



Dell EqualLogic Best Practices Series

Sizing and Best Practices for Deploying VMware View 4.5 on VMware vSphere 4.1 with Dell EqualLogic Storage

A Dell Technical Whitepaper

This document has been archived and will no longer be maintained or updated.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.

Ananda Sankaran, Chhandomay Mandal

Storage Infrastructure and Solutions

Dell Product Group

January 2011

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

© 2011 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™*, *Dell™ EqualLogic™*, *PowerEdge™* and *PowerVault™* are trademarks of Dell Inc. *Broadcom™* is a registered trademark of Broadcom Corporation. *Intel™* is a registered trademark of Intel Corporation in the U.S. and other countries. *Microsoft™*, *Windows™*, *Windows Server™*, and *Active Directory™* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Table of Contents

1	Introduction	1
1.1	Audience	1
1.2	The Rest of this Paper.....	1
2	Virtual Desktop Infrastructures	2
2.1	Benefits of VDI	2
2.2	Storage Requirements for VDI	2
2.3	Addressing VDI Storage Challenges with EqualLogic SANs	3
2.4	The VMware View Solution.....	4
2.5	Benefits of using Dell EqualLogic Hybrid Arrays for VDI.....	4
3	VMware View Infrastructure and Test Configuration	5
3.1	Test Infrastructure: Component Design Details.....	7
3.2	ESX Host Network Configuration.....	9
3.3	Virtualized Desktop Scaling: VM Host Limits on Server and Storage.....	11
3.4	VMware View Configuration	12
3.4.1	Using Linked Clones	13
4	VMware View 4.5 Characterization Tests	13
4.1	Test Results.....	14
4.2	Analysis: EqualLogic PS6000XVS SAN hosting non-persistent desktops	15
4.2.1	The login storm	15
4.2.2	The steady state task worker activity phase (after the login storm)	16
4.2.3	Host performance in the View Client ESX cluster	18
5	Sizing Guidelines for EqualLogic SANs.....	21
5.1	Using Persistent vs. Non-Persistent View Desktops	21
5.2	Calculating Capacity Requirements for Linked Clone Desktop Pools	21
5.3	Sizing Storage Capacity.....	23
5.4	Performance Considerations	24
5.4.1	When to use the EqualLogic PS6000XVS	24
5.4.2	When to use the EqualLogic PS6000XV	24
6	Best Practices.....	25
6.1	Application Layer	25

6.2	Virtual Machine and Guest OS Configuration	27
6.3	ESX Host Configuration.....	28
6.4	Host Server Resources	28
6.5	SAN Design	28
6.6	Storage	29
Appendix A	Test System Component Details	31

1 Introduction

Virtual Desktop Infrastructure (VDI) solutions are starting to gain considerable foothold. In addition to the traditional server virtualization benefits, VDI solutions can provide significant additional cost savings due to streamlined implementation and ease of management.

In VDI environments the storage infrastructure must be carefully designed and sized to meet I/O performance requirements while supporting efficient capacity utilization. The goal of this paper is to present the results of a series of VDI tests conducted by Dell Labs, and provide storage sizing guidance and best practices for designing and running **VMware® View 4.5** based virtual desktop infrastructures. Our test infrastructure included:

- **VMware® vSphere™ 4.1 Enterprise Plus**
- **Dell™ PowerEdge™ M610 blade servers**
- **Dell™ PowerConnect™ switches**
- **Dell™ EqualLogic™ SANs.**

1.1 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 VMware View based VDI solutions using Dell EqualLogic storage. It is expected that the reader has working knowledge of VMware vSphere system administration, Microsoft® Windows® server and client operating system administration, iSCSI SAN network design and Dell EqualLogic iSCSI SAN operation.

1.2 The Rest of this Paper

The rest of this paper contains the following sections:

- Section 2, **Virtual Desktop Infrastructures** on page 2
- Section 3, **VMware View Infrastructure and Test Configuration** on page 5
- Section 4, **VMware View 4.5 Characterization Tests** on page 13
- Section 5, **Sizing Guidelines for EqualLogic SANs** on page 21
- Section 6, **Best Practices** on page 25

2 Virtual Desktop Infrastructures

Desktop virtualization is emerging as an important strategy for organizations seeking to reduce the cost and complexity of managing an expanding variety of client desktops, laptops, netbooks, and mobile handheld devices. In a VMware View based VDI environment, user desktops are hosted as virtual machines in a centralized infrastructure. The user interface to the desktop virtual machine is transmitted over a network to an end-user's client device.

2.1 Benefits of VDI

In a VDI environment all desktops are centrally managed. Each desktop can be derived from a small number of "gold" base images that system administrators can easily and efficiently manage. When patches and upgrades are applied to a base image, the changes are automatically propagated to all user desktops derived from that base image. This eliminates any issues with application and device compatibility. Workstation provisioning is simpler and faster than in traditional desktop PC environments. User workloads can be (automatically) moved between resource pools in the centralized infrastructure during demand spikes. Using a VDI delivery model, IT organizations can execute smoother transitions during upgrades or new operating systems deployment periods.

The centralized desktop management model of VDI lends itself to easier data protection. This is because, from a storage point of view, a VDI desktop environment is similar to any enterprise application running in a data center. Security and compliance are other areas where VDI solutions create advantages over the traditional desktop PC computing model. In a VDI environment, it is possible to ensure that no data resides on end-point devices. The enterprise security risks that are introduced by end users locally installing applications can be significantly reduced by using a tightly-controlled centrally-managed VDI environment.

2.2 Storage Requirements for VDI

A VDI deployment must be carefully designed to ensure that it delivers the performance and scalability needed to support large scale client computing requirements inside an enterprise. When migrating to a VDI model, a lot of the components—storage, processors, memory, and networking—that supported the traditional desktop PCs are moved to a centralized data center.

A VDI deployment can place high capacity and performance demands on the underlying storage platform. For example, consolidation of large amounts of inexpensive stand-alone desktop storage into a centralized infrastructure translates into capacity requirements that the shared storage used by the VDI solution will need to meet. Performance considerations include how to estimate the aggregate volume of I/O operations per second (IOPS) generated by basic desktop client operations (system boot, logon, logoff, and application usage patterns) that VDI solution will need to support. IOPS requirements will change based on user profiles. The application I/O profile generated by a typical task worker will differ significantly from that created by a knowledge worker or power user. I/O "storm" events can be caused by large amount of simultaneous activity, such as morning logons and afternoon logoffs. The events can cause I/O spikes that place high performance demands on the storage infrastructure.

When designing storage systems for a VDI deployment you must take all of these considerations into account. The storage platform will need to meet the performance demands generated by utilization spikes and be able to cost-effectively scale to meet capacity requirements. A VDI storage infrastructure should also be virtualization aware, so that the virtualization layer can offload processor-intensive tasks (such as copying hundreds of virtual machines for desktop provisioning) to the storage layer, where it can be done more efficiently.

2.3 Addressing VDI Storage Challenges with EqualLogic SANs

Dell EqualLogic PS Series iSCSI SANs are well suited for supporting VDI deployments. They offer scalable, high-performance virtualized storage that is designed for reliability, manageability, and efficiency. EqualLogic SANs come with a range of efficiency features designed to enhance utilization and help reduce costs. For example, automated workload tiering and load balancing features help to optimize storage performance. Storage virtualization coupled with thin provisioning can provide tremendous capacity while reduced physical storage requirements. Also, thin clones¹ allow administrators to clone volumes without initially consuming any additional storage space. This feature is ideally suited for VDI environments where it may be necessary to create many copies of virtual machine images. For example, one can populate a template volume with virtual machine images, then clone that template volume as needed to provision virtual desktops. VMware View linked clones allow you to achieve similar management and utilization efficiencies.

