



Sizing and Best Practices for Deploying VMware View 5.1 on VMware vSphere 5.0 U1 with Dell EqualLogic Storage

A Dell EqualLogic Best Practices Technical White Paper

This document has been archived and will no longer be maintained or updated. For more information go to the [Storage Solutions Technical Documents page on Dell TechCenter](#) or contact support.

Storage Infrastructure and Solutions Engineering

Dell Product Group
October 2012

© 2012 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.

Dell, the Dell logo, and the Dell badge, PowerConnect™, EqualLogic™, PowerEdge™ and PowerVault™ are trademarks of Dell Inc. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Intel® is a registered trademark of Intel Corporation in the U.S. and other countries. VMware®, Virtual SMP®, vMotion®, vCenter®, and vSphere® are registered trademarks or trademarks of VMWare, Inc. in the United States or other countries. Microsoft®, Windows®, Windows Server®, Internet Explorer®, SQL Server, and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries



Table of contents

Acknowledgements.....	6
Feedback	6
Executive summary	7
Audience.....	8
1 Introduction	9
2 Desktop Virtualization with VMware View	10
2.1 VMware View components	10
2.2 Virtual desktops	10
2.3 VMware View desktop pools	10
2.4 Using linked clones	11
2.5 VMware View Storage Accelerator.....	11
3 Infrastructure and test configuration.....	13
3.1 Component design	13
3.2 iSCSI SAN configuration.....	14
3.3 EqualLogic storage array configuration	15
3.4 ESXi host network configuration	16
3.5 VMware View configuration	16
3.6 VMware View Pool configuration.....	16
3.7 Windows 7 VM configuration.....	17
4 View test methodology	18
4.1 Test objectives	18
4.2 Test approach	18
4.3 Test tools	18
4.3.1 Load generation.....	18
4.3.2 Monitoring tools	19
4.4 Test criteria	19
4.4.1 Storage capacity and I/O latency	19
4.4.2 System utilization at the hypervisor	19
4.4.3 Virtual desktop user experience.....	19
4.5 Test setup.....	20
5 Test results and analysis.....	22
5.1 Test scenarios	22



- 5.2 View Storage Accelerator/CBRC benefits..... 22
- 5.3 Boot storm I/O..... 24
- 5.4 Task worker – Login storm and steady state I/O 28
 - 5.4.1 Login storm I/O..... 28
 - 5.4.2 Steady state I/O..... 30
 - 5.4.3 Server host performance..... 31
 - 5.4.4 User experience monitoring..... 33
- 5.5 Results summary 34
- 6 Sizing guidelines for EqualLogic SANs 35
- 6.1 Capacity sizing considerations 35
 - 6.1.1 VMware View Accelerator sizing..... 36
 - 6.1.2 Base image datastore..... 36
 - 6.1.3 Replica datastore 36
 - 6.1.4 Linked clone datastores..... 36
- 6.2 Performance sizing considerations..... 37
 - 6.2.1 Boot storm sizing considerations 37
 - 6.2.2 Login storm and steady state sizing considerations 38
- 7 Best Practices..... 39
- 7.1 Application layer 39
 - 7.1.1 Implement roaming profiles and folder redirection..... 39
 - 7.1.2 Boot and login storm considerations 39
 - 7.1.3 Windows 7 master image for desktop VMs..... 39
 - 7.1.4 SSL certificate requirements..... 39
 - 7.1.5 VMware View recommendations..... 39
- 7.2 Server host layer 39
- 7.3 Network layer 40
- 7.4 Storage 40
- 8 Conclusions 42
- Appendix A VMware View 5.1 solution configuration..... 43
- Appendix B Network design and VLAN configuration..... 45
 - B.1 Management LAN configuration 45
 - B.2 VDI LAN configuration..... 45
- Appendix C ESXi host network configuration..... 47
 - C.1 vSwitch0..... 47



Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions of Dell Inc.

The team that created this whitepaper:

Sujit Somandepalli, Chhandomay Mandal, and Margaret Boeneke

We would like to thank the following Dell team members for providing significant support during development and review:

David Glynn, Ananda Sankaran, Keith Swindell, and Mark Welker

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 lower the Total Cost of Ownership (TCO) of desktop devices in all organizations, small and large. It is made even more attractive by the potential to make data more secure and improve desktop deployment flexibility.

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™ EqualLogic™ hybrid arrays – EqualLogic PS6100XS and PS6110XS – are well suited for supporting VDI deployments because they are specifically designed for VDI type workloads where a relatively small portion of the storage system can become very hot (very high read I/O rates). With high performance SSD drives and high capacity SAS hard disk drives within the same chassis, the hybrid arrays automatically move hot data from the SAS tier to the SSD tier. The automatic tiering function makes the hybrid EqualLogic 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 includes the results of a series of storage I/O performance tests and provides storage sizing guidance and best practices based on those results for designing and deploying VMware® View™ 5.1 (View) based VDI using the EqualLogic PS6100XS array as the primary storage platform. The test infrastructure included VMware View 5.1, VMware vSphere 5.0 U1, Dell PowerEdge™ blade servers, Dell Force10™ and PowerConnect™ switches, and Dell EqualLogic storage devices.

The paper demonstrates how an 830 task worker virtual desktop environment can be deployed in the chosen VDI infrastructure leveraging a single PS6100XS array. 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 (ESXi server performance and user experience as determined by Liquidware Labs Stratusphere UX). In the test environment, the PS6100XS array delivered approximately 14,500 IOPS during a boot storm with a read/write ratio of about 76% / 24%. More detailed capacity and sizing guidelines are provided in Section 6 and best practices are provided in Section 7.



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 View based VDI Solutions using Dell EqualLogic storage. It is expected that the reader has a working knowledge of the View architecture, VMware vSphere system administration, iSCSI SAN network design, and Dell EqualLogic iSCSI SAN operation.

1 Introduction

Desktop virtualization platforms like VMware View 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 storage sizing guidance and best practices based on those results for a VMware View based VDI solution.

This paper is an update to the previous VMware View paper, *Sizing and Best Practices for Deploying VMware View VDI with EqualLogic Storage*, and uses the latest generation of VMware View software and ESXi hypervisor on Dell infrastructure. Please refer to the following link for information on the previous VMware View solution document.

Sizing and Best Practices for Deploying VMware View VDI with EqualLogic Storage:

<http://en.community.dell.com/techcenter/storage/w/wiki/2642.sizing-and-best-practices-for-deploying-vmware-view-vdi-with-equallogic-storage-by-sis.aspx>

Additionally, you can also review the sizing and best practices paper for deploying virtual desktops with Dell EqualLogic Virtual Desktop Deployment Utility in a VMware View environment.

Dell EqualLogic Virtual Desktop Utility - Sizing and Best Practices:

<http://en.community.dell.com/techcenter/storage/w/wiki/3621.dell-equallogic-virtual-desktop-utility-sizing-and-best-practices-by-sis.aspx>

The test infrastructure used includes:

- VMware View 5.1
- VMware vSphere 5.0 U1 hypervisor
- Dell PowerEdge M610 blade servers
- Dell Force10 and PowerConnect switches
- Dell EqualLogic Storage

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

- Develop best practices and sizing guidelines for a View 5.1 based VDI solution deployed on EqualLogic PS series storage, Dell Force10 and PowerConnect switches, and Dell PowerEdge blade servers with ESXi 5.0 as the server virtualization platform.
- Determine how many virtual desktops can be deployed in this environment using a single PS6100XS array with acceptable user experience indicators for a task worker workload profile.
- Determine the performance impact on the storage array during peak I/O activity such as boot and login storms.

2 Desktop Virtualization with VMware View

VMware View is a VDI solution that includes a complete suite of tools for delivering desktops as a secure, managed service from a centralized infrastructure. A View infrastructure consists of many different software, network, and hardware layer components.

2.1 VMware View components

A functional list of View components used in this solution is shown below:

Client Devices	Personal computing devices used by end users to run the View Client.
View Connection Server	A software service that acts as a broker for client connections by authenticating and then directing incoming user requests to the appropriate virtual desktop, physical desktop, or terminal server.
View Client	Software that is used to access View desktops.
View Agent	A service that runs on all systems used as sources for View desktops and facilitates communication between the View Clients and View Server.
View Administrator	Web-based administration platform for View Infrastructure components.
vCenter Server	Central administration platform for configuring, provisioning, and managing VMware virtualized datacenters.
View Composer	A service running on the View Servers used to create pools of virtual desktops from a shared base image to reduce storage capacity requirements.

2.2 Virtual desktops

Virtual Desktops can be classified into two major categories: persistent and non-persistent. In a persistent desktop, a user keeps all configuration and personalization between sessions on their assigned desktop. When using persistent desktops, the administrator usually has to provision additional storage in addition to other administrative requirements like patching and upgrading of desktops.

In a non-persistent desktop environment, users are assigned virtual desktop VMs from a pool of resources during login. This type of virtual desktop does not maintain any user data or other information between sessions. At logoff, all the changes are simply discarded and the virtual desktop is returned to the original state. Patching and upgrading non-persistent desktops is as simple as making the change to the base image and redeploying the pool. Thus, these desktops are easier to manage but lack the potential for persistent user customization.

2.3 VMware View desktop pools

A Desktop Pool is a VMware term used to describe a collection of desktops that is managed as a single entity by the View Administration interface. View Desktop Pools allow administrators to group users depending on the type of service the user requires. There are two types of pools – Automated Pools and Manual Pools.

In View, an Automated Pool is a collection of VMs cloned from a base template, while a Manual Desktop pool is created by the View Manager from existing desktop sources, physical or virtual. For each desktop in the Manual Desktop pool, the administrator selects a desktop source to deliver View access to the clients.

Microsoft® Active Directory® folder redirection and roaming profile features can be used to achieve persistent desktop-like behavior on non-persistent desktops. By using these features, the administrator can design a user account so that all the configuration settings can be written to a remote profile that is stored separately from the virtual desktop image files. This reduces the need for additional management on individual virtual desktops while still providing a custom user experience. The disadvantage in using this method is the additional time required to sync the remote profiles at logon and logoff which can be substantial if many users are performing this task simultaneously. This can be improved by using a high performance file service appliance like the Dell EqualLogic FS series Unified NAS appliance.

For more information on the various types of View Pools, see: <http://pubs.vmware.com/view-50/topic/com.vmware.view.administration.doc/GUID-0A9CA985-3A78-428A-BCFB-B3E2DCCA90AD.html>

2.4 Using linked clones

Significant storage space savings and increased efficiencies in desktop VM provisioning and administration are possible when using VMware linked clones. A linked clone is a duplicate VM that shares the same base image with the original VM, but has separate differential data disks to track the differences from the original one. Each linked clone functions as an independent desktop VM with its own unique identity.

Because linked clones share the base image, they consume significantly less storage space than a set of complete independent VM images. Temporary system data and other data unique to each linked clone desktop VM are written to separate differential data storage, and these temporary data are discarded at reboot and/or user logoff. Persistent data such as user profiles and application / user data can be optionally redirected to a CIFS share. With this model, software maintenance updates, antivirus remediation, and patches need to be applied only on the base image. These base image changes will automatically take effect on all linked clones without affecting any unique user settings and data.

To configure linked clones, the administrator creates a parent VM image with the required OS, settings, and software installed. View Composer first creates a full replica of the parent VM, and then uses this replica to create linked clones. The replica can be placed on the same datastore as the linked clones or on a separate datastore.

More information on configuring linked clones can be found in the "Creating Desktop Pools" section of the VMware View Online Library here: <http://pubs.vmware.com/view-51/topic/com.vmware.view.administration.doc/GUID-0A9CA985-3A78-428A-BCFB-B3E2DCCA90AD.html>

2.5 VMware View Storage Accelerator

VMware introduced a set of new features and functionalities in View 5.1. From a storage perspective, the most important of these is *View Storage Accelerator*, which optimizes storage load and improves performance by caching common disk blocks while reading virtual desktop images. View Storage

Accelerator uses the VMware vSphere platform feature called Content Based Read Cache (CBRC) implemented inside the ESXi hypervisor.

CBRC is an in-memory cache of common blocks. When enabled for VMs used as virtual desktops, the ESXi host scans the storage disk blocks to generate digests of block contents. When the blocks are first read into the ESXi server, they are cached in the CBRC of the ESXi server. Subsequent reads of blocks with the same digest are served directly from the in-memory cache of the server. This results in significant read performance improvements, specifically for boot storms when a large number of VMs simultaneously access a large number of blocks with identical content.

The test results in Section 5.2 demonstrate the impact of View Storage Accelerator on an 800 desktop deployment with CBRC enabled and disabled. These tests confirm the benefits of CBRC in read-intensive VDI workload scenarios like boot storms. All the other results reported in this paper are obtained with CBRC enabled for virtual desktop VMs.

3 Infrastructure and test configuration

This section provides information on the test setup and configuration used for hosting View virtual desktops, including infrastructure components, networking, and storage subsystems.

3.1 Component design

The entire infrastructure and test configuration was installed in a Dell PowerEdge M1000e blade chassis because it delivers better performance per watt, increases capacity, and lowers operating costs with its integrated management tools for the blade environment. As shown in Figure 1 below, the 16 PowerEdge M610 blade servers used in this solution were divided into three ESXi clusters.

