



Citrix XenDesktop VDI with Dell Storage SC4020 All-Flash Arrays for 1,800 Persistent Desktop Users

Dell Storage Technical Solutions
August 2014

Revisions

Date	Description
June 2014	Initial publication – Damon Zaylskie, Chhandomay Mandal
August 2014	Minor edits

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT:

<http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Boom!™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of contents

Revisions.....	2
Executive summary	5
1 Introduction.....	6
1.1 Objectives.....	6
1.2 Audience.....	6
2 VDI with Compellent all-flash storage arrays.....	7
2.1 Flash-optimized data progression.....	7
2.2 Dynamic capacity.....	8
2.3 Hypervisor integration.....	8
3 Solution architecture.....	9
3.1 Software.....	9
3.1.1 Citrix XenDesktop	9
3.1.2 VDI Hypervisor platform: VMware vSphere 5.5	14
3.2 Hardware	15
3.2.1 Storage.....	15
3.2.2 Hosting environment	17
3.2.3 Networking.....	19
4 Test methodology	21
4.1 Test objectives.....	21
4.2 Test tools.....	21
4.2.1 Load generation	21
4.2.2 Monitoring tool.....	22
4.3 Test phases.....	22
4.3.1 Boot storm	22
4.3.2 Login storm	23
4.3.3 Steady state	23
4.4 Test criteria.....	23
4.4.1 Storage throughput and latency	23
4.5 Test configuration.....	24
4.6 Test results	25
4.6.1 Boot storm	25



4.6.2 Login storm	26
4.6.3 Steady state	27
5 Conclusion	28
A Test specification components	29
A.1 XenDesktop	29
A.2 vSphere Hypervisor	29
A.3 VMware vCenter	29
A.4 LoginVSI	29



Executive summary

Desktop virtualization is an important strategy for organizations seeking to reduce the cost and complexity of managing an expanding variety of client desktops, laptops, and mobile handheld devices. Virtual Desktop Infrastructure (VDI) offers an opportunity not only to reduce the operational expenses for desktop management and provisioning, but also to improve user mobility and data security.

A VDI deployment can place high performance and capacity demands on the storage platform. For example, consolidating large amounts of inexpensive stand-alone desktop storage into a centralized infrastructure can create tremendous capacity demands on centrally managed shared storage used in VDI deployments. Performance demands are determined by the number of I/O operations per second (IOPS) generated by basic desktop client operations such as system boot, logon and logoff, and by desktop usage operations from different users. Storm events such as morning logons and afternoon logoffs by many users at approximately the same time can cause I/O spikes that place high performance demands on the storage infrastructure. Moreover, IOPS generated per desktop can vary greatly depending on the user type.

A cost-effective, appropriately-sized storage platform is critical for VDI deployment success. Dell has recently introduced a new, small footprint Dell Storage SC4020 array that leverages proven Dell Compellent technology to provide an automated, self-tuning storage system capable of efficiently scaling to satisfy VDI production level workloads. Specifically, SC4020 uses Flash-optimized Data Progression to tier data across high-endurance, high-performance, write-intensive, single-level cell (SLC) SSDs and higher capacity, lower cost, read-intensive, Multi-level cell (MLC) SSDs. This significantly drives down the cost of each desktop while providing the required high performance in VDI deployments.

This paper highlights a 1,800 user persistent desktop VDI architecture using Citrix XenDesktop 7.5 with Citrix Machine Creation Services (MCS) and Citrix Personal vDisk technology. The solution leverages a single SC4020 all-flash array with both write-intensive SLC SSDs and read-intensive MLC SSDs to deliver very high performance (125,000 IOPS with less than 7 ms latency) at a low cost for each desktop. Details are provided for the storage I/O characteristics under various VDI workload scenarios such as boot and login storms as well as steady states with significant loads per user (20 IOPS for each desktop).

The presented persistent desktop VDI storage solution with SC4020 retains the end-user productivity with customized desktops while providing high performance, scalability, and minimized desktop administration overheads in a cost effective way.



1 Introduction

The Dell Storage SC4020 all-flash array is an excellent storage platform for VDI. With two controllers and 24 internal drives within a single enclosure, the SC4020 can support a large number of virtual desktops in a 2U form factor. When used in conjunction with Citrix XenDesktop and Machine Creation Services (MCS), layered on top of VMware vSphere hosts, large scale VDI deployments can be made simple and very cost-effective.

This document provides a storage reference architecture for building persistent desktop VDI deployments with light to medium user workloads from a storage perspective. The solution presented in this paper used the following software:

- Citrix XenDesktop
- Citrix Machine Creation Services
- VMware vSphere 5.5 hosts
- VMware vSphere management
- LoginVSI (for simulating real world VDI workloads)

The hardware used to host the environment consisted of:

- Dell PowerEdge M620 blade servers for VDI workload
- Dell PowerEdge M710HD blade servers for infrastructure hosting
- Dell Networking MXL Ethernet blade switches
- Dell Networking S5000 Top of Rack (ToR) switch
- Brocade M5424 Fiber-Channel blade modules
- Brocade 5100 Fibre Channel (FC) ToR switches
- Dell Storage SC4020 storage array

The hardware components are detailed in Section 3.2.

1.1 Objectives

The primary objectives of the tests conducted for this paper are:

- Determine the performance impact on the storage array during the peak VDI I/O activity such as boot and login storms.
- Demonstrate the latency at the full steady-state user load of 20 IOPS per desktop.

1.2 Audience

This document is intended for anyone who wishes to understand the performance characteristics of an SC4020 all-flash array in a VDI solution. This includes, but is not limited to, solution architects, storage architects, customers, solutions partners, and enterprise technologists.



2 VDI with Compellent all-flash storage arrays

Compellent is an enterprise storage solution designed with best-in-class capabilities and a wide range of features and functionalities. Sections 2.1 through 2.3 describe the top three reasons customers select it as the VDI storage platform of choice.