EqualLogic storage is designed to be fully VMware virtualization aware. This enables automation and optimization of key storage functions such as desktop provisioning. For example, EqualLogic SANs integrate with VMware vStorage Application Programming Interfaces (APIs) for Array Integration (VAAI). This integration enables the VMware software to offload tasks such as full copy directly to EqualLogic SANs. This reduces the time required to provision and deploy new virtual machines. The hardware-assisted locking feature in VAAI reduces boot time in a VDI environment. The EqualLogic Multipath Extension Module (MEM) for VMware vSphere provides advanced multi-path I/O functionality, such as automatic connection load balancing.

EqualLogic PS Series also includes the Host Integration Tool Kit (HIT Kit) for VMware. The HIT Kit includes Auto-Snapshot Manager/VMware Edition (ASM/VE), which can provide automatic virtual machine protection using VMware-based and SAN-based snapshots. The HIT Kit for VMware also includes a Virtual Desktop Deployment Tool that leverages thin clone technology to space-efficiently provision and deploy virtual desktops in VMware-based VDI environments.

Finally, the EqualLogic PS6000XVS and EqualLogic PS6010XVS hybrid SANs further enhance the performance and efficiency of VDI deployments by combining SSD and Serial Attached SCSI (SAS) drives within a single chassis. The on-board intelligence of the EqualLogic firmware provides automatic tiering and balancing of stored data between the SSD and SAS tiers. This device tiering within the array creates a flexible and powerful balance between performance and responsiveness provided by the low latency SSD tier and capacity provided by the SAS tier. The results we present in this paper confirm that performance and capacity balance provided by the EqualLogic XVS arrays is exceptionally well suited for large VDI environments.

¹ The thin clone feature was introduced with EqualLogic controller firmware Version 5.

2.4 The VMware View Solution

VMware View is one of the leading VDI solutions in the market today. It includes a complete suite of tools for delivering desktops as a secure, managed service from a centralized infrastructure. The components of a VMware View VDI deployment include client connection devices (desktops, laptops or thin clients), VMware View Manager for connection brokering and virtual desktop management, VMware View Composer for space-efficient provisioning of desktop virtual machines (VMs), and vSphere ESX server virtualization for hosting virtual desktops.

A VMware View infrastructure consists of many different software, network and hardware layer components. Here is a functional list of VMware View components:

Client Devices	Personal computing devices used by end users to run the View Client or View Thin Client.
View Connection Server	Authenticates users via Microsoft Active Directory and brokers client connections
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. Facilitates communication with View clients.
View Administrator	Web-based administration platform for View Infrastructure components.
vCenter Server	Central administration platform for configuring, provisioning and managing VMware virtualized datacenters.
View Portal	Web-based service for provisioning and starting Windows PC based View clients.
View Composer	A service that runs on vCenter servers. Used to create pools of linked clone VM images in order to reduce storage requirements.
View Transfer Server	Provides the data transfer service needed by Local Mode View Clients for offline operation.

Note: For detailed descriptions for each component above, see the Architecture Planning Guide section of the VMware View Online Library:

➤ <http://pubs.vmware.com/view45/ol/>

2.5 Benefits of using Dell EqualLogic Hybrid Arrays for VDI

In VMware View, it is possible to have many linked-clone desktops sharing the same base desktop image. This scenario creates a situation where the data in a relatively small portion of the storage system becomes extremely hot (very high read/write I/O rates). The EqualLogic PS6000XVS and PS6010XVS are designed to support fast access to high demand "hot" data segments. These hybrid storage arrays have sixteen drives – eight high performance SSD drives and eight high capacity 15K rpm SAS hard disk drives. When data becomes "hot" (accessed frequently) it is automatically moved from the SAS tier to the SSD tier. This automatic tiering function allows the hybrid XVS SAN very

effectively balance peak load capabilities and scale-out storage capacity requirements for VDI environments.

3 VMware View Infrastructure and Test Configuration

The core VMware View infrastructure components used in our test configuration are shown in Figure 1 below.

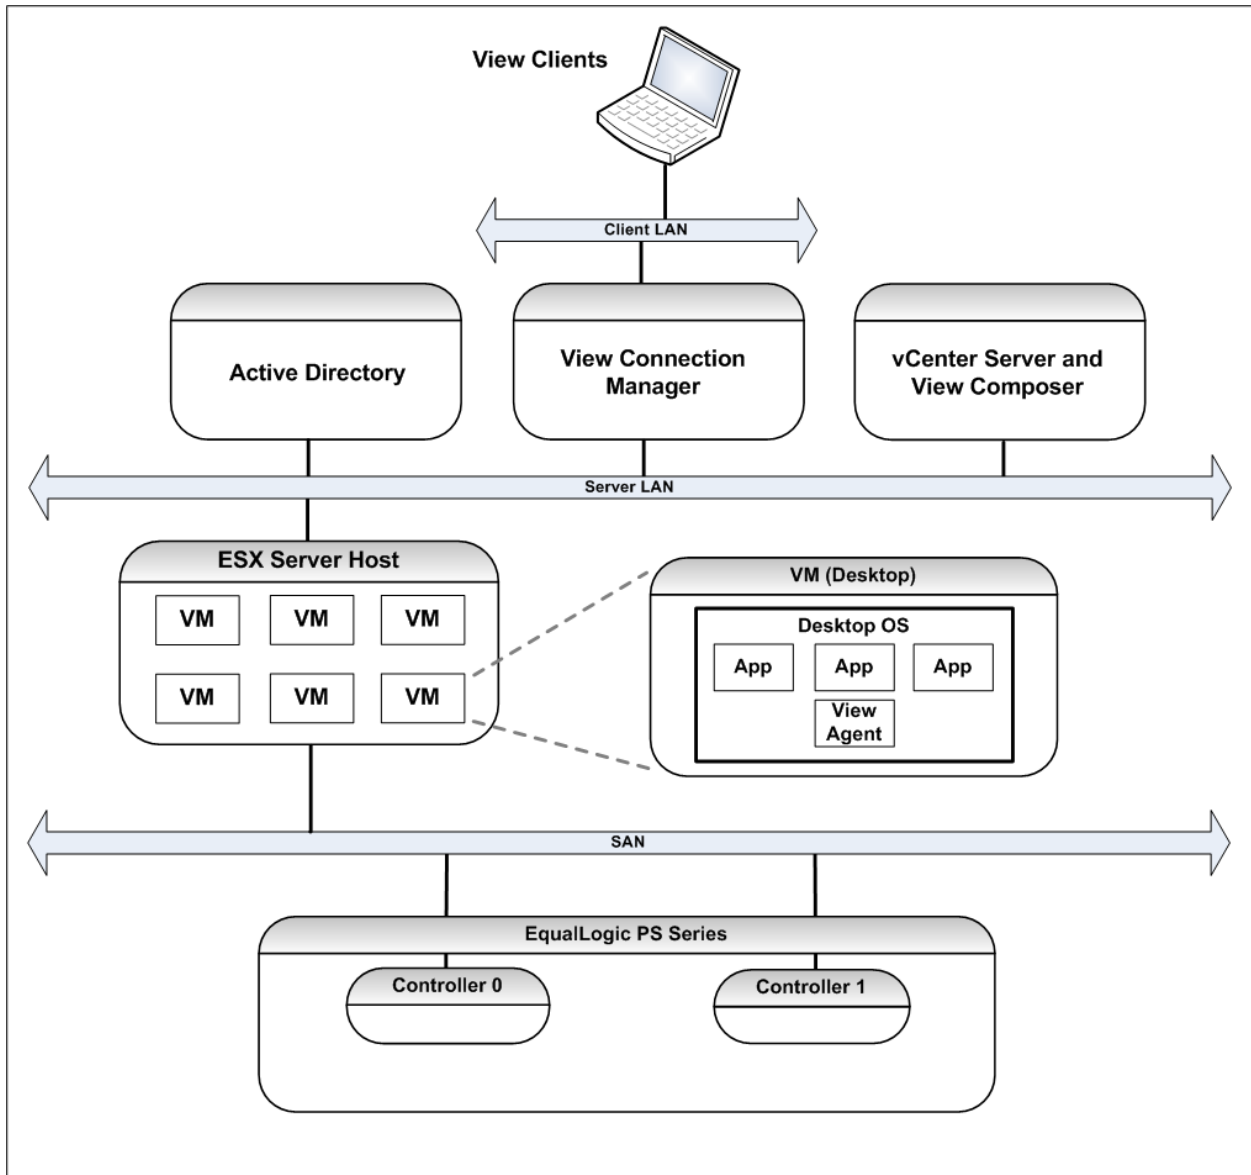


Figure 1: Test Configuration Functional Components

We added the following components to the test system configuration to simulate a realistic VDI workload:

RAWC² Controller	Provides the RAWC system configuration and management GUI for desktop workload simulation
RAWC Session Launchers	Automates launch of View VDI client sessions
Microsoft SQL Server	Provides database for vCenter and View Composer
Microsoft Exchange 2010	Email server for Outlook clients running on each desktop

We used Dell PowerEdge M610 blade servers and Dell PowerConnect M6220/M6348 Ethernet blade switches within a Dell PowerEdge M1000e Modular Blade Enclosure as the host platform for the test system components. Figure 2 shows how we distributed the virtual machines running the core View components and the additional simulation components on our M610 blade servers.




Figure 2: ESX Blade Clusters

The sixteen PowerEdge M610 Servers were subdivided into three ESX clusters. Three M610 blade servers were combined as the infrastructure cluster. The infrastructure cluster provided an aggregate computing power of 57.6GHz and 388GB of memory. Thirteen M610 E5680 blade servers were split into two different ESX clusters for hosting View client VMs: the first with 8 hosts³ and the second with 5 hosts. The two View Client VM clusters provided an aggregate computing power of 520GHz and 1248GB of memory.

² Desktop Reference Architecture Workload Code. See *Workload Simulation for Virtual Desktop Reference Architectures*, <http://www.vmware.com/files/pdf/VMware-WP-WorkloadConsiderations-WP-EN.pdf>

³ Due to View Composer limitations in View 4.5, VMware HA clusters in a View infrastructure are limited to a maximum of 8 ESX nodes. For more information, see the vSphere Clusters section of the VMware View Online Library: <http://pubs.vmware.com/view45/ol>

The size of the parent VM image used in our desktop pools was 15GB. We allocated one vCPU and 1GB of memory for each desktop VM. A linked-clone desktop pool was assigned to each ESX cluster.

3.1 Test Infrastructure: Component Design Details

Table 1 below provides an overview of the components used to build our test configuration.

Component	Purpose / Usage
Servers	
3 x Dell PowerEdge M610 Blade Servers: <ul style="list-style-type: none"> 2 x Quad Core Intel Xeon E5620 2.40 GHz CPUs and 96GB of Memory Broadcom 5709 quad-port 1GbE mezzanine card 	VMware ESX 4.1 Infrastructure Cluster (vCenter Server, domain Controllers, View Connection Manager, etc.)
13 x Dell PowerEdge M610 Blade Servers: <ul style="list-style-type: none"> 2 x Six Core Intel Xeon X5680 3.33 GHz CPUs and 96GB of Memory Broadcom 5709 quad-port 1GbE mezzanine card 	VMware ESX 4.1 View Client VM Clusters (one 8 host cluster and one 5 host cluster)
Storage Arrays	
1 x EqualLogic PS6000XVS <ul style="list-style-type: none"> 8 x 450GB 15K rpm SAS 8 x 100GB SSD 1 x EqualLogic PS6000XV <ul style="list-style-type: none"> 16 x 600GB 15K rpm SAS 	VMware View Desktop Clients: ISOs, Software, VM Template storage
1 x EqualLogic PS6000XV <ul style="list-style-type: none"> 16 x 600GB 15K rpm SAS 1 x EqualLogic PS6000XV <ul style="list-style-type: none"> 16 x 600GB 10K rpm SAS 1 x EqualLogic PS6000S <ul style="list-style-type: none"> 16 x 50GB SSD 	Infrastructure Data and VMs: - ISOs, Software, VM Templates - RAWC Server, vCenter Server, Active Directory, View Connection Brokers, Email Server
Network Switches	
Dell PowerConnect M6220	Ethernet blade switch for external and internal LAN connectivity (VMware Service Console, Server LAN)
Dell PowerConnect M6348	Ethernet blade switch for 1GbE iSCSI
Dell PowerConnect 6248	External switch for 1GbE iSCSI
Software	
VMware ESX 4.1 Enterprise Plus	Hypervisor
VMware View 4.5 Premier Edition	VMware Virtual Desktop Infrastructure Application
Windows 7 Enterprise Edition	View Client OS for characterization tests
Windows Server 2008 Standard x64	For Domain Controller,, vCenter Server, Microsoft SQL server, Exchange 2010 and other infrastructure VMs
MS Office 2007	Applications used to generate user workload
RAWC 1.2	VMware View VDI workload generation
Management and Monitoring	
EqualLogic SAN HeadQuarters 2.1	Storage Monitoring

vCenter Performance Monitor	Performance monitoring and capture at the ESX host
ESXTOP	Performance monitoring and capture at the ESX host

Table 1: Test Components

The PowerConnect M6348 switch modules and the PowerConnect 6224 top-of-rack switches were dedicated to the iSCSI SAN. The PowerConnect M6220 switch modules were dedicated to the server LAN. Figure 3 shows the overall topology and the connection paths used by each M610 blade servers (only one server is shown).




Figure 3: Server LAN and iSCSI SAN Diagram

Server LAN Configuration:

- Each PowerEdge M610 server has an onboard Broadcom 5709 dual-port 1GbE NIC.
- Dual PowerConnect M6220 switches were installed in fabric A of the blade chassis. We connected the onboard NICs to each the M6220 switches.
- The two M6220 switches were inter-connected using a 2 x 10GbE LAG.

SAN Configuration:

- Each PowerEdge M610 server included two Broadcom NetXtreme II 5709 quad port NIC mezzanine cards. We assigned one card to fabric B and the other to on fabric C.
Note: The tests for this white paper used only 2 NIC ports per card as shown in Figure 3.
- We installed dual PowerConnect M6348 switches into fabric B and fabric C on the blade server chassis. The NIC mezzanine cards connected to these switches via the blade chassis mid-plane. Note: PowerConnect M6220 switches can be used instead if only two ports per fabric are needed per server.
- We used four PowerConnect 6248 switches as our external SAN switches. We connected each of our EqualLogic storage arrays to these switches. Each port on the quad-port EqualLogic storage controller (PS6000XV or PS6000XVS) was connected to a different 6248 switch. Figure 3 shows these connection paths.
- The PowerConnect 6248 switches were configured with 10GbE SFP+ uplink modules in both module bays. As shown in Figure 3, one module was used to create a 2 x 10GbE LAG uplink to the M6348 blade switches. The other module was used to create a ring of 10GbE links between each 6248. Spanning tree protocol (STP) settings were adjusted to create a logical disconnect in the ring (if one of the other links in the loop fails, the open link is re-enabled by STP).
- Each PowerConnect M6348 switch on the chassis connected to the external SAN switches using a 2 x 10GbE Link Aggregation Group (LAG). Note: this infrastructure design could scale out with addition of a second M1000e chassis. SAN connectivity in this case would be accomplished by using a single 10GbE LAG inter-connect between the M6348 switch modules in each blade chassis and the external SAN switches. The switch modules in a second blade server chassis could then inter-connect to the same external SAN switches by using the second 10GbE LAG connection ports on each switch.

3.2 ESX Host Network Configuration

We configured two virtual switches, vSwitch0 and vSwitchiSCSI, on each ESX host as shown in Figure 4. Virtual switch configuration details:

vSwitch0	vSwitch0 provided connection paths for all Server LAN traffic . We assigned the physical adapters corresponding to the two onboard NICs (fabric A) to this switch.
vSwitchiSCSI	vSwitchiSCSI provided connection paths for all iSCSI SAN traffic . We assigned four physical adapters to this vSwitch: two connecting to fabric B and two connecting to fabric C. Four VMkernel ports were created and attached to the ESX iSCSI software initiator. Each

physical NIC uplink was exclusively assigned to each port.