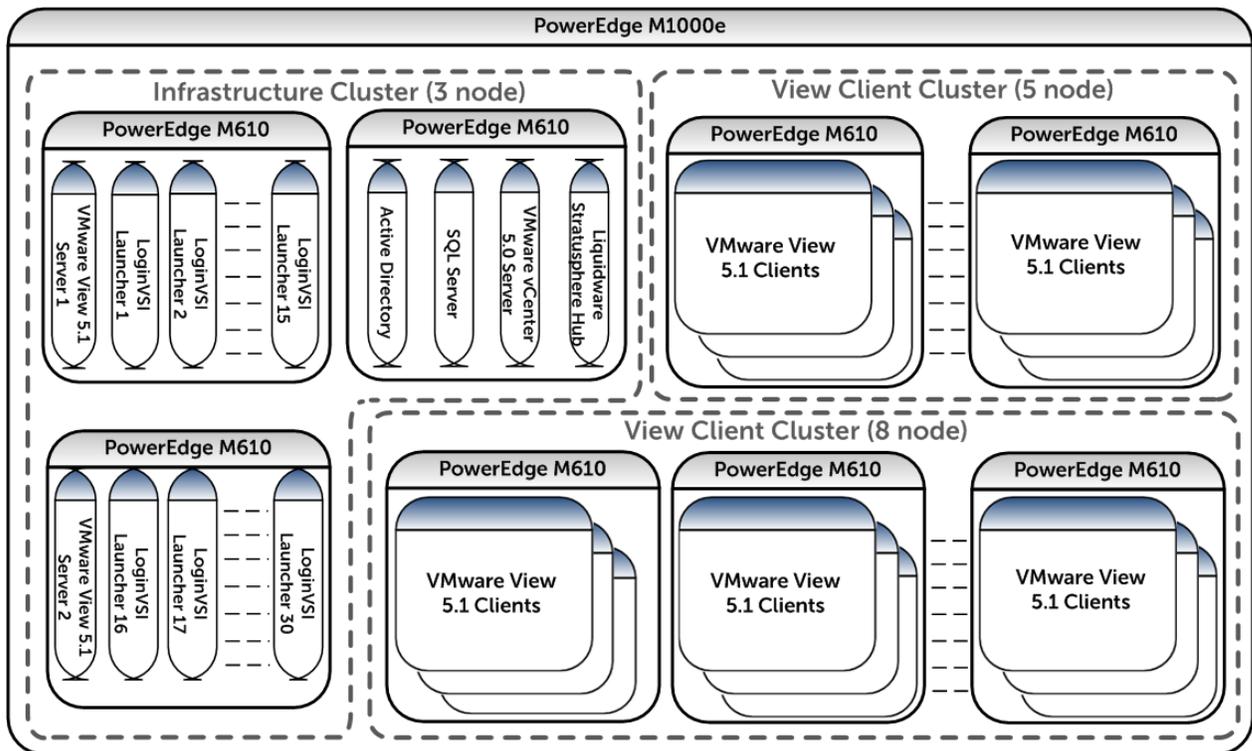


Figure 1 ESXi Blade Clusters on PowerEdge M1000e

These three clusters included:

- **Infrastructure Cluster:** Three M610 blade servers hosting virtual machines for Active Directory services, VMware vCenter 5.0 Server, VMware View 5.1 Server (primary and secondary), Windows Server 2008 R2 file server, SQL Server 2008R2, and 30 Login VSI launchers.
- **Two View Clusters:** One with eight M610 blade servers and another with five M610 blade servers, because VMware View allows a maximum of eight servers in a non-NFS cluster.

The backplane of the blade chassis consisted of the following switches:

- Two PowerConnect M6220 blade switches in fabrics A1 and A2 for connectivity to a Management LAN.
- Four PowerConnect M6348 blade switches in fabrics B1, B2, C1, and C2.

- Switches in fabrics B1 and C1 were uplinked to a pair of Dell Force10 S55 switches for dedicated access to the iSCSI SAN.
- Switches in fabrics B2 and C2 were connected to the VDI LAN and provide client connectivity.

The three physical networks that allow for segregation of different types of traffic are:

- **Management LAN:** This network provides a separate management network for all the physical ESXi hosts, network switches, and EqualLogic arrays. It also allows communication between various infrastructure components such as Microsoft Active Directory server, Microsoft SQL Server® and VMware vCenter server.
- **Dedicated iSCSI SAN:** Dedicated iSCSI storage area networking through which all the virtual desktops and other infrastructure components access the EqualLogic storage.
- **VDI LAN:** This is the network over which the clients access the virtual desktops in the View desktop pools.

For detailed information on how the Management LAN and VDI LAN were setup, please refer to Appendix B.

3.2 iSCSI SAN configuration

Figure 2 shows the network connectivity between a single M610 blade server and the storage array through the blade server chassis. The figure shows only one M610, though all 16 blades were used in the testing. The topology is identical for all the remaining blades in the chassis.

- Each PowerEdge M610 server included two Broadcom NetXtreme® II 5709s quad-port NIC mezzanine cards. One card was assigned to fabric B and the other to fabric C on the blade chassis.
- One NIC port from fabric B and one from fabric C are used to provide SAN connectivity. This also ensures connectivity in case of mezzanine card failure. Additionally, one NIC port from each of fabrics B and C are used to provide VDI LAN connectivity.
- Dual PowerConnect M6348 switches were installed in fabric B and fabric C on the blade server chassis. The NIC mezzanine cards were internally connected to each of these switches through the mid-plane on the blade server chassis.
- Two Force10 S55 switches were used for external SAN access. These switches were configured with stacking modules and 10GbE SFP+ uplink modules in the expansion bays.
- The uplink modules were used to create 2 x 10GbE LAG uplinks to the PowerConnect M6348 blade chassis switches.
- The stacking module was used to connect the two S55 switches together for failure resiliency and ease of management.
- Each EqualLogic PS6100XS array has two controllers with four iSCSI network ports and one Management port. Two ports on each storage controller were connected to one Force10 S55 switch while the other two ports were connected to the other Force10 S55 switch. The management port was connected to the PowerConnect 6248 switches on the Management LAN.

Note: For this white paper only two NIC ports per mezzanine card were used. Additional ports could be used to provide additional bandwidth.

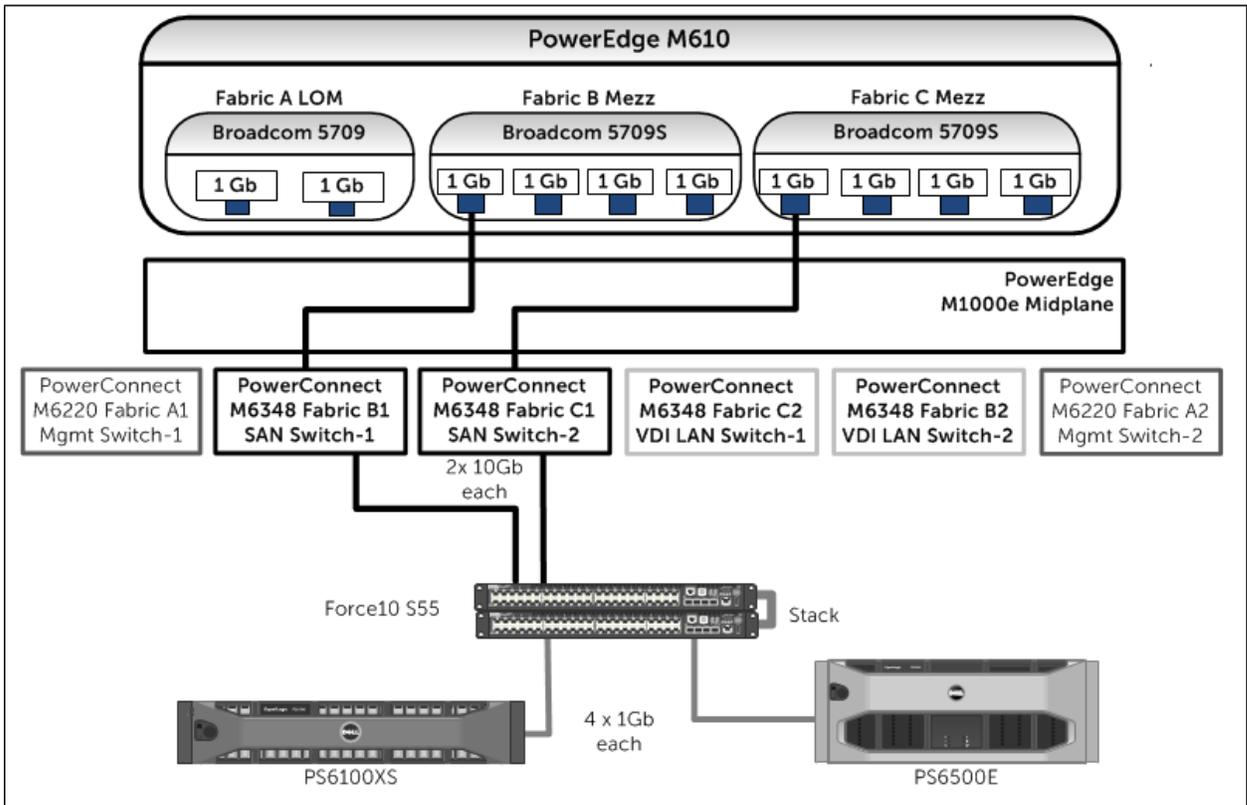


Figure 2 iSCSI SAN connectivity

3.3 EqualLogic storage array configuration

A single EqualLogic PS6100XS array hosted all the virtual desktops. The volume layout used for configuring the Replica volumes and the VDI volumes on a single PS6100XS array is shown in Table 1.

Table 1 EqualLogic volume layout on single PS6100XS array

Volume Name	Size	Purpose
VDI-BaselImages	150GB	Storage for Base image for VDI deployment
VDI-Replicas	150GB	Storage for Replica images created by VMware View
VDI-Images1	500GB	Storage for VDI Virtual Machines in View Cluster 1
VDI-Images2	500GB	Storage for VDI Virtual Machines in View Cluster 1
VDI-Images3	500GB	Storage for VDI Virtual Machines in View Cluster 1
VDI-Images4	500GB	Storage for VDI Virtual Machines in View Cluster 1
VDI-Images5	500GB	Storage for VDI Virtual Machines in View Cluster 2
VDI-Images6	500GB	Storage for VDI Virtual Machines in View Cluster 2
VDI-Images7	500GB	Storage for VDI Virtual Machines in View Cluster 2
VDI-Images8	500GB	Storage for VDI Virtual Machines in View Cluster 2

In addition to the PS6100XS array, a PS6500E array in a different pool provided a CIFS share for all the virtual desktops through a Microsoft Windows® OS-based file server. This CIFS share was used to redirect the user data and roaming profiles of the virtual desktops to a centralized location.



3.4 ESXi host network configuration

VMware ESXi 5.0 U1 hypervisor was installed on all 16 blades. The network configuration on each of those hosts is described below. Each ESXi host was configured with three virtual switches: vSwitch0, vSwitch1, and vSwitchiSCSI to separate the different types of traffic on the system.

Table 2 vSwitch configuration in ESXi

vSwitch	Description
vSwitch0	Management Network
vSwitch1	VDI LAN
vSwitchiSCSI	iSCSI SAN

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

3.5 VMware View configuration

View 5.1 was installed by following the documentation provided by VMware.

View 5.1 Documentation: <http://pubs.vmware.com/view-51/index.jsp>

Here are the specific configuration decisions used in the configuration:

- Two View servers were configured to provide load balancing and high availability.
- The View servers were installed as VMs on two separate hosts with two virtual CPUs, 10GB of RAM, and 40GB virtual hard drive.
- The first View Server was installed as a "View Standard Server" during the installation, while the second View Server was installed as a "View Replica Server".
- View Composer was installed on the same VM as the vCenter server, even though View 5.1 provides the option of installing it on a stand-alone VM. There was no difference found in performance with either option.
- Self-signed SSL certificates were applied to the VMware vCenter server VM and View 5.1 server VMs.

3.6 VMware View Pool configuration

The Add pool wizard in View was used to choose the following specific configuration options to create the virtual desktop pools.

- Virtual desktop pool type: Automated Pool
- User Assignment: Floating
- vCenter Server: View Composer linked clones
- View Composer Disks: Redirect disposable files to a non-persistent disk of size 4096MB
- Storage Optimization: Select separate datastores for replica and OS disk
- Advanced Storage Options: Use host caching
- Guest Customization: Sysprep

More information about View Pool configuration can be found in the VMware View 5.1 documentation: <http://pubs.vmware.com/view-51/topic/com.vmware.view.administration.doc/GUID-0A9CA985-3A78-428A-BCFB-B3E2DCCA90AD.html>



3.7 Windows 7 VM configuration

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

- VMware Virtual hardware v8
- One virtual CPU
- 1.5GB RAM
- 25GB virtual hard drive
- One virtual NIC connected to vSwitch3
- Windows 7 64bit OS

Additionally, the base image was customized using VMware's Optimization guide for Windows 7, available here: <http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf>



4 View test methodology

This section outlines the test objectives along with the test tools and criteria used to determine the sizing guidelines and best practices for deploying View on the EqualLogic storage platform.

4.1 Test objectives

- Develop best practices and sizing guidelines for a View 5.1 based VDI solution deployed on EqualLogic PS series storage, Dell Force10 and PowerConnect switches, and Dell PowerEdge blade servers with ESXi 5.0 as the server virtualization platform.
- Determine how many virtual desktops can be deployed in this environment using a single PS6100XS array with acceptable user experience indicators for a task worker workload profile.
- Determine the performance impact on the storage array during peak I/O activity such as boot and login storms.

4.2 Test approach

The key VDI use cases that were validated are:

- Task worker workload characterization
- Boot storms analysis
- Login storms and steady state analysis

All virtual desktops were restarted at the beginning of each login storm and steady state test to ensure they were in a consistent state. Only non-persistent desktops were used where any changes made to the desktop images are lost when the user logs off. However, changes to the user profiles were preserved by storing them on a CIFS share serviced by a Microsoft Windows Server 2008 R2 file share server. This was achieved through folder and profile redirection using Group Policy Objects (GPOs) on the Active Directory server.

4.3 Test tools

All tests were conducted using Login VSI 3.6 as the workload generator. Login VSI, developed by Login Consultants, is a widely adopted VDI benchmarking and load generation tool.

Login VSI is a unique benchmarking tool to measure the performance and scalability of centralized desktop environments such as Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI). More information can be found at: <http://www.loginvsi.com>

4.3.1 Load generation

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

- Up to five applications are open simultaneously.
- Applications include Office, IE, Word, PDF, Power Point, Excel, 7-Zip.
- Once a session is started, the medium workload repeats every 12 minutes.
- During a loop, the response time is measured every two minutes.



- Idle time is about two minutes between each 12 minute loop.
- Type rate is approximately 160ms per character.

Although Login VSI provides other workloads, the Medium workload was used in the tests because it closely resembles the workload of a task worker.

4.3.2 Monitoring tools

The following monitoring tools were used:

- Dell EqualLogic SAN Headquarters (SAN HQ) for monitoring storage array performance
- VMware vCenter Statistics for ESXi performance
- Login VSI Analyzer
- Liquidware Labs Stratusphere UX for user experience monitoring
- Custom script polling array counters for monitoring TCP retransmissions

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

4.4 Test criteria

The primary focus of the tests was to determine the maximum number of desktops which can be deployed using a single PS6100XS array in this environment while using VMware View Composer to provide Linked Clone virtual desktops in an automated pool.

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

The specific test criteria are described in the following sections.

4.4.1 Storage capacity and I/O latency

The typical industry standard latency limit for storage disk I/O is 20ms. 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 15% spare capacity on the storage array.

4.4.2 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 are:

- CPU utilization on any ESXi server should not exceed 90%
- Minimal or no memory ballooning on ESXi servers
- Total network bandwidth utilization should not exceed 90% on any one link
- TCP/IP storage network retransmissions should be less than 0.5%

4.4.3 Virtual desktop user experience

The Liquidware Labs Stratusphere UX user experience metric was used to ensure that all desktops had acceptable levels of application performance.

Liquidware Labs Stratusphere UX was used to gather data for user experience and desktop performance. Data gathered from the hosts (in VMware vCenter) and the virtual machines (software installed on the VM) is reported back to the Stratusphere Hub. The Stratusphere Hub was used to generate a comprehensive report on the Desktop Virtualization Performance Analysis. This report includes information about the host performance, virtual machine performance, and user performance. The report also provides a scatter plot that shows the performance of all the users in the system.