2.1 Flash-optimized data progression

Compellent data progression offers a complete, fully-integrated sub-LUN tiering solution that migrates data to the optimum storage tier and/or RAID level based on actual usage and performance needs with real-time intelligence and automation. Dell has recently introduced a new level of data progression optimized for flash drives that can tier data across high-endurance, high-performance, write-intensive SLC SSDs and high-capacity, low-cost, read-intensive MLC SSDs, disrupting the current flash cost model. The new enhancements include features to maximize performance of multiple SSD types, flash endurance management, new monitoring and management features as well as new default storage profiles and page sizes tailored for flash. Whereas the traditional data progression algorithms would run only once each day to optimize RAID levels and data tiering, flash-optimized data progression has the capability to move data across tiers continuously throughout the day. For details about flash-optimized data progression, see the white paper [Flash-optimized Data Progression](#). If additional capacity (that does not have the performance requirements of flash-based drives) is required, spinning disks can be added. These range from the performance-focused 15,000 RPM drives to the high-capacity 7,200 RPM drives.

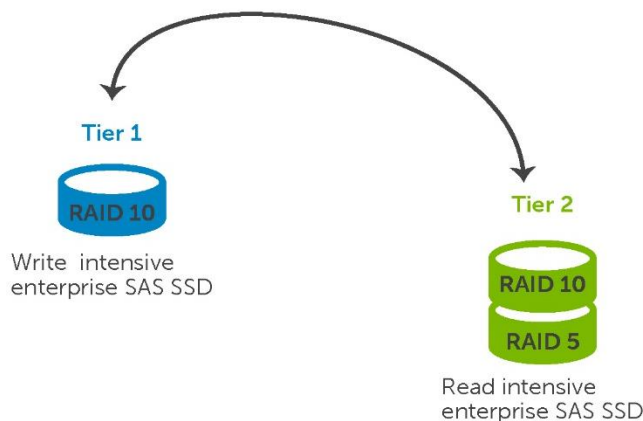


Figure 1 Tiering data with Compellent data progression

The Flash-optimized, automated, sub-LUN storage tiering is critical for a successful VDI deployment that can offer not only high performance but also a low cost. In VDI, different types of data have a different usage pattern. For example, virtual desktop VM data is much more frequently accessed than user data. Even within VM data, different blocks will see changing access patterns depending on the virtual desktop VM lifecycle. Unfortunately, the current VDI storage solutions in the market either put every piece of data (irrespective of its usage) on flash or depends on the user to separate VM data (on flash) and user data (on spinning media). While the former approach is not cost effective, the latter approach might not even be possible with a significant number of VDI use cases (for example, persistent desktops where VM data and user data are in the same drive or volume). With automated storage tiering from write-intensive SLCs to

read-intensive MLCs to spinning media, the intelligent storage layer of Compellent puts the right data on the right storage at the right time with the right cost.

2.2 Dynamic capacity

Compellent 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.

2.3 Hypervisor integration

Additional operational efficiencies are realized through integration with the hypervisor layer. Compellent seamlessly integrates with industry-leading hypervisors: VMware vSphere, Microsoft Hyper-V and Citrix XenServer. This integration delivers optimal storage performance, lower costs, and simplified VDI deployments.

As an example, the Compellent vSphere 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.

The Compellent advanced storage architecture and feature set, combined with tight hypervisor integration, provides an automated, self-tuning storage system capable of efficiently scaling to VDI production-level workloads at a very low cost.



3 Solution architecture

The overall architecture of the VDI solution used to characterize the SC4020 all-flash storage array is presented in Figure 2.

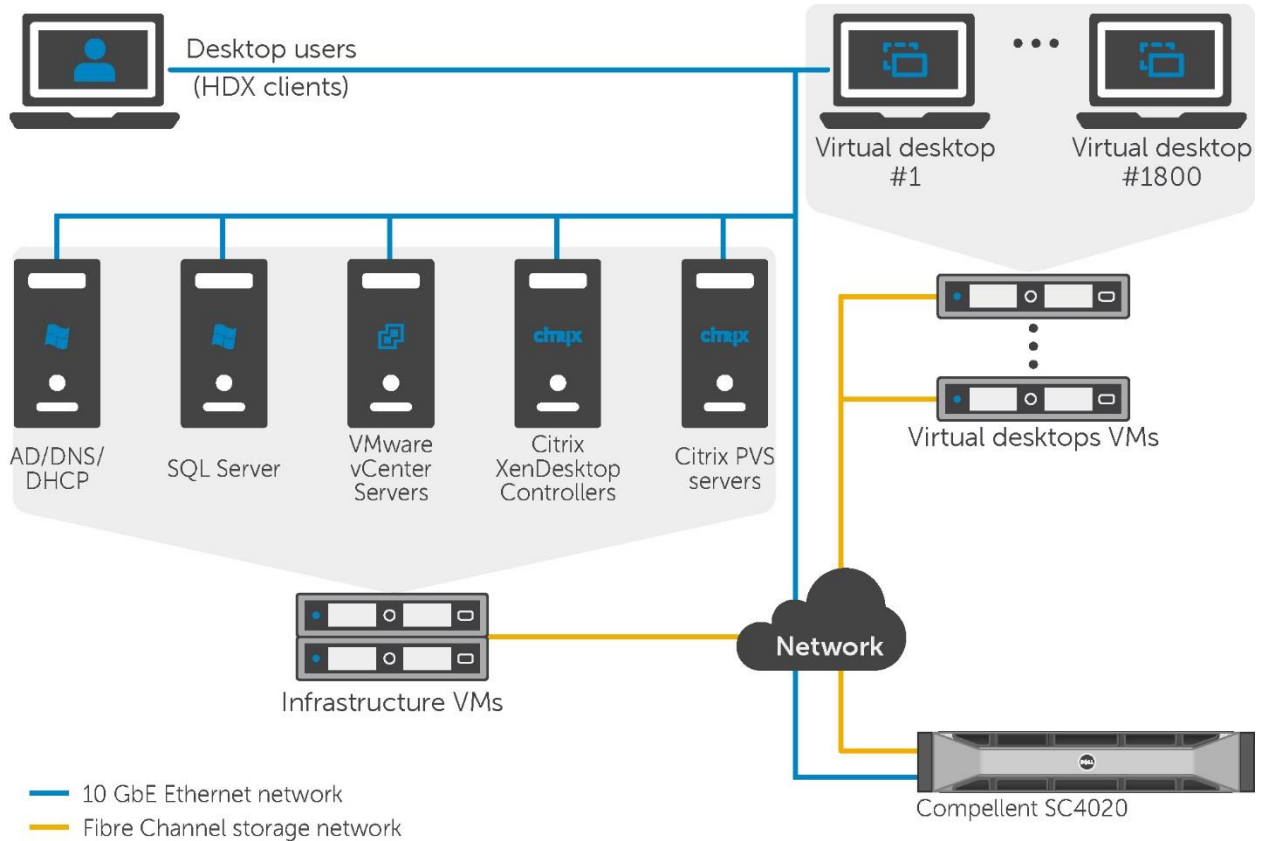


Figure 2 VDI solution architecture

3.1 Software

This section presents the software component overview of the VDI solution.