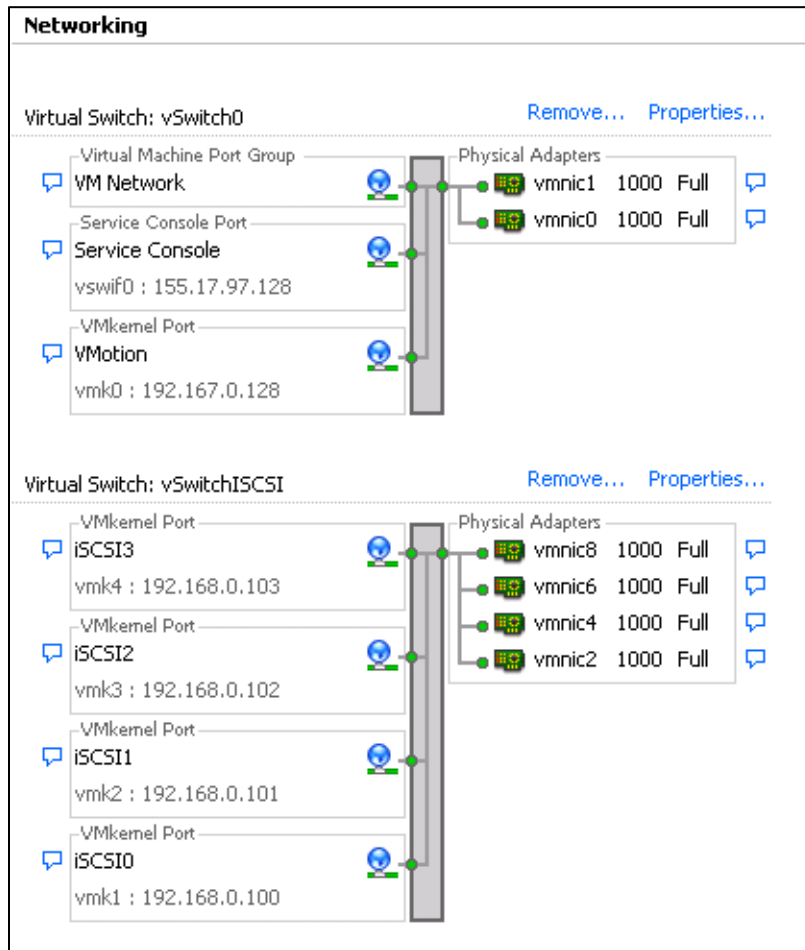


Figure 4: ESX vSwitch Configuration

We used VLANs to segregate network traffic into different classes (tagged packets) within the Server LAN. VLAN and port group assignments for the Server LAN (vSwitch0) were assigned as shown in Table 2. Figure 5 shows the logical connection paths for vSwitch0 and vSwitchiSCSI.

Port Group	VLAN	Description
Service Console	101	Service console access network
Virtual Machine Network	102	VM / Server LAN
vMotion (VMkernel)	103	vMotion network

Table 2: Server LAN (vSwitch0) Port Group Assignments

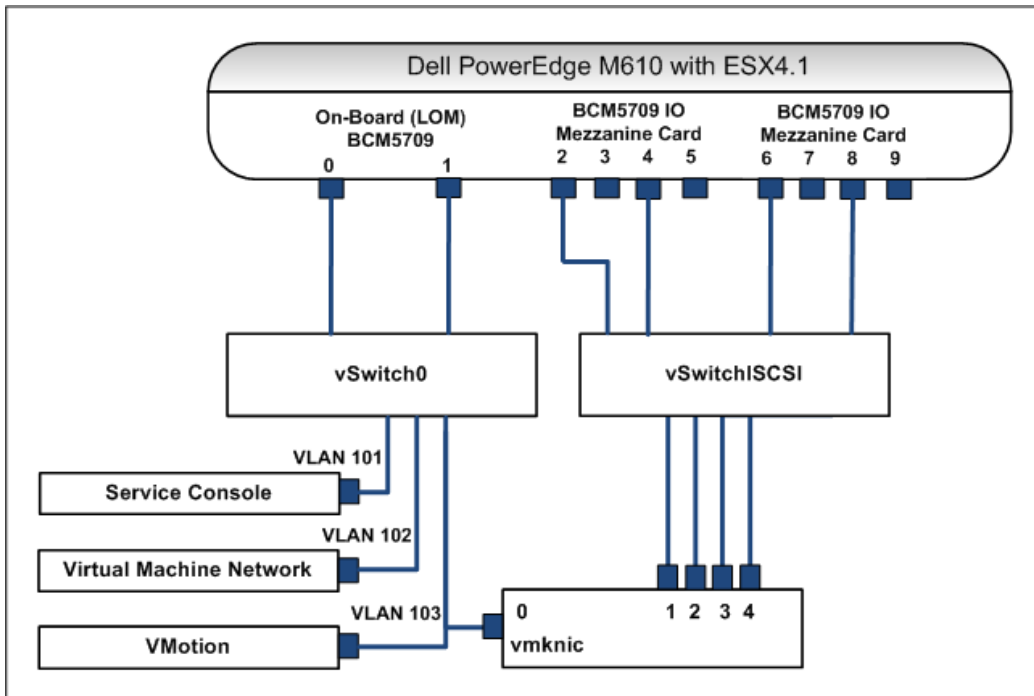


Figure 5: ESX vSwitch Connection Paths

In our configuration we used the software iSCSI initiator provided by the ESX host. To take advantage of EqualLogic-aware multi-path I/O, the EqualLogic Multipathing Extension Module (MEM) for VMware vSphere was installed on each ESX host.

Note: For detailed information on using the EqualLogic Multipathing Extension Module, see the following publication:

➤ *Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere 4.1 and PS Series SANs:*

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=9823>

3.3 Virtualized Desktop Scaling: VM Host Limits on Server and Storage

We ran a series of scale-up tests using the RAWC simulation tool to determine the maximum number of virtual desktop VMs that we could host on a single PowerEdge M610 server used in our View Client ESX clusters. We used optimal CPU, memory, network bandwidth utilization, network TCP retransmission and disk IOPS latency as the criteria for this test. The results indicated that we could optimally host up to 78 virtual desktop VMs per M610 ESX server. We could have scaled the number of VMs per ESX server to a higher level than this. However, we took a conservative strategy by including minimal or no memory ballooning as part of the criteria. Our primary focus in this case was to test scaling of EqualLogic SANs in a VDI deployment while not hitting any server performance limits during the tests. To determine the maximum number of View Client VMs the EqualLogic SAN could host, we used the following performance criteria:

- Maintain less than 20ms disk latency under any VDI workload;
- Provision, on average, a minimum of 2GB storage per linked clone delta disk.

3.4 VMware View Configuration

The VMware View configuration settings used during all performance characterization tests are shown in Table 3 below.

Entity	Description
<i>Pool Definition</i>	
Desktop pool type	Automated Pool
User assignment	<ul style="list-style-type: none"> • For persistent desktop tests: Dedicated with automatic assignment • For non-persistent desktop tests: Floating
vCenter Server (Desktop VM provisioning type)	View Composer linked clones
<i>Pool Settings</i>	
Remote Desktop Power Policy	Ensure desktops are always powered on
Default Display Protocol	PCoIP
<i>View Composer Disks</i>	
Persistent Disk	<ul style="list-style-type: none"> • Persistent Desktop tests: All changes – writes to operating system, temporary system and application data, as well as user data are saved to a persistent disk (2GB) • Non-Persistent Desktop tests: No changes are persisted – reboots (and refresh on user logoff) changed the VM to its initial state.
<i>Provisioning Settings</i>	
Provisioning	Provision all desktops up-front
Datstores	<ul style="list-style-type: none"> • PS6000XV tests: OS disks, persistent disks (if applicable) and replica disks on same datastore. Each datastore included a copy of the master replica image. Storage over-commit⁴ set to moderate for linked clones. • PS6000XVS tests: OS disks along with persistent disks (if applicable) on separate datstores from replica disks. Only the replica datastore included copy of master replica image. Storage over-commit set to aggressive for linked clones.

