

5,000 Persistent VMware View VDI Users on Dell EMC SC9000 Storage

Abstract

This reference architecture document records real-world workload performance data for a virtual desktop infrastructure (VDI) storage solution based on the Dell EMC™ SC9000 array.

August 2017

Revisions

Date	Description
December 2016	Initial release
August 2017	Upgrade from 4,000 user architecture

Acknowledgements

This paper was produced by the following members of the Dell EMC storage engineering team:

Author: Damon Zaylskie

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	4
1 Introduction.....	5
1.1 Objectives.....	5
1.2 Audience.....	5
2 Choosing VDI storage.....	6
3 Solution architecture.....	7
3.1 Environment.....	7
3.2 Software.....	7
3.2.1 VMware vSphere host.....	7
3.2.2 VMware Horizon 7.1.....	7
3.2.3 Microsoft Windows 7.....	9
3.2.4 Microsoft Windows Server 2016.....	9
3.3 Hardware.....	9
3.3.1 Storage.....	10
3.3.2 Hosting environment.....	11
3.3.3 Networking.....	13
4 Testing.....	14
4.1 Test objectives.....	14
4.2 Test tools.....	14
4.2.1 Load generation.....	14
4.2.2 Load monitoring.....	14
4.3 Test phases.....	14
4.3.1 Boot phase.....	14
4.3.2 Login phase.....	15
4.3.3 Steady-state workload.....	17
4.4 Test results.....	19
5 Conclusion.....	20
A Environment.....	21
B Technical support and resources.....	22
B.1 VMware support.....	22

Executive summary

Dell EMC SC Series storage arrays are modern, mature, and constantly evolving. With the ever-changing workload demands that organizations face today, it is essential that SC Series arrays continually advance to keep pace. The architectural changes incorporated into the Storage Center OS (SCOS) 7.2 release allows you to face these workload demands head on.

This document records real-world workload performance data for a virtual desktop infrastructure (VDI) storage solution based on the Dell EMC SC9000 array. The SC9000 is the flagship product in the SC Series family of arrays and provides an ideal solution for large-scale systems, high-end workload performance, and distributed data center environments. The test environment simulates a mid- to large-sized company with 5,000 end-user VMs running Microsoft® Windows® 7 and a workload typical of knowledge workers. The target IOPS load was 40 IOPS per virtual machine, for a total requirement of 200,000 IOPS.

The results show the SC9000 excels at delivering low latency I/O under heavy load. The maximum steady state IOPS exceeded 216,000 at just over 1.5ms of latency. The improvement over previous versions of SCOS are due to large improvements in code efficiency and improved utilization of hardware. The change in performance will vary based on customer I/O patterns, but workloads should demonstrate increased performance.

1 Introduction

Whether you are building a private or hybrid cloud, a big data analytics center, or a secure OLTP platform, the SC9000 array provides a quick, strong return on investment by optimizing, accelerating, and protecting the data that is most important to your business goals.

The SC9000 builds on the success of previous-generation SC Series arrays with a fully upgraded, more scalable hardware platform powered by the robust SCOS 7 operating system. The improvements in SCOS 7.2 also highlight the performance gains of platform maturity coupled with enhancements in code efficiency. Each build of SCOS has increased performance with no additional hardware requirements, as long as sufficient drives are attached to handle the increased throughput.

All-flash and hybrid SSD/HDD configurations can revolutionize the speed and impact of mission-critical applications. The ability to apply write- and read-intensive SSDs within the same volume, plus add an ultra-dense HDD tier at any time, gives the SC9000 flexibility to target a variety of mixed-workload performance requirements. This enables a seamless evolution from one configuration to another without having to replace drives.

1.1 Objectives

The primary objective of the testing effort is to demonstrate the performance and scalability of the SC9000 array. By providing a baseline for performance and scale, an SC9000 array can be sized to fit any midrange storage requirement, providing low latency and excellent performance for the cost. The key items tested are:

- IOP capability
- Latency under load
- Consistent performance under varying load

1.2 Audience

This reference architecture is targeted at anyone who wishes to better understand the performance and scalability of the SC9000 storage platform running the current generation of mixed-use solid-state drives. It is intended to show what the platform is capable of from a VDI workload perspective. The results will be different under other workloads. Performance under other workloads can be found on the Dell TechCenter page, [SC Series Technical Documents](#).

2 Choosing VDI storage

When choosing a VDI storage platform, there are a few important factors to consider:

RAID tiering: Solutions from some vendors require preallocation and segregation of RAID levels, holding large amounts of storage captive, which can have negative cost and performance impacts. SC Series arrays present a fully virtualized storage pool, configuring and mixing RAID levels on demand across all drives to address reads and writes differently. RAID 10 is used for new writes, while blocks not being updated are converted to read only at parity RAID 5 or 6. This provides the performance benefits of RAID 10, plus the capacity benefits of RAID 5/6. RAID acts as a flexible sub-tier for Dell EMC arrays, optimizing data placement within each drive. This maximizes performance, capacity, and administrator efficiency.

Dynamic capacity: Dynamic capacity technology makes the initial VDI capacity sizing an easy exercise. Virtual volumes can be created for thousands of virtual desktops without pre-allocating physical capacity. Actual storage capacity is consumed only when data is written to disk. As the virtual environment grows to accommodate more users, the system dynamically provisions storage from a centralized pool of unused capacity. With dynamic capacity, storage purchases can be deferred until actually required and seamlessly added to the pool without any service disruption, thereby significantly enhancing project ROI. Capacity can be grown to over one petabyte without downtime.

