%, and a desktop experience nearly indistinguishable from a traditional PC. At 500 or more users, the VDI capital expenditure was comparable to traditional PCs, while yearly support and maintenance costs were lower.

"We chose the Violin flash Memory Array because it's very well-engineered, reliable and offers high storage performance."

Gregor Waddell, Assistant Director, Anglia Ruskin University

Are You Ready to Run VDI at the Speed of Memory?

Violin's flash Memory Arrays offer unmatched benefits for your VDI deployment:

- Exceptional storage IOPS to meet peak workloads
- Achieve better than Ultrabook desktop experience
- Scalability to 1,000s of desktops – Up to 5,000 in 3U
- Enable TCO comparable to physical desktops

Contact Violin Memory today and start running your virtual desktops at the speed of memory!

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 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 terabytes 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.

Violin MEMORY

Violin Memory, Inc.
685 Clyde Ave, Mountain View, CA 94043
Ph: 1-888-9VIOLIN (984-6546)
Email: sales@vmem.com

Violin Memory EMEA
Quatro House, Lyon Way, Camberley, Surrey, GU16 7ER. UK
Ph: +44 1276 804620
Email: emeasales@vmem.com