Table 3 VMware View Configuration Settings

⁴ See *Storage Overcommit for Linked-Clone Desktops* section in the VMware View Online Library

3.4.1 Using Linked Clones

You can realize significant storage space savings and increased efficiencies in client VM provisioning and administration when using linked clones. To setup a linked clone pool, you first create a “parent” VM image with the required OS, settings and software installed. View Composer will coordinate with vCenter to create a pool of linked clones, using the parent VM image as the base image for the clones.

Each linked clone functions as an independent desktop VM with its own unique identity. Linked clones share the base image as their OS image and thus consume significantly less storage space than a complete independent VM image. Temporary system data as well as user data unique to each linked clone desktop VM is written to storage separately. This temporary data can be deleted at the end of the session. Persistent data such as user profile and application/user data can optionally be redirected to a persistent storage location assigned to each desktop VM. Thus software maintenance updates and patches need only be applied only to the base image to take effect on all linked clones, without affecting any unique user settings and data.

View composer first creates a full replica of the parent VM, and then uses that replica to create linked clones. The replica can be placed on the same datastore as the linked clones, or on a separate datastore. By placing the replica on a separate datastore, you have the opportunity to use high performance storage (SSD) for storing the parent image for the linked clones. This strategy can lead to significant I/O performance gains in View VDI infrastructures.

Note: See the “**Creating Desktop Pools**” section in the *VMware View Online Library* for more information about how to configure Linked Clones:

➤ <http://pubs.vmware.com/view45/ol/>

4 VMware View 4.5 Characterization Tests

The objective of our testing was to measure VMware View VDI I/O workloads and determine scale-out guidelines for how many virtual desktop clients can be supported by Dell EqualLogic SANs. From the storage perspective, the primary limits governing how much you can scale a virtual desktop infrastructure are the storage capacity and the maximum sustained IOPS that your storage system can provide, while staying below acceptable I/O latency limits. The VDI storage system needs to be able to sustain the required IOPS per desktop while staying within acceptable read/write I/O latency limits. The typical industry standard latency limit for I/O reads and writes is less than 20ms. This limit will result in good user application response times (assuming no other infrastructure components become a bottleneck). For capacity, we provisioned a minimum of 2GB storage on average per linked clone delta disk.

We used VMware’s VDI load generation tool named Reference Architecture Workload Code Simulator (RAWC) to simulate full scale VDI workloads while measuring system I/O performance. RAWC consists of a controller server and scripts installed on each VDI client. The tool launched View Client Sessions (via session launcher VMs) into the desktop VMs hosted in the View infrastructure. Within a desktop

VM, RAWC also automates launching of applications such as Word, Excel, Outlook, Adobe Reader and others, based on specific criteria such as typing speed, number of emails and number of pages to change. We customized the application simulation settings to achieve our desired user load profile.

Note: For more information on the Reference Architecture Workload Simulator, see the following VMware Information Guide:

- **Workload Considerations for Virtual Desktop Reference Architectures:**
www.vmware.com/go/view4rawc

In the following section, we present scale-out results for a task worker profile on different EqualLogic iSCSI SANs. We created the task worker profile by choosing MS Word, MS Excel, MS PowerPoint, MS Outlook, Adobe Reader, Internet Explorer and McAfee Anti-Virus (scan on access mode) from the set of available applications in RAWC tool. This task worker profile was configured to generate 3-5 IOPS per desktop in steady state. We also simulated login storm events by starting all client sessions within a set time frame, in order to simulate realistic login periods within an enterprise setting.

4.1 Test Results

We tested both persistent and non-persistent desktops to determine the maximum number of desktops that can be supported in EqualLogic iSCSI SANs for the chosen workload. For persistent desktops, all changes (writes to operating system, temporary system and application data as well as user data) are redirected to persistent storage. All such changes are lost when a non-persistent desktop is rebooted or refreshed at user logoff. Thus, every time the non-persistent desktop login event occurs, extra client I/O operations take place in order for Windows to properly initialize the client session. Consequently, during login storms the non-persistent desktop creates a higher processing and I/O workload test case than the persistent desktop. The highest IOPS levels we measured during our tests occurred during the non-persistent desktop login storm periods. Table 4 presents a summary of the test cases and the maximum recommended VDI client loading levels – both for non-persistent and persistent configurations – for the task worker workload.

View Client VM Storage Platform	Test Case	VDI Client Maximum (Task Worker Workload Profile)
EqualLogic PS6000XVS 8 x 100GB SSD 8 x 450GB 15K rpm SAS	A) Non-persistent desktop	1014
	B) Persistent desktop	1014
EqualLogic PS6000XV 16 x 600GB 15K rpm SAS	C) Non-persistent desktop	390
	D) Persistent desktop	780

Table 4: Test Case Summary

The results in Table 4 above show the following:

- For both persistent and non-persistent desktops, we comfortably scaled up to 1014 desktop VMs using a single PS6000XVS array for storage (test cases A and B above). Non-persistent desktops created a peak load of 8600 IOPS during the login storm. The PS6000XVS was able to meet this I/O demand with performance headroom to spare. Also, the PS6000XVS was able to store all 1014 desktop VM images while maintaining 15% spare capacity in the array, where each VM image was provisioned with at least 2GB of storage for the linked clone delta disk. These test results show that the PS6000XVS, using a combination of SSD and 15K SAS drives, provides a balance between storage performance and capacity that is exceptionally well suited for supporting VDI environments.
- For test case C (using the PS6000XV with non-persistent desktops), 390 desktops was the maximum we reached before we exceeded 20ms disk I/O latency (measured from the ESX host) during the login storm phase.
- For test case D (using the PS6000XV with persistent desktops), 780 desktops was the maximum we reached before we exceeded 20ms I/O latency (measured from the ESX host) during the login storm phase.

4.2 Analysis: EqualLogic PS6000XVS SAN hosting non-persistent desktops

In this section we present a detailed analysis of results for Test Case A in Table 4: EqualLogic PS6000XVS hosting **non-persistent desktops**. The results for Test Cases B, C and D can also be analyzed along the same lines using server and storage configuration, EqualLogic SAN Headquarters (SAN HQ), and vCenter data available from respective simulation runs.

Configuration details for Test Case A at maximum VDI client load:

- 13 ESX server hosts were utilized to scale the simulation up to 1014 concurrent VDI client sessions.
- The 13 ESX hosts were split into two ESX clusters of 8 hosts and 5 hosts each (VMware View Composer supports only 8 hosts per ESX cluster).
- The 8 host cluster was provided with 4 SAN volumes on the EqualLogic PS6000XVS. Each volume was mounted as a VMFS datastore. One of these volumes (32GB) was used for the linked-clone master replica image. The other 3 volumes (500GB each) were used for storing the client VM image differential data.
- The second 5 host ESX cluster was provided with 3 volumes: one for the linked clone master replica (32GB), and two for client VM image differential data (500GB each).

4.2.1 The login storm

Figure 6 below shows I/O data as measured by SAN HQ at the peak I/O point during the login storm phase of the 1014 desktop test. For non-persistent desktops, each login event causes Windows to recreate the user profile, read and apply GPO data, and perform other tasks associated with first time user login. As you can see in Figure 6, the login storm creates significant I/O load on the storage platform. IOPS were 8,607 while I/O latencies remained well below the 20ms limit during the login storm peak.