Hypervisor integration: Additional operational efficiencies are realized through integration with the hypervisor layer. Dell SC Series arrays tightly integrate with industry-leading hypervisors such as VMware® vSphere®, Microsoft Hyper-V®, and Citrix® XenServer®. This integration enables optimal storage performance, lower costs, and simplified VDI deployments. As an example, the Dell Storage vSphere Web Client Plug-in reduces complexity by providing integrated storage provisioning and management. Support for VMware vStorage APIs for Array Integration (VAAI) enhances vSphere server performance by offloading storage-related tasks (such as hardware-assisted locking, full copy, and block zeroing), reducing vSphere server compute overhead, network traffic, and virtual desktop deployment times.

Compression and deduplication: Compression and deduplication, used together or separately, can provide cost savings, depending on the number of base images, clone type (linked or full), number of VMs, and numerous other factors. The number of variables make it difficult to predict actual space savings, but in general, VMs deduplicate fairly well. One exception is linked clones, which are already very space efficient and typically do not benefit from deduplication, which is why many vendors express their numbers using full clones to give the impression of more space savings from deduplication.

3 Solution architecture

3.1 Environment

There are a large number of software and hardware components required to drive an environment of this scale. They all need to work together to provide a stable and seamless solution for VDI to work correctly.

3.2 Software

The solution is built on layers of software and hardware, each providing a key piece of the puzzle. The solution presented in this paper used the following software:

- VMware vSphere hypervisor 6.5
- VMware vSphere management
- VMware Horizon® 7.1
- Login VSI™ (for simulating real-world VDI workloads)
- Microsoft Windows Server® 2016 and Microsoft Windows® 7

All software components were updated to the most currently available patch levels for the tests.

3.2.1 VMware vSphere host

VMware vSphere is one of the top enterprise virtualization platforms used for building VDI and cloud infrastructures. VMware vSphere includes three major layers: virtualization, management, and interface. The virtualization layer includes infrastructure and application services. The management layer is central for configuring, provisioning, and managing virtualized environments. The interface layer includes the vSphere client and the vSphere web client.

Throughout the solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (for example, NTP, DNS, and Microsoft Active Directory®). The vCenter servers used in the solution are VMware vCenter Server Appliances v6.5. All hosts were running vSphere 6.5.

3.2.2 VMware Horizon 7.1

The VDI management platform used was VMware Horizon 7.1 with linked clones. Horizon allows a number of different cloning options, as well as managing discrete desktops. It is designed to scale to large numbers of end-points in a single interface, and provide consistency and manageability in large environments.

Linked clones were chosen for their management and small footprint. This is a very commonly used configuration because of its efficiencies and ease of use.

Horizon has several components that make up the product:

Connection Server: The Connection Server is the central hub for managing connections. Users connect using either the Horizon View client or a web browser. The Connection Server then connects the user to the proper virtual desktop session based on the assignments. Users can either be assigned a static machine (in this case, for persistence) or given a random machine in a non-persistent environment.

Composer: The Horizon View Composer server handles the tasks of provisioning virtual desktops to the proper pool using the correct base image. It also handles maintenance tasks, such as a desktop refresh or recompose. There were four Composer VMs, each registered to a connection server.

Horizon View agent: Installed on user devices, the Horizon View agent provides users with quick, secure, multi-protocol access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs. VMware Horizon provides on-demand access to Windows, web, and Software as a Service (SaaS) applications.

Management console: The Horizon View management console is a web service application that can be run from any Windows server in the environment. The management servers used were configured using VMware vCenter™ 6.5 with a central Microsoft SQL Server® 2014 instance. There were four management servers configured, each controlling 875 virtual machines.

3.2.2.1 [Linked clones](#)

All 5,000 virtual machines were created in less than two hours using parallel environments taking advantage of the high performance all-flash array.

The virtual machines were created using the linked clone method in Horizon View. This provides efficient use of space and simplifies management. By using linked clones, a replica image is used for shared data resulting in a very small footprint. After cloning, the total space consumed is less than 2GB/VM in a persistent state before user login.

After the user profiles were created, the VMs each consumed approximately 3GB of space. This includes the user customizations, page file, and changes made to customize each VM to the assigned user. Each VM runs Microsoft Office 2010, Adobe® Reader®, and a mix of Login VSI applications

Virtual desktops are assigned to the same user every time with desktop persistence. All changes made by a user are retained for future sessions. The changes are stored in a delta VMDK stored with the VM.

All the desktops are in an automated desktop pool with dedicated user assignments. Those desktops are based on a master desktop template, which is selected when the pool is created. This ensures each user gets a dedicated machine they can customize.

Horizon View also supports creating non-persistent environments to streamline support and troubleshooting. These desktops automatically revert to the default configuration upon user logoff.

3.2.3 Microsoft Windows 7

Each of the desktops was a Windows 7 Enterprise client. Windows 7 32-bit was chosen for this configuration. All patches were applied to the desktop image before the base replica was created.

Some Windows components were disabled to reduce resource requirements, in accordance with best practices. In each desktop, indexing was disabled, the frequency of Windows updates was limited, and pagefile and temporary files were redirected to a separate volume. All disk traffic was captured during the test runs.

3.2.4 Microsoft Windows Server 2016

For all Windows Server based functions, Windows Server 2016 was used, with all current patches. The VM hardware configuration varied based on the roles and is documented in the appendices.

3.3 Hardware

A fully functional VDI environment for 5,000 virtual machines requires a large amount of hardware with significant compute and memory resources. The act of virtualizing does realize great efficiencies but cannot eliminate the requirement to process data. The act of virtualizing collapses the environment into dense clusters of resources.

The hardware used to host the environment consisted of the following:

- Dell PowerEdge™ blade servers for VDI workload
- Dell PowerEdge M1000e blade enclosures
- Dell EMC Networking MXL Ethernet blade switches
- Dell EMC Networking S5000 top-of-rack (ToR) Ethernet switch
- Brocade® M5424 Fibre Channel (FC) blade modules
- Brocade 6505 FC ToR switches
- Dell EMC SC9000 storage array

For the tests performed in this document, all hardware was updated to the newest, generally available firmware from the manufacturer.

