



Dell EqualLogic Best Practices Series

Sizing and Best Practices for Deploying Citrix XenDesktop on VMware vSphere with Dell EqualLogic Storage

A Dell Technical Whitepaper

Storage Infrastructure and Solutions Engineering

Dell Product Group
January 2012

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.

© 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. Adobe®, Acrobat®, Flash®, Reader®, and Shockwave® are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Broadcom® is a registered trademark of Broadcom Corporation. Citrix® and XenDesktop® are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. Intel® is a registered trademark of Intel Corporation in the U.S. and other countries. Liquidware Labs is a trademark of Liquidware Labs, Inc. Login Consultants is a trademark or registered trademark of Login Consultants in the Netherlands and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter®, and vSphere® are registered trademarks or trademarks of VMWare, Inc. in the United States or other countries. Microsoft®, Windows®, Windows Server®, Internet Explorer®, 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
2	Virtual Desktop Infrastructures	2
2.1	Addressing VDI storage challenges with EqualLogic SANs	2
3	Citrix XenDesktop solution infrastructure.....	4
3.1	Test infrastructure: Component design details	4
3.2	XenDesktop VDI delivery using provisioning servers	7
3.3	Provisioning server components	8
3.4	EqualLogic storage array configuration	9
3.5	Citrix XenDesktop networking	10
3.5.1	Network architecture.....	10
3.5.2	Network characteristics	12
3.5.3	vSwitch and VLAN configuration.....	12
4	Citrix XenDesktop test methodology	14
4.1	Test objectives	14
4.2	Test approach	14
4.3	Test tools	14
4.3.1	Load Generation.....	15
4.3.2	Liquidware Labs Stratusphere UX.....	15
4.3.3	Monitoring tools	15
4.4	Test criteria	15
4.4.1	Storage capacity and I/O latency	15
4.4.2	System resource utilization on the hypervisor infrastructure	16
4.4.3	Virtual desktop user experience	16
5	Citrix XenDesktop test results and analysis	17
5.1	Test Scenarios	17
5.2	vDisk I/O Characteristics.....	18
5.3	CIFS share I/O Characteristics	19
5.4	Task worker – 100% Pre-boot and login storm	20
5.4.1	Storage layer: EqualLogic array performance	20
5.4.2	Login storm I/O	23
5.4.3	Hypervisor layer: ESXi host performance.....	24

5.4.4	Citrix infrastructure layer: PVS and DDC performance	24
5.4.5	User Experience: application performance	25
5.4.6	Write cache utilization on virtual desktops.....	25
5.5	Task worker – 20% Pre-boot.....	25
5.5.1	Storage layer: EqualLogic array performance.....	25
5.6	Virtual desktop scaling with two arrays.....	26
5.6.1	Storage layer: EqualLogic array performance.....	26
5.6.2	Login Storm I/O	27
5.7	Boot storm.....	28
5.7.1	Storage layer: EqualLogic array performance.....	28
5.8	Results summary	30
6	Sizing guidelines for EqualLogic SANs	32
6.1	Storage capacity considerations	32
6.1.1	vDisk capacity requirements	32
6.1.2	Write cache capacity requirements	33
6.2	Storage performance considerations.....	33
6.3	Sizing Calculations	34
7	Best practices.....	36
7.1	Desktop profiles and I/O storms	36
7.1.1	Implement roaming profiles and folder redirection	36
7.1.2	Use XenDesktop Power management feature.....	36
7.1.3	Boot and login storm considerations	36
7.2	ESXi host configuration.....	36
7.3	Network configuration	37
7.4	Citrix infrastructure	37
7.4.1	Read-Only vDisk.....	37
7.4.2	Write cache recommendations	38
7.4.3	Optimizing the Windows 7 master image.....	38
7.4.4	Allocate more RAM on PVS servers.....	38
7.5	Storage	38
7.5.1	Volumes and datastores.....	38
7.5.2	Storage array characteristics	38
8	Conclusions.....	40
Appendix A	Citrix XenDesktop solution configurations.....	41

Appendix B	Folder redirection and roaming profiles.....	44
Appendix C	Network Design and VLAN configuration.....	45
	VDI Network.....	45
	Infrastructure network.....	45
Appendix D	Liquidware Labs Stratusphere UX.....	47

Acknowledgements

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

The team that created this whitepaper:

Chidambara Shashikiran, Chhandomay Mandal, and Margaret Boeneke

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

Brian Drew, David Glynn, Ananda Sankaran, 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 SIFeedback@Dell.com.



