

Dell EMC SC Series: 3,000 VMware Horizon Linked Clone VDI Users

Dell EMC SC7020F all-flash array

Abstract

This document provides a reference architecture for running 3,000 VMware Horizon linked clone VDI users on a Dell EMC SC7020F all-flash array.

October 2019

Revisions

Date	Description
October 2019	Initial release

Acknowledgements

This paper was produced by the following:

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.

Copyright © 2019 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. [10/4/2019] [Reference Architecture] [4735-RA-V]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
Audience	5
1 VDI design considerations.....	6
1.1 Linked clone pools.....	6
1.2 Full clones.....	6
1.3 Instant clones.....	6
1.4 Persistent and non-persistent desktops	6
1.5 Storage pools.....	7
1.5.1 RAID configurations.....	7
1.6 Validating the storage design	7
1.6.1 Validating the I/O path	7
2 SC7020 overview	9
2.1 Intuitive management	9
2.2 Storage Center OS	9
2.3 Self-optimizing, automated storage.....	9
2.4 Space savings in the data center	9
2.5 Data center automation with leading technology partner integrations	10
2.6 Flexibility for now and in the future.....	10
2.7 Robust SAN with feature-rich software	10
2.8 Eliminate the expense of managing separate systems.....	11
2.9 SC Series expansion enclosures.....	11
3 VDI test configuration	12
3.1 Block storage configuration	12
3.1.1 Pools.....	12
3.1.2 Volumes.....	12
3.1.3 VMware datastore.....	12
3.2 Improving VM performance	13
3.2.1 Memory.....	13
3.2.2 Host-based caching.....	13
3.2.3 VMware OS optimization	13
4 Testing results	14

Table of contents

4.1	VDI performance data	14
4.1.1	Boot phase.....	15
4.1.2	Login	15
	Summary.....	19
A	Environment	20
A.1	Horizon View.....	20
A.2	VMware hosts	20
A.3	VMware vCenter	20
A.4	Login VSI	21
A.5	Windows desktops.....	21
A.6	Hardware components	21
A.7	Logical diagram	22
B	Technical support and resources	23
B.1	VDI resources	23

Executive summary

The Dell EMC™ SC Series is an excellent platform for clustered storage. This document describes the performance and integration characteristics of an SC Series array running a VMware® Horizon® workload. This example has 3,000 users at 40 IOPS per user, which represents a heavy workload that is well above a typical user environment in terms of disk traffic.

The very low latency, high-throughput load generated shows the performance capability of a Dell EMC SC7020F array, which is capable of supporting a substantial number of users on a single array. The SC7020F model tested for this document is an all-flash array designed to achieve maximum performance from a dense storage design.

Dell EMC SC7000 series arrays are block storage devices that provide an exceptional number of IOPS per dollar. They are designed to support midrange customers looking for excellent performance in a block storage device.

Audience

This document is intended for storage administrators, virtualization users, architects, partners, and anyone responsible for configuring SC Series systems. Some familiarity with Dell EMC storage systems is assumed.

We welcome your feedback along with any recommendations for improving this document. Send comments to StorageSolutionsFeedback@dell.com.

1 VDI design considerations

VDI workloads are some of the most I/O-intensive workloads supported by a storage array. The exact I/O load for VDI deployments is determined by the number of concurrent users and applications. The capacity requirements for VDI tend to be small compared to the throughput requirements, which is especially true for linked clone deployments.

1.1 Linked clone pools

Linked clones are commonly used in VDI deployments due to their ease of management, small footprint, and consistency. Since a base image is used for deployment, all desktops are consistent. When upgrades are required, a new image is deployed to upgrade virtual machines (VMs) quickly and easily.

One of the tradeoffs of linked clones is reduced customization without additional tools. To retain the ease of management and reduced footprint, all VMs are mirrored to a base image. While persistent desktops can be enabled, this branches the VM images and makes VM replacement more challenging.

For a large number of identical desktops, linked clones are the most efficient method and one of the most common.

1.2 Full clones

VMware Horizon does support managed full clones. These have a much larger footprint and require individual updating, or redeployment, upon update. They also require far more storage space. This can be offset by deduplication, but this has a performance impact and typically cannot match the capacity savings of linked clones.

1.3 Instant clones

VMware Horizon 7 has support for Instant Clone Technology. This feature creates VMs as needed from the base image using vmFork technology. This creates a machine on demand from in-memory data, and writes the data to disk as a background process. This process is fast and easy to manage but requires additional components and configuration for a persistent user experience. This is becoming more common in VDI and should be considered for certain scenarios.