There are many sizing tools used to gauge the processor and memory requirements for virtual machines. They are outside the scope of this document. See the [VMware community sizing guidelines](#) and [VMWare Windows 7 sizing guidelines](#). While the guidelines state 128 VMs per LUN as a practical limit, The SC9000 supports higher VM counts with no performance impact. In this scenario, 167 VMs were created per LUN. The sizing guidelines also state 26+ IOPS per user. By achieving 40 IOPS per user, the SC9000 far exceeds the performance requirements of demanding workloads.

3.3.1 Storage

The storage used was an [SC9000 array](#) with mixed-use solid-state drives. The mixed-use drives provide excellent performance with a very attractive price point. They are also low-power devices relative to spinning drives and generate far less heat. The storage density is also very good.

3.3.1.1 SC9000 overview

Technical highlights of the SC9000 include:

Array configurations: All-flash, hybrid, or HDD arrays.

Auto-tiering method: Policy-based migration based on real-time data usage, customizable 512KB–4MB page size.

Auto-tiering structure: Up to three primary (media-based) tiers total, with up to two SSD tiers (write- and read-intensive SSDs).

RAID support: RAID 0, 1, 5, 6, RAID 10, and RAID 10 DM (dual mirror); any combination of RAID levels can exist on a single array.

RAID tiering: Auto-provisions and dynamically restripes multiple RAID levels on the same tier; no need to pre-allocate RAID groups.

Thin provisioning: Active by default on all volumes, operates at full performance across all features.

Thin snapshots: Records changes only; snapshots auto-migrate to lower-cost storage.

Intelligent deduplication and compression: Selectable option per volume on SSD and HDD tiers. Compression-only option also available.

Standard replication: Synchronous/asynchronous through FC or iSCSI.

Federated multi-array systems: Live Migrate (included in base product) enables host-transparent data movement among arrays. Volume Advisor continually monitors federation to suggest optimal data placement based on customizable capacity or performance load balancing policies. Dell Storage Manager (DSM) provides data-center-level viewpoint and control.

Business continuity: Live Volume auto-failover, auto-repair, VMware Metro Stretch Cluster support, and Microsoft Azure® Site Recovery.

3.3.1.2 Storage test configuration

The tests were conducted using Storage Center OS (SCOS) 7.2 with system write cache disabled. There are considerable improvements to performance in SCOS 7.2 over previous versions. These gains will be available to all platforms supporting this release of firmware.

The drives used were 1.92TB drives, providing usable space of 1.75TB per drive. The system contained 60 drives with 3 dedicated spares. This provided 99TB of raw space. In this case, a smaller drive could be used as long as it provides the same performance per drive. This system is sized for capacity growth and to allow snapshots to be kept. The target for this test is performance not capacity, and final requirements will dictate the drive capacity used.

All storage hardware was configured to Dell best practices with vSphere. For more information, see the document, [Dell EMC SC Series Best Practices with VMware vSphere 5.x–6.x](#).

3.3.2 Hosting environment

The environment running the VMs makes up the majority of the hardware required. The power, rack space, cooling, and management requirements are much larger than the storage components. This is normal based on the processing power and memory requirements of 5,000 virtual machines.

Table 1 Host configuration

Component	Description
Host	PowerEdge M620, two 8-core Intel® Xeon® Processor E5-2680, 320GB RAM
Operating system	vSphere 6.5
Fibre Channel	QLogic® QME2572 HBA, Brocade M5424 chassis-based switch
Network	Dell EMC Networking MXL 10Gb dual-port Ethernet module
Virtual desktop configuration	2 vCPUs, 2GB RAM, VMXNet3, 30GB volume

Table 2 Virtual machine configurations

Component	Description
Infrastructure VMs	4 vCPUs, 8GB RAM, VMXNet3 vNIC, Login VSI 4.1.25 agent
Operating system	Windows Server 2016
vNIC	VMXNet3
Applications	Microsoft Active Directory®, DNS, DHCP, vCenter Servers, Horizon Connection and Composer servers, Microsoft SQL Server®, DSM

Table 3 Infrastructure servers

Component	Description
Operating System	Windows Server 2016
vNIC	VMXNet3
Applications	Active Directory, DNS, DHCP, vCenter Servers, Horizon Connection and Composer servers, SQL Server, DSM

Table 4 Storage

Component	Description
SC9000	Dual storage controllers
16Gb x 4-port Fibre Channel	QLogic 16Gb 4-port Fibre Channel adapters, 2 per controller head
SCOS 7.2	Storage firmware
Drives	60 x 1.92TB: 57 active, 3 spares

