VDI with Citrix XenDesktop and VMware vSphere on Dell EqualLogic PS6210XS Hybrid Storage Arrays

A Dell Storage Reference Architecture

Dell Storage Engineering
December 2014

A Dell Reference Architecture

Revisions

Date	Description
December 2014	Initial release

Table of contents

Acknowledgements

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to SISfeedback@Dell.com.

SISfeedback@Dell.com

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 to not only reduce the operational expenses for desktop management and provisioning but also to improve user mobility and data security.

A VDI deployment can place high capacity and performance 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. There may also be situations like an unexpected power shutdown which requires booting of all the virtual desktops at the same time. This boot storm creates significantly higher IOPS on the underlying storage platform than the regular business hour usage.

The Dell PS Series hybrid arrays are well suited for supporting VDI deployments because they are specifically designed for VDI type workloads where the most frequently accessed data – the active data set – which needs very high I/O rates, resides in a relatively smaller portion of the storage capacity. With high performance SSD drives and high capacity SAS hard disk drives within the same chassis, the hybrid arrays automatically move frequently accessed data from the SAS tier to the SSD tier. The automatic tiering function makes the hybrid PS Series SAN a cost efficient option for VDI environments where the peak load in an I/O storm from hundreds of virtual desktops is typically concentrated on the relatively small-sized base images.

This paper demonstrates how a 2,000 standard user virtual desktop environment can be deployed in a Citrix XenDesktop 7.5 platform with Citrix Provisioning Services 7.1 (PVS) VDI infrastructure leveraging a *single* Dell EqualLogic PS6210XS hybrid array. The test infrastructure also included Dell Networking switches and Dell PowerEdge servers running VMware vSphere 5.5 as hypervisor. Details are provided for the storage I/O characteristics under various VDI workload scenarios like boot and login storms along with performance characteristics throughout the VDI stack.

In the test environment, the single PS6210XS array successfully hosted 2,000 desktops with satisfactory performance results across all layers of the stack, including the user layer, hypervisor layer, and storage layer. As an example, at the peak of the 2,000 user login storm, the array delivered approximately 13,200 IOPS with less than 1 ms write latency with a 99% write VDI I/O workload. The very low latency at the storage layer ensured excellent end-user desktop application response times. The detailed tests results are available in Section 6. Best practices are provided in Section 7.

1 Introduction

Desktop virtualization platforms like Citrix XenDesktop 7.5 can provide organizations with significant cost savings, streamlined implementation, and ease of desktop management. In order to achieve these VDI benefits and to ensure optimal user experience, storage infrastructure design and sizing considerations need to be addressed carefully.

The goal of this paper is to present the results of a series of storage I/O performance tests and provide the optimal virtual desktop density for a PVS based XenDesktop VDI solution leveraging a Dell EqualLogic PS6210XS hybrid storage array.

For more information on sizing guidelines, see the paper *Sizing and Best Practices for Deploying VMware View 5.2 on VMware vSphere 5.1 with Dell EqualLogic Storage* at:
http://en.community.dell.com/techcenter/extras/m/white_papers/20438105/download.aspx

1.1 Objectives

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

- Develop best practices for a XenDesktop, PVS and vSphere based VDI solution deployed using the PS6210XS hybrid array
- Determine how many virtual desktops can be deployed in this environment using a single PS6210XS storage array with acceptable user experience indicators for a standard user workload profile
- Determine the performance impact on the storage array of peak I/O activity such as boot and login storms

The test infrastructure used for the tests included:

- Citrix XenDesktop 7.5, Provisioning Services 7.1
- VMware vSphere 5.5 hypervisor
- Dell PowerEdge M620 blade servers
- Dell Networking MXL switches and S6000 switches
- Dell EqualLogic PS6210XS storage array

1.2 Audience

This paper is intended for solution architects, storage network engineers, system administrators, and IT managers who need to understand how to design, properly size, and deploy PVS based VDI Solutions using Dell PS Series storage. It is expected that the reader has a working knowledge of the XenDesktop architecture, vSphere system administration, iSCSI SAN network design, and Dell PS Series iSCSI SAN operation.

2 VDI with Dell Storage PS Series hybrid storage arrays

Dell Storage PS Series hybrid storage arrays combine both solid-state drives (SSDs) and traditional hard disk drives (HDDs) in the same chassis and include features such as automated data tiering and centralized management tools to support efficient VDI deployments.

Combining SSDs and HDDs in a single Dell Storage PS Series hybrid array enclosure helps optimize performance, capacity, and scalability for the multi-tiered workloads that characterize VDI deployments. Intelligent PS Series software with automated tiering capabilities places the frequently accessed base desktop virtual machine image on low-latency, high-performance SSD storage to help ensure maximum performance. Less frequently accessed data, such as end-user application files, are placed on low-cost, capacity-optimized HDDs, enabling organizations to meet VDI performance demands without costly overprovisioning.

Administrators can take advantage of the virtualized, scale-out Dell Storage PS Series hybrid array storage architecture to roll out VDI deployments in phases as their needs grow and without disruption in availability or performance. Additionally, robust management and reliability features help ensure extremely high availability and rapid recoverability.

Dell Storage PS Series hybrid arrays like PS6210XS are designed to offer VDI-optimized storage. Combined high-performance SSD capacity and capacity-optimized SAS drives provide a hot-to-warm data ratio that is optimized for high-performance virtual desktop environments. SSD capacity is available for frequently accessed desktop virtual machine data, while end-user application data can be stored on capacity-optimized drives in either of these arrays or in preexisting capacity located on other Dell Storage PS Series arrays. In addition to automatic tiering within the hybrid arrays, PS Series software can automatically tier data across arrays based on utilization rates when multiple arrays are deployed.

> For more information on Dell Storage PS Series hybrid array load balancer, see *EqualLogic PS Series Architecture: Hybrid Array Load Balancer* at:
> http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20349614/download.aspx

Figure 1 below shows the 2,000 standard user VDI architecture for XenDesktop-based environments with 22 PowerEdge M620 servers, eight MXL switches, and one PS6210XS hybrid array for virtual desktops and one PS6110E array for user data. Out of the 22 blade servers, 20 were used to host virtual desktop virtual machines (VMs), and the remaining two were used for hosting infrastructure VMs. Information about how this reference architecture was developed ensuring the optimal utilization across the servers, switches, and storage arrays is detailed in the rest of this paper.
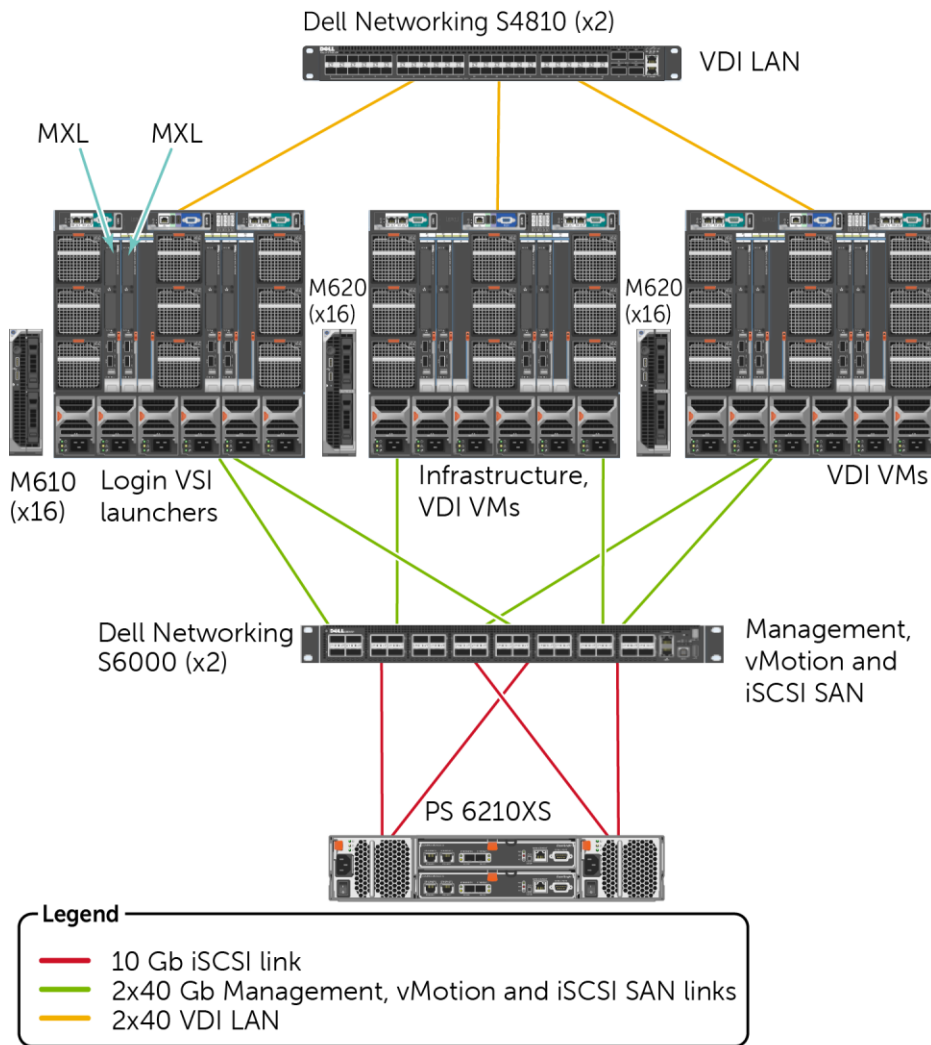
Figure 1   Compute, storage, and network layout in the PowerEdge M1000e blade chassis for the solution

