

3,500 Persistent VMware View VDI Users on Dell EMC SC7020 Storage

Dell EMC Engineering
July 2017

Revisions

Date	Description
December 2016	Initial release
July 2017	Updated for Dell™ Storage Center OS (SCOS) 7.2

Acknowledgements

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.

© 2016–2017 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.....	5
1 Introduction.....	6
1.1 Objectives.....	6
1.2 Audience.....	6
2 Choosing VDI storage	7
3 Solution architecture.....	8
3.1 Environment.....	8
3.2 Software.....	8
3.2.1 Software components.....	8
3.2.2 VMware vSphere host	8
3.2.3 VMware Horizon 7.1	8
3.2.4 Microsoft Windows 7	10
3.2.5 Microsoft Windows Server 2016.....	10
3.3 Hardware	10
3.3.1 Hardware components	10
3.3.2 Storage	10
3.3.3 Hosting environment.....	12
3.3.4 Networking.....	14
4 Testing.....	15
4.1 Test objectives.....	15
4.2 Test tools	15
4.2.1 Load generation.....	15
4.2.2 Load monitoring.....	15
4.3 Test phases	16
4.3.1 Boot phase.....	16
4.3.2 Login phase	17
4.3.3 Steady-state workload	20
4.4 Test results	21
5 Conclusion.....	23
A Environment	24

A.1	Horizon View.....	24
A.2	VMware hosts.....	24
A.3	VMware vCenter.....	24
A.4	Login VSI.....	25
A.5	Windows desktops.....	25
B	Additional resources.....	26
B.1	Technical support and resources	26
B.2	VMware support.....	26

Executive summary

This document provides a storage template for a 3,500-user persistent VDI environment running on a Dell EMC™ SC7020 storage array. It highlights the SC7020 with the release of Dell Storage Center OS (SCOS) 7.2. This release includes performance enhancements which increased the IOP capability and decreased IOP latency over previous versions by as much as 35 percent. Customers can upgrade to this release to take advantage of these performance gains.

The performance tests simulated 3,500 end-user VMs running Microsoft® Windows® 7 and a workload simulating knowledge workers. The target I/O workload, measured in I/O operations per second (IOPS), was 20 IOPS per virtual machine. This workload was chosen to show the ability of an SC7020 array to support the majority of small- to mid-sized workloads, even with demanding storage requirements. Even at 74,000 IOPS, the latency stayed below .7 ms. VDI environments are extremely heavy on writes, and this simulation shows the performance of an SC7020 can deliver with heavy write workloads.

Note that while the storage provided excellent performance, an appropriately sized compute environment is required as well.

1 Introduction

The SC7000 series is a Dell EMC midrange storage solution capable of supporting a broad array of workloads. It is dynamic, scalable, easy to manage, and supports a large number of client configurations, making it an excellent platform for VDI workloads. It is also a dense storage solution combining drives and controllers in a single unit, and supports a broad spectrum of solid-state and spinning drive options for great flexibility in storage sizing and performance.

For VDI workloads, solid-state drives (SSD) are a great choice. They provide very good performance in a very small form-factor, greatly increase performance density, and have an excellent cost/performance ratio. The new generation of high-performance read-intensive SSDs deliver much better performance and capacity per dollar. The cost per IOP is better than write-intensive drives, while providing a considerably better cost per GB.

For this solution, drive performance and cost are the primary factors. The capacity footprint per VM is small but a large number of read/write operations per second are required for a good user experience. VDI is one of the most storage-intensive workloads today, and latency is the most critical measure of storage performance. The SC7020 provides all of this in a small-footprint solution.