3.1.1 Citrix XenDesktop

The solution is based on Citrix XenDesktop 7.5 because it provides a complete end-to-end solution delivering Microsoft Windows virtual desktops to users on a wide variety of endpoint devices. Virtual desktops are dynamically assembled on demand, providing users with pristine and personalized desktops each time they log on.



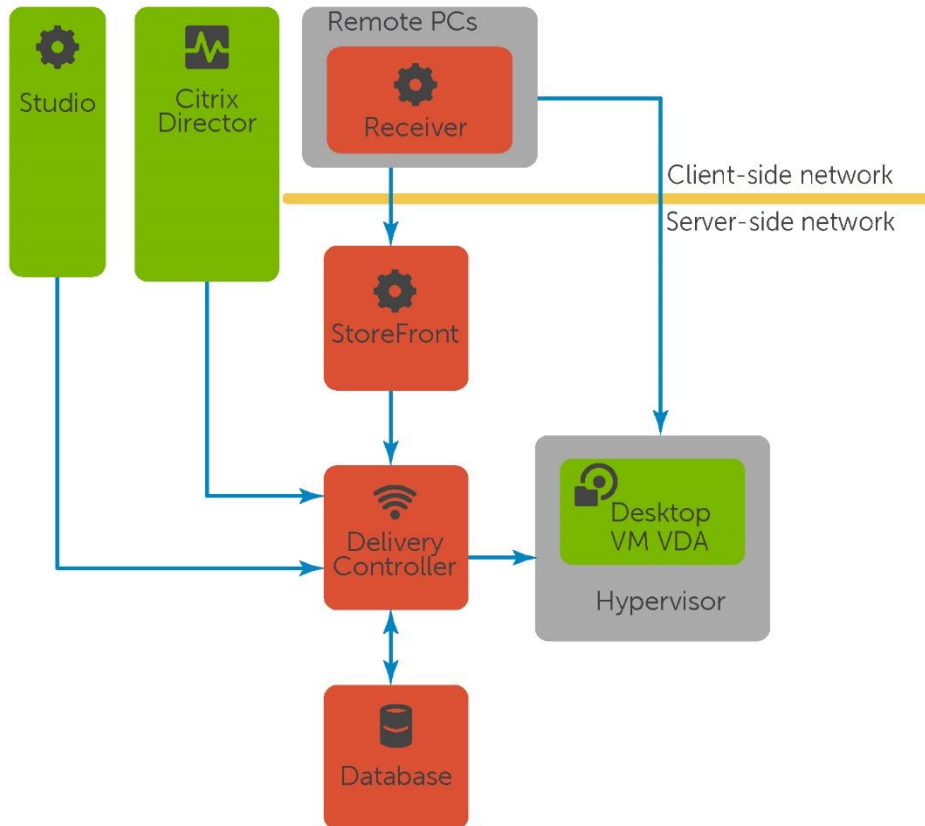


Figure 3 XenDesktop architecture

Citrix XenDesktop provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure.

The core XenDesktop components include:

Studio: The management console that enables configuration and management of a deployment; it eliminates the need for separate management consoles for managing delivery of applications and desktops. Studio provides various wizards that provide guidance through the process of setting up an environment, creating workloads to host applications and desktops, and assigning applications and desktops to users.

Director: A web-based tool that enables IT support teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users. It is also possible to view and interact with a user sessions using Microsoft Remote Assistance.

Receiver: Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user devices including smartphones, tablets, and PCs. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications.



Delivery Controller (DC): Installed on servers in the data center, the controller authenticates users, manages the assembly of users virtual desktop environments, and brokers connections between users and their virtual desktops.

StoreFront: StoreFront authenticates users to sites hosting resources and manages stores of desktops and user applications.

License Server: The Citrix License Server is an essential component at any Citrix-based solution. Every Citrix product environment must have at least one shared or dedicated license server. License servers are computers that are either partly or completely dedicated to storing and managing licenses. Citrix products request licenses from a license server when users attempt to connect.

Machine Creation Services (MCS): A collection of services that work together to create virtual servers and desktops from a master image on demand, optimizing storage utilization and providing a pristine virtual machine to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.

Virtual Delivery Agent (VDA): The Virtual Delivery Agent is a transparent plugin that is installed on every virtual desktop or XenApp host (RDSH) and enables the direct connection between the virtual desktop and users endpoint devices.

3.1.1.1 Machine Creation Services

Citrix Machine Creation Services is the native provisioning mechanism within Citrix XenDesktop for virtual desktop image creation and management. Machine Creation Services uses the hypervisor APIs to create, start, stop, and delete virtual desktop images. Desktop images are organized in a Machine Catalog and within that catalog, there are a number of options available to create and deploy virtual desktops:

Random: Virtual desktops are assigned randomly as users connect. When they logoff, the desktop is reset to its original state and made free for another user to login and use. Any changes made by the user are discarded at log off.

Static: Virtual desktops are assigned to the same user every time with three options for how to handle changes made to the desktop: Store on local vDisk, Personal vDisk, or discarded on user log off.