3 Solution architecture

The overall architecture of the VDI solution used to characterize the PS6210XS hybrid storage array is shown in Figure 2.
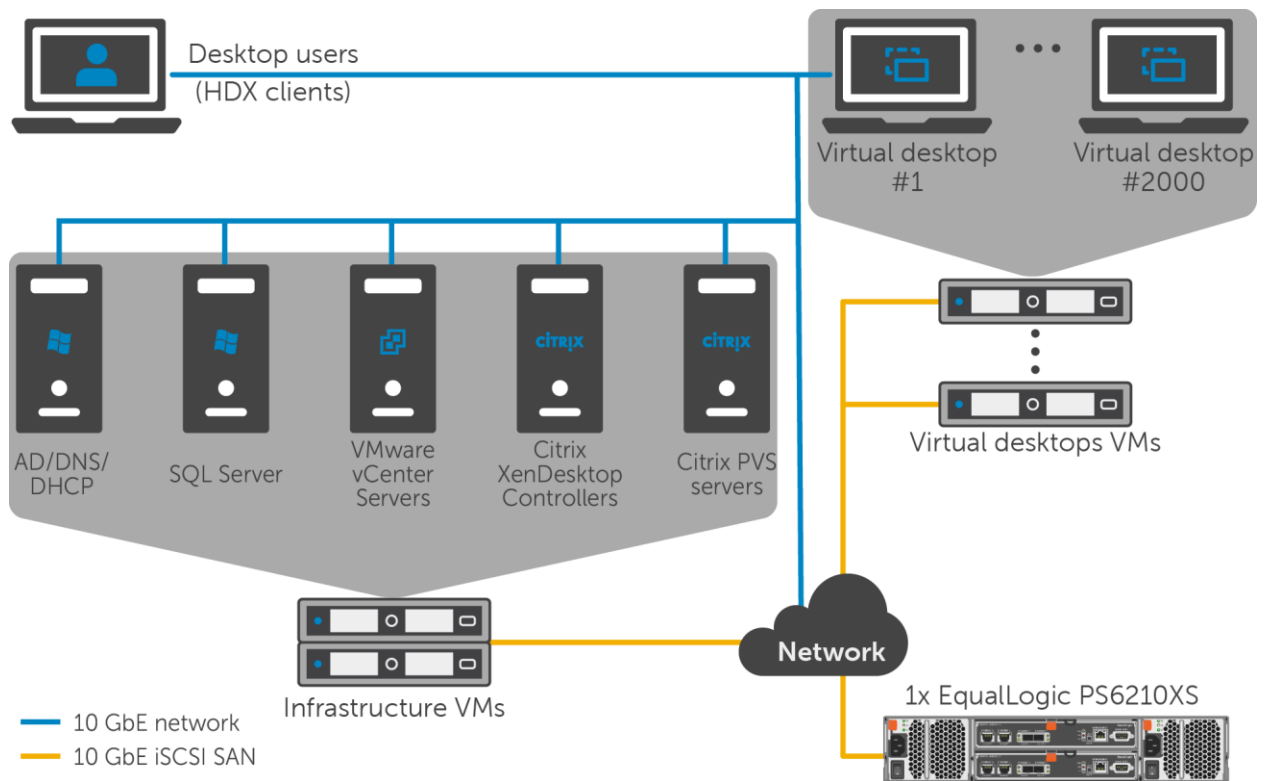
Figure 2   VDI solution architecture

3.1 Software

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

3.1.1 Citrix XenDesktop

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:

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.

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 session using Microsoft Remote Assistance.

License Server: The Citrix License Server is an essential component of 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 VM to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.

Provisioning Services (PVS): The PVS infrastructure is based on software-streaming technology. This technology allows computers to be provisioned and re-provisioned in real-time from a single shared-disk image. This means administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is done on the master image.

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

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

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 guide you through the process of setting up an environment, creating workloads to host applications and desktops, and assigning applications and desktops to users.

Virtual Delivery Agent (VDA): A transparent plugin that is installed on every virtual desktop or XenApp host (Remote Desktop Session Host) that enables the direct connection between the virtual desktops and users' endpoint devices.

3.1.1.1 PVS

Citrix XenDesktop provides PVS as an alternative provisioning mechanism to the Citrix MCS for virtual desktop image creation and management. PVS uses streaming to share a single base vDisk image instead of copying images to VMs. PVS can be used to deliver shared vDisk images to physical or VMs. PVS enables real-time provisioning and re-provisioning of shared vDisk images, which allows administrators to completely eliminate the need to manage and patch individual systems.

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.

Using PVS, vDisk images are configured in Standard Image mode, read-only, or Private Image mode, read/write. A vDisk in Standard Image mode allows multiple desktops to boot from it simultaneously, greatly reducing the number of images that must be maintained and the amount of storage that is otherwise required (non-persistent). Private Image mode vDisks are equivalent to dedicated hard disks and can only be used by one target device at a time (persistent). The Provisioning Server runs on a virtual instance of Windows 2012 R2 on the Infrastructure Servers.

3.1.1.2   PVS Write Cache

PVS delivery of standard images relies on write-caches to store any writes made by the target OS. The most common write-cache implementation places write-cache on the target machine's storage. Independent of the physical or virtual nature of the target machine, this storage has to be allocated and formatted to be usable.

While there are four possible locations for storage of the write cache in PVS, it is recommended to place the PVS write cache either in the target compute host's RAM with overflow enabled or use a hard disk on the target VM. For the tests used in this paper, a hard disk on the target VM was used.

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 were used (NTP, DNS, Active Directory, and others).

4 Infrastructure and test configuration

This section provides information on the test configuration used for hosting XenDesktop virtual desktops, including infrastructure components, networking, and storage sub-systems.

4.1 Host design considerations

The entire infrastructure and test configuration was installed in two Dell PowerEdge M1000e blade chassis. The 22 PowerEdge M620s used in this solution were divided into two vSphere clusters.
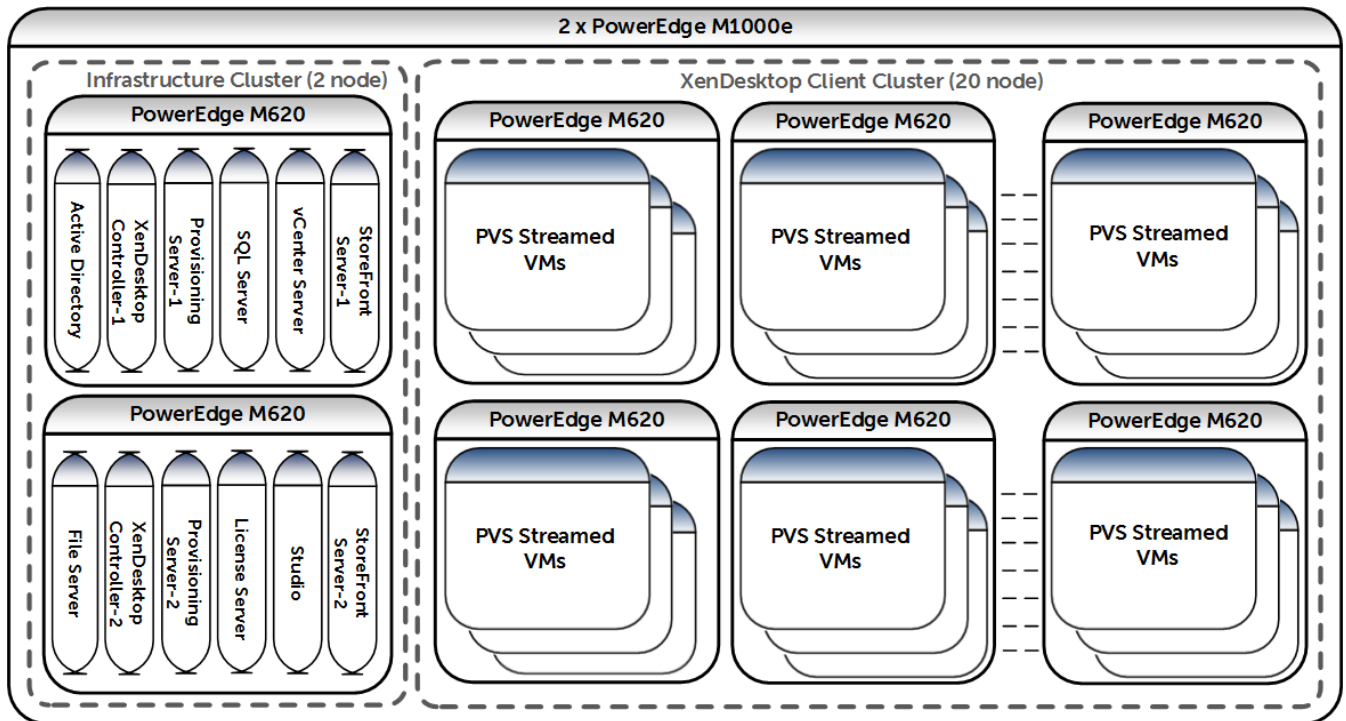
Figure 3   vSphere blade clusters on PowerEdge M1000e

The clusters included:

- **Infrastructure Cluster:** Two M620 blade servers hosting VMs for Active Directory services, VMware vCenter 5.5 server, Citrix XenDesktop Delivery Controllers 7.5 (primary, secondary and tertiary), XenDesktop PVS 7.1 server (primary, secondary and tertiary), XenDesktop StoreFront Server, Citrix License Server, XenDesktop Studio, Citrix Director, Microsoft Windows Server 2012 R2 based file server, and SQL Server 2012 R2.
- **XenDesktop Client Cluster:** 20 M620 blade servers hosting the streamed virtual desktops.

In addition to the above servers, eight Dell PowerEdge M610 blade servers in a separate M1000e blade chassis were used for VDI load generation purposes.

4.2 Network design considerations

Figure 4 below shows the network layout of one of the 22 PowerEdge M620 blade servers with vSphere 5.5 installed.
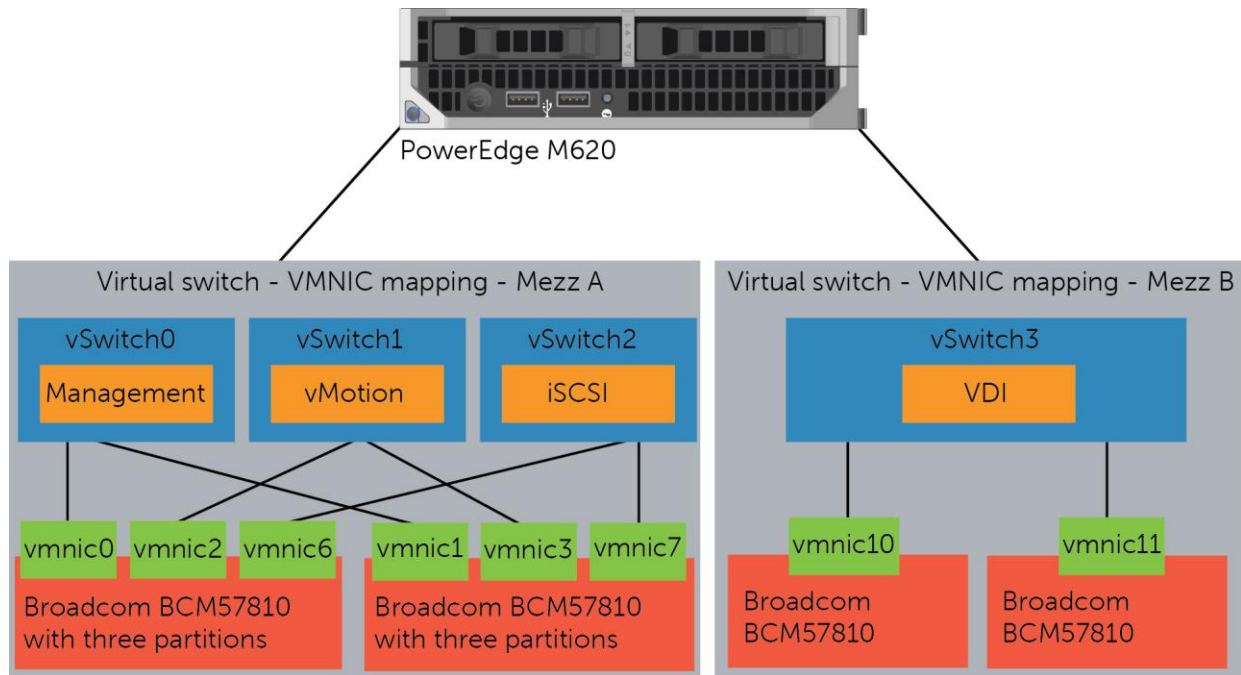
Figure 4   vSphere host network configuration

Each M1000e blade chassis where the VDI solution was deployed consisted of the following switches:

- Two MXL blade switches in Fabric A for connectivity to the Management LAN, iSCSI SAN, and the vMotion LAN
- Two MXL blade switches in Fabric B for connectivity to the dedicated VDI LAN

The MXL blade switches used to provide connectivity to the Management LAN, iSCSI SAN, and vMotion LAN were connected to two Dell Networking S6000 switches to provide connectivity to the PS6210XS storage array and connectivity between the blade chassis.