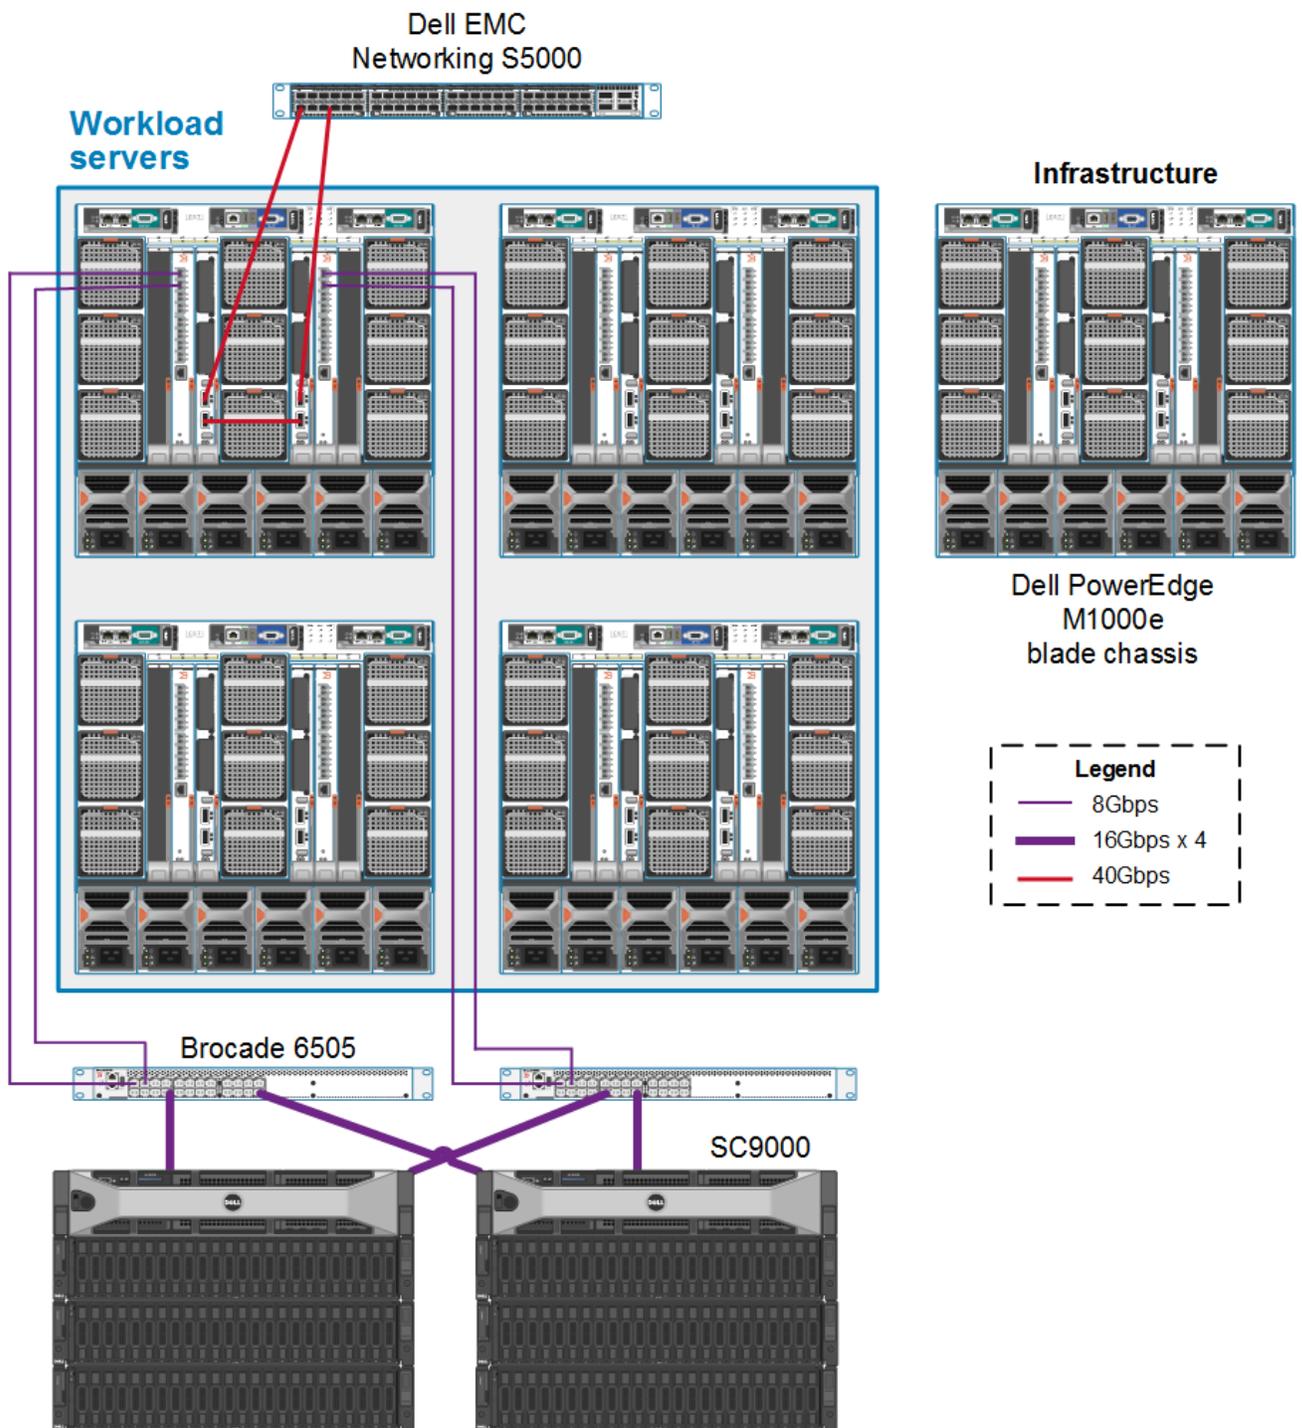


Figure 1 Hosting environment

A dedicated server was used for the Login VSI control and file-sharing server. This is due to the **very** high I/O requirements of the centralized share. At times, the throughput exceeded 10Gb/sec from the file share. The Login VSI share server was a PowerEdge M620, with 144GB of RAM and 2 x 10Gbps Intel Ethernet modules.

3.3.3 Networking

The network architecture consists of 10Gbps dual-port network adapters in each server trunked to a 10/40Gbps Dell EMC Networking MXL dual-switch stack in each blade chassis. Each MXL stack pair is cross-connected with a 40Gbps connection and trunked with two 40Gbps QSFP+ connections to a ToR Dell EMC Networking S5000 switch.