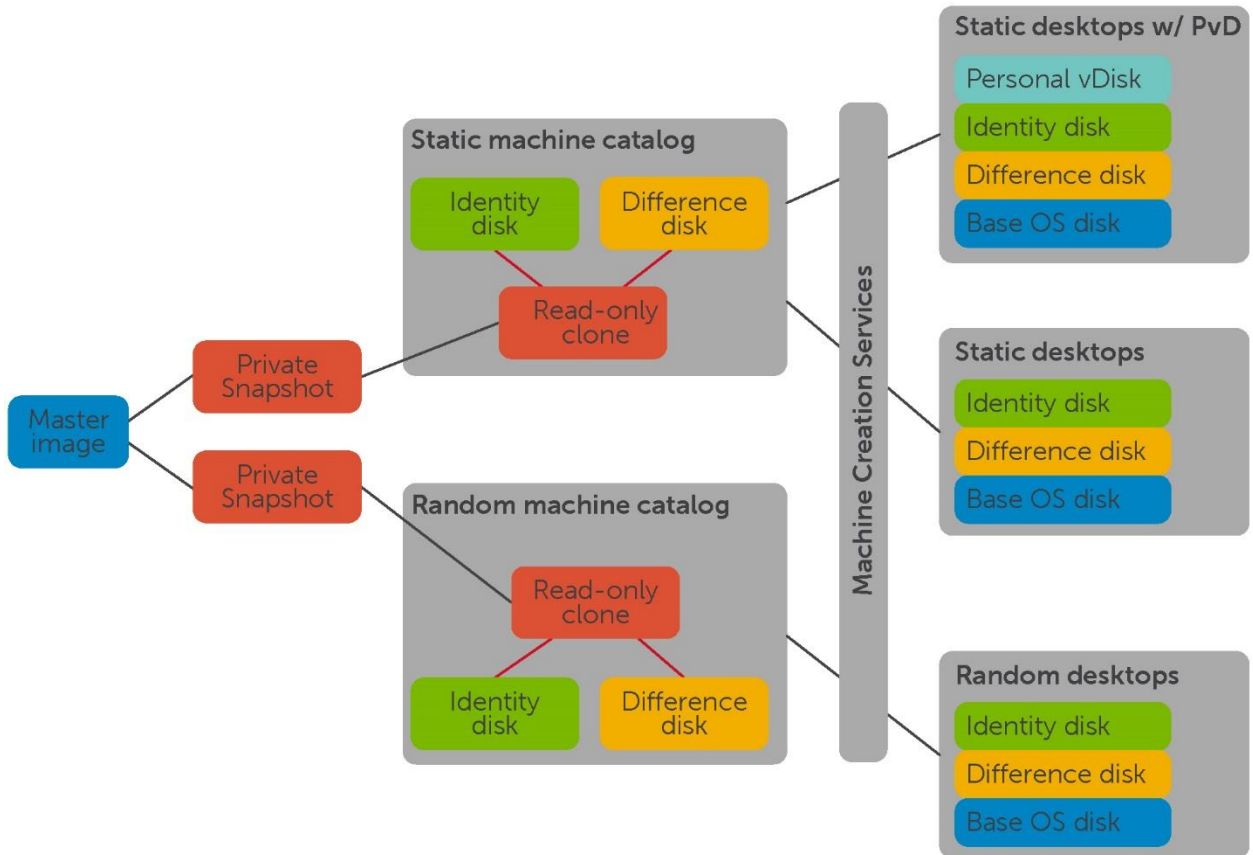


Figure 4 Virtual desktop creation with MCS

All the desktops in a random or static catalog are based off a master desktop template which is selected during the catalog creation process. MCS then takes snapshots of the master template and layers two additional virtual disks on top: an Identity vDisk and a Difference vDisk. The Identity vDisk includes all the specific desktop identity information such as host names and passwords. The Difference vDisk is where all the writes and changes to the desktop are stored. These Identity and Difference vDisks for each desktop are stored on the same data store as their related clone.

While traditionally used for small to medium sized XenDesktop deployments, MCS can bring along with it some substantial shared storage cost savings because of the snapshot/identity/difference disk methodology. The disk space requirements of the identity and difference disks when layered on top of a master image snapshot, is far less than that of a dedicated desktop architecture.

3.1.1.2 Citrix Personal vDisk technology

Citrix Personal vDisk is a high-performance enterprise workspace virtualization solution that is built in to Citrix XenDesktop to provide user customization and personalization benefits of a persistent desktop image, with the storage savings and performance of a single/shared image.

With Citrix Personal vDisk, each user receives personal storage in the form of a layered vDisk that enables them to personalize their desktop environment and make it persistent.

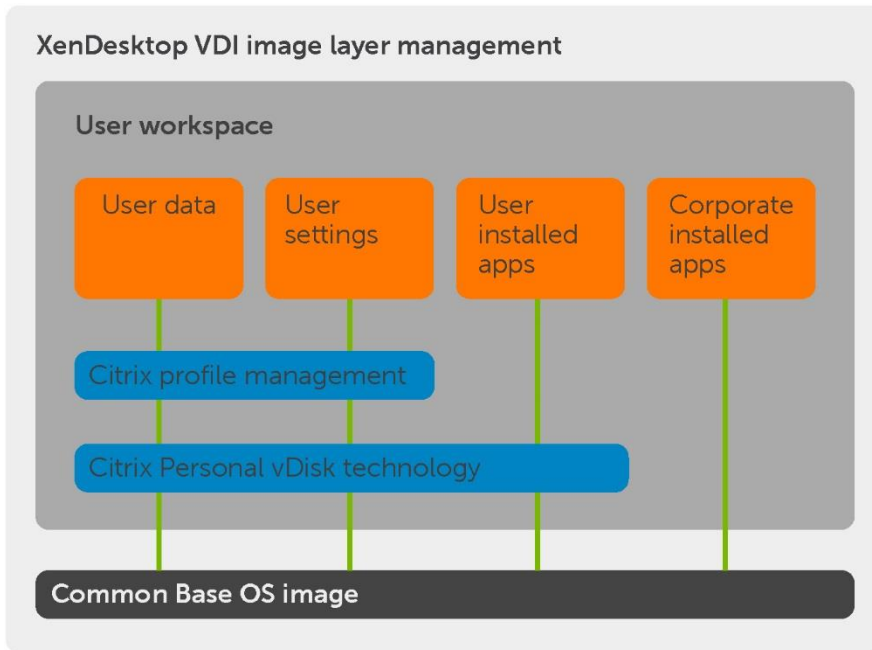


Figure 5 Persistent desktops with Citrix Personal vDisk

In addition, this vDisk stores user or departmental applications as well as data or settings the VDI administrator chooses to store. Personal vDisk provides the following benefits to XenDesktop:

- Persistent personalization of user profiles, settings and data
- Enables deployment and management of user installed and entitlement based applications
- Fully compatible with Application delivery solutions such as Microsoft SCCM, App-V and Citrix XenApp
- 100% persistence with VDI pooled Storage management
- Near Zero management overhead



3.1.1.3 Citrix profile manager

Citrix profile management is a component of the XenDesktop suite that manages user profiles and minimizes many of the issues associated with traditional Windows roaming profiles in an environment where users may have their profile open on multiple devices at the same time. The profile management toolset has two components. The first is the profile management agent installed on devices (or virtual desktops) where the user profiles are managed by the toolset. The second component is a group policy administrative template that is imported into a group policy assigned to an organizational unit within active directory. It contains the devices where the user profiles are managed.