Liquidware Labs Stratusphere UX can generate a variety of reports that compare and validate user experience metrics. More information about these reports, including sample reports can be found here: <http://www.liquidwarelabs.com/products/stratusphere-ux-validation-reports.asp>

The criteria for the scatter plot is shown below.

VDI UX Definition			
Machine UX Indicators (Weight)	Good	Fair	Poor
Login Delay (20%)	< 15.000 s	15.000 - 60.000 s	> 60.000 s
Application Load Time (20%)	< 10.000 s	10.000 - 30.000 s	> 60.000 s
CPU Queue Length (20%)	< 3.000	3.000 - 6.000	> 6.000
Page Faults (0%)	< 2000.000	2000.000 - 10000.000	> 10000.000
Non-Responding Apps (20%)	< 2.000	2.000 - 3.000	> 3.000
I/O Experience Indicators (Weight)	Good	Fair	Poor
Disk IOPS (25%)	< 25.000	25.000 - 50.000	> 50.000
Disk Queue Length (25%)	< 1.000	1.000 - 3.000	> 3.000
Network Latency Roundtrip (25%)	< 150.000 ms	150.000 - 300.000 ms	> 300.000 ms
Failed Connections (25%)	< 5.000	5.000 - 15.000	> 15.000

Figure 3 Liquidware criteria for the scatter plot

4.5 Test setup

Two virtual desktops pools were configured using the VMware View 5.1 Administrator. Each pool was built from a Windows 7 base image. The Windows 7 configuration information is available in Section 3.7.

Desktop pool properties:

- Two View desktop pools
- Pool 1 with 512 virtual desktops (64 virtual desktops per host, eight hosts), spread across four 500 GB volumes (VDI-Images 1 – 4)
- Pool 2 with 320 virtual desktops (64 virtual desktops per host, five hosts), spread across four 500GB volumes (VDI-Images 5 – 8)
- Base images are stored on a separate 150GB volume (VDI-BaseImages)
- Replica images are stored on a separate 150GB volume (VDI-Replicas)
- Storage over-commit for all volumes was set to conservative
- View Composer disposable disk size was 4096 MB
- Disposable disk drive letter was set to auto

- Host caching was enabled for all hosts

A total of 832 virtual desktops were split into two pools of 512 and 320 desktops each. These numbers were chosen so that the number of virtual desktops hosted by each physical host was identical.

The test began by finding the maximum number of virtual desktops that can be hosted on a single ESXi host and then continued by adding hosts until the storage array limitations were met.

For example, a single ESXi host was able to handle up to 96 virtual desktops without showing any memory ballooning. However, as more hosts were added, the storage array limit had to be taken into consideration. The number of virtual desktops per host was reduced as the number of hosts was increased so that all the remaining 13 physical hosts in the blade chassis could be used with an even distribution of the workload.

5 Test results and analysis

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

5.1 Test scenarios

- **View Storage Accelerator/CBRC impact**

As noted in Section 2.5, View Storage Accelerator is a key technology introduced in VMware View 5.1. This feature uses CBRC at the ESXi host to reduce read IOPS on the shared SAN storage. Tests were conducted with an 800 desktop task worker VDI deployment with and without CBRC to determine CBRC impact in a VDI environment.

- **Boot storm**

Boot storm represents 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.

- **Task worker – Login storm and Steady state**

In this test, all 832 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 workload to simulate the Medium workload for a task worker profile.

5.2 View Storage Accelerator/CBRC benefits

In order to better understand the impact of CBRC, the array performance was tested on an 800 desktop VDI deployment. Using the same test methodology described earlier, the IOPS load was tested on the EqualLogic PS6100XS array with and without CBRC. When CBRC was disabled, the 800 task worker workload created 2565 read IOPS and 4250 write IOPS for a total of 6815 IOPS at the steady state. With CBRC enabled, the same workload created 811 read IOPS and 4674 write IOPS for a total of 5485 IOPS at steady state (see Figure 4). This 68% reduction in read IOPS with CBRC enabled is a significant improvement. When looking at the IOPS load at volume level, the replica volume (base image) read IOPS decreased from 1763 with CBRC disabled to 18 with CBRC enabled. Thus, overall read IOPS reduction was almost entirely because all the reads to the base image were served out of the host cache when CBRC was enabled.

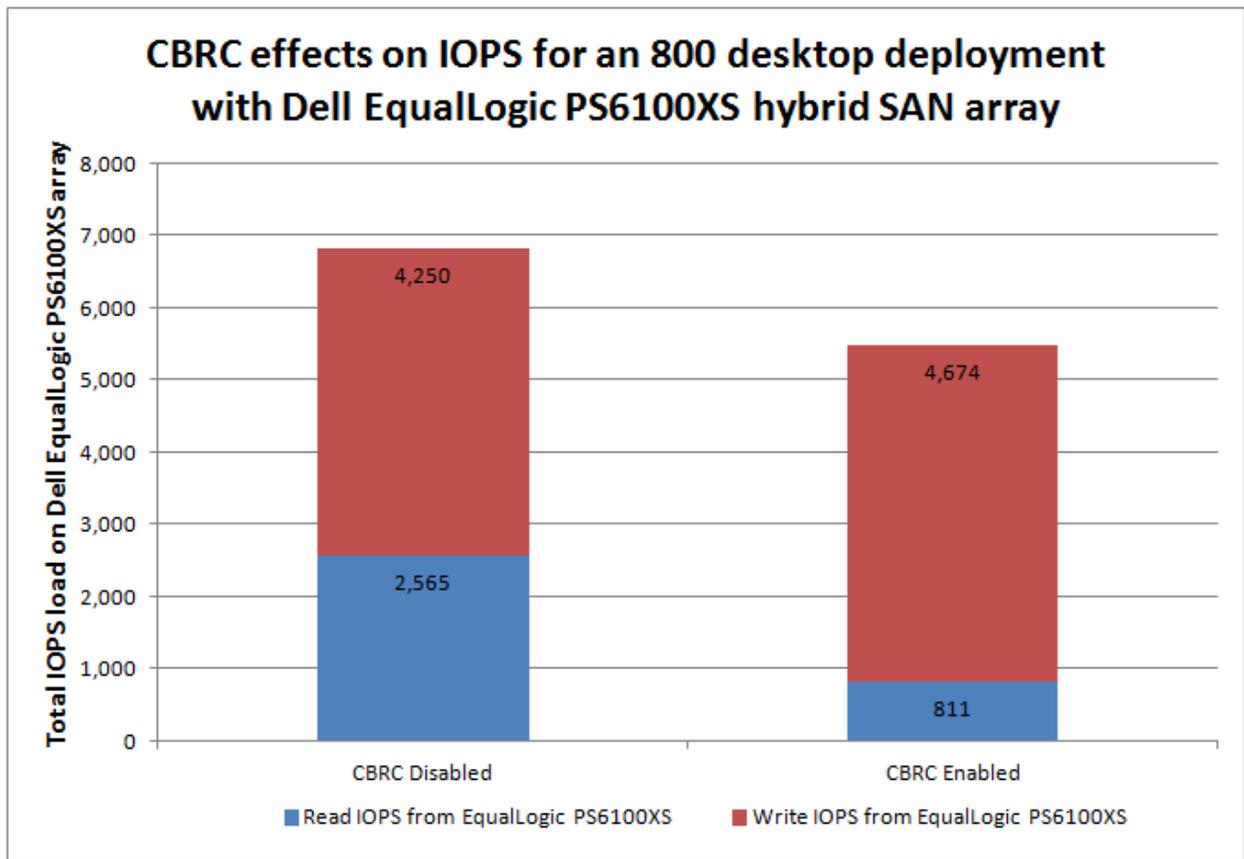


Figure 4 CBRC effects on VDI workload

The following analysis shows the CBRC benefits for the enterprise in terms of desktop density and cost for each desktop. It is a sample analysis, and exact benefits will vary based on the actual VDI workload, array performance characteristics, and price.

Table 3 below shows the IOPS impact on a single desktop that creates 10 IOPS in a steady state with a 30%:70% read:write mix without CBRC. The test shows a 68% read IOPS reduction rate when CBRC is enabled, so the same desktop creates eight IOPS in a steady state with a 12%:88% read:write mix with CBRC.

Table 3 CBRC effect on a single desktop IOPS

	Read IOPS	Write IOPS	Total IOPS	Read %	Write %
CBRC Disabled	3	7	10	30%	70%
CBRC Enabled	1	7	8	12%	88%

Using an example storage array that delivers a maximum of 7300 IOPS under 20ms latency with a 30%:70% read:write mix; the same array delivers 6700 IOPS when the corresponding read:write mix is 12%:88%. In the example, also assume that the array costs \$50,000. Table 4 shows the desktop density for the array with and without CBRC.



Table 4 CBRC effect on desktop density on a storage array

	Read percentage	Write percentage	Total array IOPS (Sample)	IOPS/desktop	Total desktops
CBRC Disabled	30%	70%	7300	10	730
CBRC Enabled	12%	88%	6700	8	838

As Table 5 shows, in this example, CBRC improves desktop density by 15% and cost for each desktop by 13%.

Table 5 CBRC improvements

	Desktops per array	Cost for each desktop
CBRC Disabled	730	\$68
CBRC Enabled	838	\$60
CBRC Improvement	15%	13%

5.3 Boot storm I/O

In the boot storm test, the virtual desktops were reset simultaneously from the VMware vSphere client to simulate a boot storm. Figure 5 shows the storage characteristics during the boot storm – the PS6100XS array delivered nearly 14,500 IOPS (14-18 IOPS per VM) under the peak load during this test.

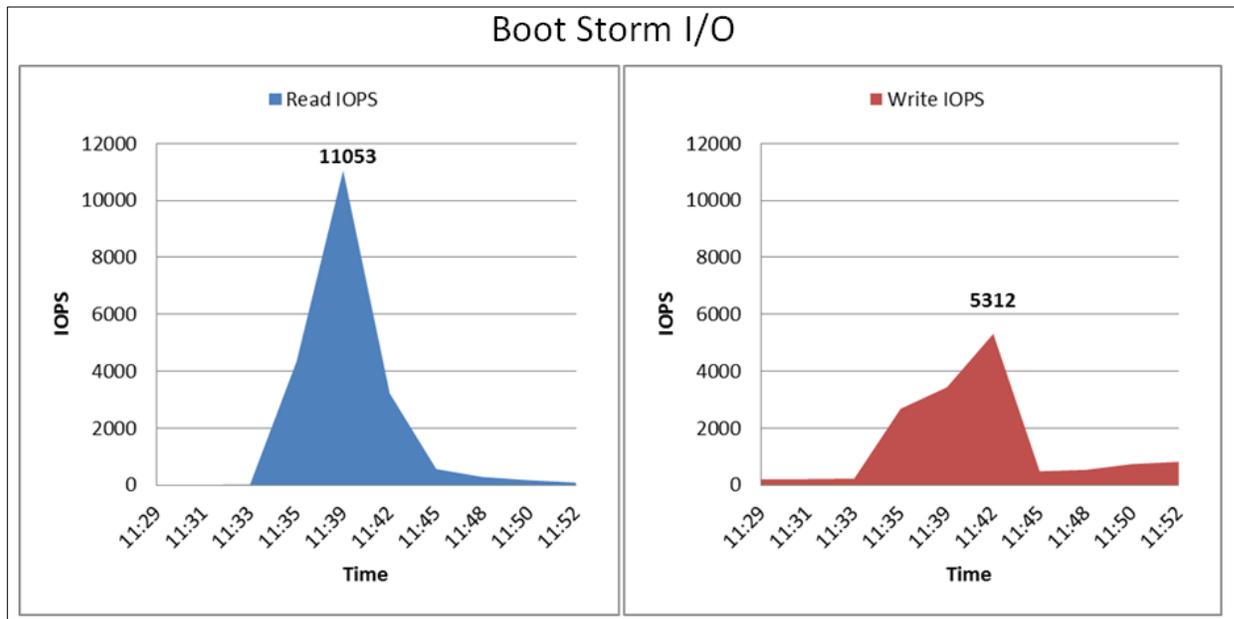


Figure 5 SAN HQ data showing PS6100XS array IOPS during boot storm

The spike seen in Figure 5 was primarily due to read operations, because the boot process of the virtual desktops creates many simultaneous reads to the replica image. The read and write ratios during the boot storm were about 76% reads to 24% writes.



On the replica image volume nearly 100% reads and almost negligible writes were observed.

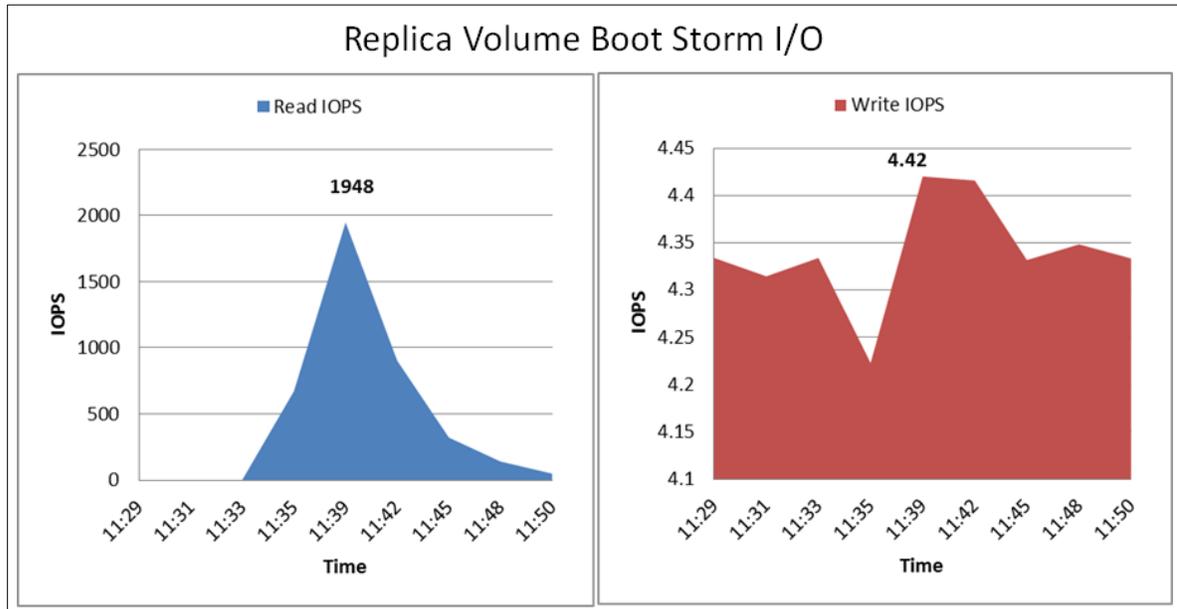


Figure 6 Read and write IOPS on the replica volume

The read and write ratios on the remaining volumes that hosted the virtual machines were nearly 53% reads and 47% writes. The IOPS on one such volume are shown in Figure 7 below. The remaining volumes that hosted the VMs had similar IOPS as shown in the figure.