4 Testing

4.1 Test objectives

The series of tests run for this reference architecture are designed to capture the performance of the SC9000 controllers. The tests were scaled to run within the limits of the SC9000 controllers and highlight their performance.

4.2 Test tools

4.2.1 Load generation

The load-generation tools used were created by Login VSI consultants. These are End User Computing standard tools for consistent load generation and allow Dell EMC to compare performance results across platforms and the industry by leveraging a standardized workload.

4.2.2 Load monitoring

Dell Storage Manager (DSM) was used to chart and report on all aspects of the underlying storage. DSM can share real-time I/O, report on historical performance or capacity utilization, as well as provide a single pane of glass for multiple storage arrays. DSM is also the control plane for replication on SC Series arrays.

4.3 Test phases

4.3.1 Boot phase

Figure 2 shows the storage load as the virtual machines are turned on as quickly as possible. In order to achieve faster boot times, PowerCLI was used to connect to each of the hosts and power on virtual machines as quickly as possible. This may be required to boot an environment after a complete outage, whether controlled or unplanned. The figure shows the rapid boot times available on Dell EMC all-flash storage.

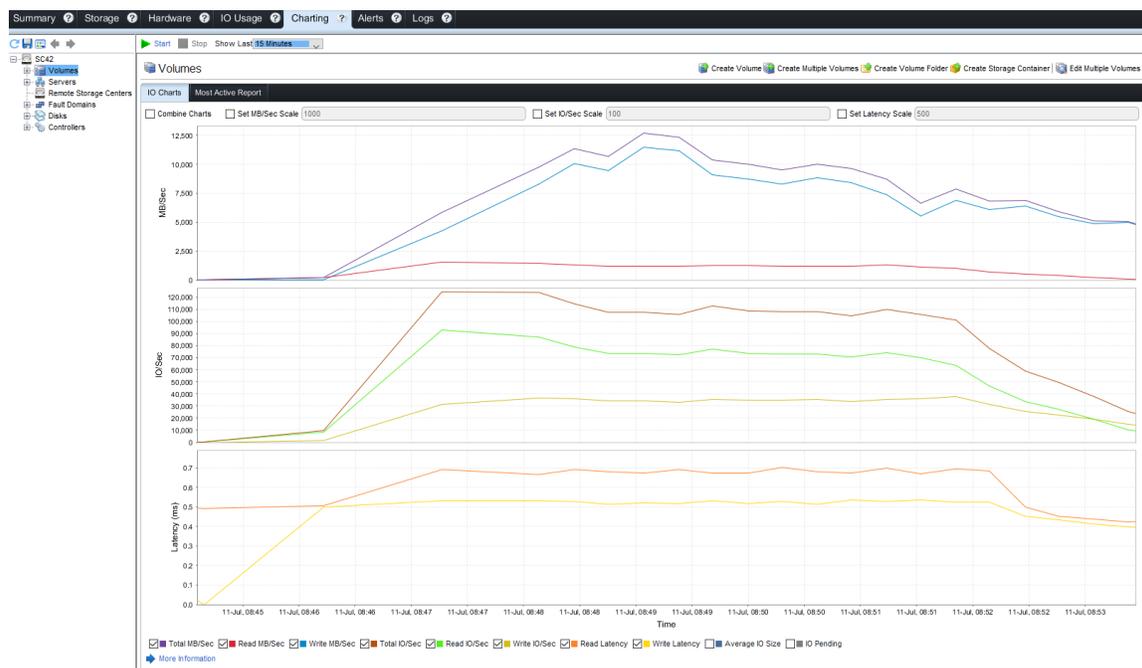


Figure 2 Boot storm

As the chart shows, almost 125,000 IOPS at 12.5GB/Sec were generated during the time it took to boot all 5,000 VMs. The time required to boot all VMs to the login screen was less than 6 minutes. This workload is uncommon because most customers very rarely reboot an entire environment.

4.3.2 Login phase

This test is a more important scenario because it demonstrates the expected performance during mass user logins, a workload which is a daily occurrence. Each morning, users log in in waves, generating large amounts of data. The chart transitions directly into the next phase which is steady state. Figure 3 shows the login phase.



Figure 3 Login phase

This chart shows the steady increase in load as more and more users log in to the system. As seen in the chart, the SC9000 provides excellent performance with predictable latency as the load increases.

Figure 4 shows the time to login and the number of sessions successfully launched.