In order to further optimize the profile management, folders within the user profile that can be used to store data are redirected the users home drive. The folder redirection will be managed using group policy objects within Active Directory. The following folders are redirected:

- Contacts
- Downloads
- Favorites
- Links
- My Documents
- Searches
- Start Menu
- Windows
- My Music
- My Pictures
- My Videos
- Desktop

3.1.2 VDI Hypervisor platform: VMware vSphere 5.5

VMware vSphere 5.5 is the enterprise virtualization platform used for building VDI and cloud infrastructures. VMware vSphere 5.5 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 (NTP, DNS, Active Directory, etc.). The vCenter 5.5 VM used in the solution is a single Windows Server VM or vCenter 5.5 virtual appliance, residing on a host in the management layer. SQL server is a core component of the Windows version of vCenter and is hosted on another VM that also resides in the management tier. As a best practice, install additional XenDesktop components in a distributed architecture, one role per server VM.



3.2 Hardware

This section presents the details of the VDI solution hardware used to determine the performance of the SC4020 all-flash storage array for the XenDesktop-based persistent desktop VDI deployments including boot storms, login storms and steady states. Therefore, sufficient headroom was left in the other components ensuring none of these other components in the VDI stack was the bottleneck.

3.2.1 Storage

The storage used to host the virtual desktops was a SC4020 array running Storage Center Operating System (SCOS) 6.5. All eight front-end ports were 8 Gb FC and all four back-end ports were 6 Gb serial-attached SCSI (SAS). The array used the internal enclosure for SSDs. The write-intensive, high-performance drives were 400 GB SLC drives composed of 17 active drives with one hot spare. The five active and one spare read-intensive, high-capacity drives were 1.6TB MLC drives, providing capacity for user persistence. The use of dual drive types provides a combination of high-performance and high-capacity drives in a single array, improving the cost and capacity of the array. Table 1 summarizes the storage hardware configuration.

Table 1 Storage hardware components

Storage role	Type	Qty.	Description
Controllers	SC4020	1	System Center Operating System (SCOS) 6.5 includes dual hot swap controllers
Enclosures	Internal	1	24 bay – 2.5" disk drive enclosure
Ports	FC - 8 Gbps	8	Front end host connectivity
	SAS - 6 Gbps	4	Back end drive connectivity
Drives	400 GB SLC drives	18	17 Active with 1 hot spare
	1.6 TB MLC Drive	6	5 Active with 1 hot spare

The volumes created to host the virtual desktops took advantage of the Compellent Dynamic Capacity technology for efficient capacity allocation. Table 2 details the volume layout on the SC4020 all-flash array hosting the virtual desktops.

Table 2 Volumes for hosting virtual desktops

Name	Type	Qty.	Size	Description
VDI Volumes	VMFS	40	2 TB	Each Dynamic Capacity (thin provisioned) volume stores 45 virtual machine

For each persistent desktop VM, an allocation of 30GB was made for its Personal vDisk (see Table 6). Compellent Dynamic Capacity technology was leveraged, and repeated testing over long periods of time



consistently showed low actual capacity utilization. The Compellent Flash-optimized Data Progression takes care of moving data automatically across tiers depending on the data access patterns and usage. Figure 6 shows the actual utilization as well as data placement after Flash-Optimized Data Progression tiered data across SLCs and MLCs. This is the most advantageous location, as determined by the system, based on block access, snapshot activity, and performance optimization.