The MXL blade switches used for providing connectivity to the dedicated VDI LAN were connected to two Dell Networking S4810 switches.

Data Center Bridging (DCB) was used to provide converged LAN and SAN access on the S6000 / MXL switches, PS Group and ESXi servers.

Network Interface Card (NIC) Partitioning (NPAR) was used to divide the physical NICs in Fabric A into multiple logical NICs. This allows for dynamic allocation of bandwidth for the different partitions and this helps reduce the total cost of ownership for the solution. In these tests, the iSCSI partition was given access to 50% of the available 10 Gb bandwidth and the remaining bandwidth was shared among the remaining partitions.

More information on NIC partitioning with Broadcom NICs is available in the white paper *Enhancing Scalability through Network Interface Card Partitioning*, available here: http://www.dell.com/downloads/global/products/pedge/en/Dell-Broadcom-NPAR-White-Paper.pdf

The following partitions were created on the Broadcom NICs on Fabric A.

Table 1   NPAR partitions

Partition name	Purpose
Management	Management access to the vSphere hosts and infrastructure VMs
vMotion	VMware vMotion access to provide live migrations with zero downtime as well as load balancing on vSphere hosts
iSCSI	Access to Dell Storage iSCSI SAN where the VDI VMs are stored

The networks were segregated at the S6000 and MXL switches using DCB and VLANs to separate different types of traffic, specifically:

- **Management LAN (VLAN 10):** This network provided a separate management network for all the physical vSphere hosts. It also allowed communication between various infrastructure components such as Microsoft Active Directory Server, Microsoft SQL Server, and VMware vCenter Server.
- **VMware vMotion LAN (VLAN 20):** This is the network over which the VMs were migrated to provide high availability (HA) and load balancing between the hosts. Since the 12 M620s were divided into two clusters, only hosts in the same cluster could provide HA and load balancing.
- **VDI LAN (VLAN 30):** This is the network over which the virtual desktops were streamed to the user's thin clients. The connectivity to the existing network of the client was provided by uplink modules on the Force10 MXL switches.
- **iSCSI SAN (VLAN 172):** This is the network over which all the virtual desktops and the infrastructure services accessed the Dell storage arrays.

Appendix B has extended information on the vSwitch configuration for each vSphere host.

4.3   Separation of user data and virtual desktop data

Typically, user data can be maintained on a separate file share through roaming profiles and folder redirections. In this approach, a desktop, which is comprised of the shared read-only base image and the individual differential data from temporary changes, is always stateless. When a user logs in, the personal data is mapped to the desktop and all the changes to the data are maintained on the file share over the network. When the user logs off, the personal data is no longer on the desktop and the desktop can be put back in the original state with the temporary data discarded.

This approach has two benefits. First, the advantages of the non-persistent desktop deployment model can be leveraged while using the persistent desktop model. Second, the performance needs of VM data and user data are distinctly different, with the former needing very high performance to handle the I/O storms. The VM data can be placed in a high-performance storage array, while the user data can be served from the file shares running on a capacity-oriented storage array.

In our test environment, we used a separate storage array to host the volumes for infrastructure and user data. Each user was assigned about 2 GB of user data space on a 4 TB volume that could support up to 2,000 users.

4.4 Dell Storage PS Series storage array configuration

PS6210XS storage arrays hosted all the virtual desktops while an additional PS6110E array was used to host the infrastructure VMs and the user data volumes used in this solution. The volume layout used for the infrastructure functions including user data is shown in Table 2.

Table 2   Volume layout for hosting infrastructure components and user data

Volume name	Size	Purpose
Infra-1	500 GB	Storage for Active Directory, SQL Server, vCenter Server, XenDesktop Controller-1 & 3, Provisioning Services Server-1, Director, StoreFront Server-1, KMS Server
Infra-2	500 GB	Storage for File server, XenDesktop Controller-2, Provisioning Services Server-2 & 3, StoreFront Server-2, Studio,
UserSpace	4 TB	Storage for User profiles and folder redirection space (Average, 2 GB per user)

In addition to the infrastructure volumes, the storage array also provided shared storage for hosting the virtual desktops. The volume layout used for configuring the base vDisk image and VDI volumes on the array is shown in Table 3.

Table 3   Dell Storage PS6210XS layout for volumes hosting virtual desktops

Volume name	Size	Purpose
PVS-Master-1	50 GB	Storage for base vDisk image for VDI deployment
PVS-Master-2	50 GB	Storage for base vDisk image for VDI deployment
PVS-Master-3	50 GB	Storage for base vDisk image for VDI deployment.
VDI-Images1 through 20	500 GB each	Storage for VDI VMs in VDI Cluster. Each volume hosted 100 desktops.

4.5   vSphere host network configuration

VMware vSphere 5.5 hypervisor was installed on all 22 blades. The network configuration on each of those hosts is described below. Each vSphere host was configured with four virtual switches, vSwitch0, vSwitch1, vSwitch2, and vSwitch3, to separate the different types of traffic on the system.

Table 4   vSwitch configuration in vSphere hosts

vSwitch	Description	Virtual NICs used
vSwitch0	Management Network	vmnic0, vmnic1
vSwitch1	iSCSI SAN	vmnic6, vmnic7
vSwitch2	vMotion LAN	vmnic2, vmnic3
vSwitch3	VDI LAN	vmnic10, vmnic11

For additional information on individual vSwitch configuration, refer to Appendix B.

4.6 XenDesktop configuration

XenDesktop 7.5 was installed by following the documentation provided by Citrix.

XenDesktop 7.5 installation documentation:
http://support.citrix.com/proddocs/topic/xenapp-xendesktop/cds-xenapp-xendesktop-75-landing.html

Table 5   Specific configuration used in the tests

Purpose	Count	Type	Memory	CPU
XenDesktop Controllers	3	VM	8 GB	8
PVS Servers	3	VM	12 GB	8
StoreFront Servers	2	VM	8 GB	8
License Server	1	VM	4 GB	4
Citrix Studio	1	VM	4 GB	4
Citrix Director	1	VM	4 GB	4

4.7 Windows 7 VM configuration

Following the guidelines from Citrix, VMware and Login VSI, the Windows 7 base image was generated based on a generic base VM with the following properties:

- VMware Virtual Hardware version 8
- One virtual CPU
- 3 GB RAM
- 25 GB virtual hard drive
- One virtual NIC connected to the VDI network
- Windows 7 64 bit OS

Additionally, the base image was customized using the *Citrix Windows 7 Optimization Guide*, available here: http://support.citrix.com/servlet/KbServlet/download/25161-102-648285/XD%20-%20Windows%207%20Optimization%20Guide.pdf

5 XenDesktop test methodology

This section outlines the test objectives along with the test tools and criteria used to determine the optimal virtual desktop density and best practices for deploying Citrix XenDesktop on PS6210XS arrays.

5.1 Test objectives

As noted in Section 1.1, the test objectives were:

- Develop best practices and optimal virtual desktop densities for a XenDesktop based VDI solution deployed with PS6210XS hybrid arrays
- Determine how many virtual desktops can be deployed in this environment using a single PS6210XS storage array with acceptable user experience indicators for a standard user workload profile
- Determine the performance impact on the storage array of peak I/O activity such as boot and login storms

5.2 Test tools

All tests were conducted using Login VSI 4.0 as the workload generator and user experience analyzer tool. Login VSI is a benchmarking tool to measure the performance and scalability of centralized desktop environments such as Server Based Computing (SBC) and VDI.

> **Note:** More information can be found at the Login VSI website: http://www.loginvsi.com

5.2.1 Load generation

The "Medium" workload from Login VSI was used to simulate the Standard user workload. The characteristics of the Medium workload are:

- Up to five applications are open simultaneously.
- Applications include Microsoft Internet Explorer, Microsoft Word, Microsoft Excel, Microsoft PowerPoint, PDF reader, 7-Zip compression software, Movie player and FreeMind.
- Once a session is started, the Medium workload repeats approximately every 48 minutes.
- During a loop, the response time is measured every 3-4 minutes.
- Idle time is about two minutes in each 48 minute loop.
- Type rate is approximately 160 ms per character.

Although Login VSI provides other workloads, the Medium workload was used in the testing because it closely resembles the workload of a Standard user.

> More information on the classifications of the user types can be found in Appendix D.

5.2.2 Monitoring tools

The following monitoring tools were used:

- Dell SAN Headquarters (SAN HQ) for monitoring storage array performance
- VMware vCenter statistics for vSphere performance
- Login VSI Analyzer for end-user performance statistics

Detailed performance metrics were captured from the storage arrays, hypervisors, virtual desktops, and the load generators during the tests.

5.3 Test criteria

The primary focus of the tests was to determine the maximum number of desktops which can be deployed on a single PS6210XS hybrid storage array in this environment while using PVS.

VDI configurations involve many components at different layers, like application layer, hypervisor, server, network, and storage. As a result, multiple metrics need to be captured at different layers to ensure that the environment is healthy and performing optimally and appropriately for all users.

The specific test criteria are described in the following sections.