Figure 6: Login storm I/O performance detail: PS6000XVS hosting 1014 non-persistent desktops

Figure 7 shows the I/O performance measured at the individual disk drives within the storage system at the same login storm peak I/O point. The data in the table below the chart shows that during peak load approximately 90% of the array IOPS were served by the SSD drives. During login storm most of the read I/O is targeted to the replica image. As this happens, the replica image data becomes “hot.” The PS6000XVS array automatically moves the hot data to the SSD drives. SSD drives provide significantly better random read I/O performance as compared with the 15K SAS drives. Automatic movement of the hot data to the SSD tier within the array significantly increased the overall I/O performance of the array during login storm.




Figure 7: PS6000XVS disk drive IOPS during non-persistent desktop login storm

4.2.2 The steady state task worker activity phase (after the login storm)

After all of the desktop sessions have launched and completed their login phase there was a period of steady state application activity. Figure 8 shows I/O data as measured by SAN HQ at a peak I/O point during the steady state phase of the 1014 desktop test. The load created 3600 IOPS (approximately 3.5 IOPS per user desktop) and the array handled the load while staying below 5ms latency on average.




Figure 8: Steady state I/O performance detail: PS6000XVS hosting 1014 desktops

Figure 9 shows the I/O performance measured at the individual disk drives within the storage system at the same measurement point as in Figure 8. The data table below the chart shows that approximately 85% of the array IOPS were served by the SSD drives and the remaining 15% served by the 15K SAS drives in the array. The increased I/O load on the 15K SAS drives are due to the fact that the steady state phase of the test has a high write percentage (25% read and 75% writes in the test workload) and a significant portion of this I/O is targeted at the differential data disks of the individual desktop VMs.




Figure 9: PS6000XVS disk drive IOPS during non-persistent desktop steady state phase

4.2.3 Host performance in the View Client ESX cluster

During the test, we measured CPU, memory, network and disk performance on all of the View Client cluster ESX hosts⁵. The performance of one of the ESX hosts is presented here. It is representative of all the ESX hosts under test. The results shown in the figures in this section were captured using VMware vCenter.

Figure 10 shows the CPU utilization for one of the ESX hosts during the test. The CPU usage ramped up during the login phase as new client sessions were launched. Overall, the CPU utilization of the ESX host averaged less than 50% during the test. Figure 11 shows that server memory utilization stayed less than 90% during the test run. Memory ballooning was minimal to none on the ESX host. As seen in the figure, transparent memory sharing between VMs started in the latter part of the test period. This led to lower memory usage than the first part of the test. The increased memory usage toward the end of the test was caused by client activities during logouts.




Figure 10: CPU utilization for one ESX View Client Cluster Host during 1014 non-persistent desktop test

⁵ Dell PowerEdge M610 blade servers configured with dual Six Core Intel Xeon X5680 CPUs running at 3.33 GHz, with 96GB of installed memory.




Figure 11: Memory utilization for one ESX View Client Cluster Host during 1014 non-persistent desktop test

Figure 12 shows disk latencies for read and write operations, as measured from the ESX host, averaged for the seven volumes under test. As shown in the figure, disk latency stayed well within 20ms on average. Figure 13 shows the iSCSI read and write network data rates during the same test period. The network utilization was low, with ESX hosts using approximately 35% of the total available bandwidth for iSCSI SAN traffic.




Figure 12: Disk latency for one ESX View Client Cluster Host during 1014 non-persistent desktop test




Figure 13: iSCSI network utilization for one ESX View Client Cluster Host during 1014 non-persistent desktop test

5 Sizing Guidelines for EqualLogic SANs

Client desktop usage in enterprise environments follows typical patterns or phases. For example, at the beginning of the workday most employees login into their desktops within a relatively small time frame. During this time "login storms" can be expected. After the login storm, periods of high and low steady state application activities will occur. For example, high user activity on their workstations during morning/afternoon hours and low activity during break hours. These phases of activity cause different I/O patterns and loads in the VDI environment SAN.

5.1 Using Persistent vs. Non-Persistent View Desktops

Persistent desktops use persistent disk storage to store all changes – writes to the OS, temporary system and application data as well as user data. For non-persistent desktops, no changes to disk persist between client login sessions. Because of this, when a user logs into a non-persistent desktop, tasks such as user profile creation, reading and applying GPO data etc. generate significant amounts of I/O. The results in Section 4 demonstrate that when using non-persistent desktops, there can be significant spikes in SAN I/O load in large VDI environments during login storm periods.

When considering I/O performance requirements for a VDI SAN you must take into account the difference in I/O load patterns between persistent and non-persistent desktops. Our test results confirm that using persistent desktops minimizes the incremental I/O load effect on the VDI SAN caused by login storm periods. When using persistent desktops we were able to support a higher number of clients before we reached SAN I/O performance limits.

5.2 Calculating Capacity Requirements for Linked Clone Desktop Pools

There are many variables you need to consider when calculating storage requirements for desktop pools in VMware View VDI environments. In this section we provide an overview of considerations and inputs necessary for the calculation. We also provide an example calculation for a 1000 linked-clone desktop environment. You should consult the VMware View documentation before calculating requirements for your environment.

See the "*Storage Sizing for Linked-Clone Desktop Pools*" section in the *VMware View Online Library* for detailed guidance & formulas:

➤ <http://pubs.vmware.com/view45/ol/>

Considerations for calculating desktop pool size:

- You should allocate space for two replica images: one image for the linked-clone base image, and a second image for use by VMware View Composer for recompose activities. The replica image can be hosted in multiple datastores. Each datastore used for hosting a replica should allocate space for two replica images.
- The space allocated for each replica image should be the same as the size of the parent image. (Note: you can place the replica images in a different datastore than the one used for the desktop VM images.)
- For each desktop VM:

- o Allocate space equal to the memory size of the VM for swap usage
- o Allocate space for linked clone delta files. The size of the delta file allocation is determined by the amount of change that occurs between the desktop VM and the base image. Determine the maximum change expected (percent of base image size) and use that value for the delta file allocation. (Note: The setting for storage over-commit level⁶ will also affect the capacity calculation.)
- o Allocate at least 100MB for VM log files (if VM logging is enabled)
- o If using persistent desktops then you need to allocate additional space for persistent storage. (Note: you can place persistent storage for desktops in a different datastore than the one used for the desktop VM images.)

The following formula may be used for calculating the datastore capacity requirement:

$$\text{Capacity (GB)} = (\text{NumVM} * (\text{VMMem} + \text{Log} + \text{RepSize} * (\text{C}\%) + \text{PDisk}) + 2 * \text{RepSize} * \text{NumDS}) * \text{OF}\%$$

Table 5 describes the formula components and provides a sample calculation. We assume a parent VM replica image size of 15GB.

Variable	Description	Non-Persistent: 1000 desktops	Persistent: 1000 desktops
RepSize	Size of the linked clone master replica image. We allocated space for two replica images: one for the linked clone master replica, and a second for use by View Composer ⁷ .	2 * 15GB	
NumDS	Number of datastores (dedicated for images or shared with linked clones) hosting the replica image	2	
NumVM	Number of View Client desktop VM images	1000	
VMMem	Amount of memory provisioned to the desktop VMs	1GB	
C%	The maximum expected percent change in the linked clone images from the master replica image (ie. the growth rate of the linked clone delta disk). See note (1) below.	7.5%	4.0%
PDisk	For persistent desktops, this is the maximum expected size of the persistent storage used by the desktop VM. For non-	0	0.5GB

⁶ See *Storage Overcommit for Linked-Clone Desktops* section in the VMware View Online Library

⁷ View Composer was not actually used in our test configurations.

	persistent desktops PDisk=0.		
Log	Allocated for VM log files (if VM logging is enabled)	0 (no logging)	
OF%	Overhead factor.	15%	
Calculation: $(1000 * (1\text{GB} + 15\text{GB} * (0.075)) + 2 * 2 * (15\text{GB})) * 1.15$ $(1000 * (1\text{GB} + 15\text{GB} * (0.05) + 0.5\text{GB}) + 2 * 2 * (15\text{GB})) * 1.15$		2513GB	2484GB
(1) For non-persistent desktops, this will include OS image changes, user profile data and application data. The linked clone delta disk for non-persistent desktops will incur a higher change rate (%C) as compared to persistent desktops. This is because persistent desktops redirect user profile and user data to a separate disk.			