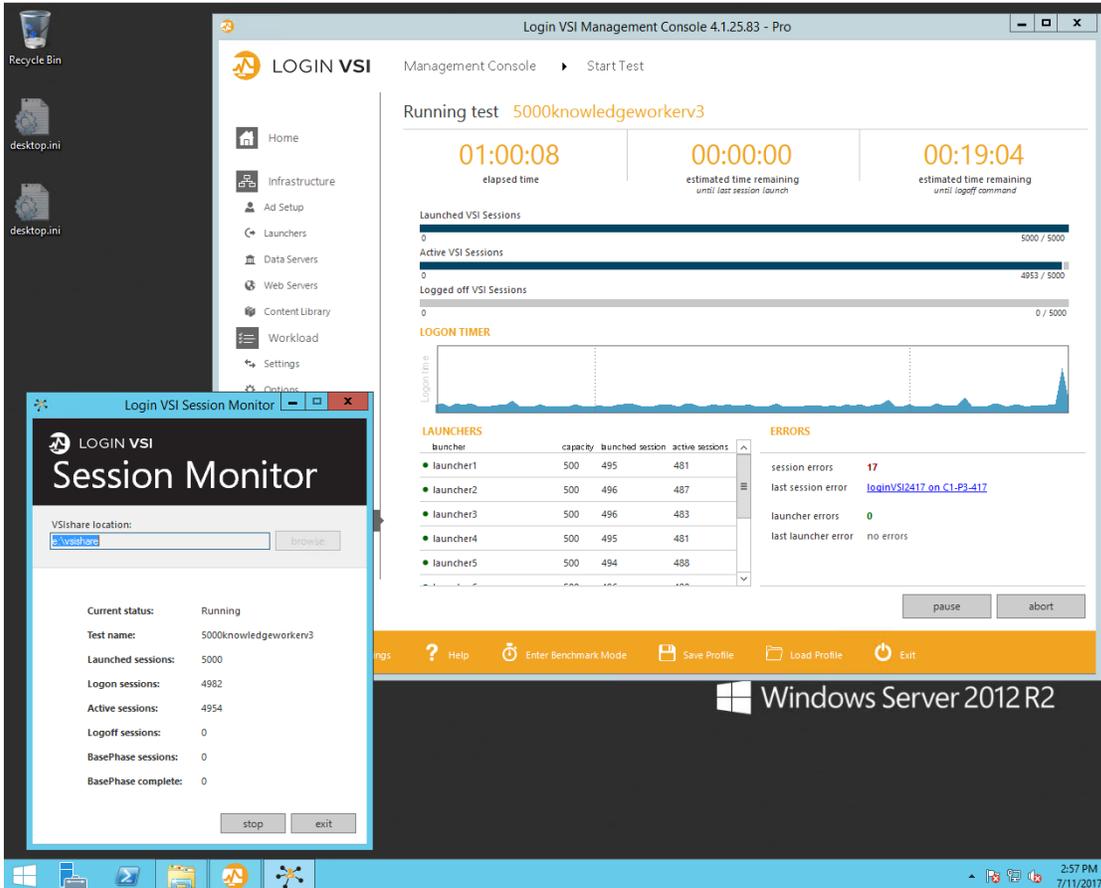


Figure 4 Login time and sessions launched

Figure 5 shows the peak load generated.

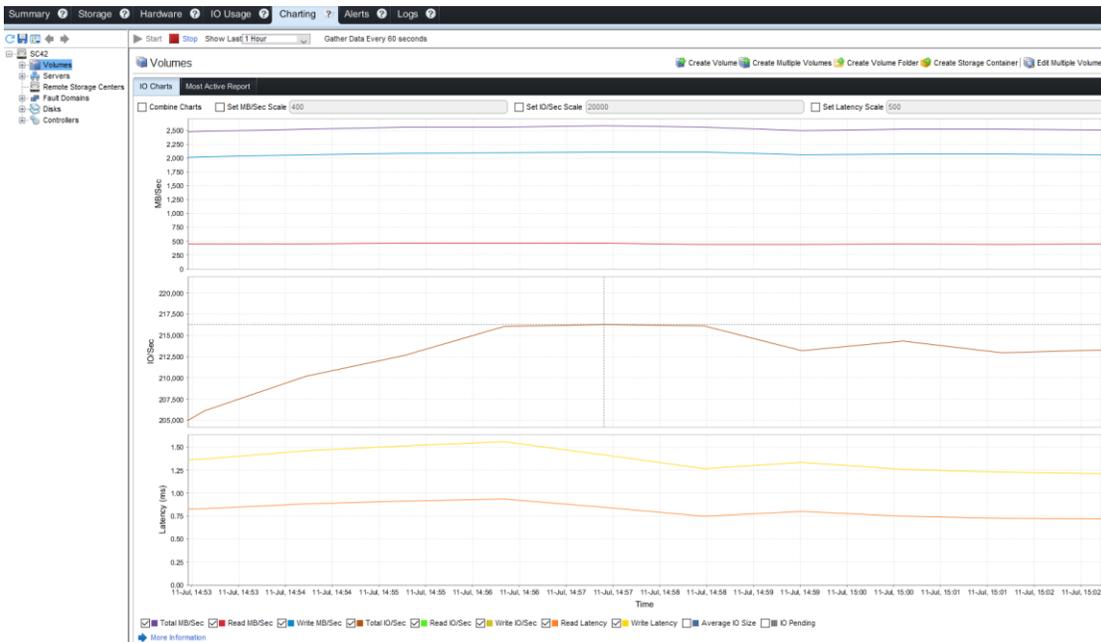


Figure 5 Peak login load

The logins generated 216,000 IOPS, with latency just over 1.5ms. This level of performance allows an excellent end-user experience, provided the compute and memory are properly sized. A storage latency this low ensures the storage will not be the gating factor in the user experience.

4.3.3 Steady-state workload

The steady-state workload used for these tests was targeted for 40 IOPS per user, to simulate a heavier-than-normal workload per user.

Figure 6 shows the load generated.