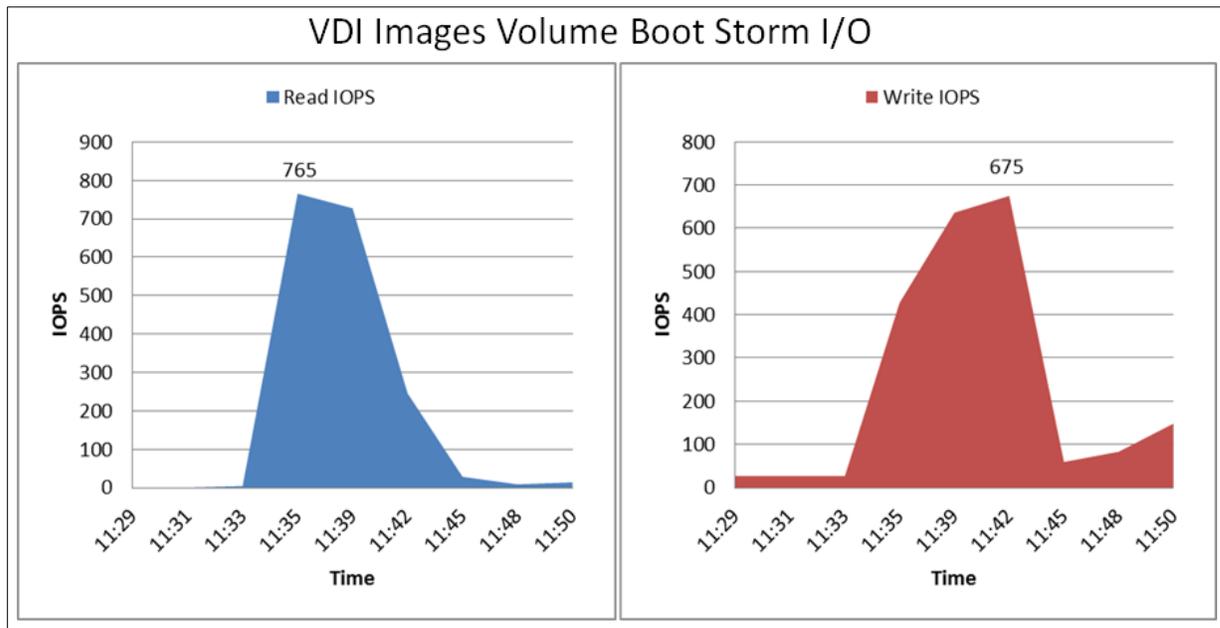


Figure 7 Read and write IOPS on one of the VDI-Images volumes

The replica volume only hosts a copy of the base image that all the virtual machines are linked to and due to this almost all the IO operations on this volume are reads. The volumes on which the VDI VMs are hosted have a mix of reads and writes. This is because all the changes made by the OS are written to these volumes.



The ESXi hosts did not show any bottle necks with respect to meeting the CPU and memory resource demands during the boot operation. There were no CPU utilization alerts, and no memory ballooning was observed. Other infrastructure services like the View Servers and Active Directory servers did not show any exceptional load during the boot storm.

The table below shows that most of the operations during the boot storm were handled by the SSDs in the hybrid array. This ensured that the boot storm lasted for a short time and all the desktops were available quickly. The SSD drives also acted as the accelerated write cache on the storage array and were able to handle the write I/O load.

Table 6 Disk usage on PS6100XS during boot storm

Member	Disk	Description	Average IOPS	Read I/O Rate	Write I/O Rate	Status
View51-6100XS-2	0	SSD 400GB SAS	2,423.90	25.2 MB/sec	12.6 MB/sec	online
View51-6100XS-2	1	SSD 400GB SAS	2,437.60	25.2 MB/sec	12.6 MB/sec	online
View51-6100XS-2	2	SSD 400GB SAS	2,406.80	25.2 MB/sec	12.6 MB/sec	online
View51-6100XS-2	3	SSD 400GB SAS	2,397.40	25.3 MB/sec	12.6 MB/sec	online
View51-6100XS-2	4	SSD 400GB SAS	2,424.50	25.3 MB/sec	12.6 MB/sec	online
View51-6100XS-2	5	SSD 400GB SAS	2,432.30	25.3 MB/sec	12.5 MB/sec	online
View51-6100XS-2	6	SSD 400GB SAS	2,417.30	25.3 MB/sec	12.5 MB/sec	online
View51-6100XS-2	7	10K 600GB SAS	274.80	2.0 MB/sec	678.1 KB/sec	online
View51-6100XS-2	8	10K 600GB SAS	200.90	2.0 MB/sec	864.6 KB/sec	online
View51-6100XS-2	9	10K 600GB SAS	203.40	2.0 MB/sec	791.1 KB/sec	online
View51-6100XS-2	10	10K 600GB SAS	197.40	1.8 MB/sec	683.8 KB/sec	online
View51-6100XS-2	11	10K 600GB SAS	227.90	1.9 MB/sec	712.0 KB/sec	online
View51-6100XS-2	12	10K 600GB SAS	285.20	2.1 MB/sec	915.4 KB/sec	online
View51-6100XS-2	13	10K 600GB SAS	0.00	0 KB/sec	0 KB/sec	spare
View51-6100XS-2	14	10K 600GB SAS	204.40	1.7 MB/sec	683.8 KB/sec	online
View51-6100XS-2	15	10K 600GB SAS	192.00	1.7 MB/sec	802.4 KB/sec	online
View51-6100XS-2	16	10K 600GB SAS	197.70	1.7 MB/sec	695.1 KB/sec	online
View51-6100XS-2	17	10K 600GB SAS	208.20	1.9 MB/sec	858.9 KB/sec	online
View51-6100XS-2	18	10K 600GB SAS	205.30	1.7 MB/sec	712.0 KB/sec	online
View51-6100XS-2	19	10K 600GB SAS	210.80	1.9 MB/sec	864.6 KB/sec	online
View51-6100XS-2	20	10K 600GB SAS	206.30	1.7 MB/sec	678.1 KB/sec	online
View51-6100XS-2	21	10K 600GB SAS	204.40	1.8 MB/sec	808.1 KB/sec	online
View51-6100XS-2	22	10K 600GB SAS	225.40	1.8 MB/sec	695.1 KB/sec	online
View51-6100XS-2	23	10K 600GB SAS	278.20	2.0 MB/sec	915.4 KB/sec	online

As Figure 8 shows, storage network utilization was well within the available bandwidth. The peak network utilization during boot storm reached approximately 30% of the network bandwidth, and then it gradually declined once all the machines were booted up. There were also no retransmissions on the iSCSI SAN.

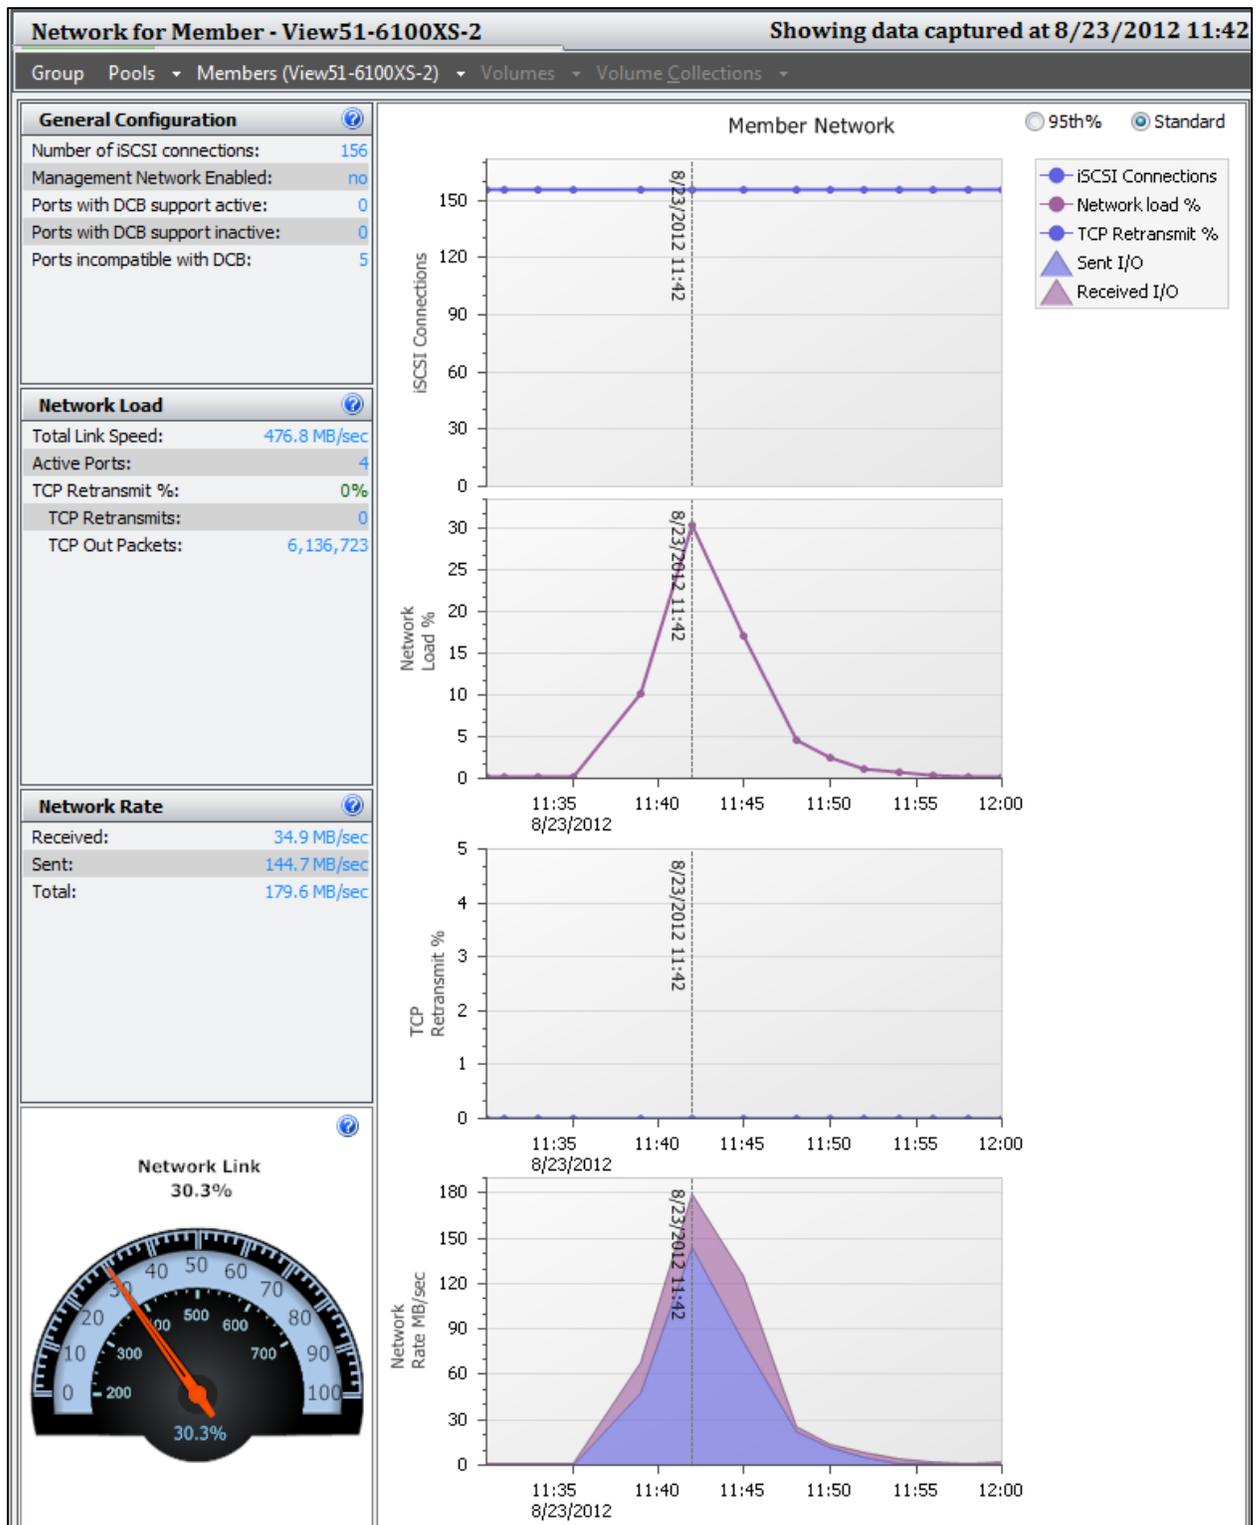


Figure 8 SAN HQ data showing network performance during boot storm

These results show that the EqualLogic PS6100XS can handle a heavy random I/O load like a boot storm with no issues. Typically boot storms rare, but they can occur after a power failure or a major system outage.



5.4 Task worker – Login storm and steady state I/O

The following sections provide results from the login storm and steady state testing.

5.4.1 Login storm I/O

Login VSI was programmed to launch 832 virtual desktops over a period of 20 minutes after pre-booting the virtual desktops. The peak IOPS observed during the login storm was about 6200 IOPS (6-8 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
- Operating system services on the virtual desktop
- First launch of applications

Once a virtual desktop has achieved a steady state after the user login, the Windows OS has cached applications in memory and does not need to access storage each time the application is launched. This is why IOPS tend to be lower during the steady state.

Figure 9 below shows the IOPS and latency observed during the login storm.