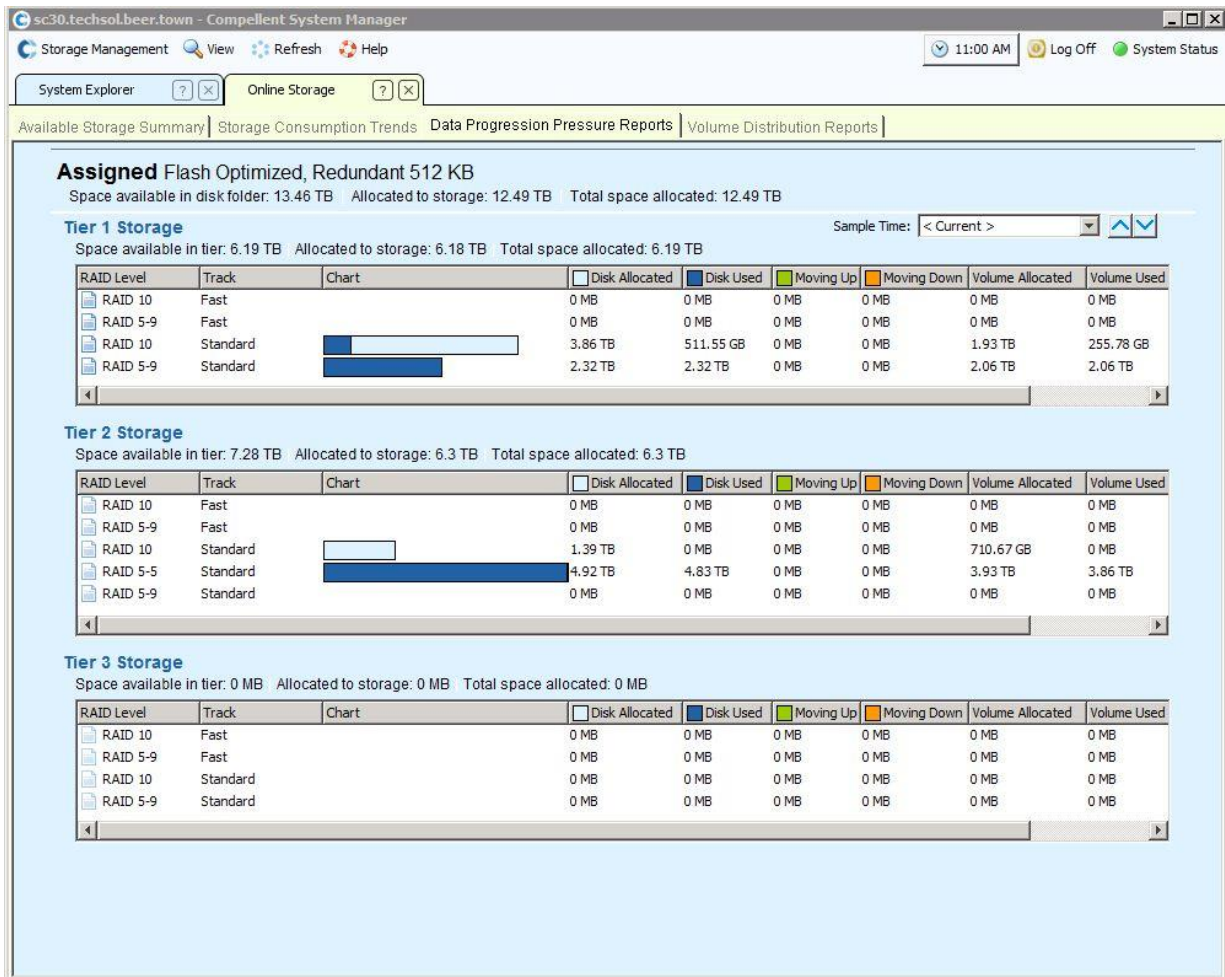


Figure 6 Data Progression

In addition to the array hosting virtual desktops, there was a Compellent SC8000 storage array used for the storage needed for infrastructure. This Compellent SC8000 storage array used a combination of SSDs and HDDs. The role of this infrastructure storage is purely to provide storage for domain controllers, XenDesktop servers, VMware vCenter and related server roles, and LoginVSI launchers (for VDI desktop load generation). In an actual deployment, this array could be part of the existing storage platform or could be combined with the VDI storage array depending on actual requirements.

All storage hardware was configured to Dell Compellent best practices with vSphere 5.x. For more information on these Best Practices, see *Dell Compellent Best Practices with VMware vSphere 5.x* at



<http://www.dell.com/learn/us/en/19/business~solutions~whitepapers~en/documents~compellent-best-practices-vmware-vsphere-5x.pdf>

Table 3 VMWare vSwitch configurations

Item	VLAN	Ports/Server	NIC Connectivity	NIC Teaming mode
vSwitch Cluster 1	31	120	Dual 10 Gb NICs, LACP trunked	Load Balancing, Route based on IP Hash
vSwitch Cluster 2	41	120	Dual 10 Gb NICs, LACP trunked	Load Balancing, Route based on IP Hash
Infrastructure Chassis	550	120	Dual 10 Gb NICs, LACP trunked	Load Balancing, Route based on IP Hash

Each VDI workload chassis is a member of a distinct VLAN, to control the number of devices per VLAN. This allows scaling of the number of VMs per VLAN up to 4,000 per VLAN, based on the vSwitch configuration and using S5000 switches, which have a limit of 16,000 address resolution protocol (ARP) entries.

3.2.2 Hosting environment

The virtual desktop hosting environment ran vSphere 5.5.

All hosts were blade chassis based, providing full hardware remote control and power management, as well as simplifying hardware updates, switch management, and standardizing hardware from a replacement perspective. Chassis based servers also simplify cluster management, with a chassis based networking backplane providing high-performance, redundant connectivity. Figure 7 shows the physical networking diagram.






Figure 7 Physical connectivity



Table 4 Host and desktop virtual machine configuration

Component	Description
Server	M620, 320 GB RAM, 2x8 core E5-2680 @ 2.7 Gz
Operating system	vSphere 5.5
FC	QLogic QME2572 HBA, Brocade M5424 IO module
Network	Dell Networking MXL 40Gb/10Gb
Virtual desktop hardware	2vCPUs, 2 GB RAM, VMXNET3, 30 GB cache disk

Table 5 Infrastructure virtual machine configuration

Component	Description
Operating system	Windows Server 2008 R2
vNIC Driver version	VMXNET3
Applications	Active Directory Domain Controllers, DNS, DHCP, vSphere servers, Enterprise Manager server, Liquidware Labs Stratusphere UX, XenDesktop, Microsoft SQL

Table 6 Persistence configuration

Component	Description
XenDesktop	7.5
Personal vDisk	30GB

3.2.3 Networking

Table 7 Network switches

Component	Description
ToR switch	Dell Networking S5000
Quantity	1
Firmware	9.0
Applications	ToR switch, VLAN management, VLAN routing



Table 8 Infrastructure Blade Chassis network switches

Component	Description
Chassis switch	Dell Networking MXL 10 Gb Ethernet
Quantity	2
Firmware	9.0
Application	Chassis based switching for Launcher hosts and PVS

Table 9 Workload Blade Chassis network switches

Component	Description
Chassis switch	Dell Networking MXL
Quantity	4 (2 per blade chassis)
Firmware	8.3.16
Applications	Chassis based switching, server connectivity

Table 10 Blade Chassis FC Switches

Component	Description
Chassis switch	Brocade 5424
Quantity	6 (2 per blade chassis)
Firmware	7.1.0.a
Applications	Chassis based switching, storage connectivity

Table 11 ToR FC switches

Component	Description
FC switch	Brocade 5100
Quantity	2
Firmware	7.1.0a
Applications	FC fabric, chassis and storage connectivity



4 Test methodology

This section outlines the testing process, criteria, tools, and monitoring methods. The criteria defined here can be used to compare a customer configuration or desired architecture.

4.1 Test objectives

As stated previously, the primary objectives of the tests conducted for this paper are:

- Determine the performance impact on the storage array during the peak VDI I/O activity such as boot and login storms
- Demonstrate the latency at the full steady state user load of 20 IOPS per desktop

This document provides a storage architecture reference point for building large-scale persistent desktop VDI deployments with medium to heavy workloads from a storage perspective.