1.4 Persistent and non-persistent desktops

Linked clone desktops can be created in a persistent manner. This allows users to save customizations to their desktops and keep files on the desktops of their VMs. This provides users a more customized environment but increases the cost of management. Persistent desktops cannot be replaced without a method to retain user files and settings.

Non-persistent desktops have a far lower management cost. If a desktop becomes corrupt, there typically is no intervention required to repair it. When a user logs out, the desktop is reset to the base image automatically. This ensures users will always have a clean desktop to work with. However, the tradeoff is an increased storage load. Each time a user logs out, the changes are discarded and a replacement change file must be created. The resulting increase in I/O should be accounted for.

1.5 Storage pools

In general, it is recommended to use fewer storage pools within SC Series arrays because this reduces complexity and increases flexibility, and usually improves performance. Dell EMC recommends using a single disk pool for hosting VDI desktop volumes. This provides better performance by leveraging the aggregate I/O bandwidth of all disks to service I/O requests from the VMware ESXi™ hosts. A single drive pool also allows all hosts to share the volumes enabling VMware vSphere Storage vMotion®. This enables hosts to go into maintenance mode and VMs to be moved without taking an outage. This greatly simplifies host updates and patches.

A single storage pool also allows all storage to be available for all services. Since snapshots and replication use the same pool, having all storage available in a single pool means all services have access to the entire storage array. If there are requirements for latency or bandwidth thresholds to be defined, QoS and latency priorities can be defined per volume.

1.5.1 RAID configurations

By default, SC Series storage uses the RAID 10 protection level when creating new pages within a storage pool. The I/O requirements for VDI desktops need to be clearly defined to size the storage correctly. As data becomes less active, it is moved to RAID 5 or 6, depending on the activity level.

The small footprint and high I/O density typically allows a smaller drive size, reducing drive rebuild times. This means RAID 5 is preferred over RAID 6 in most cases. The choice of RAID 5 or 6 is based on drive size and rebuild time requirements.

1.6 Validating the storage design

VDI workloads can be challenging to simulate. They consist of hundreds to thousands of virtual machines, making it difficult to recreate the load. There are few tools that can accurately simulate this.

Login VSI is the tool used for this VDI design. It allows controlling a large number of virtual machines in a consistent manner.

To accurately simulate the storage load, the following needs to be performed:

- Create the base image to match the intended configuration
- Simulate the correct user count
- Run the applications as a user would
- Determine the optimum number of volumes based on recovery requirements and management overhead
- Create a baseline for the number of I/Os and GBs per user to determine the storage requirements

1.6.1 Validating the I/O path

VDI workloads tend to be fairly small block. The average I/O size is between 24k and 32k. The number of I/Os generated are fairly high due to the number of users per volume.

Spreading users across more volumes typically does not affect SC Series array performance. Using 20 to 30 volumes can support 3,000 users. For a smaller number of users, fewer volumes can be used. More volumes require additional management effort, but spread the I/O across more LUNs for better performance. This also results in smaller volumes for recovery or replication purposes.

VMware guidelines typically suggest using fewer users per volumes. With snapshot-based protection, the number of VMs is less of a concern because recovery times change very little based on a larger user count per volume.

Because of the small I/O size of VDI workloads, bandwidth is typically not a limiting factor. With modern 16Gb, or with higher Fibre Channel or 10Gb+ Ethernet architectures, the number of ports is less of an issue. Four ports per controller and two controllers per storage appliance can easily handle 3,000 VDI users.

2 SC7020 overview

2.1 Intuitive management

The SC7020 storage array has the control points to streamline day-to-day tasks, smooth major migrations, and keep all data fully available and optimized. Such features include convenient web management from a mobile device using Unisphere™ for SC Series, cloud-based analytics and monitoring with CloudIQ, and robust client-side software with Dell™ Storage Manager (DSM).

2.2 Storage Center OS

Dell Storage Center OS (SCOS) gives diverse environments the ideal access point to deploy uniquely adaptable, future-ready storage solutions. Capabilities include the following:

- Over 10:1 capacity savings with intelligent deduplication and compression
- Federated data control with Live Migrate
- Seamless business continuity with Live Volume auto-failover between arrays
- Federated management and cross-platform replication between SC Series and PS Series arrays
- Quality-of-Service (QoS), VMware vSphere Virtual Volumes™ (vVols) support, iSCSI Datacenter Bridging (DCB)

See the product specification sheets and product documentation on Dell.com for additional features and specific platform support.