5.3.1 Storage capacity and I/O latency

The typical industry standard latency limit for storage disk I/O is 20 ms. Maintaining this limit ensures good user application response times when there are no other bottlenecks at the infrastructure layer. In addition to this, it is also recommended to maintain a 10% spare capacity on the storage array for optimal performance.

5.3.2 System utilization at the hypervisor

Even though the primary focus of these tests was storage characterization, additional metrics at the hypervisor infrastructure layer were defined to ensure solution consistency. These were:

- CPU utilization on any vSphere server should not exceed 85%
- Minimal memory ballooning on the VMs
- Total network bandwidth utilization should not exceed 90% on any one link
- TCP/IP storage network retransmissions should be less than 0.5%

5.3.3 Virtual desktop user experience

Login VSI Analyzer was used to gather metrics on the user experience at the virtual desktop layer to ensure that all the desktops had acceptable levels of application performance. See Appendix C for details about Login VSI's VSImax (Dynamic) parameter results.

5.4 Test configuration

A single virtual desktop catalog was configured using the Citrix Studio administrator interface. Each catalog was built from a Windows 7 base image. The Windows 7 configuration information is available in Section 4.7.

Desktop pool properties:

- One XenDesktop catalog
- Catalog with 2,000 desktops, spread across 20 500 GB volumes (VDI-Images1 through 20, 100 desktops per volume)
- 2,000 desktops were deployed across 20 hosts (100 desktops per host)
- Shared vDisk Images are stored on a separate volume and replicated across the three PVS servers (PVS-Master-1, PVS-Master-2, PVS-Master-3)

6 Test results and analysis

This section presents the results from the different XenDesktop VDI characterization tests and the key findings from each test. The standard user workload represents a majority of the VDI users in the industry today, and the testing was focused on this workload profile.

6.1 Test scenarios

The following tests were conducted to gather results and analysis on the solution stack.

1. **Boot storm:** Boot storms represent the worst-case scenario where many virtual desktops are powered on at the same time and they all contend for the system resources simultaneously. This test was used to evaluate if the storage array hosting the desktops was capable of handling huge spikes in storage I/O without causing significant impact on other services.
2. **Login storm:** Login storms also represent a heavy I/O scenario where many users are logging into their virtual desktops at the beginning of a work day or a shift. In this test, all the desktops were pre-booted and left in an idle state for more than 20 minutes to let their I/O settle before running the Login VSI Medium workload to simulate users logging in to their virtual desktops.
3. **Steady state workload for standard users:** Once the login storm for the previous test was completed, the Login VSI Medium workload was allowed to run for at least one hour to simulate the real world scenario of users performing their daily tasks. The VSImax (Dynamic) parameter from Login VSI evaluated the end-user experience of simulated users working on their virtual desktops throughout these tests.

6.2 2,000 standard-user desktop tests

The following sections provide results from the boot storm, login storm, and steady state testing for a Standard user on a single PS6210XS.

6.2.1 Boot storm I/O

To simulate a boot storm, the virtual desktops were reset simultaneously from the VMware vSphere client. Figure 5 shows the I/O pattern on the pool comprised of the single PS6210XS array.
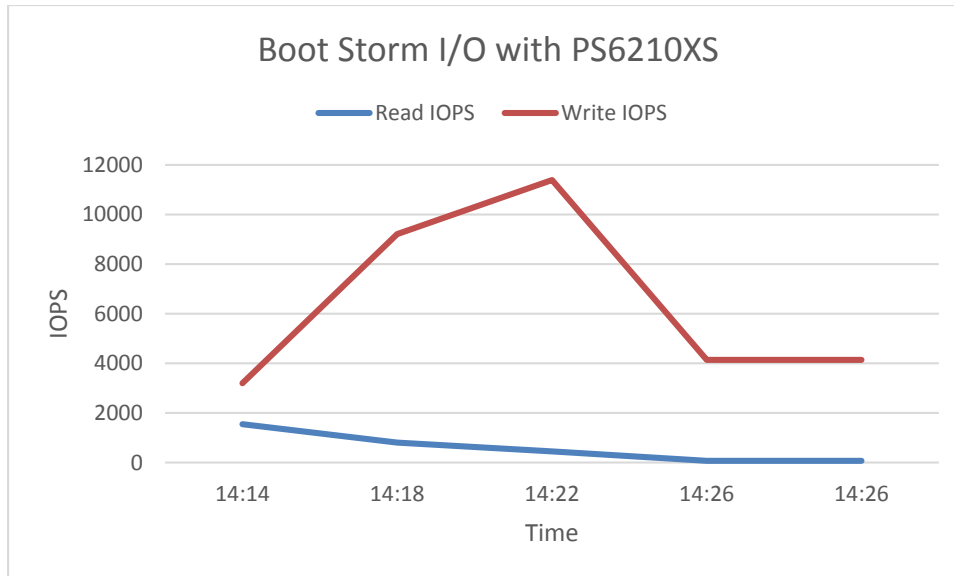
Figure 5   SAN HQ data showing boot storm performance with PS6210XS array

With 2,000 desktops, the boot storm generated about 11,800 total IOPS with a majority of them being write operations. All desktops were available in less than 20 minutes. There are very few read I/O operations due to the fact that the PVS servers are using their RAM as a cache to stream the master image vDisk and no longer need to read from the storage array. This also means that that the only remaining operations are writes that occur from the boot operations on the virtual desktops. The read-write ratio for the pool was 8% reads to 92% writes.

6.2.2   Login storm I/O

Login VSI was programmed to launch 2,000 virtual desktops over a period of about 50 minutes after pre-booting the virtual desktops. The peak IOPS observed during the login storm was about 13,200 IOPS (8-12 IOPS per VM).

Login storms generate significantly more write IOPS than a boot storm or steady state due to multiple factors including:

- User profile activity
- Starting operating system services on the desktop
- First launch of applications

Once a virtual desktop has achieved a steady state after user login, the Windows 7 OS has cached applications in memory and does not need to access storage each time the application is launched. This leads to lower IOPS during the steady state. Figure 6 below shows the IOPS and latency observed during the login storm.
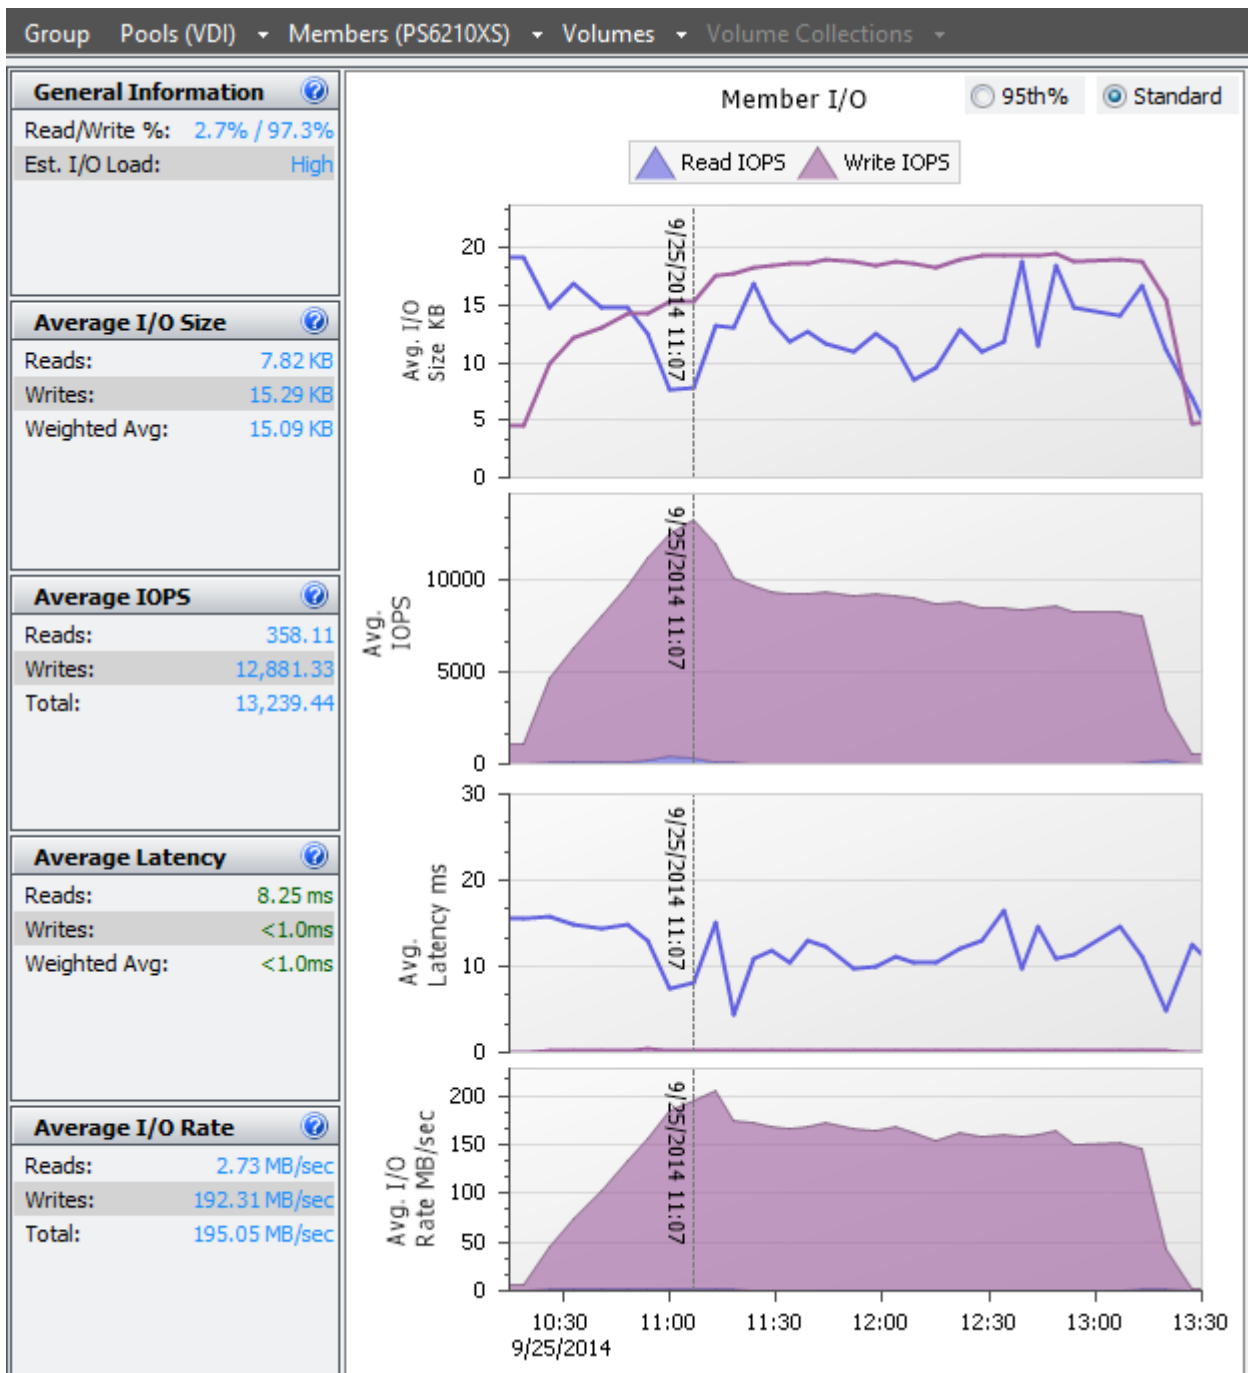
Figure 6   SAN HQ data showing login storm I/O with PS6210XS array

The maximum latency seen in this test is well below the 20 ms criteria (write latency of less than 1 ms with 97% write workload), and the array was able to sustain about 13,200 total IOPS for the login storm. This shows that the array is able to support the 2,000 desktops with no issues.