4.2 Test tools

In this project, LoginVSI was used to control the test. LoginVSI offers a standardized workload generator capable of simulating various user types. In this case, it was used to generate 20 IOPS per VM to simulate a medium user.

XenDesktop was used to throttle the virtual machine boot rate. For the purposes of this test, a variable rate of machines was booted per minute. Since XenDesktop is normally used to boot the virtual machines before they are actually needed using a scheduler, the boot storm phase was purely an IOPS/latency exercise.

Since the workload was generated across a large number of machines in a consistent manner, LoginVSI was used to control the login rate for the workstations. By testing varying rates of machines per minute that login, the maximum login rate per minute that the environment can handle was determined.

4.2.1 Load generation

The load generation was done with by aLoginVSI. The following configuration was tested:

Table 12 Load generation configuration

Workload	LoginVSI
Medium User	20 IOPS



The characteristics of the LoginVSI medium workload, designed to run on two vCPUs per desktop VM in LoginVSI 4.0, were:

- Once a session began, the workload repeated (looped) every 48 minutes.
- The loop was divided in four segments; each consecutive LoginVSI user logon started at different segments. This ensured that all elements in the workload were equally used throughout the test.
- During each loop, the response time was measured every three to four minutes.
- The medium workload opened up to five applications simultaneously.
- The keyboard type rate was 160 ms for each character.
- Approximately two minutes of idle time was included to simulate real-world users.

In every loop, each user opened and used:

- Microsoft Outlook to browse messages
- Microsoft Internet Explorer to browse different webpages; a YouTube style video (480p movie trailer) was opened three times in every loop
- Microsoft Word to measure response time in one instance, in another instance to review and edit a document
- Doro PDF Printer & Acrobat Reader to print the Word document and review in a PDF file
- Microsoft Excel to open a very large randomized worksheet
- Microsoft PowerPoint to review and edit a presentation
- FreeMind, a Java based Mind Mapping application

4.2.2 Monitoring tool

The performance and latency measuring were done with Dell Compellent Enterprise Manager. This tool provided consolidated performance data across the entire environment from a single console.

4.3 Test phases

A VDI environment has multiple types of load, each with a distinct profile. This poses a challenge from a sizing and cost perspective because the need to support high boot loads needs to be balanced against the overall cost of the solution.

4.3.1 Boot storm

In the boot storm test phase, all of the virtual machines were turned on in rapid succession. The boot storm was measured from the time the first machine was turned on until the time the last machine was available for login.

This resulted in a nine-minute window to power on all machines and reach the login screen. This is a very acceptable boot time for 1,800 virtual machines.



4.3.2 Login storm

The login storm phase tested the storage array performance when 1,800 users were logged into the virtual machines. A high login rate was used to generate substantial I/O load during this login storm phase. The login time for 1,800 users was just under 10 minutes. At the 10-minute mark, all virtual machines were displaying the desktop.

4.3.3 Steady state

The steady state performance was a primary focus of this test and used to achieve the lowest possible latency during the steady state phase of the test with 1,800 users. The majority of the usage of a VDI solution is providing users with centrally managed workstations. The most important result from this solution was that the best possible user experience was determined.

The steady state latency with all 1,800 users logged in, generating 20 IOPs per user, was an average of 2 ms of latency.

4.4 Test criteria

Two primary test criteria were monitored during the testing to show the performance of the system under load.

4.4.1 Storage throughput and latency

This test maintained under 5 ms of volume latency to highlight the performance characteristics of SSDs and flash-optimized data progression offered by Compellent. With a steady state latency of under 2 ms, the Dell Compellent SC4020 offers excellent performance, and will provide an excellent platform for a VDI solution.



4.5 Test configuration

The virtual machines are configured in two catalogs with XenDesktop, with a single delivery group. Each catalog contains 900 virtual machines. By using a single delivery group, a single instance of LoginVSI can control the entire group of virtual machines.

Table 13 Testing details

Item	Count	Description
Machine Catalogs	2	Logical collections of machines for management purposes, each catalog uses a single configuration, replicated to each storage resource
XenDesktop Servers	3	Desktop Controllers and web access controllers
Virtual machines	1,800	The virtual machines being tested
VM Datastores	40	The volumes for virtual machine storage and virtual machine disk cache
VM Workload hosts	32	vSphere 5.5 hosts
VM Infrastructure hosts	16	vSphere 5.5 hosts, used to run the LoginVSI launchers, which initiate RDP sessions
LoginVSI Launchers	100	The machines initiating the remote connections to the virtual machines

Table 14 LoginVSI configuration

Configuration option	Setting
Launchers	100
Connection	Citrix Receiver Storefront initiated
User count	1,800
Workload	Medium
Run time	4 hours



4.6 Test results

The following sections detail the test results in each of the test phases.

4.6.1 Boot storm

Boot storm is the least common scenario because it only happens if an entire system is restarted, such as if it is taken down for maintenance, there is a power failure or another catastrophic error occurs.

The maximum I/O achieved was just over 125,000 IOPS at less than 7 ms average latency. A single vCenter server was used to manage all 1,800 desktops (done for simplicity, may not be appropriate for all environments). This test resulted in a boot time of less than 9 minutes for all 1,800 virtual machines from restart to the login screen.




Figure 8 Boot storm I/O



4.6.2 Login storm

A login storm is a more common scenario. Because it occurs every time the workday starts, it is a more important test.

Login storm I/O chart in Figure 9 shows the entire login to steady-state transition and highlights the maximum I/O achieved during this test. During the login storm, where all 1,800 users were logged in within 10 minutes, the maximum I/O was 49,000 IOPS with 2.5ms average latency. Even in this aggressive login rate, the SC4020 all-flash array delivered excellent performance.