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

1.1 Objectives

The primary objective of the testing effort is to demonstrate the performance and scalability of the SC7000 series array. By providing a baseline for performance and scale, an SC7000 series array can be sized to fit any midrange storage requirement, providing low latency and excellent performance per cost. The key items tested are:

- IOP capability
- Latency under load
- Impact of compression and deduplication under load
- Design baseline for customers and partners

1.2 Audience

This reference architecture is targeted at anyone who wishes to understand the performance and scalability of the SC7020 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 in assets posted to 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 pre-allocation 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 pages 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 SC Series 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 they are 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. SC Series arrays tightly integrate with industry-leading hypervisors such as VMware® vSphere® and Microsoft Hyper-V®. 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 large 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 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

This solution is built on layers of software and hardware. Each component provides a key piece of the puzzle.

3.2.1 Software components

The solution presented in this paper uses the following software:

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

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

3.2.2 VMware vSphere host

VMware vSphere is the top enterprise virtualization platform 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.

3.2.3 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 three management servers configured, each controlling 875 virtual machines.

3.2.3.1 Linked clones

All 3,500 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 2 GB/VM in a persistent state before user login.

After the user profiles were created, the VMs each consumed approximately 2.8GB 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 workloads simulated by Login VSI.

The virtual desktops were permanently assigned to specific users, creating a persistent environment. All changes made by a user are retained for future sessions. The changes are stored in a delta VMDK stored with the VM.

All 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.

3.2.4 Microsoft Windows 7

Each of the desktops was a Windows 7 Enterprise client. Windows 7 32-bit was chosen for this configuration. All Windows and Office updates 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.

3.2.5 Microsoft Windows Server 2016

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

3.3 Hardware

A fully functional VDI environment for 3,500 virtual machines requires a large amount of hardware with significant compute and memory resources. The act of virtualizing does realize great efficiencies but it cannot eliminate the requirement to process data. The act of virtualizing collapses the environment into dense clusters of resources.

3.3.1 Hardware components

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) switch
- Brocade® M5424 Fibre Channel (FC) blade modules
- Brocade 6505 FC ToR switches
- Dell EMC SC7020 storage array

There are many sizing tools to gauge the processor and memory requirements for virtual machines. They are outside the scope of this document. See the [VMware community sizing guidelines](#).

3.3.2 Storage

The storage used was an [SC7020 array](#) with read-intensive solid-state drives. The high performance read-intensive drives provide excellent performance at 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.

Technical highlights of the SC7020 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/or 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.

The drives used were 1.6 TB drives, providing usable space of 1.46TB per drive. The system contained 30 drives with two dedicated spares. This provided 40.75TB 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 allows 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 Storage SC Series Best Practices with VMware vSphere 5.x-6.x](#).

3.3.3 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 3,500 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 MXL10Gb dual-port Ethernet module
Virtual desktop configuration	2 vCPUs, 2GB RAM, VMXNet3, 30GB volume