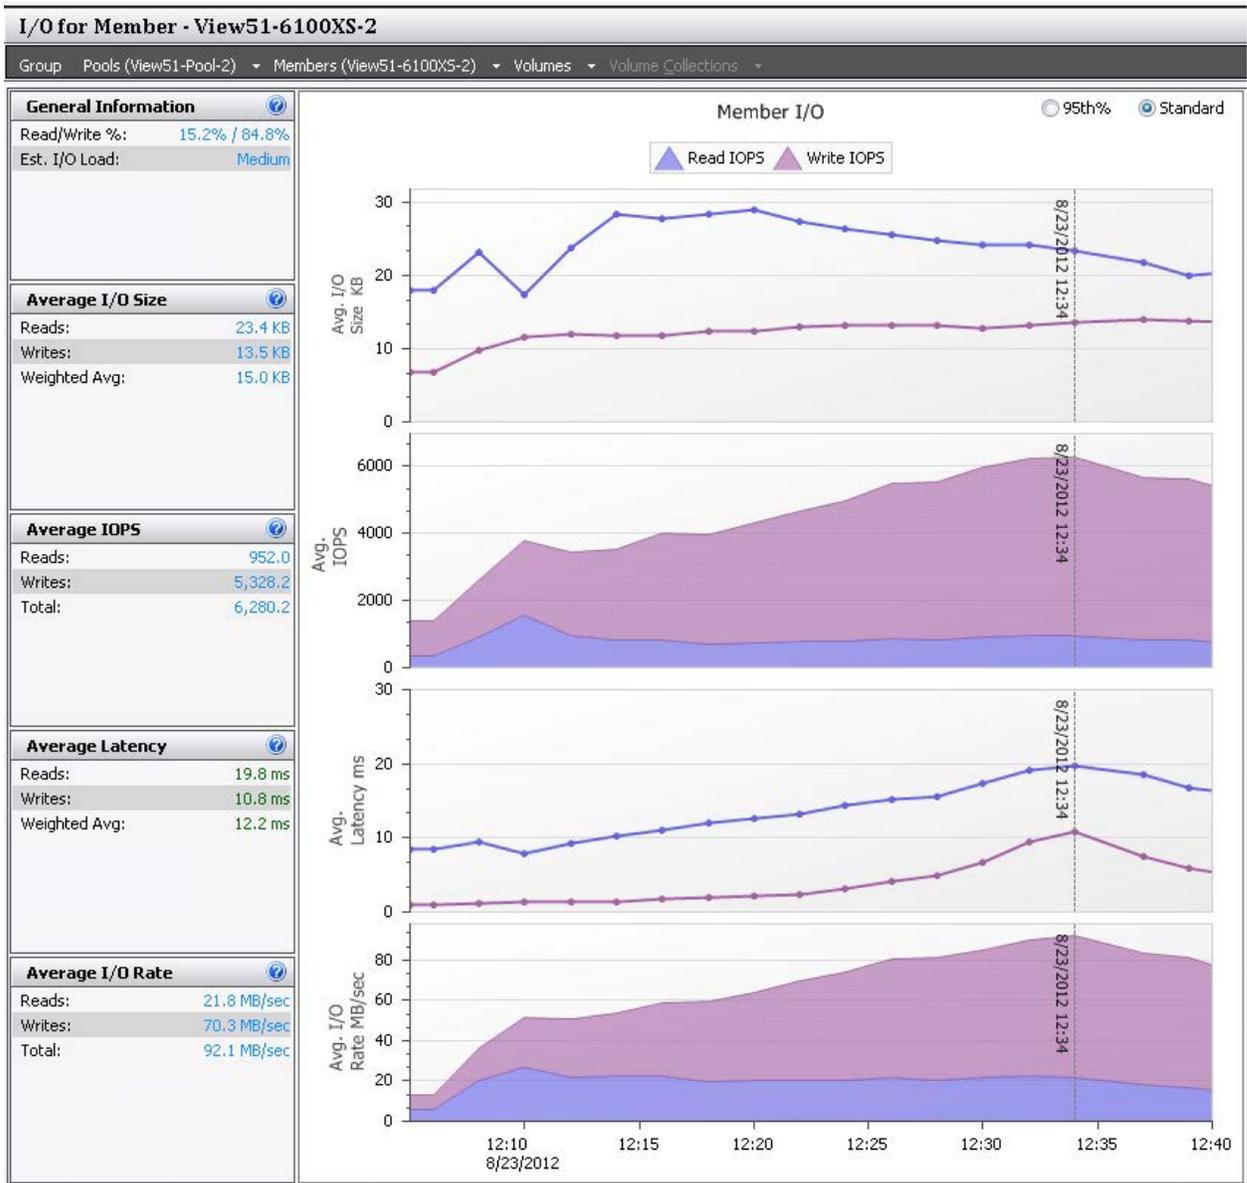


Figure 9 SAN HQ data showing login storm I/O

The peak latency seen on the storage array is less than 20ms, and the storage array is able to handle the 832 users logging in over such a short duration with no performance issues. The following table shows the overall usage of the disks in the array during a login storm as collected by SAN HQ.



Table 7 Disk usage on EqualLogic PS6100XS array during login storm

Member	Disk	Description	Average IOPS	Read I/O Rate	Write I/O Rate	Status
View51-6100XS-2	0	SSD 400GB SAS	1,952.80	19.4 MB/sec	21.5 MB/sec	online
View51-6100XS-2	1	SSD 400GB SAS	1,939.90	19.4 MB/sec	21.7 MB/sec	online
View51-6100XS-2	2	SSD 400GB SAS	1,934.60	19.5 MB/sec	21.6 MB/sec	online
View51-6100XS-2	3	SSD 400GB SAS	1,947.00	19.5 MB/sec	21.5 MB/sec	online
View51-6100XS-2	4	SSD 400GB SAS	1,950.20	19.4 MB/sec	21.5 MB/sec	online
View51-6100XS-2	5	SSD 400GB SAS	1,953.10	19.5 MB/sec	21.5 MB/sec	online
View51-6100XS-2	6	SSD 400GB SAS	1,957.50	19.6 MB/sec	21.5 MB/sec	online
View51-6100XS-2	7	10K 600GB SAS	147.30	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	8	10K 600GB SAS	147.40	2.1 MB/sec	2.2 MB/sec	online
View51-6100XS-2	9	10K 600GB SAS	145.00	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	10	10K 600GB SAS	145.70	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	11	10K 600GB SAS	147.90	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	12	10K 600GB SAS	148.30	2.1 MB/sec	2.2 MB/sec	online
View51-6100XS-2	13	10K 600GB SAS	0.00	0 KB/sec	0 KB/sec	spare
View51-6100XS-2	14	10K 600GB SAS	141.90	1.9 MB/sec	2.1 MB/sec	online
View51-6100XS-2	15	10K 600GB SAS	144.50	2.0 MB/sec	2.2 MB/sec	online
View51-6100XS-2	16	10K 600GB SAS	139.20	1.9 MB/sec	2.1 MB/sec	online
View51-6100XS-2	17	10K 600GB SAS	140.10	2.0 MB/sec	2.2 MB/sec	online
View51-6100XS-2	18	10K 600GB SAS	141.20	1.9 MB/sec	2.1 MB/sec	online
View51-6100XS-2	19	10K 600GB SAS	137.70	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	20	10K 600GB SAS	144.90	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	21	10K 600GB SAS	142.40	2.1 MB/sec	2.2 MB/sec	online
View51-6100XS-2	22	10K 600GB SAS	142.20	2.0 MB/sec	2.1 MB/sec	online
View51-6100XS-2	23	10K 600GB SAS	144.30	2.1 MB/sec	2.2 MB/sec	online

Table 7 clearly shows that most of the login storm I/O is handled by the SSD drives and therefore the array is able to provide the best possible performance.

5.4.2 Steady state I/O

The total IOPS on the PS6100XS array averaged during the steady state with all the users logged in was around 5400 (6.4 IOPS per VM). Of this, the read IOPS accounted for about 719 (approximately 14% of the total I/O load) and the remaining 4700 IOPS or 86% were write IOPS. Read and write latencies were also well below 20ms throughout the test.

All changes that occur on the desktop (including temporary OS writes such as memory paging) are being written to disk. The I/O pattern is primarily writes due to this activity. Once the desktops are booted and in a steady state, the read I/O becomes minimal due to CBRC caching on the ESXi hosts.

During the steady state, there is minimal activity on the replica volume and most of the activity is seen on the VDI-Images volumes that host the virtual machines.

Figure 10 shows the performance of the array during the steady state test.



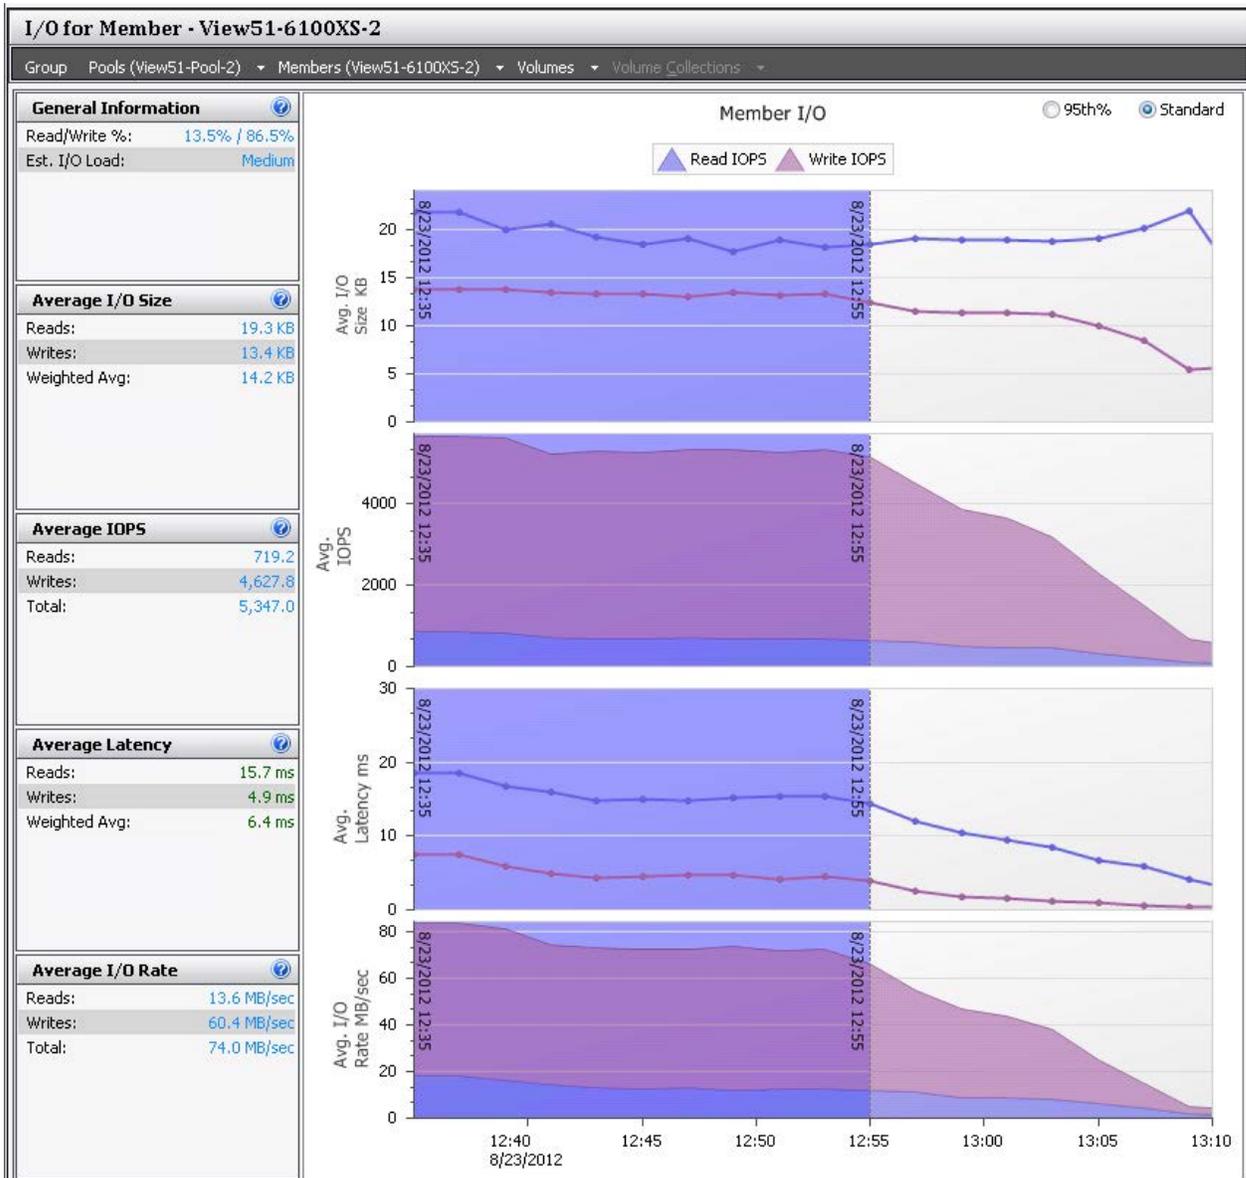


Figure 10 SAN HQ data showing steady state I/O

5.4.3 Server host performance

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

Statistics for the ESXi hosts were captured using VMware vCenter server. The figures below show the CPU, memory, network, and storage adapter utilization of one ESXi server hosting the virtual desktops.

The key observations from the statistics were:

- CPU utilization was well below 80% throughout the test.
- Active memory usage was about 35% and there was no memory ballooning observed.



- Network utilization was about 10MBps which included an average of 5MBps usage on the iSCSI SAN and very minimal traffic on management and VDI LANs.
- Average Read and Write latencies at the storage adapter level matched the observed latencies in SAN HQ.