SIFeedback@Dell.com

1 Introduction

Virtual Desktop Infrastructure (VDI) products such as Citrix® XenDesktop® 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 conducted by Dell™ Labs and provide storage sizing guidance and best practices based on those results for a Citrix XenDesktop based VDI solution.

Our test infrastructure included:

- Citrix XenDesktop 5.0
- VMware® vSphere™ 4.1 Hypervisor™
- Dell PowerEdge™ M610 blade servers
- Dell PowerConnect™ switches
- Dell EqualLogic™ SANs

The primary objectives of our testing were:

- Determine how many virtual desktops can be deployed using a single Dell EqualLogic PS6010XVS virtualized iSCSI SAN array with acceptable user experience indicators for task workers.
- Determine the performance impact on the storage array during peak I/O activity such as boot and login storms.
- Determine how to scale the number of virtual desktops by adding additional storage array members.

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

2 Virtual Desktop Infrastructures

VDI helps IT organizations simplify administration and reduce costs while enhancing security and regulatory compliance, increasing IT flexibility and business agility, and strengthening business continuity and disaster recovery. However, a VDI deployment must be carefully designed to ensure that it delivers the performance and scalability needed to support an enterprise-wide client community. When moving from a distributed standalone desktop model, all components—storage, processors, memory, and networking—are moved to a centralized data center for all the users, making appropriate design and sizing critical to success.

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 expensive shared storage used in VDI deployments. Performance demands include 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. Typically in a VDI environment, the IOPS generated by virtual desktops are random and can have significant impact on the storage. 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 require booting of all the virtual desktops at the same time. This kind of I/O storm (boot storm) creates significantly higher IOPS on the underlying storage array.

To be successful, storage designs for a VDI deployment must take these demands into account. In particular, the storage platform should have not only the high performance and scalability required to handle large I/O and system resources utilization spikes, but also the ability to cost-effectively handle large 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.

2.1 Addressing VDI storage challenges with EqualLogic SANs

Dell EqualLogic PS Series iSCSI SANs are well suited for supporting VDI deployments because they offer scalable, high-performance, virtualized storage designed for reliability, manageability, and efficiency. EqualLogic SANs come with a range of efficiency features to enhance utilization and help reduce costs. For example, automated workload tiering and load balancing help optimize storage performance. Storage virtualization coupled with thin provisioning can provide increased capacity with reduced physical storage requirements. EqualLogic Multipath Extension Module (MEM) for VMware vSphere provides advanced multi-path I/O functionalities such as automatic connection load balancing for highly available access to the EqualLogic SAN in a VDI environment. EqualLogic also offers additional features such as thin clone technology, comprehensive integration with VMware vStorage Application Programming Interfaces (APIs) for Array Integration (VAI), and VMware virtualization awareness that enable you to create more efficient deployments.

In a Citrix XenDesktop VDI solution, many virtual desktops can share the same base desktop image. This scenario causes the data in a relatively small portion of the storage system to become extremely hot (very high read/write I/O rates). The EqualLogic hybrid arrays – EqualLogic PS6010XVS and PS6100XS – are designed to support fast access to high demand hot data segments. These hybrid SANs have high performance SSD drives as well as high capacity SAS hard disk drives. When data

becomes hot, it is automatically moved from the SAS tier to the SSD tier. This automatic tiering function makes the hybrid EqualLogic SAN a very cost efficient option for VDI environments where the peak load from hundreds of desktops during login storm is concentrated on the relatively small capacity base image volume (the hot data blocks).