Table 2 Virtual machine configurations

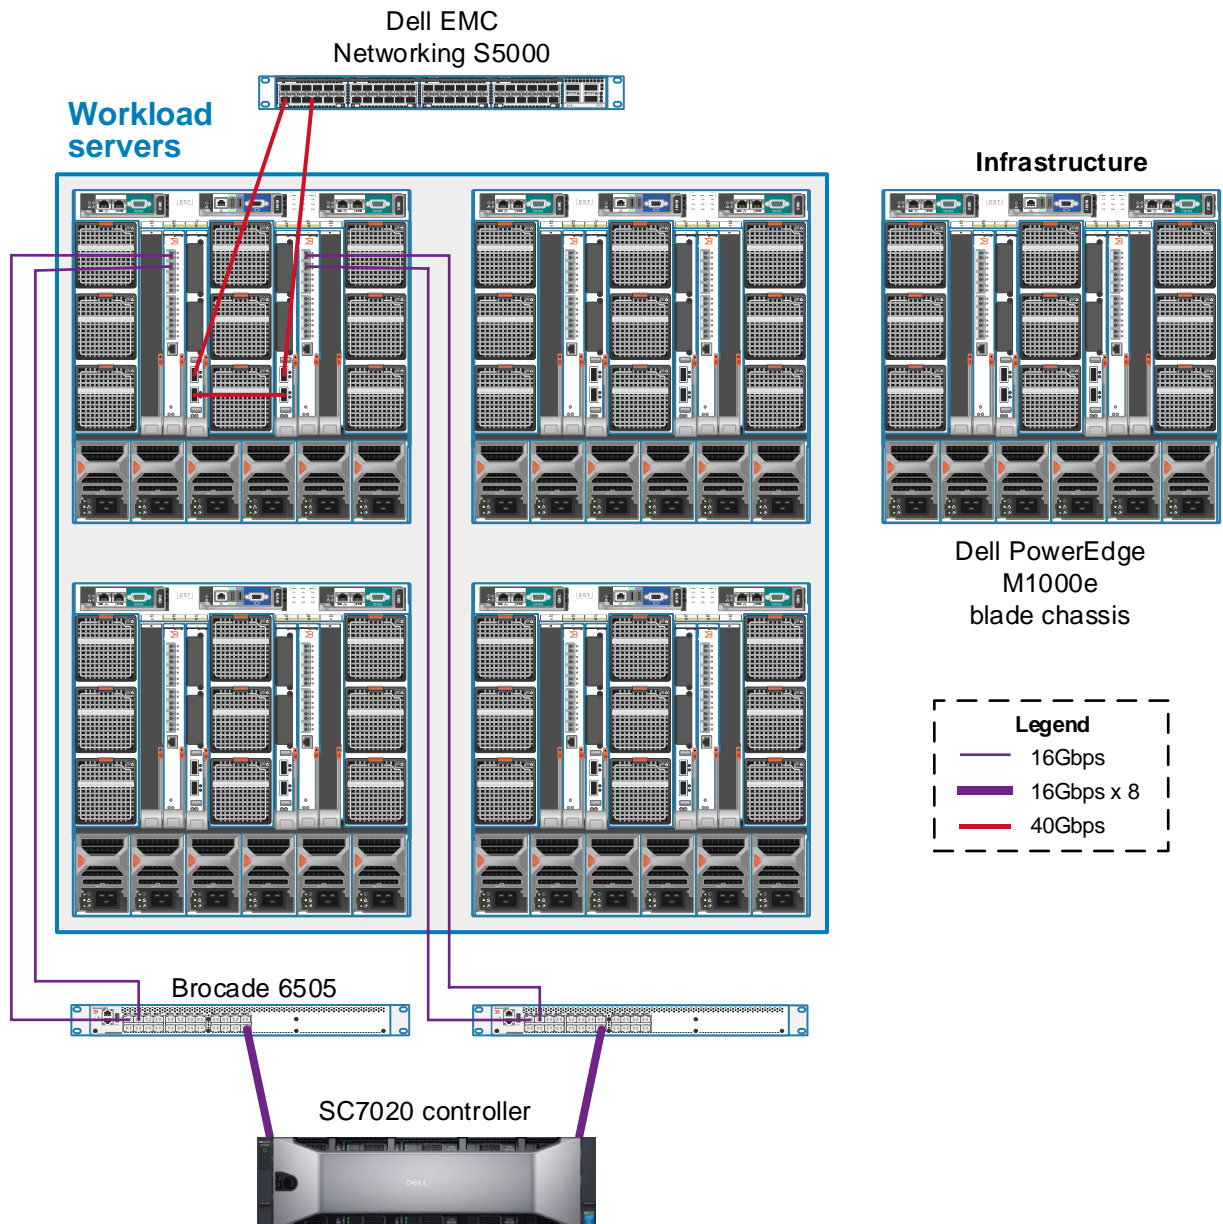
Component	Description
Infrastructure VMs	4 vCPUs, 8GB RAM, VMXNet3 vNIC, Login VSI 4.1.25 agent
Operating system	Windows Server 2016
vNIC	VMXNet3
Applications	Microsoft Active Directory®, DNS, DHCP, vCenter Servers, Horizon Connection and Composer servers, Microsoft SQL Server®, Dell Storage Manager (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
SC7020	Dual storage controllers, 30 drive bays
16Gb 8-port Fibre Channel	QLogic 16 Gb 4-port Fibre Channel adapters, 1 per controller head
SCOS 7.2	Storage firmware
Drives	30 x 1.92TB: 28 active, 2 spares



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 4 Gb/sec from the file share. The Login VSI share server was a PowerEdge M620, with 144GB of RAM and 2 x 10 Gb/sec Intel Ethernet modules.