2.3 Self-optimizing, automated storage

Based on actual use and application performance requirements, SC Series Data Progression automatically optimizes drive tiers and RAID levels for performance and cost savings. By default, all new data is written to fastest tier 1 drives at RAID 10 performance levels, then converted to economical RAID 5/6 for subsequent reads.

As data ages, it is moved to less expensive storage until it becomes active again. All-flash and hybrid arrays can leverage the unique attributes of write- and read-intensive SSDs in combination. Data is placed in the ideal location, providing the perfect combination of high IOPS and cost savings.

2.4 Space savings in the data center

SC Series storage delivers unprecedented levels of platform efficiency, along with comprehensive data reduction, as well as the following benefits:

- Conservation of rack space for future growth
- Smaller footprint, which benefits power and cooling expenses
- Intelligent deduplication and compression for all storage tiers and drive types

Operating costs are reduced with advanced features like volume-level QoS, vVols support, and departmental charge-back reporting, allowing instant visibility and control of a multi-terabyte, multi-location environment.

2.5 Data center automation with leading technology partner integrations

SC Series arrays feature integration points with VMware, Microsoft, Oracle, OpenStack, IBM, CommVault, Symantec, Foglight, and more. The VMware vSphere plug-in, vCenter™ SRM adapter, and VAAI support help manage storage and virtual machines together.

Application Protection Manager for Oracle , which was enhanced for the Linux® environment, allows creating application-consistent snapshots (replays) for Oracle® databases. Combined with SC Series snapshots, it ensures reliable applications, and consistent backup and recovery of Oracle and Microsoft Volume Shadow Copy (VSS) enabled applications (Exchange, SQL and Hyper-V) and VMware VMs.

2.6 Flexibility for now and in the future

SC7020 combines more capacity and greater flexibility together in one system to simplify management and lower costs. The SC7020 array offers the following:

- Scalable IOPs with block data
- Over 10:1 capacity savings with intelligent deduplication and compression
- Over 60 percent reduction in rack space required
- Optimized performance at an affordable price
- Support for Fibre Channel and iSCSI
- Easy management with Dell Storage Manager (DSM)

The SC7020 array enables consolidating multiple types of workloads into a single array, and helps achieve agile performance and a rapid return on investment. SC7020 storage lowers the cost per terabyte, for all array types, including hybrid or all flash options, and uses your current architecture.

Expand with your business: Grow when and how business dictates with seamless scalability up to 500 drives.

Future-proofed data: Support various array options including flash, hybrid, multi-protocol connectivity support, data protection, ProSupport Powered by Copilot Services, and Dell EMC Deployment Services.

Optimized resources: Seamlessly manage storage with a single pane of glass across the federated family of SC Series and PS Series.

2.7 Robust SAN with feature-rich software

The SC7020 offers a full suite of features to meet a variety of workloads. The automated intelligence of Data Progression optimizes drive tiers and RAID levels, while Live Volume maintains seamless business continuity through auto-failover between arrays. Features include the following:

- Federated data control with Live Migrate
- Guaranteed performance with QoS, application by application
- Data protection with snapshots (replays)
- Dynamic Capacity, which separates storage allocation from utilization, optimizing capacity
- Unified management of block and file storage and cross-platform replication between SC Series and PS Series
- vVols and iSCSI DCB support

2.8 Eliminate the expense of managing separate systems

DSM provides comprehensive monitoring and management of all local and remote SC Series arrays. Easy-to-use data replication and migration between platforms gives PS Series customers flexibility to add SC arrays to their environment when desired.

DSM streamlines SAN administration and reduces operating costs with advanced features like volume-level QoS, vVols support, and departmental charge-back reporting. PowerShell and REST management APIs allow integrating the SC7020 into business workflows.

2.9 SC Series expansion enclosures

Supporting up to 500 total drives and 3 PB raw capacity in a single array, the SC7020 array offers five supported expansion enclosures that can be mixed and matched depending on specific performance, cost, or density needs.

Options are available for ultra-dense enclosures (84 drives in 5U) with the SC280 and high-performance 12Gb SAS SC400 and SC420 enclosures. Adding a new enclosure is simple — the new drives can be automatically assimilated into the existing virtualized pool or set up as their own standalone pool. Data-in-place upgrades make it easy to change controllers and other key components, or migrate from legacy SC Series arrays.