6.2.3 Steady state I/O

During the steady state for 2,000 desktops, the array was able to easily sustain about 8,700 IOPS that the Login VSI load generator created. The read/write ratio was <1% reads (34 IOPS) and 99% writes (8,636 IOPS). The latency was well below 20 ms (in fact, less than 1 ms for this 99% write workload) indicating that the arrays is able to sustain the 2,000 desktops without any issues.

Figure 7 below shows the performance of the storage pool with single PS6210XS hybrid array during a steady state of 2,000 desktops.
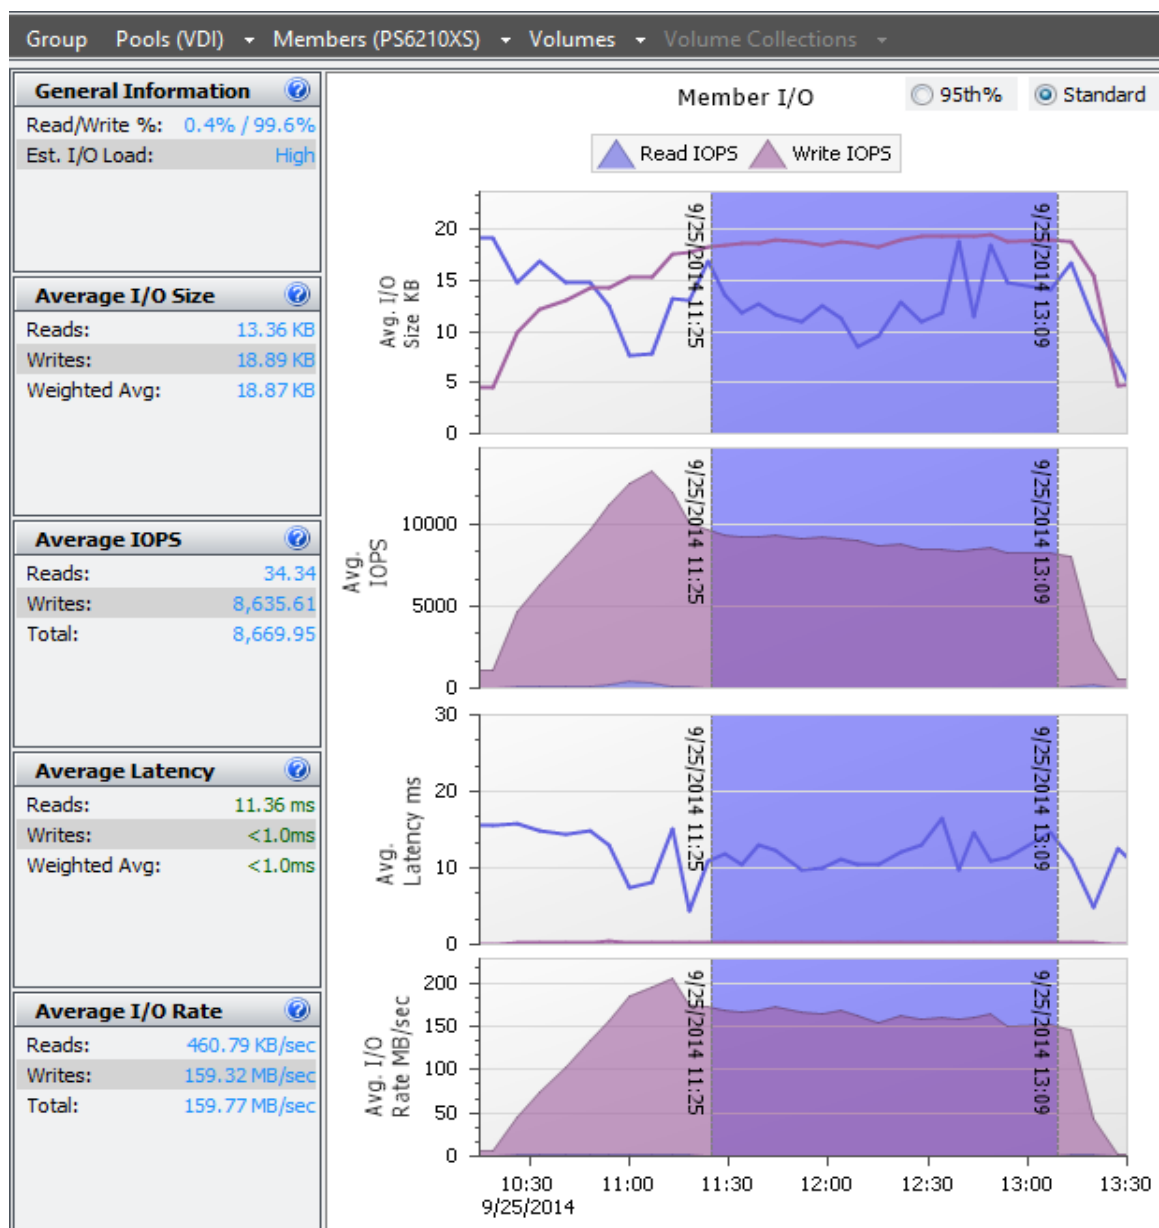
Figure 7   SAN HQ data showing steady state I/O for PS6210XS array

6.3 Server host performance

During the login storm and steady state of the test, the vSphere host CPU, memory, network, and storage performance was measured on all the servers hosting the virtual desktops. The performance of one such vSphere server is presented here. The other vSphere servers had similar performance characteristics.

Statistics for the vSphere hosts were captured using VMware vCenter Server. The figures below show the CPU, memory, and network utilization for boot storm, login storm and steady state of one of the vSphere servers hosting the virtual desktops. The results shown below are for a test run with 2,000 desktops in the desktop pool.

Figure 8   Average CPU performance per core on one vSphere host during login storm and steady state

Figure 9   Memory usage during login storm and steady state

Figure 10   Overall network performance during login storm and steady state
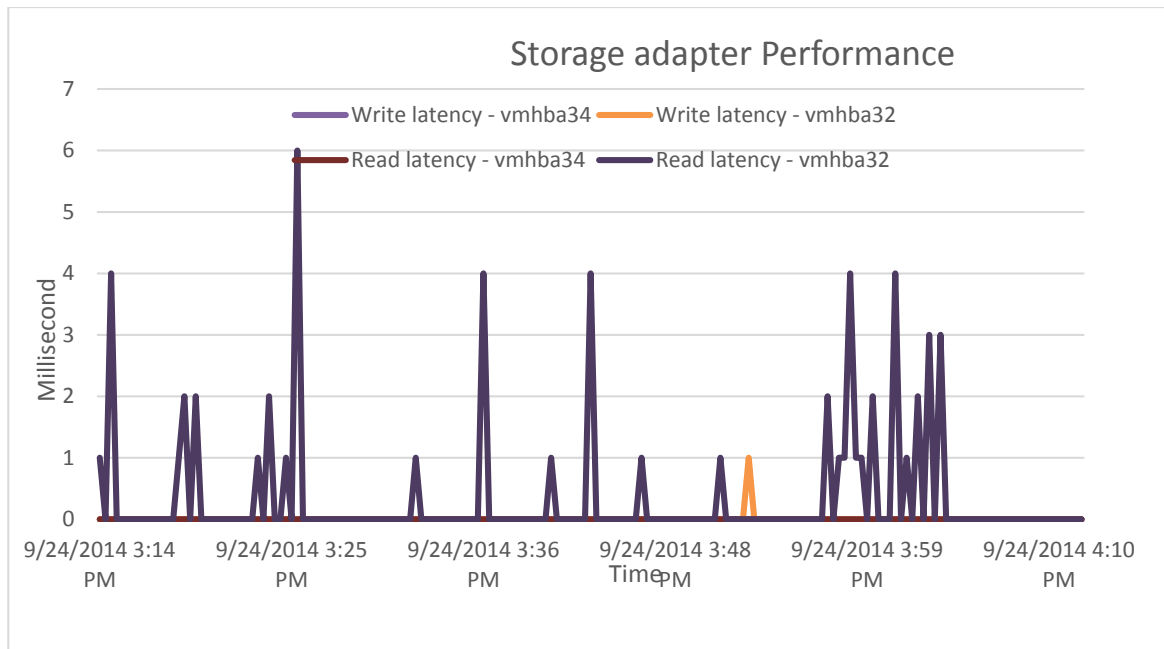
Figure 11   Storage adapter performance during login storm and steady state

The key observations from the statistics were:

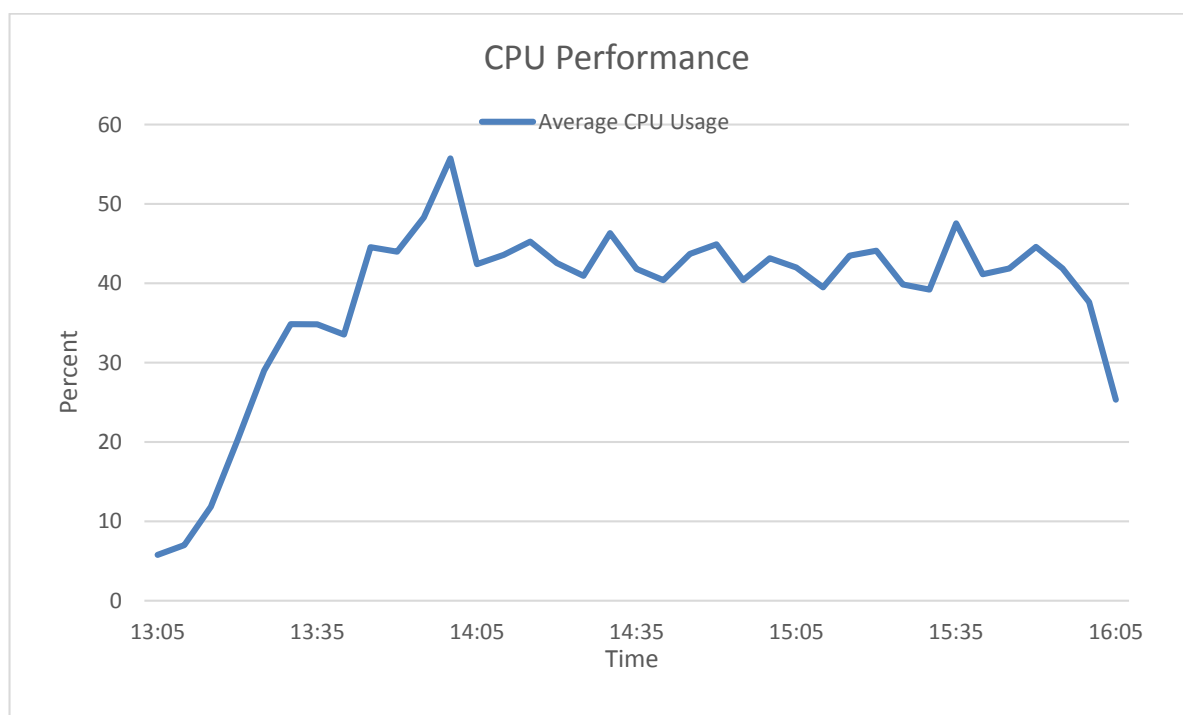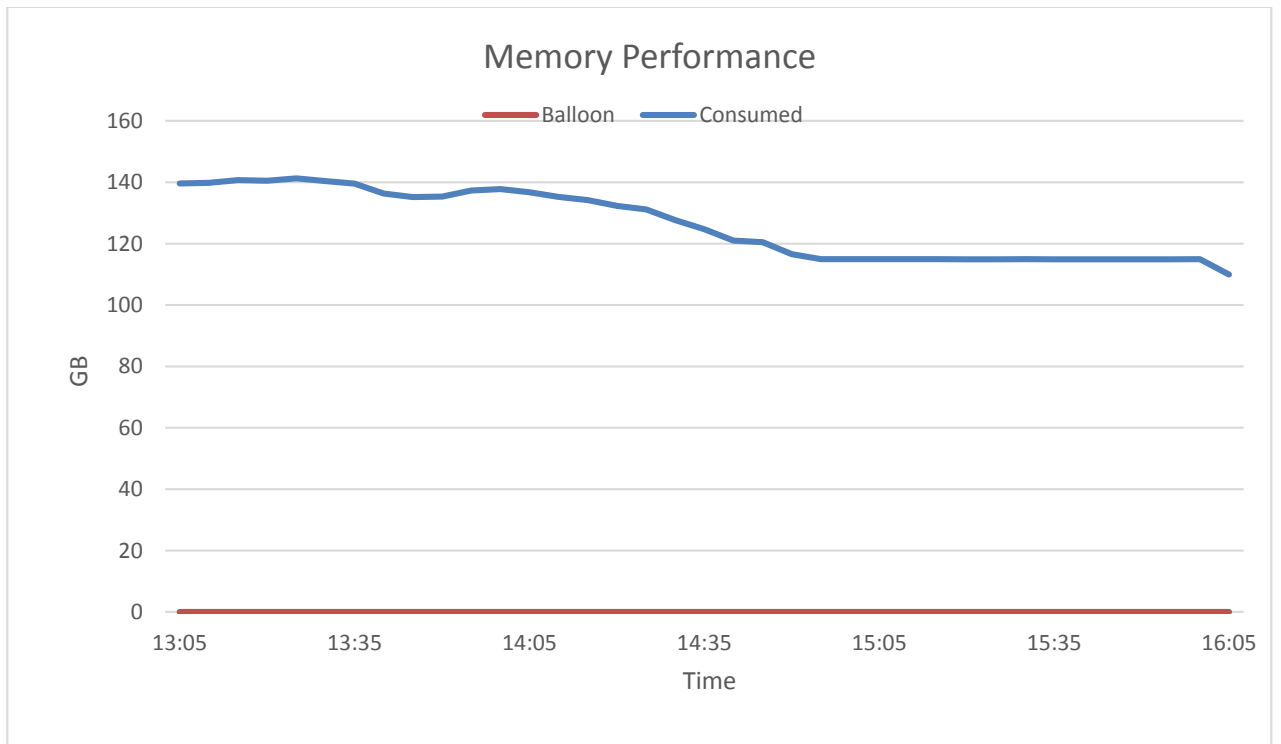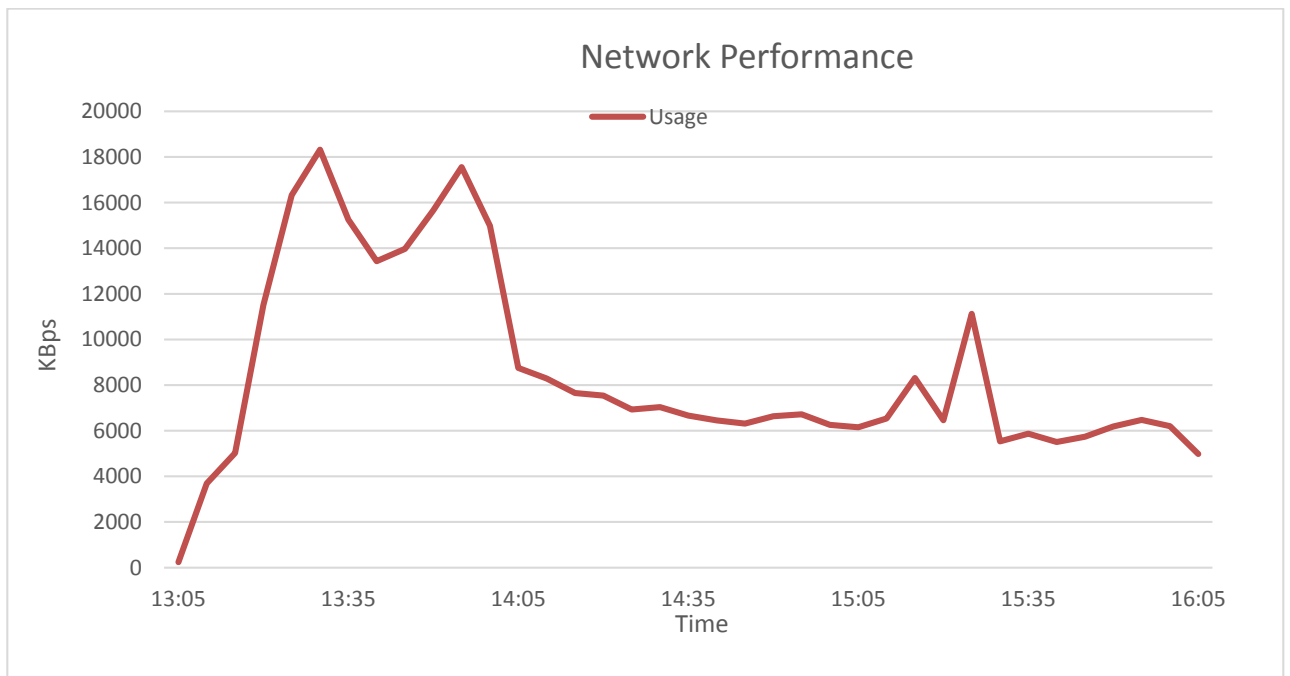- CPU utilization was below the 85% threshold throughout the test.
- Active memory usage was about 80% during the boost storm and about 60% during login storm and steady state. There was minimal or no memory ballooning observed.
- Network utilization was about 45% which included all the networks including iSCSI SAN, VDI LAN, Management LAN and vMotion LAN.
- Average read and write latencies at the storage adapter level were very close to the observed latencies in SAN HQ.

6.4   Results summary

The key observations from the test results are listed below.

- The PS6210XS array is able to host 2,000 virtual desktops and support a standard user type of I/O activity.
- The VDI I/O was mostly write-intensive I/O with more than 98% writes and less than 2% reads.
- None of the system resources on the vSphere servers hosting the virtual desktops reached maximum utilization levels at any time.
- With 2,000 desktops booting simultaneously, the array observed nearly 12,100 IOPS and all desktops were available within 30 minutes of the storm.
- When 2,000 users were logged in within 50 minutes, the PS6210XS storage array easily handled this aggressive login storm with low latency – it delivered approximately 13,200 IOPS with less than 1 ms of write latency with a 99% VDI I/O workload.

7 Best practices

This section provides the best practices derived from the testing and analysis in Section 6.

7.1 Application layer

This section provides best practices for the implementation of VDI using Citrix XenDesktop.

7.1.1 Implement roaming profiles and folder redirection

It is highly recommended that all users in the VDI environment be configured with roaming profiles and folder redirection. This preserves user profiles and user data across boots while using non-persistent virtual desktops.

It is also recommended to use a high performance file service to provide the profile and folder redirection. A separate array may be used to host these volumes for best performance.

7.1.2 Boot and login storm considerations

To avoid I/O bursts due to boot storms, it is recommended that all desktops be pre-booted, preferably with the boots staggered over time before users begin login at the start of a shift or a workday.

It is important to size storage based on the IOPS needs of boot and login storms. The storage subsystem should be designed to handle these storms in addition to the steady state IOPS. Sizing a storage system only on the steady state IOPS is not recommended because this can cause degraded user experience and performance during a boot or login storm.

7.1.3 Windows 7 master image for desktop VMs

It is recommended that the operating system be customized in order to provide the best performance in a VDI environment. This includes disabling some services which may not be required. This can improve performance for the end user. VMware has a specific set of recommendations and settings for Windows 7 that allow for faster logins, quicker screen refreshes, and generally better performance.

> The Citrix recommendations for Windows 7 image customization can be found here:
> http://support.citrix.com/servlet/KbServlet/download/25161-102-648285/XD%20-%20Windows%207%20Optimization%20Guide.pdf

7.1.4 Citrix XenDesktop recommendations

It is recommended to have redundancy and load balancing for all Citrix services, namely, the XenDesktop Controllers, Provisioning Servers, and StoreFront Servers. Ideally these should be running on different hosts to provide high availability.