Table 5: Sample Capacity Sizing Calculation for Linked Clone Desktop Pool

5.3 Sizing Storage Capacity

We used two EqualLogic array models in our test configuration: the PS6000XVS (hybrid 100GB SSD + 450GB SAS drives) and the PS6000XV (600GB SAS drives). The PS6000XVS was configured using the default accelerated RAID 6. The PS6000XV was configured using RAID 10. Table 6 shows storage capacities for these two array models using the configured RAID policies.

Model	Disk Drive Size	Storage Capacity ^(c) (GB)
PS6000XVS^(a) 8 x SSD + 8 x SAS 15K Drives	100GB SSD 450GB SAS 15K	2850 (accelerated RAID 6)
PS6000XV^(b) 16 x SAS 15K Drives	300GB	2100 (RAID 10)
	450GB	3150 (RAID 10)
	600GB	4200 (RAID 10)
(a) Accelerated RAID 6 is the only RAID policy supported by the PS6000XVS. (b) The PS6000XV supports the following RAID policies: 5, 6, 10 and 50. For the purposes of this test RAID 10 was selected to maximize I/O performance for random I/O. (c) Useable capacity will be slightly less (~10%) than these values due to space required for peer storage array management functionality.		

Table 6: Storage capacity for PS6000XVS and PS6000XV

Note that the capacity calculations shown in Table 5 in the previous section are both within the storage capacity limits for the PS6000XVS shown in Table 6.

Note: Raw storage capacity data for all EqualLogic SAN array models is available in the Dell *EqualLogic Configuration Guide*, available here:

➤ <http://www.delltechcenter.com/page/EqualLogic+Configuration+Guide>

5.4 Performance Considerations

In general, the design considerations specific to VDI performance focus on supporting the workloads generated by a typical desktop session lifecycle (boot, login, work activities and logoff). Quantifying the average I/O workload generated by the user profile (task worker, knowledge worker etc.) is important for performance and sizing considerations. The storage platform used in a VDI deployment should be able to sustain the maximum IOPS these workloads create within acceptable latency limits.

Login storm is a worst case I/O workload pattern, especially for non-persistent desktops. In our testing we confirmed that login storm periods generated much higher storage system performance demands than after all logins were complete and users were in a steady state pattern of application usage. When using persistent desktops, the login storm effect can be negligible (except for the first time login of all desktops).

From our tests we measured I/O patterns for the task worker user profile during the simulated login storm period as follows:

- Login Storm (non-persistent desktops):
 - 50/50 read/write ratio
 - 36KB read and 10KB write block size
 - 8 - 13 IOPS per desktop
- Steady State (persistent and non-persistent desktops)
 - 25/75 read/write ratio
 - 18KB read and 10KB write block size
 - 3 - 4 IOPS per desktop (from application events in the task worker user profile)

5.4.1 When to use the EqualLogic PS6000XVS

The I/O performance measured at the maximum load (presented in Section 4.2) indicates that the PS6000XVS could host significantly more desktop VM images and still stay within I/O latency limits. You should consider using the EqualLogic PS6000XVS for VDI environments that will have high IOPS requirements (larger VDI environments running non-persistent desktops). You should also consider using the PS6000XVS in persistent desktop environments that will need to support a large number of users generating heavy application I/O. In general, you should consider using the PS6000XVS in any environment where the I/O performance boost provided by automatic movement of stored data between the SSD and SAS tiers will help to meet workload performance requirements.

5.4.2 When to use the EqualLogic PS6000XV

The test results in Section 4 show that the PS6000XV was able to host 390 non-persistent task worker profile desktops while staying within I/O latency limits. Based on this result we recommend the PS6000XV for small VDI environments running non-persistent desktops with moderate I/O performance requirements. The test results in Section 4 also show that the PS6000XV was able to host 780 persistent task worker profile desktops while staying within I/O latency limits. Thus the PS6000XV may also be a good fit for larger VDI environments where login storm effects are not a concern.

6 Best Practices

Each of the component layers in the system stack shown in Figure 14 requires careful design and configuration to ensure optimal system performance.

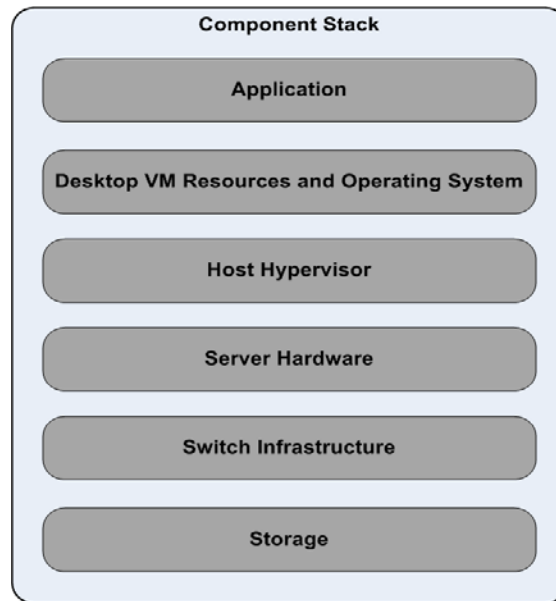


Figure 14: Component Stack

Within each layer of the stack you must consider each of the following design goals:

- Availability** Ensuring high availability is a critical design guideline for each component in the solution. Redundant sub-components and connection paths must be maintained to avoid single points of failure.
- Scalability** Scalability is also critical design goal. Your solution architecture should be able to scale to meet performance and capacity needs.
- Performance** Performance of each component in the stack must be balanced such that all performance requirements for the VDI system are met (examples: I/O latency, user response times).

We present a top-down discussion summarizing design and configuration best practices related to each of the layers in the component stack in the following sections.

6.1 Application Layer

There are many considerations at the application layer that you need to focus on.

User Data Considerations

We recommend that you do not store user data in the virtual desktop's local drive. Separate file servers should be used for storing user data. Specifically, group policies should be used to configure a user's virtual desktop profile to leverage roaming user profiles and folder redirection. Roaming profiles allow

users to have their desktop and application settings follow them between different virtual desktops. Folders like *Desktop*, *Documents*, *Pictures*, *Music*, *Videos*, *Favorites*, *Contacts*, *Downloads* and *Links* should be redirected to a user's home directory in a network file share. By following these practices, login speeds will improve and a reset of user's profile will not cause document data loss.

Anti-Virus (AV) Optimizations

A VMware View non-persistent desktop can resume its original state with refresh on logoff or reboot, so On Access Scanning (OAS), which is typically enabled for inbound (write operations) and outbound (read operations) should be sufficient. For persistent desktops, On Demand Scanning (ODS) can also be included, but ODS events should be staggered or randomized to avoid an anti-virus scan I/O storm. Additional AV best practices include:

- Install a virus scanner. You should not need to install any extra packages provided by the virus scanner vendor other than the core virus scanning features.
- Include the Network Access Control (NAC) management agent if necessary.
- Use random or staggered scan scheduling.
- Configure scanning exclusion lists to exclude files like *pagefile.sys*, **.pst*, *spool*, *database* and encrypted files.

Note: See the following VMware View Technical Note for more information: ***Anti-Virus Deployment for VMware View***

➤ <http://www.vmware.com/files/pdf/VMware-View-AntiVirusDeployment-WP-en.pdf>

Application Delivery Considerations

For VDI environments, you should categorize applications by usage, and leverage application virtualization as appropriate. For example, the set of applications needed for all users (Microsoft Office, anti-virus, etc.) should be included in the base image. Second and third tier applications could be delivered via application virtualization.