See the [SC Series expansion enclosure data sheet](#) for more details.

3 VDI test configuration

This section describes the VDI test configuration used for this document. There are many options for running VDI workloads on SC Series hardware. A linked clone configuration was chosen due to the high adoption rate among Dell EMC customers.

3.1 Block storage configuration

Block storage is used for this VDI deployment. NFS can be used for VM storage since ESXi supports NFS-based datastores, but there was no benefit to using file-based storage for these tests.

3.1.1 Pools

The SC Series architecture is built around pools of disks. The system is designed to efficiently use all available storage, without reserving disk storage for specific functions. All snapshots and volumes share the entire pool of storage. This ensures storage capacity and performance are used as efficiently as possible. Any new storage added is immediately available for volumes, replications, and snapshots.

The storage is intelligently striped at the page level to improve performance and efficiency. This is done without user interaction. The system determines what data can be stored in a more efficient RAID type, and what data requires high-performance RAID.

SC Series storage typically performs best with a single disk folder. This allows all volumes to share the performance of all available drives, and makes the most efficient use of the space. If disk contention may be an issue, QoS can be employed to control resource allocation, or multiple disk pools can be created.

A disk folder can contain disks of multiple tiers of performance and capacity. This enables data to be placed on the most effective tier of storage, whether for best performance or cheapest capacity.

A single disk folder is created and the recommended storage profile is used. This is the configuration most environments would use and provides the best performance in most scenarios.

3.1.2 Volumes

The volume size used is 4 TB. This allows VMs to grow as needed and VMware Horizon to detect the volume space for VM growth. Horizon monitors drive capacity available and prevents VM placement on a volume that it deems to be too full. This is a safety feature to prevent over-allocation to the point of failure. There are settings to allow more aggressive VM placement, but these require administrators to closely monitor volume utilization to ensure volumes do not fill up, or to have an automated process to grow volumes before they reach capacity.

SC Series arrays can dynamically expand volumes quickly and easily if needed. This ability can prevent Horizon from halting VM creation if a volume is approaching capacity.

3.1.3 VMware datastore

The VMFS 6 file system was used for all datastores. This is the default for VMware ESXi 6.7. The block size was 1 MB. This allows Horizon to dynamically reclaim blocks as they are released from VMs. Automatic space reclamations keep the volumes and disk pool from consuming excess space when VMs are torn down.

3.2 Improving VM performance

3.2.1 Memory

Virtual machine performance is greatly affected by the quantity of memory allocated to each VM. It is the most effective way to improve performance in virtual machines, but it can be difficult to scale to large environments because of the aggregated quantity. For the purposes of this scenario, 3 GB per VM is allocated.

Decreasing allocated memory results in more storage load, and this impact on storage should be tested. Creating a baseline for a single VM gives a good indication of storage load per VM. This provides a starting point for I/O load, but many other factors affect the total.

3.2.2 Host-based caching

The host configuration can change the load generated to the storage. One of the options is content-based read cache, enabled through Horizon View Storage Acceleration. This option uses host RAM to accelerate reads during boot events. This can reduce volume reads and improve boot times for virtual machines. The tradeoff is increased host memory load. On hosts that are memory-constrained, this may not be an option. For those that can spare memory, this can improve the boot times of the virtual machines when large numbers of machines are requesting the same data.

3.2.3 VMware OS optimization

VMware offers a free tool to optimize virtual machines. It is a profile-based optimization tool that can greatly reduce the resource requirements for virtual machines.

Login VSI has created a profile for VM optimization that reduces memory, disk, and CPU requirements to allow a greater number of virtual machines per host. It can have a negative impact on some applications so testing is required. The overall impact is to considerably reduce host requirements for virtual machines allowing a much greater number per host.

The Windows® 10 Login VSI profile was used to optimize tested machines. Find more information in the Login VSI article [Optimizing desktop images with VMware OSOT](#). Login VSI provides guidance on how the profile changes the behavior of virtual machines when applied.

Testing should be done once an optimization profile has been applied to ensure users are still able to use the desktops. The modifications favor performance and may require tuning to allow all applications to function correctly.

Testing for VDI load requires many interconnected systems and processes. A test harness needs to connect to all of the VMs and generate a realistic load. The load needs to perform at a level deemed acceptable to users.

4 Testing results

There are two phases that a VDI environment spends significant time and resources supporting: the login and steady state phases. These are the primary concerns for administrators since they happen every work day. These two phases must perform well for an environment to be user friendly.