> Configuration of SQL database for mirroring in a Citrix XenDesktop environment:
> http://support.citrix.com/article/CTX127359

7.2 Server host layer

The vSphere servers hosting the infrastructure service providers and the virtual desktops are recommended to be configured as follows:

- Follow VMware and Dell best practices for installing and configuring vSphere.
- Install and configure EqualLogic Multipathing Extension Module (MEM) for vSphere 5.5 to get the best performance from the storage array.
- Separate virtual switches to segregate iSCSI SAN traffic, VDI traffic, vMotion traffic, and Management network traffic.
- Each network path should be assigned to a minimum of two physical NICs for high availability.

> VMware KB article on best practices for installing vSphere 5.5: http://kb.vmware.com/kb/2052329
>
> Installing and configuring the Dell EqualLogic MEM for VMware vSphere 5: http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19991633.aspx

7.3 Network layer

It is recommended that at least two physical NICs on each server be dedicated to the VDI network or separate the NICS used for the Provisining services.

Use VLANs to segregate different types of network traffic on the same physical network. In this case, it is recommended to separate the infrastructure, vMotion, and VDI LAN traffic into separate VLANs.

Do not use VLANs to segregate iSCSI SAN traffic. It is required that SAN traffic be on a separate physical network to provide the best performance. The recommended way to provide a converged fabric for iSCSI SAN traffic is through the use of Data Center Bridging (DCB) technologies which allow the iSCSI SAN to have lossless end-to-end connections while guaranteeing a minimum bandwidth.

Virtual switches in vSphere have a default limit of 120 ports. If the number of virtual desktops on each host exceeds the available ports, vSwitch properties should be changed to support the required number of virtual desktops. This change requires a reboot of the host vSphere server.

On iSCSI SAN switches, Spanning tree should be disabled on switch ports connected to end devices for server and storage ports. The PortFast setting should be enabled in the switch configuration for these ports. Jumbo frames and Flow control (if the NICs support it) should be enabled for all components of the iSCSI network.

> More information on configuring Dell Networking switches for use with Dell Storage iSCSI SANs is available here:
> http://en.community.dell.com/techcenter/storage/w/wiki/4250.switch-configuration-guides-by-sis.aspx

7.4   Storage

For VDI environments, it is recommended to use the Dell Storage PS Series hybrid arrays, which consist of SSD drives and 10,000 RPM SAS drives within a single chassis. These hybrid arrays automatically move hot data to the SSD tier, which improves performance in VDI environments in a cost-effective way.

It is recommended to have separate volumes for base images and virtual desktops. This aids in better manageability of the volumes, easier performance monitoring, and allows for easy future growth.

It is recommended to use a separate high performance file service to provide file shares for roaming profiles and user shares.

8   Conclusions

The paper demonstrates how a 2,000 standard user virtual desktop environment can be deployed using a Citrix XenDesktop VDI platform leveraging a single Dell Storage PS6210XS hybrid array. The storage I/O characteristics under various VDI workload scenarios (boot storm, login storm and steady state) along with performance characteristics throughout the VDI stack including vSphere server performance and user experience demonstrate the optimal configuration of the PS6210XS hybrid array for this VDI deployment.

This testing of the Dell Storage PS Series storage platform showed that the PS6210XS hybrid array can easily support very fast access to high-demand data in a VDI workload environment. It can easily handle high IOPS spikes that occur during boot and login storms. With PS Series hybrid arrays automatically tiering data between SSDs and HDDs, the solution can sustain large IOPS needs with a relatively small amount of SSD capacity thereby cost-effectively boosting the storage performance in a VDI deployment.

With the ability to support a large number of desktops in a single array, the PS6210XS can help organizations that are limited by the cost-per-desktop hurdle start their VDI deployment plans much sooner. Moreover, the peer storage architecture of the PS Series arrays achieves linear scalability of both controller throughput and storage capacity as new arrays are added to accommodate VDI environment growth. This linear scaling in both performance and capacity keeps the storage-per-virtual-desktop cost low even in very large VDI deployments.

In addition to improved performance and flexible scalability, the PS6210XS hybrid arrays help reduce VDI deployment costs by handling workloads on highly cost-effective storage media, enhance VDI performance by placing the vDisks and Write Cache volumes on low-latency SSDs as needed, and simplify VDI deployments through automated and intelligent data tiering. Combined with the ease of administration of PS Series arrays, these benefits reduce the total cost for VDI in many organizations and make PS6210XS a great storage platform for desktop virtualization.

A Citrix XenDesktop 7.5 solution configuration

Table 6   Hardware components

Solution configuration - Hardware components:		Description
Blade enclosure	3 x Dell PowerEdge M1000e blade chassis	All components, (storage, networking and compute), for this solution are held within this blade chassis.
Virtual Desktops	20 x Dell PowerEdge M620 Servers: • BIOS Version: 2.27 • 2 x Intel® Xeon® E5-2650 2.00 GHz Processors • 196 GB RAM • 2 x 300 GB 10K SAS internal disk drives • 1 x Dual-port Broadcom NetXtreme II 57810 10GbE NIC (LAN on motherboard) • 1 x Dual-port Broadcom NetXtreme II 57810 10GbE NIC	vSphere 5.5 with Dell EqualLogic MEM 1.2.0 installed on all M620 servers Windows 7 (64 Bit) VMs with 3 GB RAM
Infrastructure Servers	2 x Dell PowerEdge M620 servers: • BIOS Version: 2.27 • 2 x Intel Xeon E5-2650 2.00 GHz Processors • 128 GB RAM, • 2 x 300 GB 10K SAS internal disk drives • 1 x Dual-port Broadcom NetXtreme II 57810 10GbE NIC (LAN on motherboard) • 1 x Dual-port Broadcom NetXtreme II 57810 10 GbE NIC	vSphere 5.5 with Dell EqualLogic MEM 1.2.0 installed on all M620 servers VMs are loaded with Microsoft Windows Server 2012 R2. • Microsoft Active Directory • Citrix XenDesktop Delivery Controller 7.5 Server 1, 2 • Citrix StoreFront 2.5 Server 1,2 • Citrix Provisioning Services 7.1 Server 1,2 • Citrix Licensing Server 11.11 Server • VMware vCenter Server 5.5 • Microsoft SQL Server • Microsoft File Server
Login VSI Launchers	8 x Dell PowerEdge M610 servers • BIOS Version: 2.7.0 • 2 x Intel Xeon E-5520 2.26 GHz Processors • 96 GB RAM • 2 x 146 GB 10K SAS internal disk drives • 1 x Dual-port Broadcom NetXtreme II 57711 10 GbE NIC	vSphere 5.5 installed on all M610 servers VMs are loaded with Microsoft Windows Server 2012R2 64bit Each VM supports 20 Login VSI launcher sessions.

Network	• Management, iSCSI SAN 8 x Dell Force10 MXL 10Gb Ethernet Switch (DCB enabled) 2 x Dell Force10 S6000 40Gb Ethernet Switch (DCB Enabled) • VDI LAN 4 x Dell Force10 MXL 10Gb Ethernet Switch 2 x Dell Force10 S4810 10Gb Ethernet Switch	VLT LAG interconnect for HA and load balancing Firmware: 9.5.0.0
Storage	1 x Dell EqualLogic PS6210XS Hybrid Storage Array: • 7 x 800 GB SSD • 17 x 1.2 TB 10K SAS disks • Dual 10GbE controllers • Firmware: 7.0.7 1 x Dell EqualLogic PS6110E: • 24 x 500 GB 7.2K NL-SAS disks • Dual 10GbE controllers • Firmware: 7.0.7	Virtual Desktops are all stored on the PS6210XS storage arrays. User Data and Infrastructure VMs are stored on the PS6110E storage array
Performance monitoring	• SAN HQ • vCenter Performance monitoring • Login VSI Analyzer	Performance monitoring on Dell PS Series storage arrays Performance monitoring and capture at the vSphere host User experience monitoring on the virtual desktop

Table 7   Software components

Solution configuration - Software components:	Description / Version
Citrix XenDesktop Delivery Controller	7.5.0.4531
Citrix Provisioning Services	7.1.0.4022
Citrix Receiver	4.1.0.56461
Citrix Studio	7.5.0.4531
Citrix License Server	11.11.1 build-13012
Citrix StoreFront	2.5.0.29
VMware vSphere Hypervisor	5.5.0 build-1331820
VMware vCenter Server	5.5.0 build-1623101

EqualLogic Multipathing Extension Module	1.2.0.365964
SAN Headquarters	3.0.0.7189
Microsoft Windows 7 Enterprise x64 SP1	VDI Clients for characterization tests,
Microsoft Windows Server 2012 R2 x64	VMs for hosting Citrix XenDesktop Components, VMware vCenter Server, MS-SQL server, and other infrastructure VMs
Login VSI	Version 4.0.11.500, Load generator for VDI clients and user experience analysis
Microsoft Office 2007 Standard	Service Pack 3 (12.0.6607.1000)

B   vSphere host network configuration

Each vSphere host was configured with four virtual switches - vSwitch0, vSwitch1, vSwitch2, and vSwitch3.

B.1   vSwitch0

vSwitch0 provides connection paths for all management LAN traffic. The first partition of the physical adapters from the two on board NICs (Fabric A) was assigned to this switch.
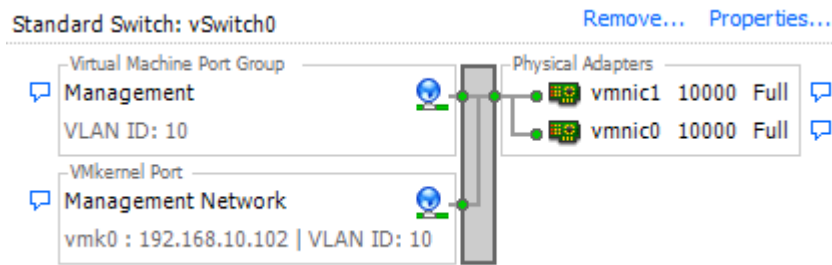

Figure 12   vSwitch0 – Management LAN

B.2   vSwitch1

Two partitions from the physical adapters in Fabric A were assigned to this virtual switch. This vSwitch carries the traffic required to provide VMware vMotion services.
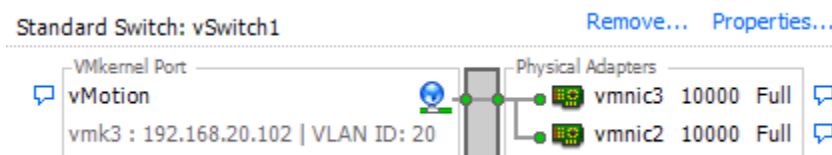