Boot Storm Considerations

Boot storms can have a large performance impact in a VDI environment as simultaneous booting of a large number of VMs create a huge IOPS load on the underlying storage systems. To avoid this, organizations deploying VDI solutions should pre-boot the environment by starting the required number of VMs well ahead of the time when the first set of users arrives at the beginning of a workday. With technologies such as hardware-assisted locking in VAAI and automatic workload tiering available in PS6000XVS, organizations can also pre-boot a certain percentage of the VMs ahead of the time and then turn on the rest of the VMs on an on-demand basis without affecting performance. This strategy will lead to lower power consumption and hence improved TCO.

Login/Logoff Storm Considerations

Login and Logoff storms are unavoidable as they depend solely on user arrival and departure times. Our test results show that it is very important to consider the performance impact of login/logoff storms, especially when using non-persistent desktops. The VDI infrastructure (and the underlying

storage platform) should be designed to handle the IO requirements during these periods. Organizations should study the arrival rate (number of user connections per minute) and the corresponding IOPS requirements when designing the VDI solution. Our test results showed that using the EqualLogic PS6000XVS (hybrid SSD/SAS storage array with automatic hot-data migration between device tiers) enabled us to support much higher peak I/O workloads while still supporting VM storage density requirements.

6.2 Virtual Machine and Guest OS Configuration

There are many guidelines that need to be considered when optimizing how you configure and host "Golden Master" VMware View images on ESX. Please refer to the following VMware publications for more details:

- *VMware View Optimization Guide for Windows 7:*
<http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf>
- *Storage Considerations for VMware View:*
http://www.vmware.com/files/pdf/view_storage_considerations.pdf
- *Best Practices for Creating VMware View Golden Master Images:*
http://www.ituda.com/Docs/VMware/Best_Practices_For_Creating_VMware_View_Golden_Master_Images.pdf

The guidelines break out into the two levels:

Optimizing the Windows OS

There are many guest desktop OS optimizations that can enhance performance, resource utilization, reliability and manageability of VDI environments. A lot, if not all, of the default settings and behaviors in any desktop OS are set for traditional standalone PC operation. Some of these default settings and behaviors can be quite detrimental to performance and user experience in a VDI environment. Proper care should be taken to set VDI-friendly values for guest OS system parameters and services. In our case, Windows® 7 optimizations for VMware View were performed following VMware's recommendations. The optimizations we performed in our test environment:

- Disabled the following: Offline Files, AutoUpdate, Background Defragmentation, Last Access Timestamp, Hibernate, Bug Check Memory Dump, Move to Recycle Bin
- Removed unnecessary Windows 7 components, such as: Tablet PC Components, Windows Fax and Scan, Windows Gadget Platform, XPS Services and Viewer, Remote Differential Compression
- Disabled unnecessary services, such as: Themes, Windows Defender, Windows Search, Windows Update, Superfetch, Diagnostic Policy Service, Background Intelligent Transfer Service, Base Filtering Engine, WLAN Autoconfig
- Disabled schedules for unnecessary tasks, such as: ScheduledDefrag, RegIdleBackup, SystemRestore, WinSAT, DiskDiagnostic\DataCollector

Optimizing the VM within ESX

The following steps are best practices for supporting a VDI deployment:

- Disable virtual machine logging
- Disable serial and parallel ports for guest desktop OS

- Align guest desktop OS vmdk image files to 64K block boundaries (Windows 7 performs this automatically)

6.3 ESX Host Configuration

The following steps are recommended for configuring the ESX host:

- Install the EqualLogic Multipathing Extension Module (MEM) for vSphere 4.1. This is supported with both the ESX software iSCSI initiator or any supported iSCSI hardware initiator on the host.
- You should configure separate virtual switches for VM network traffic and iSCSI storage traffic on the ESX hosts
- Jumbo frames should be enabled on the iSCSI vSwitch and the vmkernel ports. At least two server NICs dedicated for iSCSI traffic need to be configured as uplink NICs to the iSCSI vSwitch

Note: Detailed steps for configuring MEM and the ESX host initiator with EqualLogic storage are provided in the following document:

- *Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere 4.1 and PS Series SANs:*

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=9823>

6.4 Host Server Resources

It is important that you properly configure your servers to provide the optimal number of NICs for your VDI I/O workload. You could consider the following:

- We recommend that you configure your servers to provide dedicated NICs for VDI client/server I/O and separate dedicated NIC/iSCSI hardware initiators for iSCSI SAN I/O.
- To provide for redundancy, you should provision at least two NICs per server for each path.
- Additional NICs/HBAs should be provisioned for iSCSI SAN I/O as needed, based on throughput and IOPS requirements.
- You should also configure servers with appropriately sized CPU and memory resources for the VDI workload.
- Dell and VMware provide additional resources to assist in planning and server sizing in support VDI workloads. See: www.dell.com/flexiblecomputing for more details.

6.5 SAN Design

We recommend the following SAN design best practices:

- You should design redundant SAN component architectures. This includes the NICs on the servers and switches for the storage network (including server blade chassis switches and external switches).

- You should also make sure that the server NIC ports and storage NIC ports are connected in a way such that any single component failure in the SAN will not disable access to any storage array volumes.
- Flow control should be enabled on both the server NICs and the switch ports connecting to server and storage ports.
- We recommend you enable jumbo frames on the server NIC ports and the switch ports. Within ESX you also need to set the jumbo frame MTU value on the virtual switches and the virtual NICs used by the guest VM OS.
- You should disable spanning tree on switch ports connecting to end devices like server ports and storage ports. Portfast should be enabled for these ports.

Note: General recommendations for EqualLogic PS Series array network configuration and performance is provided in the following document.

- *PS Series Array Network Performance Guidelines:*

<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>

6.6 Storage

Storage recommendations for capacity and performance are presented in Section 5. Additional recommendations:

- You should deploy multiple, separate volumes (maximum: 500GB) for hosting VMware View linked clone desktops. Doing so will support isolation for manageability and performance.
- When using the PS6000XVS, we recommending using a separate volume (from the ones used for hosting linked clone desktops) for storing the linked clone master replica images. You should not host replica images on volumes used for linked clones. Keep the number of base images to a minimum (ideally one per XVS array). For example, in one of our tests we had one 32GB volume where one base image was stored, and four 500GB volumes that were used for linked clone delta space. Having a large number of clones linked to the same base image will make the base image an I/O "hot spot." By isolating the base image to a separate volume, we ensure quick movement of that volume data to the SSD tier in the XVS array.
- When using the PS6000XVS, first create the volumes where the linked clone master replica images will be stored. Once that is done, create the volumes for storing linked clone differential data.
- Hosting replica images on separate volumes is one of the new features View 4.5 supports⁸. Earlier versions of View required creation of a replica image on each datastore used for

➤ ⁸ See *Storing View Composer Replicas and Linked Clones on Separate Datastores* in the VMware View Online Library: <http://pubs.vmware.com/view45/ol/>

linked clones. With View 4.5, only one replica volume datastore needs to be specified per View pool and ESX cluster.

- When using persistent desktops, if user data needs to be backed up then you should host persistent disks assigned for each VM on a separate VMFS datastore. This datastore can be backed up independently to protect user data with its own SLA requirements.

Appendix A Test System Component Details

The tests were conducted using the following component firmware levels.

Component	Firmware Version
EqualLogic Arrays	v5.0.2
Network Switches <ul style="list-style-type: none">- PowerConnect 6248- PowerConnect M6220- PowerConnect M6248	3.2.07 3.1.5.2 3.1.5.2
Broadcom BCM5709 Quadport 1GigE IO Mezzanine	5.2.7, A10
M610 Server BIOS	2.2.3
M1000e Chassis Firmware <ul style="list-style-type: none">- Chassis Management Controller (CMC)- Integrated Dell Remote Access Controller (iDRAC6)- iKVM	2.30 3.02 01.00.01.01



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