The boot phase is far less common and only happens in specific circumstances. It is far less important because it normally occurs outside business hours and has very little user impact.

The environment used for the testing is documented in appendix A.

4.1 VDI performance data

The profile used is the default knowledge worker profile. It covers the majority of the core applications deployed and generates a balanced workload. The following information from the [Login VSI website](#) lists details of the workload.

The knowledge worker profile is designed for 2(v)CPU environments. This is a well-balanced intensive workload that stresses the system smoothly, resulting in higher CPU, RAM, and I/O usage.

The knowledge worker workload uses the following applications:

- Adobe® Reader
- Freemind/Java
- Microsoft Internet Explorer
- Microsoft Excel
- Microsoft Outlook
- Microsoft PowerPoint
- Microsoft Word
- Photo Viewer

This profile also uses native Windows applications (Notepad and 7-Zip) to complete the print and zip actions used by the workload meta functions.

By using a standardized profile, the workload generated can easily be duplicated or verified in any environment.

4.1.1 Boot phase

The time required to boot up all VMs is less critical since it rarely happens in a production environment. It would typically only happen during a complete system outage, or a complete image replacement for all VMs.

The performance during the boot phase is shown in Figure 1.




Figure 1 Performance during boot phase

The total IOPS achieved are just under 300,000, with an average latency of under 800 microseconds.

4.1.2 Login

The login cycle is set to one login per second until all VMs are logged in. Once VMs are logged in, they start the workload, so the login load overlays the steady-state workload until all users are running the workload.

The following sections detail the various phases of the test run.

4.1.2.1 Login ramp

Figure 2 shows how the load increases as users log in. This shows a fairly predictable increase in load and latency as the user activity increases.




Figure 2 Performance during login phase

The maximum load during the login phase reaches 175,000 IOPS. The average I/O size during this phase is 16KB, and the average latency is just under 800 microseconds.

4.1.2.2 Steady workload

This section details the performance once the environment achieves a steady workload state. This is the most important phase of the testing and the most common production state. The user experience is defined by how the storage performs in this state. Figure 3 shows the performance level of an SC7020 array.



Figure 3 Performance during steady state

Under full user workload, the I/O load generated is 164,000 IOPS. This equates to just over 54 IOPS per user. Even at this load, the latency is only 800 microseconds on average.

This is a very heavy user load and shows how well an SC7020 array can perform under a heavy load.

4.1.2.3 Login VSI test results

The Login VSI test results detail how well the environment performs under user load. The tests are written to simulate a variety of real-world user applications in a realistic work pattern. The summary shown in Figure 4 combines the results of the applications into an overall user experience chart.