3.3.4 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 and data-reduction capabilities of the SC7000 series controllers. The tests were scaled to run within the limits of the SC7020 controllers and highlight their performance.

The data-reduction components tested are not specific to the SC7000 series arrays. They are common to several SC Series platforms. The numbers shown in the test phase section are indicative of what can be expected from an SC Series system with data reduction enabled. The performance numbers will vary but the reduction numbers should be consistent.

4.2 Test tools

4.2.1 Load generation

The load generation tools used were created by Login VSI consultants. These are the 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 chart real-time I/O statistics, 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 pane for replication on SC Series arrays.

4.3 Test phases

4.3.1 Boot phase

Figure 1 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 vSphere 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.




Figure 1 Boot storm

As the chart shows, over 150,000 IOPS were generated during the time it took to boot all 3,500 VMs. The time required to boot all VMs to the login screen was less than 5 minutes. This workload is uncommon because most customers very rarely reboot an entire environment. Even under this heavy load the latency was just .75ms.

4.3.2 Login phase

This test is a much 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 2 shows the login phase.




Figure 2 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 SC7020 provides excellent performance with predictable latency as the load increases.

Figure 3 shows the peak load generated.




Figure 3 Peak login load

This chart shows just under 74,000 IOPS generated at just under .7 ms of write latency. Latency this low ensures the storage will not be the gating factor in the user experience.

The SC Series array optimizes for write performance by utilizing multiple RAID levels simultaneously. This provides both speed and performance.

Figure 4 shows the number of sessions successfully launched.