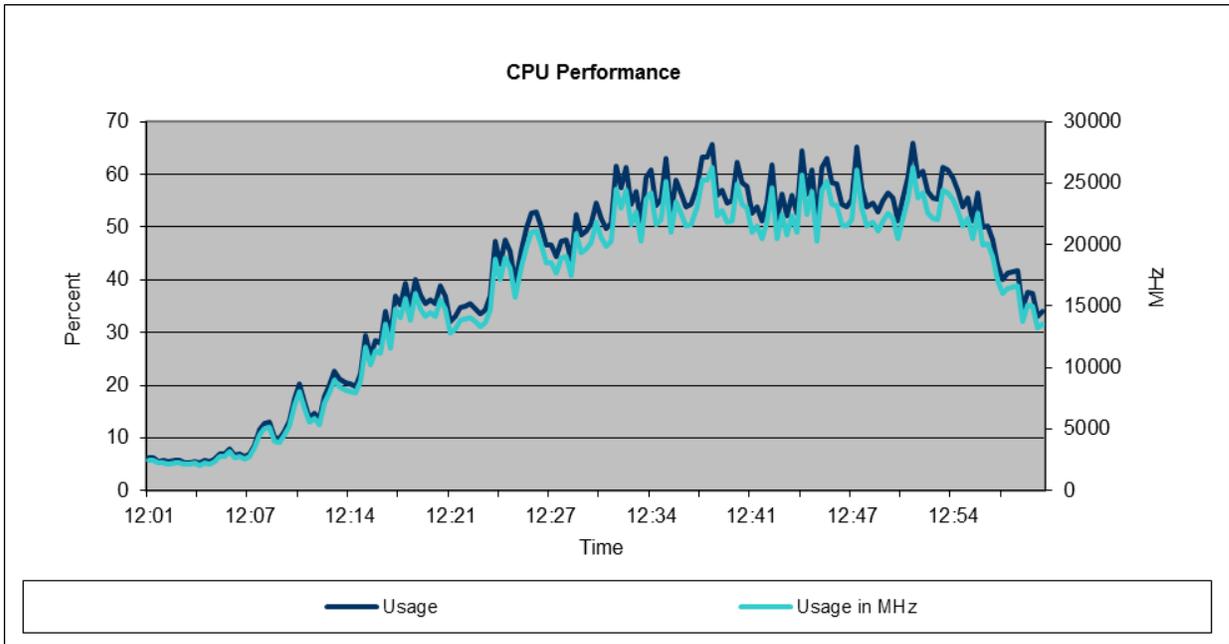


Figure 11 ESXi host CPU utilization

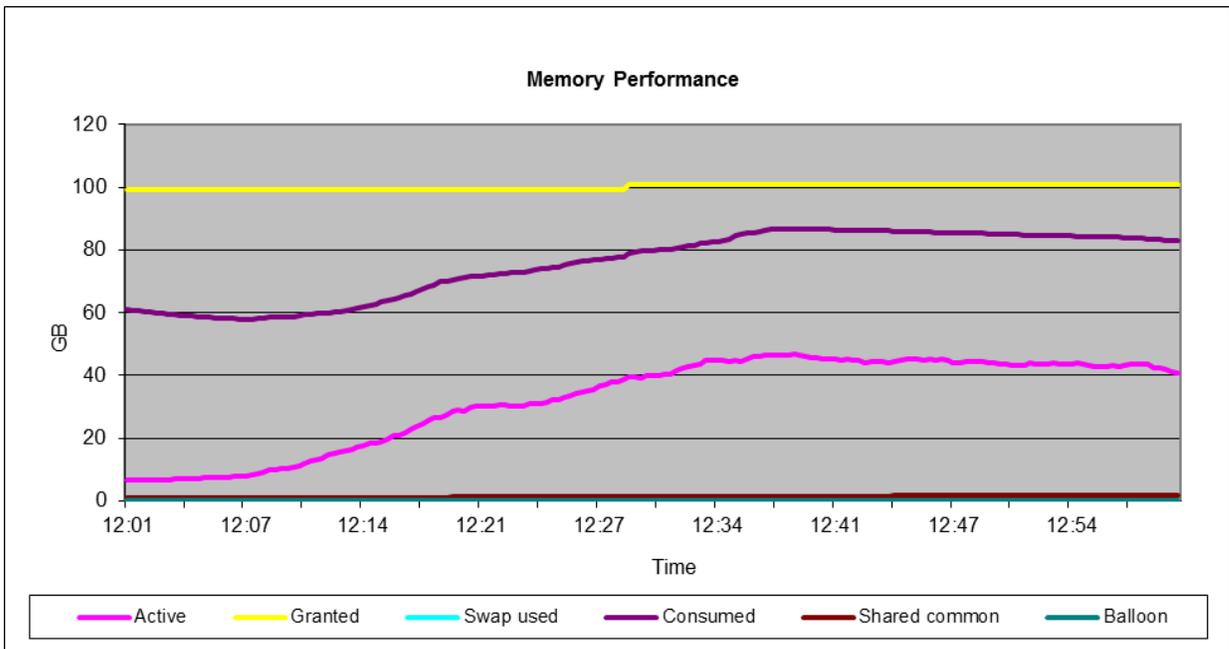


Figure 12 ESXi host memory utilization



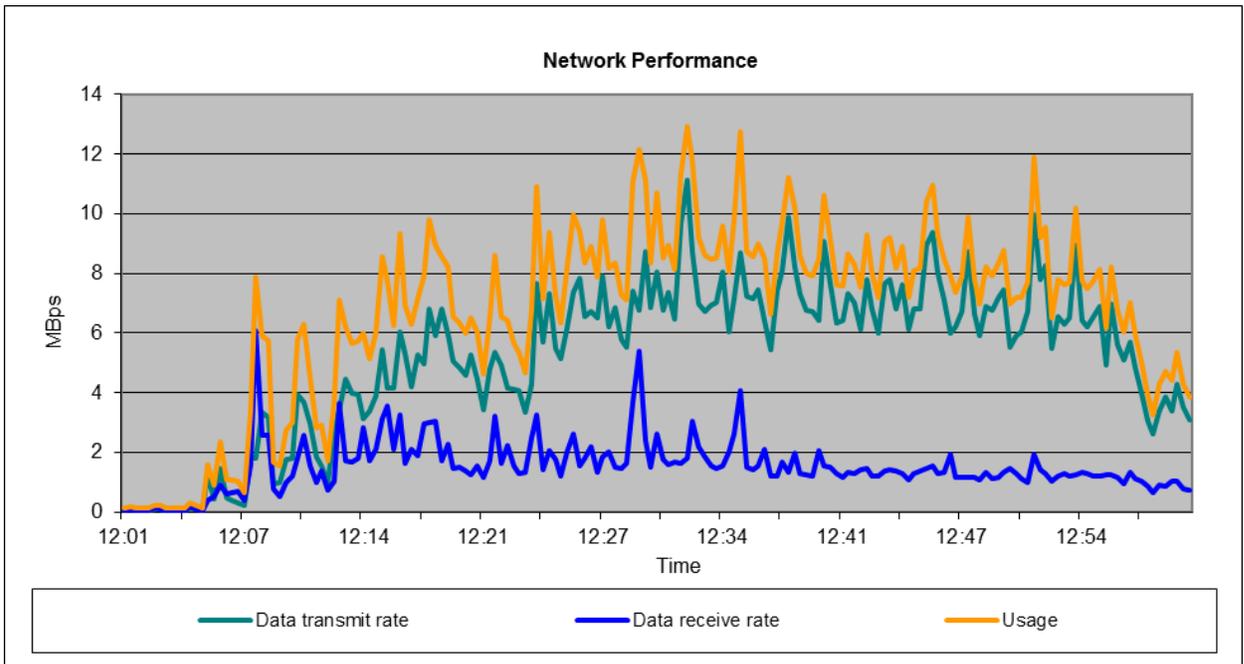


Figure 13 Combined ESXi host network utilization

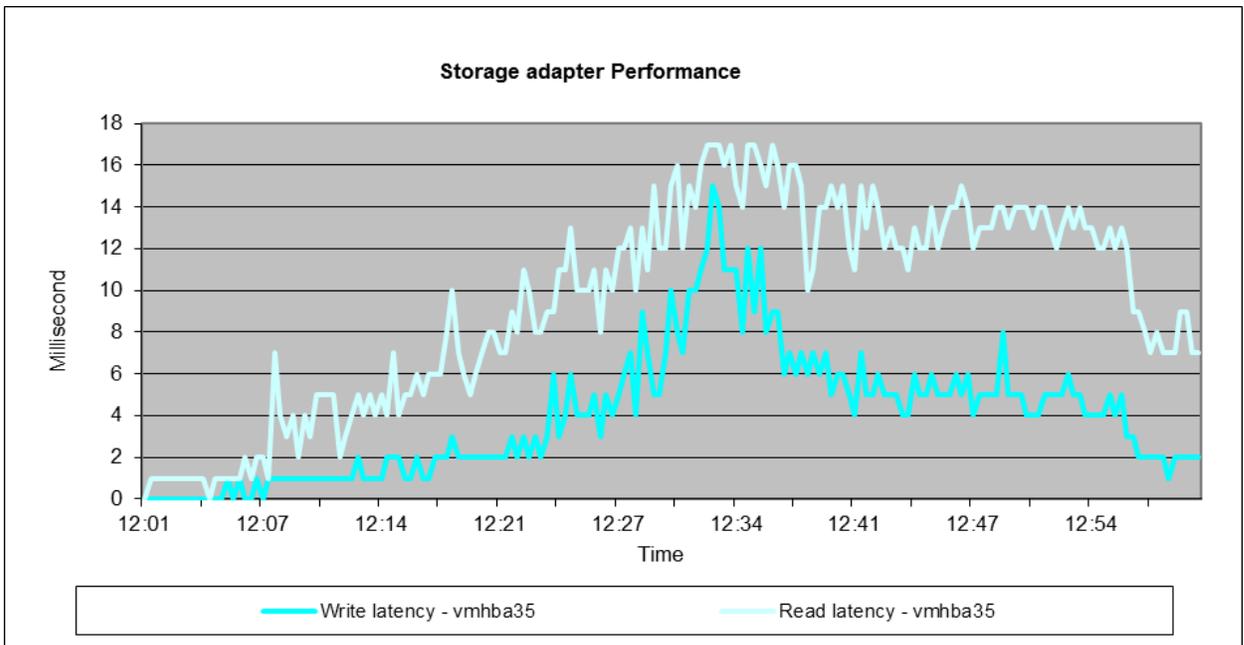


Figure 14 ESXi storage adapter read and write latencies

5.4.4 User experience monitoring

Based on the criteria in Section 4.4.3, the scatter plot for the user experience is shown below. As Figure 15 shows, virtually all the users are in the Good category.

This shows that the EqualLogic PS6100XS solution is capable of easily providing adequate user experience for all the users in the VDI environment.



Note that Liquidware Labs Stratusphere UX registered 788 of the total 832 desktops for the scatter plot in Figure 15, but the Good performance of the remaining ones were verified in other reports.

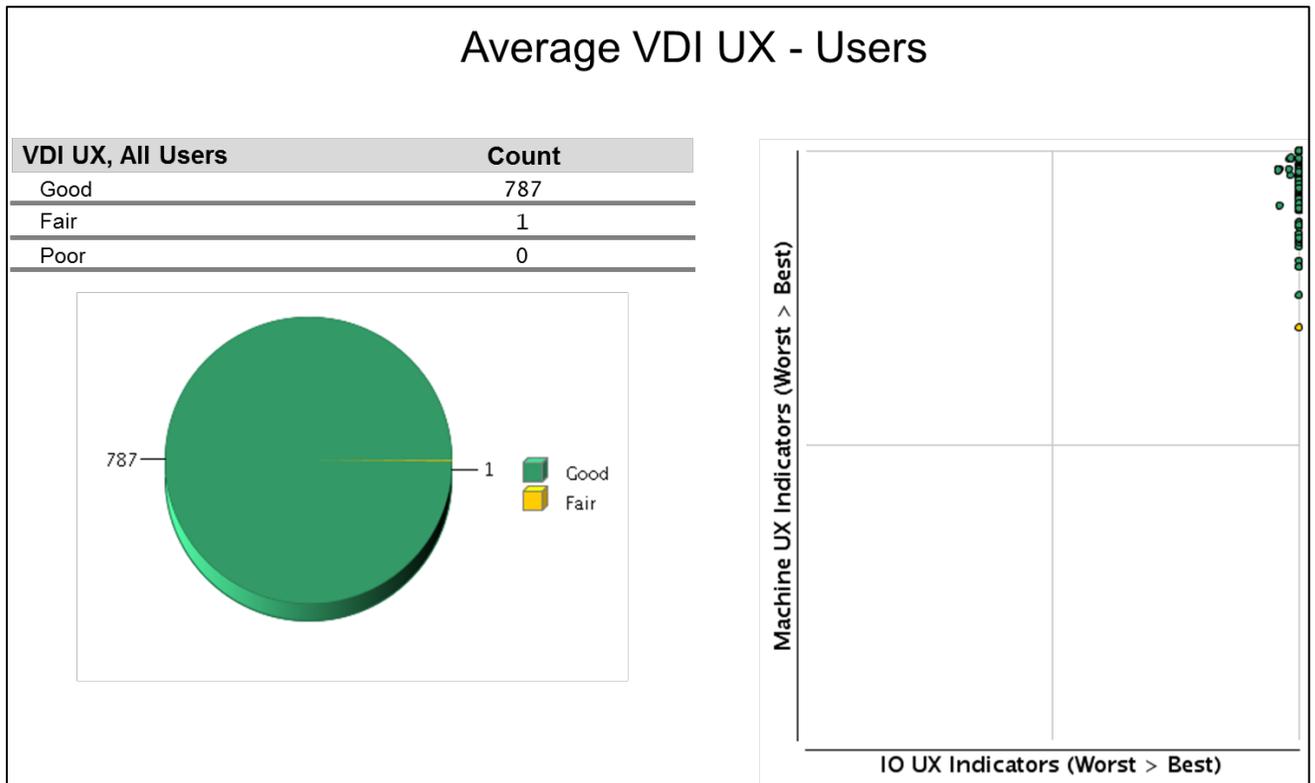


Figure 15 Average VDI user experience for all users

5.5 Results summary

The key observations from the test results are listed below.

- A single PS6100XS array was able to host 832 virtual desktops and support the task worker type of I/O activity.
- The VDI I/O was mostly write-intensive I/O with more than 85% writes and less than 15% reads.
- None of the system resources on the ESXi servers hosting the virtual desktops reached maximum utilization levels at any time.
- During a boot storm simulation, nearly 14,500 IOPs were observed, and all desktops were available in View within 20 minutes of the storm.
- The login storm for 832 desktops lasted around 30 minutes and had most of the I/O served by the SSD drives on the array.
- The Stratusphere UX tool showed that all 832 desktops were in the Good category in the scatter plots.



6 Sizing guidelines for EqualLogic SANs

The storage array selected for the solution should be able to handle various I/O patterns that occur throughout the day for a VDI solution. These include the login storm at the beginning of a shift or a work day when employees login to their virtual desktops in a relatively short period of time. Once they are logged in, the virtual desktops reach a steady state where they generate predictable IOPS as the employees go about their work day. The same storage array needs to handle recovery situations due to unexpected events such as power outages which might cause boot and login storms.

A good way to deploy the storage infrastructure for VDI is to understand the VDI workload characteristics including the performance characteristics of the applications being deployed. These include:

- Capacity requirements
- Performance requirements
 - IOPS
 - Average Disk latency
 - Read and Write ratios
 - Type of I/O (Sequential or Random)

The following sections summarize the key capacity and performance guidelines for a VMware View based VDI solution while using PS6100XS. These guidelines are based on the test results and key findings from the previous sections.

6.1 Capacity sizing considerations

There are a number of variables that need to be considered for calculating the storage requirements for linked clone desktop pools in VMware View environments. In this section you will find an overview of considerations and inputs necessary for the calculations. There is also an example calculation for 832 linked clone desktop environment. You should consult the VMware View documentation before calculating the requirements for your environment.

The storage capacity requirements for linked clone desktops using View Composer primarily depend on the following factors:

- Number of base image copies that need to be created and maintained
- Number of replica images that will be stored
- OS installed
- Additional applications that need to be installed
- The type of I/O activity for each virtual desktop
- Additional space required for linked clone delta files. The delta file size depends on the amount of change that occurs between the desktop VM and the base image.

In the configuration separate datastores were used for base images (OS disks), replica disks, and persistent disks. You can use the following formulas to calculate the sizing requirements for the respective datastores. All calculations below assume a size of 25GB for the base OS image.

6.1.1 VMware View Accelerator sizing

Enabling VMware View Accelerator turns on CBRC on the selected ESXi hosts.

CBRC works by creating a digest file for each VMDK on the VM and stores the hash information about VMDK blocks with the VM itself. The size of this digest file is between 5 to 10MB for each GB of the VMDK size. This means that for a 25GB Windows VM replica used in the testing, there was about 125MB of storage space used for the digest file. This digest file is loaded into memory when it is accessed for the first time.