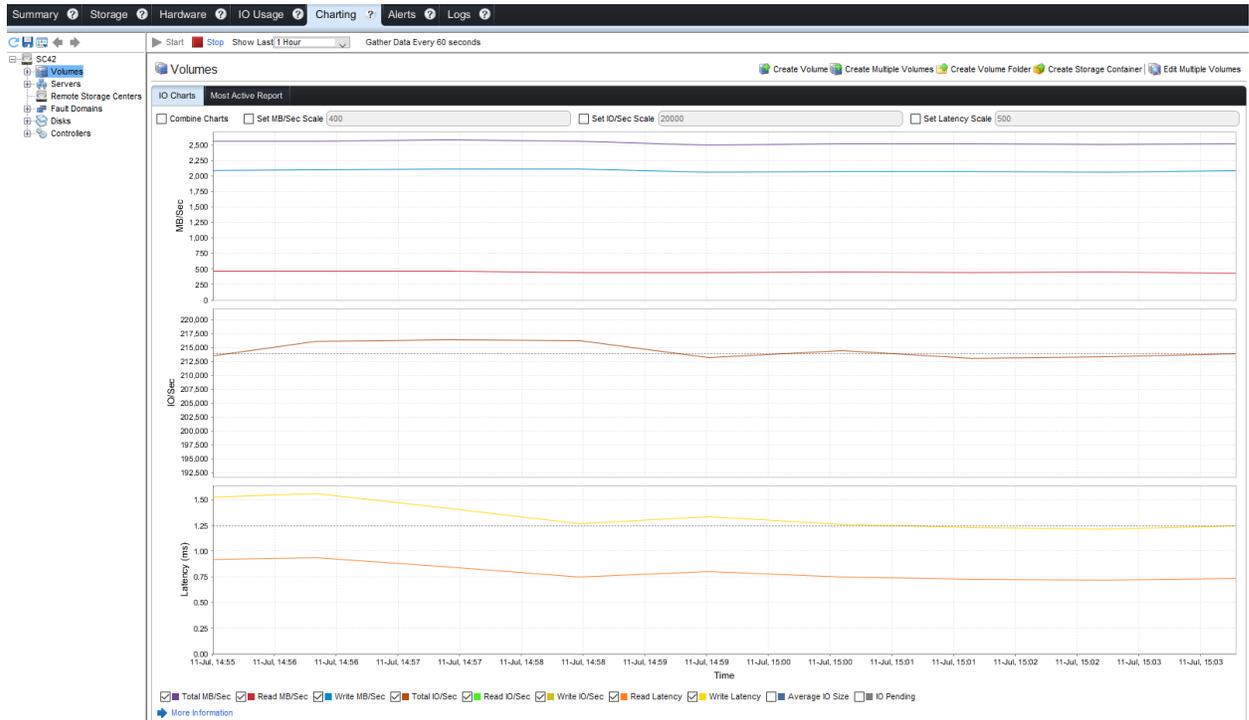


Figure 6 Steady-state workload

Figure 7 shows more detail for the steady-state load.

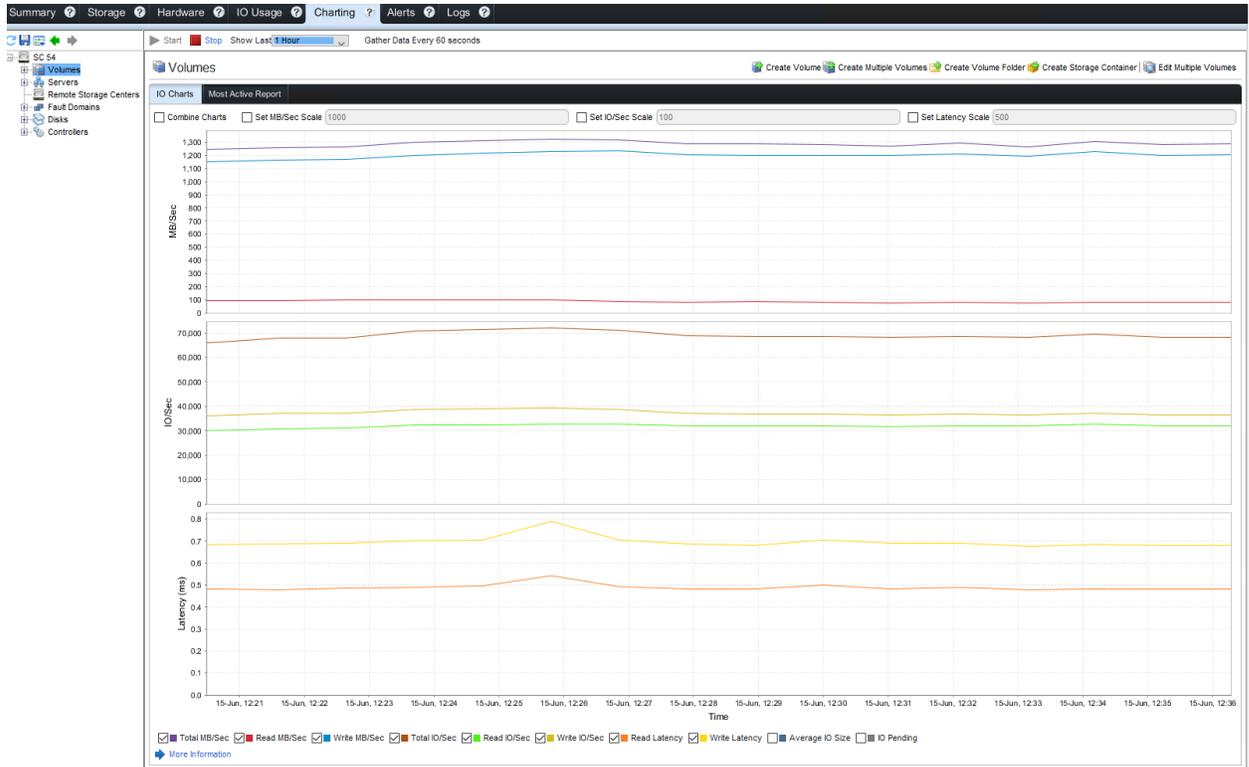


Figure 7 Steady-state load detail

As Figure 7 shows, the SC9000 handled over 40 IOPS per user while maintaining just 1.25ms of write latency. Being able to sustain such a low latency at over 40 IOPS per user ensures the SC9000 can provide fast, consistent performance for medium to large organizations.

4.4 Test results

Figure 8 shows the output from the Login VSI test tool, charting the user login and application response times. As indicated in the chart, the threshold is 1,873ms, with an average latency of just 1,456ms. This solution delivered a latency of 78 percent of the threshold.