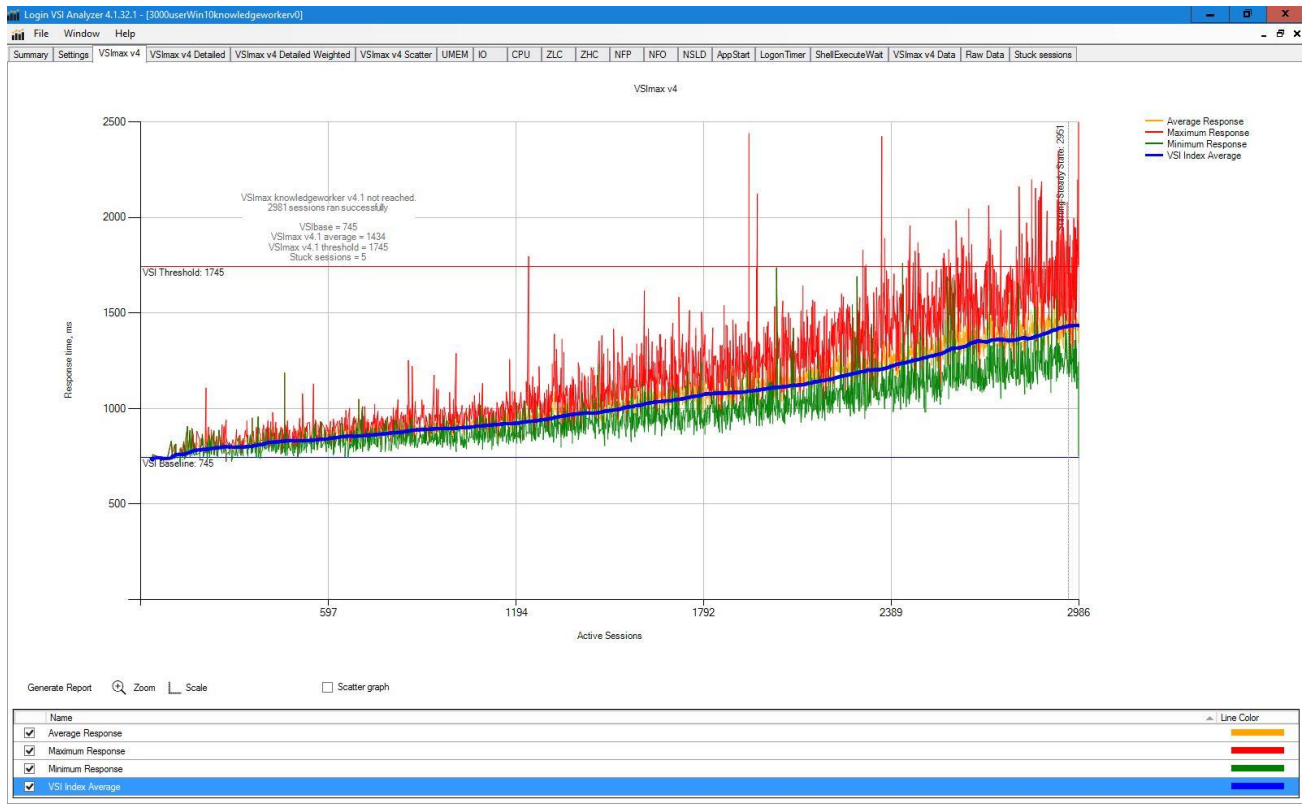


Figure 4 Testing summary

Summary

The testing described in this document shows that in a VDI deployment, the SC7020 array has performance to spare even under a heavy load. The additional performance available allows maintenance to be run or additional users to be added without compromising performance.

The user load of 3,000 VMs at 54 IOPS per user is conservative for this platform. There is considerable performance headroom available, allowing additional users to be added, or other workload types to run on the same array. This flexibility allows storage arrays to have multiple uses and not be dedicated to only VDI.

A Environment

A.1 Horizon View

Table 1 Horizon View configuration

Component	Description
Horizon View	7.8
Role	Virtual machine assignment, management, web interface, connection broker, Composer
Servers	Four 4-core servers, 16 GB RAM each, 2 Composers, 2 Connection servers
Cabling	10Gbps VMXNet3

A.2 VMware hosts

Table 2 VMware host hardware configuration

Component	Description
VMware host	6.7u2
Role	End-user compute node
Hardware	Two 18-core processors, 576 GB RAM
Storage	200 GB SSD for boot/log

A.3 VMware vCenter

Table 3 vCenter server configuration

Component	Description
VMware vCenter appliance	6.7, 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 4 Login VSI configuration

Component	Description
Login VSI	4.1.32
Role	Workload generation and application performance reporting

A.5 Windows desktops

Table 5 Virtual desktop configuration

Component	Description
Client desktops	Windows 10
Memory	3 GB
Processor	2 vCPUs
Applications	Microsoft Office 2016, Adobe Reader, WinZip®, Freemind, Doro PDF Writer

A.6 Hardware components

Component	Description
Fibre Channel switches	Two Brocade® 6510 switches for dual fabric
Network switch	Two Dell EMC Networking S5148F switches
Servers	Dell EMC PowerEdge™ R640, two Intel® Xeon® Gold 6154 CPU @ 3.00 GHz, 576 GB RAM
Network adapters	Dual 10Gb Ethernet adapters built in
Fiber-channel adapters	QLogic® 2692 dual port 16Gb

A.7 Logical diagram




Figure 5 VDI logical diagram

Component	Description
SC7020	Storage controller, dual controllers
16 ports x 16Gb Fibre Channel	QLogic 16 Gb 4-port Fibre Channel adapters, 2 per controller head
Firmware 7.4	SC Series firmware
Drives	48 x 1.92 TB

B Technical support and resources

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

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

B.1 VDI resources

There is a large amount of SQL Server information available online. While not complete, the following list contains several useful links.

Table 6 Additional VDI resources

Description	Link
Dell VDI Storage Solutions	https://www.dell.com/en-us/glossary/vdi-storage.htm
VMware Horizon home page	https://www.vmware.com/products/horizon.html
Dell Technologies VDI	Dell Client Computing resources