When memory overcommit is used to assign more RAM to VMs than there is available memory in the host, it is important to note that enabling CBRC can create a digest of significant size. When CBRC is enabled, the digest file increases the memory utilized in a host and could cause increased memory ballooning and impact the overall performance of the host server.

In the test setup, the base image was 25GB and the replica image had a digest of 125MB. Each VM had a non-persistent disk size of 4GB which created a digest file of 32MB. If a server hosted 64VMs and a replica disk, the total memory required for CBRC (assuming the maximum 2048MB is used for CBRC) would be:

$$2048MB + 125MB + (64 \times 32MB) = 4221MB$$

It is important the server has enough memory capacity to cache this data in RAM for optimal performance.

6.1.2 Base image datastore

This datastore stores the base images copies that need to be created and maintained for the virtual desktops.

$$\text{Capacity needed for OS datastore (GB)} = \text{Base Image Size} \times (2 \times \text{VM Memory}) \times \text{Number Of Base Images}$$

For example:

Two base images were stored in the OS disks, so the capacity required for the OS disks datastore is:

$$\text{Capacity (GB)} = 25 \times (2 \times 1.5) \times 2 = 150GB$$

6.1.3 Replica datastore

This datastore is used to host the replica disk images that are created from the base images during the deployment of the linked clone virtual desktops.

The space required for the replica images is identical to the space required for the OS images and the same formula from above can be used.

6.1.4 Linked clone datastores

These datastores are used to store the VDI VM images and the disposable disks for all the virtual desktops created using linked clones. The capacity required to store these virtual desktops depends on the amount of space reserved for the linked clone delta files and the aggressiveness of the storage overcommit used while creating the desktop pool.

More information on storage overcommit for linked clone desktops can be found here: <http://pubs.vmware.com/view-51/topic/com.vmware.view.administration.doc/GUID-5CE57B8D-3127-44DB-BA51-18C96CE29096.html>

Capacity needed for Linked clone datastores (GB)

$$= \text{Number of VMs} \times (2 \times \text{VMMemory}) \times \text{Number of Datastores} \times \text{Overcommit factor}$$

For example: To host 832 desktops with a conservative storage overcommit, the capacity required is

$$\text{Capacity (GB)} = 832 \times (2 \times 1.5\text{GB}) \times 8 \times 0.25 = 4992\text{GB}$$

Table 8 shows the storage overcommit level for hosting the same desktops with different options chosen.

Table 8 Capacity requirements for different storage overcommit factors.

Option	Storage overcommit level
None	No overcommit
Conservative	4 times the size of datastore. Default level.
Moderate	7 times the size of the datastore.
Aggressive	15 times the size of the datastore

For more information on the sizing formulas used in this section, see: <http://pubs.vmware.com/view-51/topic/com.vmware.view.administration.doc/GUID-9704B4D2-D592-4EF6-946E-61C4FF3D0BB1.html>

6.2 Performance sizing considerations

The storage platform used in VDI deployment should be able to sustain the maximum IOPS that boot and login storms create within acceptable latency limits.

6.2.1 Boot storm sizing considerations

A boot storm is the worst case storage I/O workload pattern for virtual desktops and is rarely encountered. In the testing, it was confirmed that boot storms generated much higher storage system performance demands than a login storm or steady state application usage.

In the test environment, a single PS6100XS array delivered nearly 14,500 IOPS during a boot storm with a read write ratio of about 76% / 24%. This is because during a boot storm, most of the operations are reads from the VDI volumes and the replica disk. As noted in Section 5.3, this boot storm was created by resetting all the virtual desktops simultaneously from VMware vCenter. It took approximately eight minutes to boot the 832 desktops and the boot storm I/O lasted for about 20 minutes. In a production environment, virtual desktop VMs should be booted before the start of a work day or a work shift, when possible, thereby reducing the boot storm impact.

It is recommended to stagger the boot storms by restarting only a few VMs at a time, for example, 200. This greatly reduces the load on the storage and allows you to size your solution to handle the IOPS generated by those desktops without causing performance issues on the already running VMs.

6.2.2 Login storm and steady state sizing considerations

In the testing, logging in all 832 users over a period of 30 minutes generated about 6200 IOPS which is about 7-10 IOPS per desktop. The maximum read/write ratio seen during the login storm was about 15% / 85%. During the login storm most of the operations are writes because of the changes to the OS and page files being written to the disk.

During the steady state, after all the users logged in, the combined IOPS generated was about 5400 IOPS which translates to about 6-8 IOPS per desktop. The read/write ratio seen here is still write heavy and nearly identical to the login storm – at about 14% / 86%.

The IOPS generated during the login storm were higher than those seen during the steady state. Quantifying the average I/O workload generated is important for performance and sizing considerations. The storage platform should be able to sustain the maximum IOPS these workloads create within acceptable latency limits.

Based on the test results, I/O characteristics observed for the task worker type of workload profile during the lifecycle of a virtual desktop are summarized as shown in Table 9.

Table 9 I/O characteristics during the login and steady state of a VM

I/O type	IOPS per desktop	Member read / write ratio	I/O block size
Boot Storm	14-18	76% / 24%	Reads 10KB /Writes 10KB
Steady State	6-8	13% / 87%	Reads 19KB /Writes 13KB
Login Storm	7-10	15% / 85%	Reads 23KB /Writes 13KB

7 Best Practices

7.1 Application layer

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 still using non-persistent virtual desktops. .

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 beginning of a shift or 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 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 which can improve performance. 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 VMware recommendations for Windows 7 image customization can be found here:
<http://www.vmware.com/resources/techresources/10157>

7.1.4 SSL certificate requirements

It is recommended that all the servers use either signed SSL certificates or use self-signed SSL certificates to provide the best security for your infrastructure. VMware View servers and View Composer servers require valid SSL certificates to communicate with each other and to work correctly.

7.1.5 VMware View recommendations

Depending on the actual applications and the actual usage of the virtual desktop, it is recommended to adjust the Adobe Flash settings for the remote sessions based on VMware View recommendations.

Enabling View Storage Accelerator is recommended to get the best performance from the storage layer and the best user experience on the virtual desktops. Additionally, use blackout times to ensure that the cache is not regenerated when the users are active on the system as this process puts additional load on the ESXi host and the storage platform causing a performance drop that affects the user experience.

7.2 Server host layer

The ESXi 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 ESXi.
- Install and configure EqualLogic Multipathing Extension Module (MEM) for vSphere 5.0 to get the best performance from the storage array.
- Separate virtual switches to segregate iSCSI SAN traffic, VDI traffic, and Management network traffic.
- Each network path should be assigned to a minimum of two physical NICs, preferably on different hardware paths, for high availability.

VMware KB article on best practices for installing ESXi 5.0: <http://kb.vmware.com/kb/2005099>

Installing and configuring the Dell EqualLogic MEM for VMware vSphere 5:
<http://www.equallogic.com/WorkArea/DownloadAsset.aspx?id=10798>

7.3 Network layer

It is recommended that at least two physical NICs on each server be dedicated to each of the following logical networks:

- VDI Network
- Management network for Infrastructure and vMotion services
- iSCSI SAN

This allows the solution to work even in the event of the failure of a single blade switch, rack switch, or individual NIC port on the storage array or on the ESXi host server.

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

Virtual Switches in ESXi 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 ESXi 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.

7.4 Storage

It is recommended to use the EqualLogic PS6100XS hybrid arrays, which consist of seven SSD drives and seventeen 10,000 RPM SAS drives within a single chassis, for VDI environments. These hybrid arrays automatically move hot data to the SSD tier, which improves performance in VDI environments.

It is very important to ensure that the SAN switches on the blade chassis are uplinked to dedicated external SAN switches. This will ensure future upgradeability by providing the ability to connect additional storage arrays without any downtime to the existing infrastructure.

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

8 Conclusions

Testing of the EqualLogic storage platform showed that the EqualLogic PS6100XS arrays can easily support very fast access to high-demand data in a multi-tiered VDI workload environment. The ability to support a large number of desktops (more than 830 task worker desktops) means that the EqualLogic PS6100XS can help organizations that are limited by the cost-per-desktop hurdle start their VDI deployment plans much sooner. Additionally, the peer storage architecture of the EqualLogic arrays achieves linear scalability of both controller throughput and storage capacity as new arrays are added with the 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.

The benefits of deploying EqualLogic PS6100XS arrays with VMware View 5.1 are also compelling. These benefits include helping reduce VDI deployment costs by handling workloads on highly cost-effective storage media, enhanced VDI performance by placing read-sensitive replica images on low-latency SSDs, and easier VDI deployment through automated and intelligent on-board tiering software. Combined with the ease of administration of EqualLogic arrays, these benefits reduce the TCO for VDI in many organizations.

Appendix A VMware View 5.1 solution configuration

Solution configuration - Hardware components:		Description
Virtual Desktops	13 x Dell PowerEdge M610 Servers: <ul style="list-style-type: none"> • BIOS Version: 6.1.0 • 2 x Hexa Core Intel® Xeon® X5680 3.33Ghz Processors • 96 GB RAM • 2 x 146GB 10K SAS internal disk drives • 1 x Dual-port Broadcom 5709 1GbE NIC (LAN on motherboard) • 2 x Broadcom NetXtreme II 5709s 1GbE NIC, Quad-Port 	ESXi 5.0 U1 with Dell EqualLogic MEM 1.1 installed on all M610 servers. Windows 7 (32 Bit) VMs with 1.5GB RAM.
VMware View Servers Login VSI Launchers [VDI Workload Generator] Infrastructure Servers	3 x Dell PowerEdge M610 servers: <ul style="list-style-type: none"> • BIOS Version: 6.1.0 • 2 x Hexa Core Intel® Xeon® X5680 3.33 Ghz Processors • 96 GB RAM, • 2 x 146GB 10K SAS internal disk drives • 1 x Quad-port Broadcom 5709 1GbE NIC (LAN on motherboard) • 2 x Broadcom NetXtreme II 5709s 1GbE NIC, Quad-Port 	ESXi 5.0 U1 with Dell EqualLogic MEM 1.1 installed on all M610 servers. VMs are loaded with Microsoft Windows 2008 R2. <ul style="list-style-type: none"> • Active Directory • VMware View 5.1 Server 1 • VMware View 5.2 Server 2 • VMware vCenter Server • SQL Server • 30 Login VSI Launchers • Liquidware Stratusphere UX virtual appliance
Network	<ul style="list-style-type: none"> • 2 x Dell PowerConnect M6220 1Gb Ethernet Switch, stacked (Mgmt) • Firmware: 4.2.2.3, A30 • 2 x Dell PowerConnect M6348 1Gb Ethernet Switch, stacked (VDI) • Firmware:4.2.2.3, A21 • 2 x Dell PowerConnect M6348 1Gb Ethernet Switch, (iSCSI) • Firmware:4.2.2.3, A21 • 2 x Dell Force10 S55 1Gb Ethernet Switch, stacked (iSCSI) • Firmware: 8.3.5.2 • 2 x Dell PowerConnect 6248 1Gb Ethernet Switch (VDI) • Firmware: 3.3.3.3, A27 	4 x 10GbE uplinks to the external Force10 S55 switch from PowerConnect M6348 switches on the blade chassis. The Force10 S55 switches are used for the dedicated iSCSI SAN. They are stacked using 2 x 12GbE stacking ports on the rear.
Storage	1 x Dell EqualLogic PS6100XS: <ul style="list-style-type: none"> • 7 x 400GB SSD • 17 x 600GB 10K SAS disks • Dual 4 port 1GbE controllers 	Pools of desktops are stored on the PS6100XS array.

	<ul style="list-style-type: none"> • Firmware: 5.2.4 • 1 x Dell EqualLogic PS6500E: • 48 x 1TB 7.2K SATA disks • Dual 4 port 1GbE controllers • Firmware: 5.2.4 	User data is presented to all virtual desktops using a CIFS share from a Windows based file server. All infrastructure VMs are stored on the PS6500E
Performance Monitoring	<ul style="list-style-type: none"> • SAN HQ – 2.2.0 • vCenter Performance monitoring • Liquidware Stratusphere UX – 5.0.0 	Performance monitoring on EqualLogic arrays. Performance monitoring and capture at the ESXi host. User experience monitoring on the virtual desktop.

Solution Configuration - Software Components:	Description / Version
VMware View Server	5.1.0 build-704644
VMware View Composer	3.0.0 build-691993
VMware View Agent	5.1.0 build-704644
VMware View Client 64bit.	5.1.0 build-704644
VMware ESXi Hypervisor	5.0 U1 build-702118
VMware vCenter Server	5.0.0 build-623373
EqualLogic Multipathing Extension Module	1.1.0
Microsoft SQL Server 2008R2 Enterprise x64	10.50.1600.1
Microsoft Windows 7 Enterprise x64 SP1	VDI Clients for characterization tests
Microsoft Windows Server 2008 Enterprise R2 x64 SP1	VMs for hosting VMware View, vCenter Server, MS-SQL server, Login VSI launchers and other infrastructure VMs.
Login VSI	Version 3.6, Load generator for VDI clients
Liquidware Stratusphere UX	Version 5.0.0
Microsoft Office 2007 Standard	Service Pack 3 (12.0.6607.1000)



Appendix B Network design and VLAN configuration

B.1 Management LAN configuration

- Each PowerEdge M610 server has an onboard Broadcom 5709 dual port 1GbE NIC.
- Dual PowerConnect M6220 switches are installed in fabric A of the blade chassis. The onboard LOM NICs are connected to each of the M6220 switches.

The two PowerConnect M6220 switches are interconnected using 2 x 10GbE stacking interconnects.