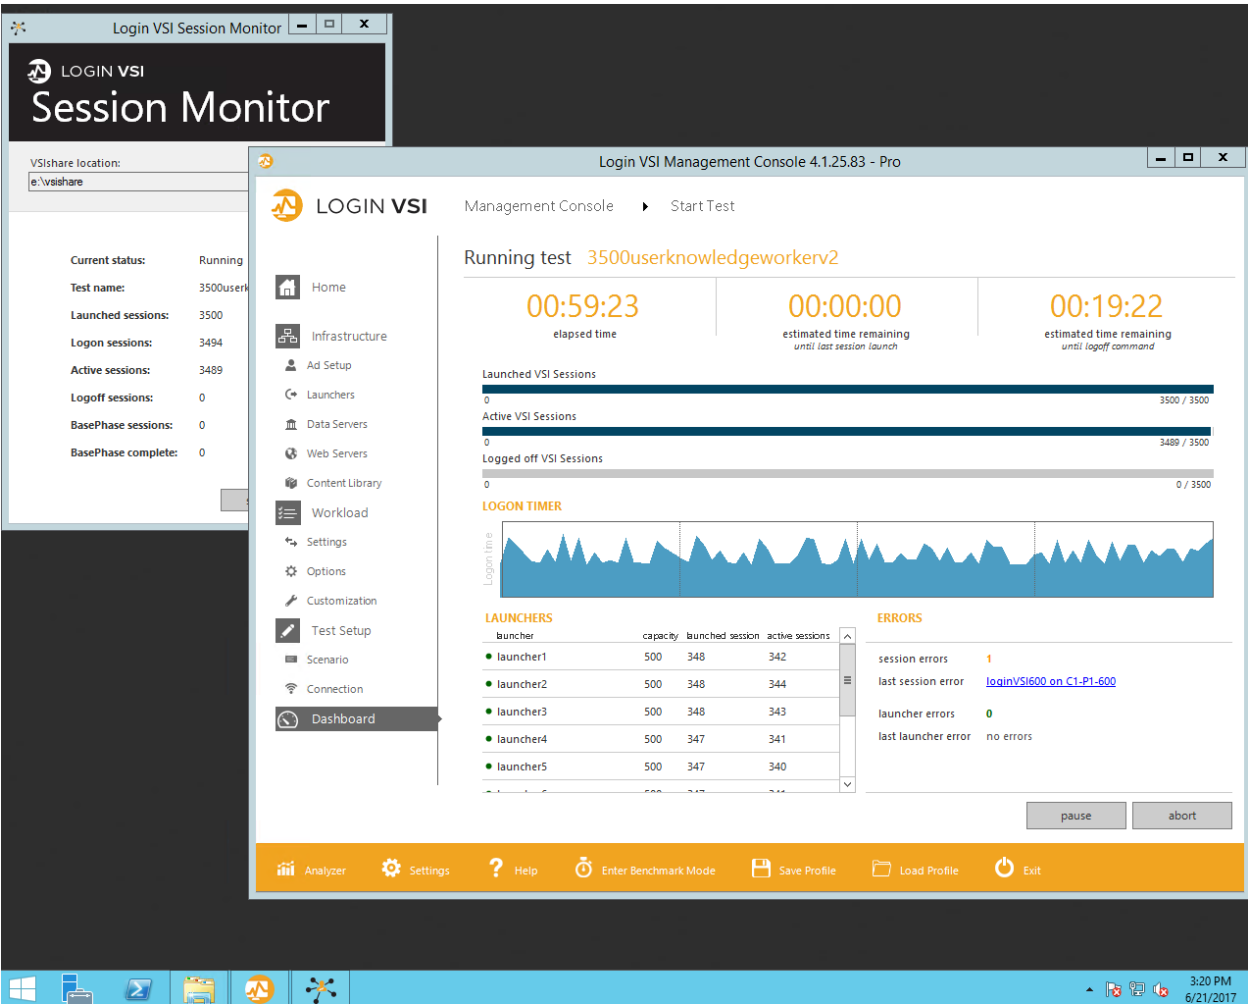


Figure 4 Sessions launched

4.3.3 Steady-state workload

The steady -state workload used for these tests was targeted for 20 IOPS per user to simulate a heavier-than-normal workload per user.

Figure 54 shows the load generated.




Figure 5 Steady-state workload

As Figure 4 shows, the constant workload generated 70,000 IOPS at under .7ms. Being able to sustain such a low latency at almost 20 IOPS per user proves that the SC7020 can provide fast, consistent performance for small to medium businesses.

4.4 Test results

Figure 6 shows the full workload from start to finish.