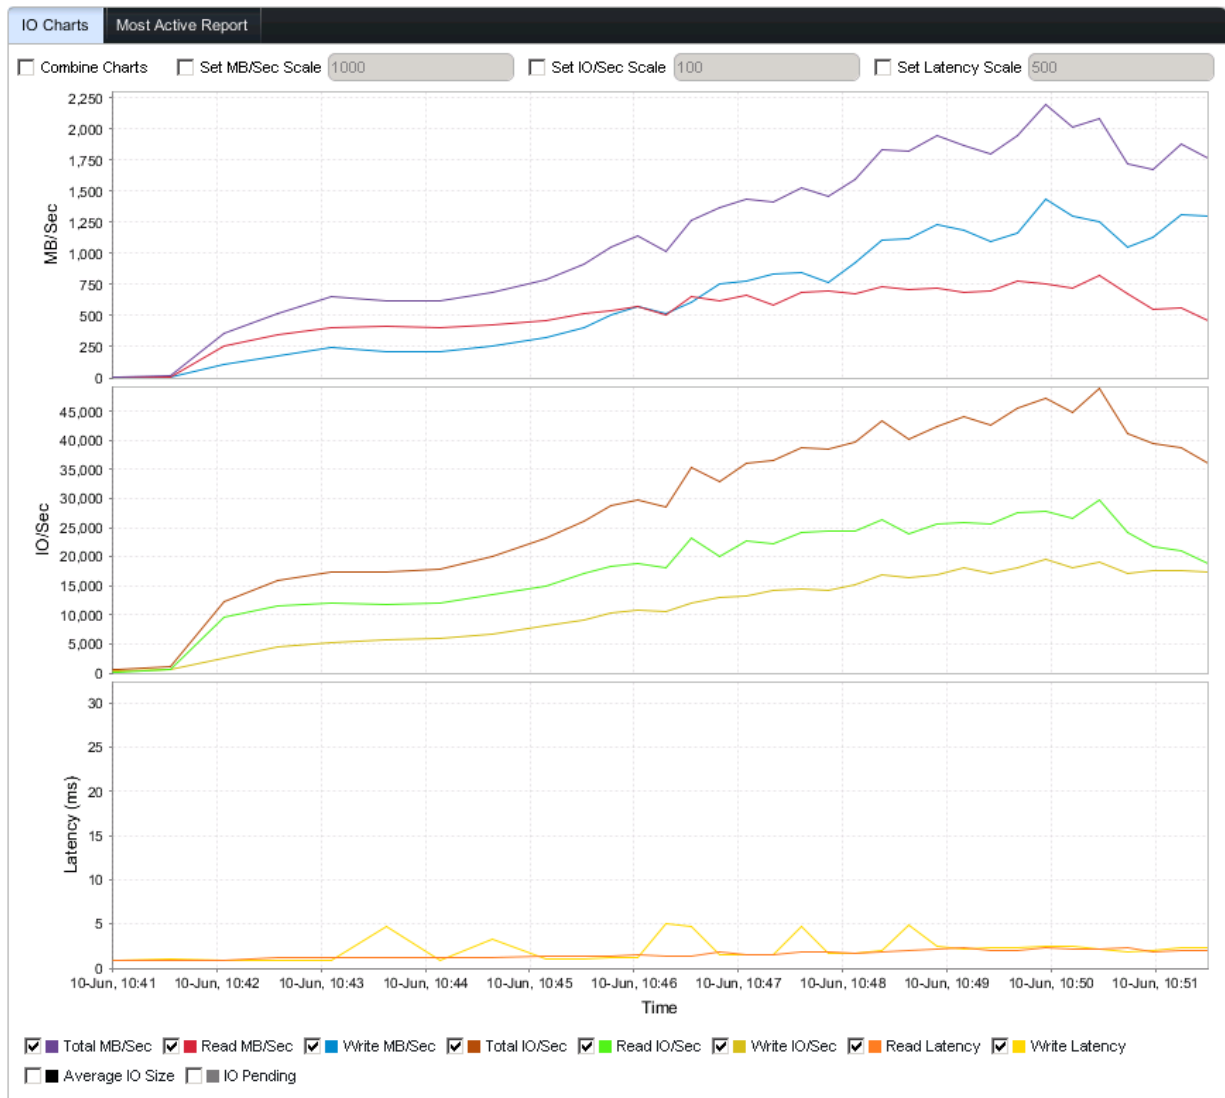


Figure 9 Login storm I/O chart



4.6.3 Steady state

A steady state phase is important because the majority of time for a system is spent in steady state. The performance during this phase of the test is the most critical because it determines the user experience. The steady state workload had a read to write ratio of 30% to 70%. Figure 10 shows the performance during the steady state phase of the test. The SC4020 delivered 40,000 IOPS for the 1,800 user persistent desktop VDI environment. The average write latency (as well as average read latency) during this test was just under 2 ms. This equates to an excellent user experience during the steady state phase.



Figure 10 Steady state I/O



5 Conclusion

This paper demonstrates how an 1,800 user persistent desktop XenDesktop VDI environment can be deployed with a single SC4020 all-flash array comprising a single 2U device, providing excellent performance in an extremely space efficient footprint. The storage I/O characteristics under various VDI workload scenarios (boot storm, login storm and steady state) demonstrate the effectiveness of SC4020 all-flash arrays as a storage building block for desktop virtualization deployments.

The testing of the SC4020 all-flash array storage platform showed that it can easily support medium VDI workloads. For example, it delivered 125,000 IOPS with less than 7 ms latency at boot storm and 40,000 IOPS with less than 2 ms latency at 20 IOPS/desktop for this 1,800-user deployment.

With the ability to support a large number of desktops with a small footprint, the SC4020 all-flash array can help organizations that are limited by the cost-per-desktop hurdle start VDI deployment plans. With its superb performance, advanced storage architecture and feature set like flash-optimized data progression, combined with a seamless hypervisor integration, SC4020 all-flash arrays provide an excellent platform for VDI infrastructure.



A Test specification components

A.1 XenDesktop

Table 15 XenDesktop server configurations

Component	Description
XenDesktop	7.5 with Machine Creation Services
Application	Virtual machine assignment management, Web interface for remote connections,
Servers	4 cores, 12 GB of RAM
Storage	40 GB SAN data volume

A.2 vSphere Hypervisor

Table 16 vSphere server configuration

Component	Description
vSphere	5.5
Application	End-User Virtual Machine host
Server	2x8 core processor, 320 GB RAM
Storage	200 GB SSD local boot volume, SAN based VMFS volumes

A.3 VMware vCenter

Table 17 VMWare vCenter configuration

Component	Description
VMware vCenter	5.5
Application	Hypervisor management, centralized management

A.4 LoginVSI

Table 18 LoginVSI details

Component	Description
LoginVSI	4.0.11
Application	Load generation test harness, controls all virtual machine sessions used for load generation