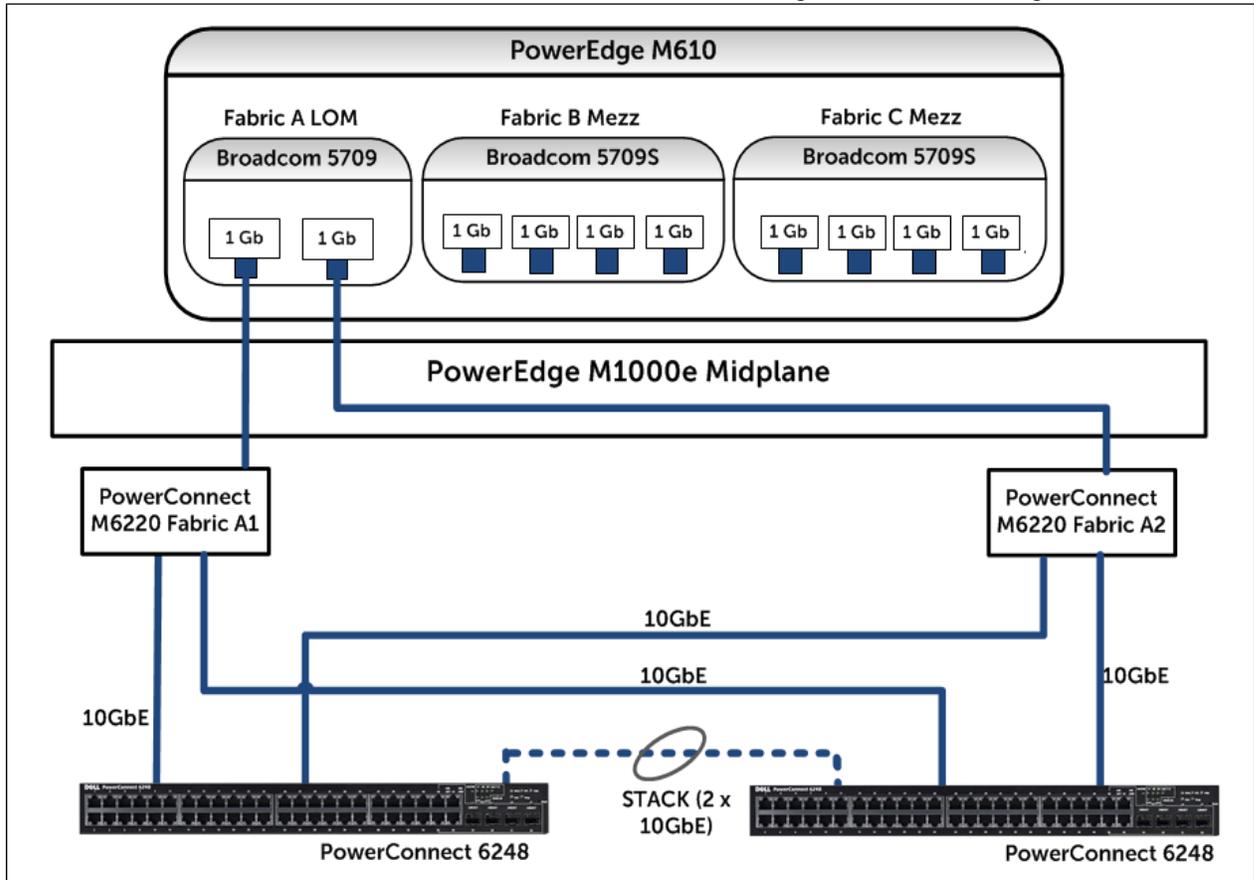


Figure 16 Management LAN connectivity

B.2 VDI LAN configuration

Users wanting to access the Virtual Desktops hosted by View connect on this LAN.

These switches may be uplinked to external switches to provide connectivity to the rest of the organization.

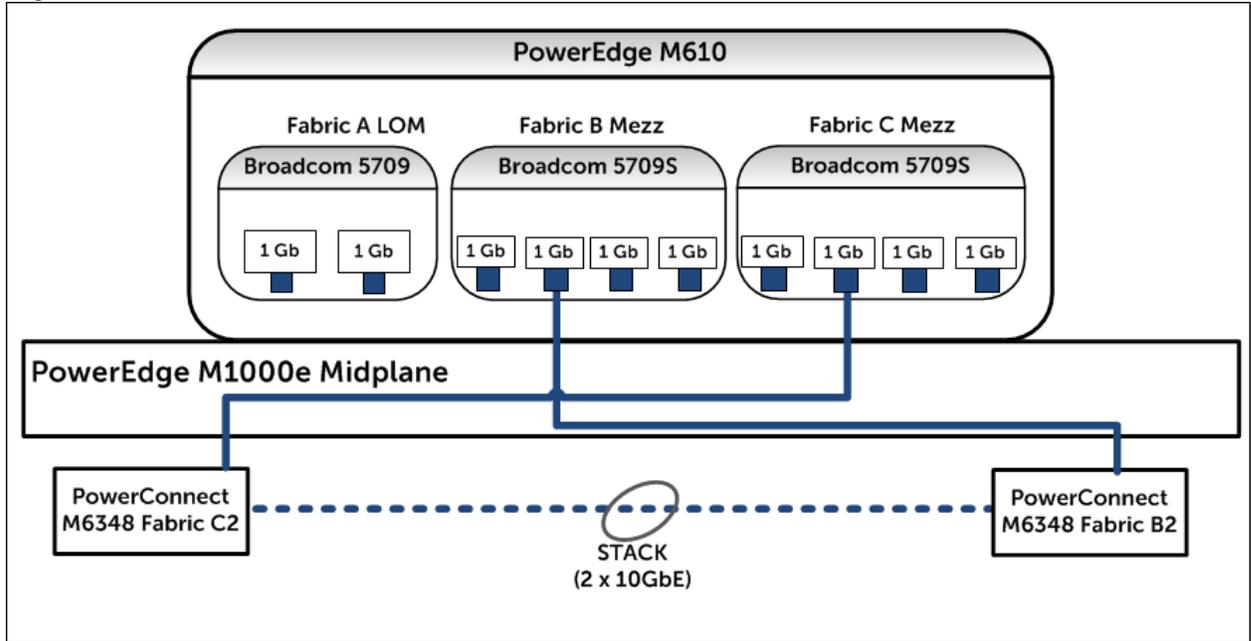


Figure 17 VDI LAN connectivity

Appendix C ESXi host network configuration

Each ESXi host was configured with three virtual switches, vSwitch0, vSwitch1, vSwitchISCSI.

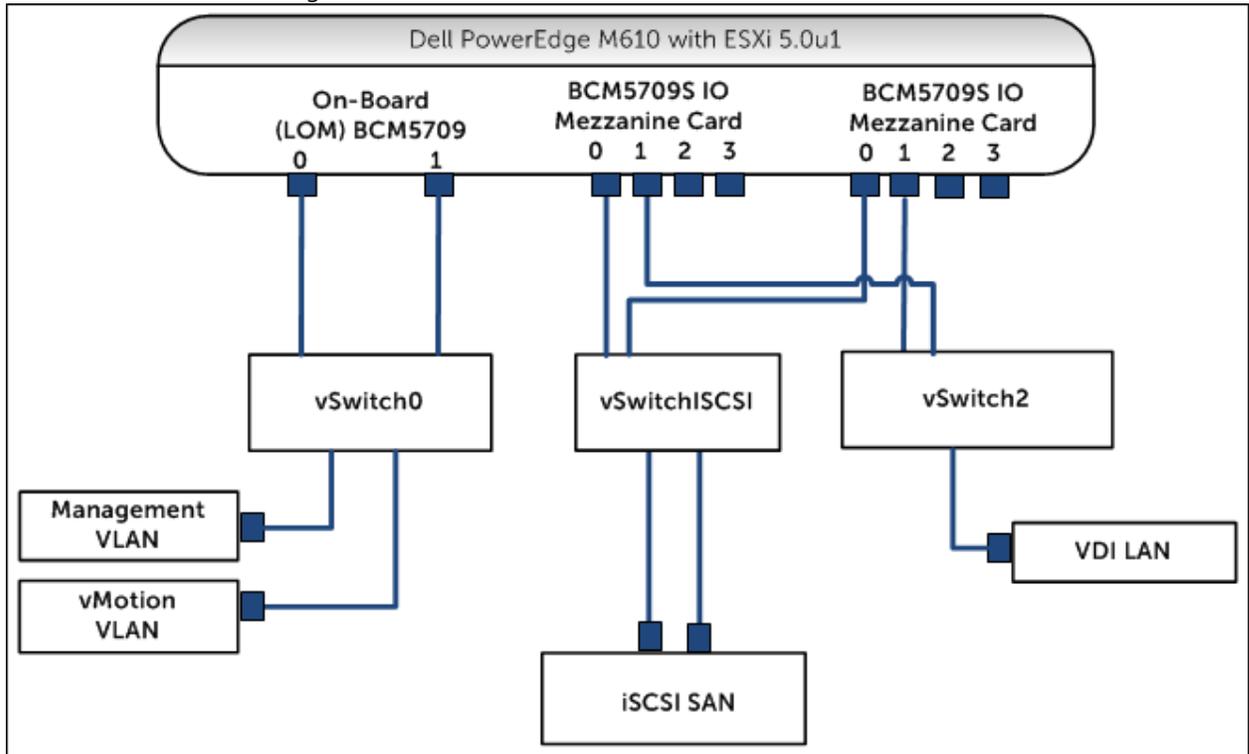


Figure 18 ESXi vSwitch logical connection paths

C.1 vSwitch0

vSwitch0 provides connection paths for all management LAN traffic. The physical adapters from the two onboard NICs (fabric A) were assigned to this switch.

VLANs are used to segregate network traffic into different classes (tagged packets) within this LAN. They are shown in Table 10.

Table 10 Management network VLANs

Port Group	VLAN ID	Description
Management Network	102	vCenter Management
vMotion	103	vMotion Network



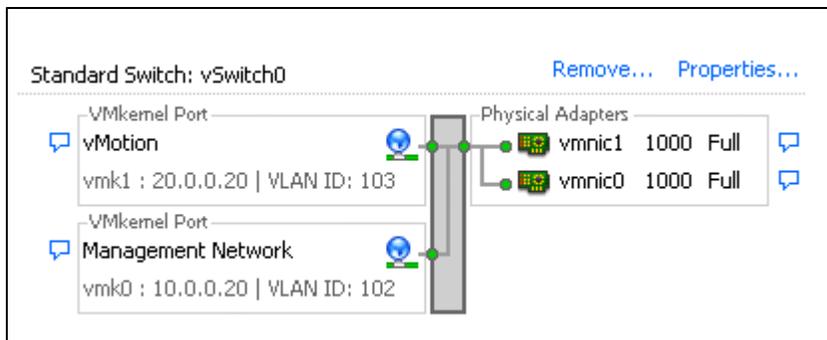


Figure 19 vSwitch0

C.2 vSwitchISCSI

This virtual switch provided paths for all the iSCSI SAN traffic. Two physical adapters were assigned, one from each of the mezzanine cards on fabric B and fabric C.

In the configuration, the software iSCSI initiator provided by the ESXi host was used. To take advantage of EqualLogic-aware multi-path I/O, the EqualLogic MEM for VMware vSphere was installed on each ESXi host.

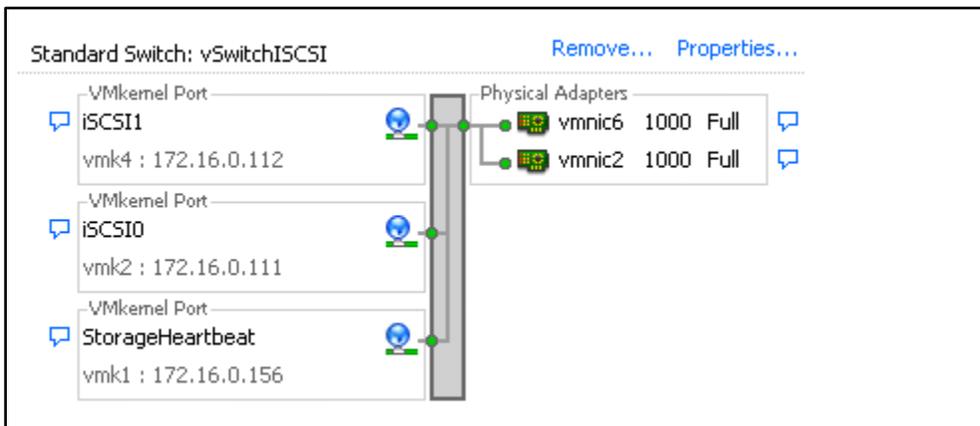


Figure 20 vSwitchISCSI

C.3 vSwitch1

Two physical adapters, one each from the mezzanine cards on fabric B and fabric C are assigned to this virtual switch. This vSwitch carries all traffic for the VDI LAN.

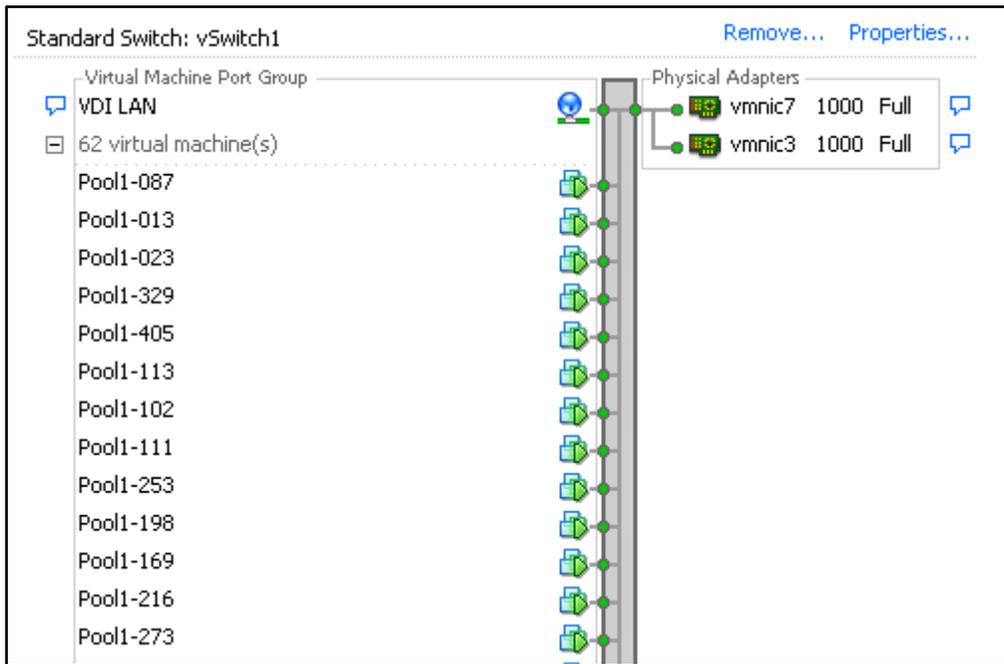


Figure 21 vSwitch1



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 for the purpose of sharing knowledge, best practices, and information about Dell products and your installations.

Referenced or recommended Dell publications:

- Dell EqualLogic PS Series Network Performance Guidelines:
<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>
- Dell EqualLogic PS series arrays – Scalability and Growth in Virtual Environments:
<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19992296.aspx>
- Installing and configuring the Dell EqualLogic MEM for VMware vSphere 5 (EqualLogic Support login required) : <http://www.equallogic.com/WorkArea/DownloadAsset.aspx?id=10798>
- Dell EqualLogic PS Series Architecture: Load Balancers:
<http://www.equallogic.com/WorkArea/DownloadAsset.aspx?id=10752>
- Dell Force10 S-Series S55 2-Switch Reference SAN:
<http://en.community.dell.com/techcenter/storage/w/wiki/3590.san-switch-reference-architectures-by-sis.aspx>

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

- VMware View 5.1 Documentation: <http://pubs.vmware.com/view-51/index.jsp>
- VMware View 5.1 Release notes: <http://www.vmware.com/support/view51/doc/view-51-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 ESXi 5.0:
<http://kb.vmware.com/kb/2005099>
- 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>



This white paper is for informational purposes only. The content is provided as is, without express or implied warranties of any kind.