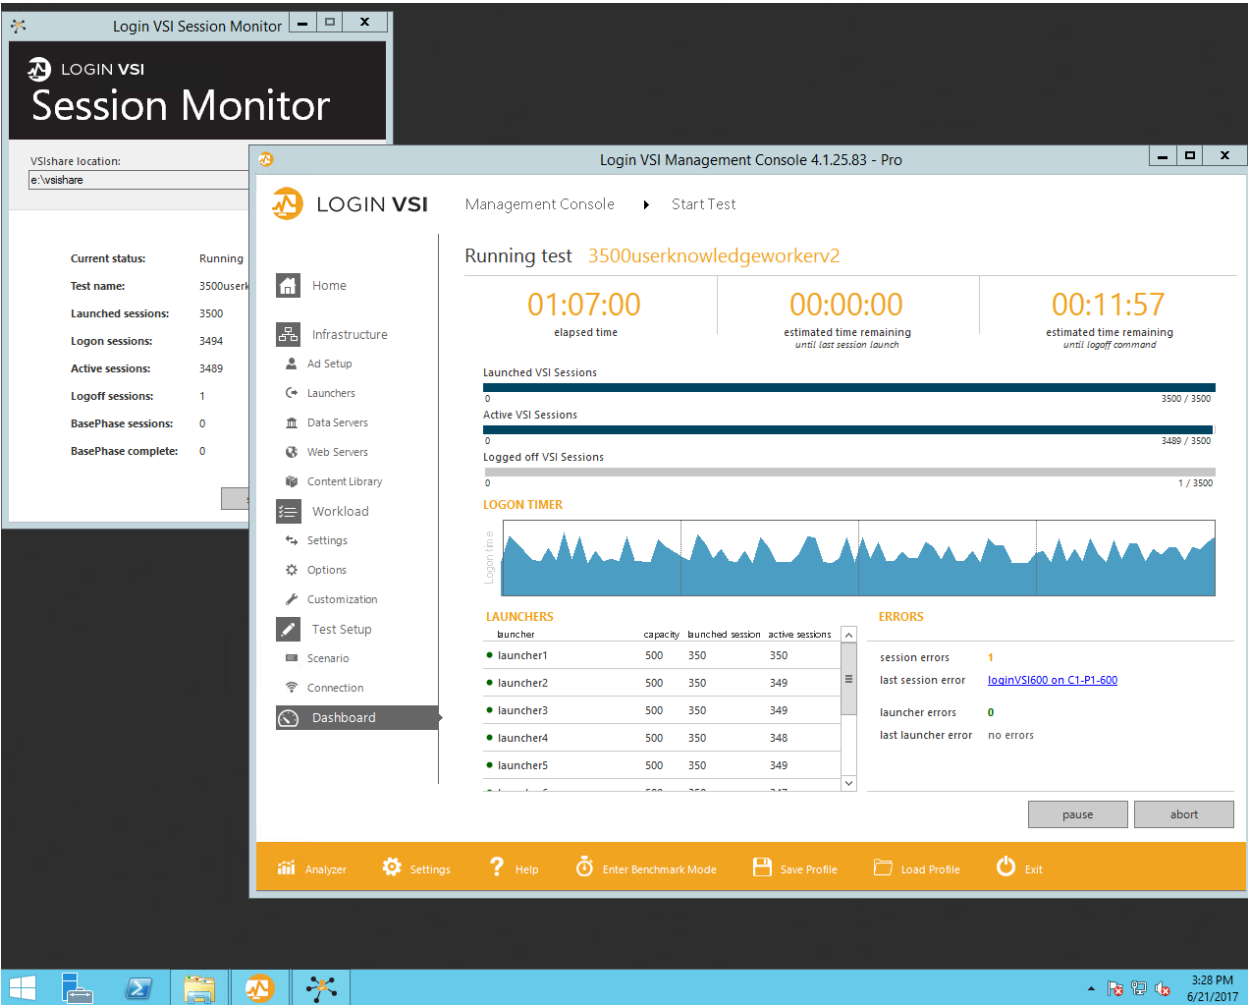


Figure 6 Login VSI workload

With 3,500 sessions configured, the response times are shown in Figure 7. This chart is the output from the Login VSI test tool, showing the login and application response times. As indicated in the chart, the threshold is 1,993 ms with an average latency provided of just 1,492 ms. The response time is an average of time to complete given tasks, such as open a new document, zip a file, or print an open file.




Figure 7 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 SC7020 storage array to be an excellent platform for up to 3,500 VDI knowledge worker users. It provides excellent performance and scalability, with up to 2PB of usable, configurable storage in a reliable and stable platform.

A Environment

A.1 Horizon View

Table 5 Horizing 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

A.2 VMware hosts

Table 6 VMware host hardware configuration

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

A.3 VMware vCenter

Table 7 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

A.4 Login VSI

Table 8 Login VSI configuration

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

A.5 Windows desktops

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 Additional resources

B.1 Technical support and resources

Dell.com/Support is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell 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.2 VMware support

For VMware support, see the following resources:

- VMware.com
- [Education and training](#)
- [Online documentation](#)
- [VMware communities](#)