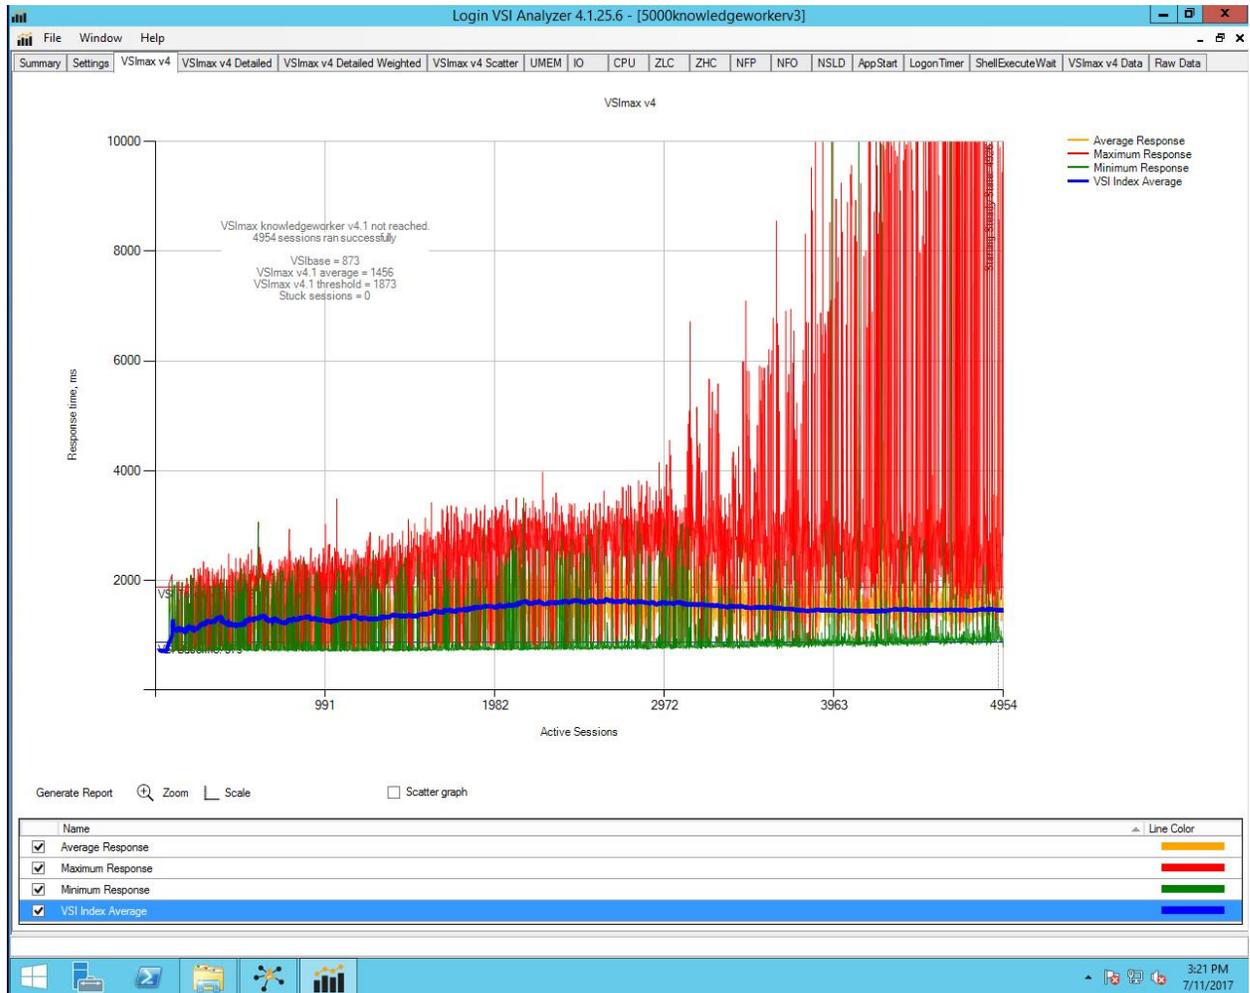


Figure 8 Login VSI results

5 Conclusion

The true test of the workload is how well the system responds to user input. Storage is just one gear in the machine enabling users to run their workloads. The testing described in this document shows how well the storage performs under a full workload.

The data show the SC9000 storage array to be an excellent platform for 5,000 heavy power users. It provides excellent performance and scalability, with up to 2PB of usable, configurable storage in a reliable and stable platform.

A Environment

Table 5 Horizon View configuration

Component	Description
Horizon View	7.1
Role	Virtual machine assignment, management, web interface, connection broker
Servers	Four 4-core servers, 12GB RAM
Cabling	10Gbps VMXNet3

Table 6 VMware host hardware configuration

Component	Description
VMware host	6.5
Role	End-user compute node
Hardware	Two 8-core processors, 320GB RAM
Storage	200GB SSD for boot/log

Table 7 VMware vCenter server configuration

Component	Description
VMware vCenter appliance	6.5, two virtual vCenter appliances for end-user VMs, one appliance for infrastructure
Role	Management of all hosts, DRM load balancing, host monitoring

Table 8 Login VSI configuration

Component	Description
Login VSI	4.1.25
Role	Workload generation and application performance reporting

Table 9 Virtual desktop configuration

Component	Description
Client desktops	Windows 7
Memory	2GB
Processor	2 vCPUs
Applications	Office 2010, Adobe Reader, WinZip®, Freemind, Doro PDF Writer

B Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Dell TechCenter](https://www.dell.com/techcenter) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware, and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC storage platforms.

B.1 VMware support

For VMware support, see the following resources:

- [VMware.com](https://www.vmware.com)
- [Education and training](#)
- [Online documentation](#)
- [VMware communities](#)