Figure 13   vSwitch1 – vMotion LAN

B.3   vSwitch2

This virtual switch provided paths for all the iSCSI SAN traffic. Two partitions of the physical adapters from the two on board NICS (Fabric A) were assigned to this virtual switch.
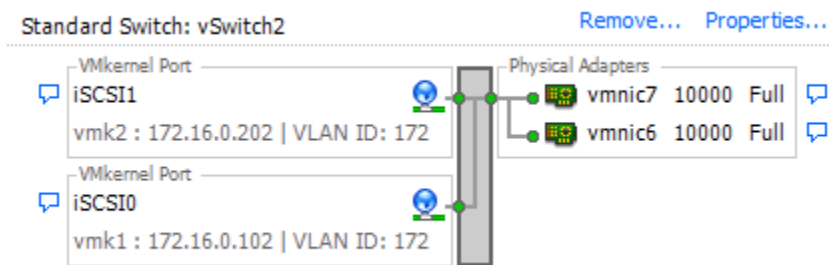
Figure 14   vSwitch2 – iSCSI virtual switch

B.4   vSwitch3

Two 10 Gb physical adapters in Fabric B are assigned to this virtual switch. This vSwitch carried all the traffic for the VDI LAN.
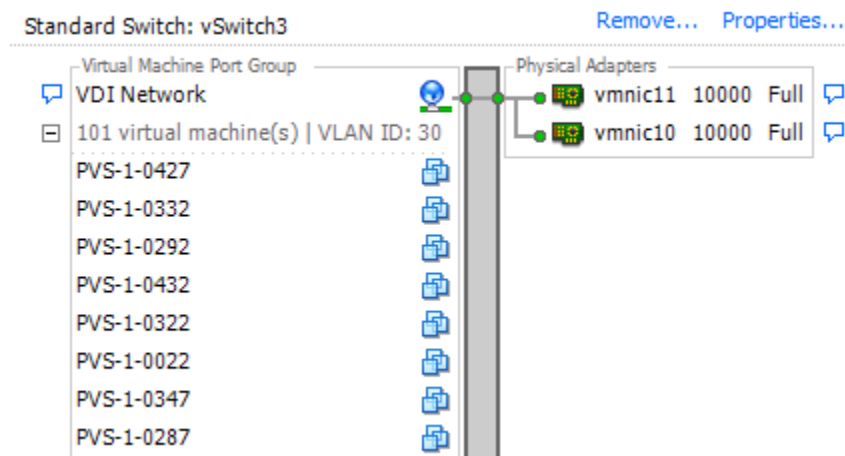

Figure 15   vSwitch3 – VDI LAN

C   User experience monitoring with Login VSI

Login VSI uses the VSImax parameter to determine the maximum number of sessions that can be obtained from a deployed solution.

The calculation methodology used in VSImax is available here:
http://www.loginvsi.com/documentation/index.php?title=Analyzing_Results

Figure 16 below shows the response time as experienced by each user as calculated by Login VSI. The graph clearly shows that the VSImax (Dynamic) value has not been reached and there were no hung or inactive sessions. This means that the array is capable of supporting the 2,000 desktops and all users have acceptable response times for running applications.
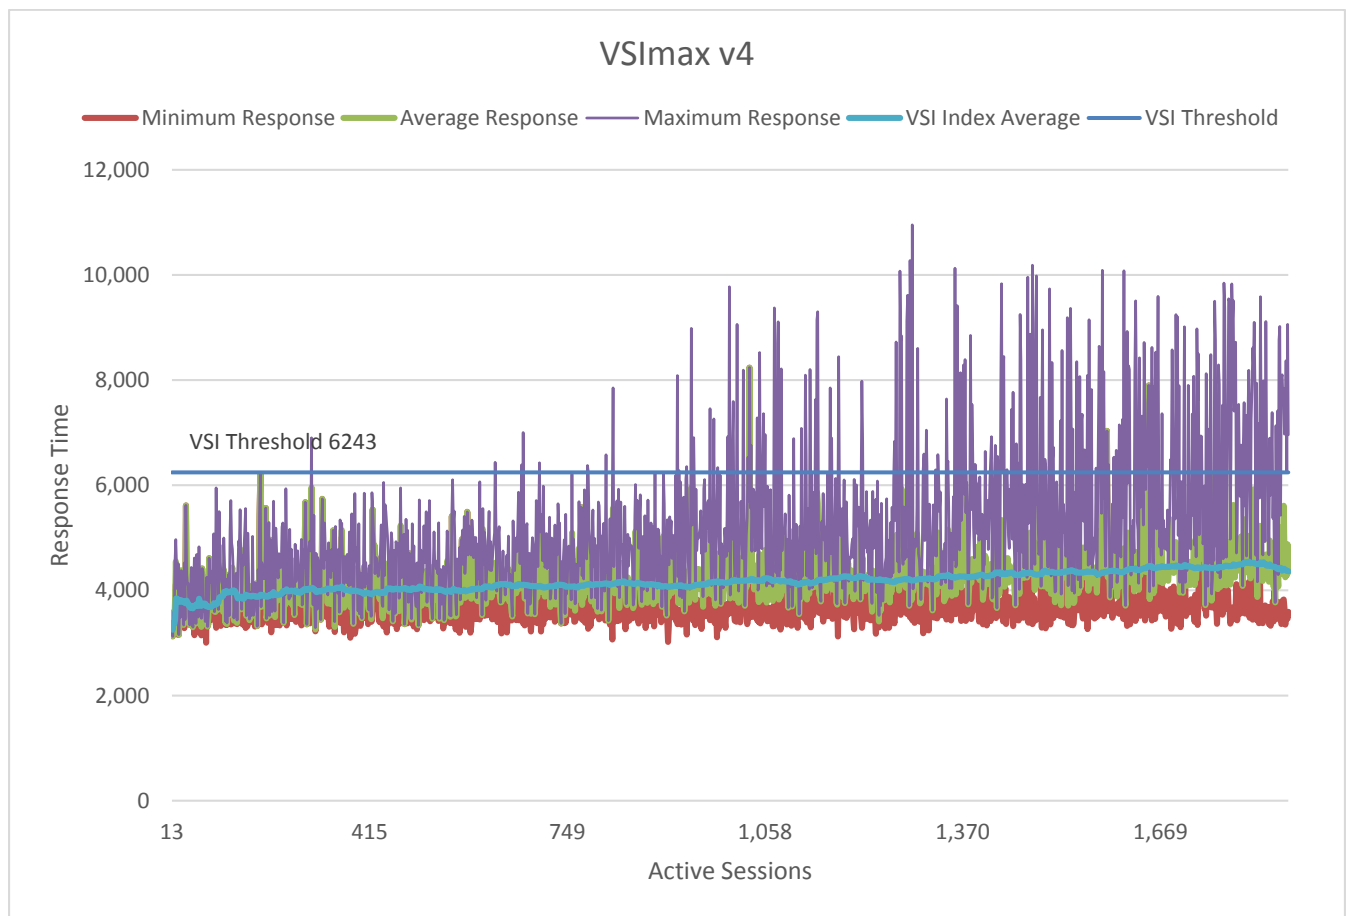

Figure 16   Login VSI generated graph showing average response time per user and VSImax for 2,000 VMs

D Workload characterizations

The Client Cloud Computing group at Dell defines three primary types of workloads for a VDI solution.

Basic workload: This workload profile consists of simple application tasks. Typically a repetitive application use profile with a non-personalized desktop image. Sample use cases may be kiosks or call-centers which do not require a personalized desktop environment and where the application stack is static. The workload requirements for a Basic user are the lowest in terms of CPU, memory, network, and disk I/O and allow the greatest density and scalability of the infrastructure.

Standard workload: This workload profile consists of email, typical office productivity applications, and web browsing for research/training. There is minimal image personalization required in a Standard user workload profile. The workload requirement for a Standard user is moderate and most closely matches the majority of office worker profiles in terms of CPU, memory, network, and disk I/O and will allow moderate density and scalability of the infrastructure.

Premium workload: This workload profile represents an advanced knowledge worker. All office applications are configured and used. The user is typically working with moderate-to-large file sizes and may include some graphics creation or editing. The Premium user requires extensive image personalization for shortcuts, macros, and menu layouts. The workload requirements for a Premium user are heavier than typical office workers in terms of CPU, memory, network, and disk I/O and allow limited density and scalability of the infrastructure.

The three types of user workloads are summarized in Table 8 below.

Table 8   User workload characterizations

User workload	VM memory allocation	vCPUs assigned	Approximate IOPS	User data disk space	Simultaneous Applications	Login VSI session idle time
Basic	1 GB	1	7-8	5 GB	2	105s
Standard	1.5 GB	1	8-9	5 GB	5	120s
Premium	2.5 GB	1	9-10	5 GB	8	120s

More information about the user workload types can be found in Appendix A.3 of the document *Introduction to the Dell DVS Enterprise Solution* here:
http://i.dell.com/sites/doccontent/business/solutions/engineering-docs/en/Documents/dvs-enterprise-6020.pdf

Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees to share knowledge, best practices, and information about Dell products and your installations.

Referenced or recommended Dell publications:

- Sizing and Best Practices for Deploying VMware View 5.2 on VMware vSphere 5.1 with Dell EqualLogic Storage:
  http://en.community.dell.com/techcenter/extras/m/white_papers/20438105/download.aspx
- Dell EqualLogic Configuration Guide:
  http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19852516/download.aspx
- Dell EqualLogic PS Series arrays – Scalability and Growth in Virtual Environments:
  http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19992296/download.aspx
- Installing and configuring the Dell EqualLogic MEM for VMware vSphere 5:
  http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19991633/download.aspx
- Dell EqualLogic PS Series Architecture: Load Balancers:
  http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19949521/download.aspx
- Dell Networking MXL 10/40GbE Blade Switch Configuration Guide:
  http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20279157/download.aspx

The following VMware publications are referenced in this document or are recommended sources of additional information:

- VMware Horizon View 5.2 Documentation:
  http://pubs.vmware.com/view-52/index.jsp
- VMware Horizon View 5.2 Release notes:
  http://www.vmware.com/support/view52/doc/horizon-view-52-release-notes.html
- VMware View Optimization Guide for Windows 7:
  http://www.vmware.com/resources/techresources/10157
- VMware KB article on best practices for installing vSphere 5.5:
  http://kb.vmware.com/kb/2052329
- VMware blog entry on Optimizing Storage with View Storage Accelerator:
  http://blogs.vmware.com/euc/2012/05/optimizing-storage-with-view-storage-accelerator.html
- VMware blog entry on View Storage Accelerator – In Practice:
  http://blogs.vmware.com/euc/2012/05/view-storage-accelerator-in-practice.html

For EqualLogic best practices white papers, reference architectures, and sizing guidelines for enterprise applications and SANs, refer to Storage Infrastructure and Solutions Team Publications at:

- http://dell.to/sM4hJT