Finally, the VDI workloads are also very write intensive, and the system dedicated write-cache in EqualLogic hybrid arrays that leverages SSD drives for this purpose provides a significant performance improvement for VDI environments.

3 Citrix XenDesktop solution infrastructure

Citrix XenDesktop is a comprehensive desktop virtualization solution that includes the capabilities required to deliver desktops, applications, and data securely to every user in an enterprise. The solution incorporates a variety of components to accomplish these various capabilities.

3.1 Test infrastructure: Component design details

The core Citrix XenDesktop infrastructure components used in our test configuration are shown in Figure 1. Storage connectivity details are abstracted in the figure.




Figure 1 Citrix XenDesktop core infrastructure components

Overview of the core infrastructure components (see Figure 1):

- We used two provisioning servers (PVS) and two Desktop Delivery Controllers (DDC) for streaming the master disk image to virtual desktops. Two servers were used for high availability and better performance.
- PVS and DDC servers were virtualized and hosted on two separate ESXi servers for high availability.

- vDisk was configured using standard mode and it was made read-only so that the two provisioning servers can access it simultaneously.
- Any modifications made to the streamed image by the virtual desktop were stored in the write cache. Each virtual desktop consumed 3 GB of storage capacity – 2 GB of write cache (that includes 1 GB page file) and 1 GB of RAM. The “Cache on Device Hard Drive” option provided by Citrix Provisioning Services was used to mount the write cache.
- A CIFS share through a Microsoft® Windows® file server virtual machine was created using a 2 TB volume hosted on a Dell EqualLogic PS6510E array. All the user folders and user shares that contain user or application data from virtual desktops were stored on the CIFS share.
- A total of 32 Windows Server 2008 R2 virtual machines were hosted on three ESXi servers for Login VSI load generation. Each virtual machine was used as a Login VSI launcher and each launcher invoked 40 sessions.

The core infrastructure components are listed below.

Table 1 Citrix XenDesktop core infrastructure components

Citrix XenDesktop infrastructure	
Component	Description
Citrix Web Interface	Provides the user interface to the XenDesktop VDI environment.
Citrix License servers	Manages Citrix licences for XenDesktop environment.
Active Directory	Common namespace and secured access to all the servers and VMs in the environment.
DHCP	Used to provide IP address for virtual desktops.
DNS	IP address and host name resolution for the entire XenDesktop environment.
Microsoft® SQL Server®	Microsoft SQL Server database was used to create databases for Citrix PVS, DDC, and VMware® vCenter® server.
Citrix Provisioning Services Servers (PVS)	Creates and provisions virtual desktops from a single desktop image on demand. Provides great flexibility, simplifies management of desktop images, and ensures that each user gets a pristine virtual desktop every time they log on.
Desktop Delivery Controllers (DDC)	Manages the virtual desktop requirements of the user. Authenticates users by interacting with Active Directory. Manages the connectivity between users and their virtual desktops.
Virtual Desktop Agent (VDA)	Installed on virtual desktops. Manages the direct connection between virtual desktops and client devices.
Independent Computing Architecture (ICA)	A proprietary protocol designed by Citrix. Used for communication between client devices and virtual desktops.
Virtual desktop storage configuration	
Component	Description
PVS vDisk	The disk image files located on provisioning servers to boot the virtual desktops. These are configured on shared EqualLogic storage hosted by PVS servers.
PVS write cache	Any writes made to the desktop OS are redirected to a temporary area called the write cache. The write cache is configured as a virtual desktop hard drive, which is physically located on the shared EqualLogic storage via ESXi datastores.
User Data (CIFS Share)	The user data for each virtual desktop is mounted from a NAS file server using shared storage.
Hypervisor	

Component	Description
VMware ESXi Enterprise Plus	The smaller footprint version of ESXi that does not include the ESXi service console.
VMware vCenter	Centralized management interface for VMware vSphere environment.
Login VSI Workload Generator	
Component	Description
Login VSI	A third-party benchmarking tool from Login Consultants that is used to simulate a real-world XenDesktop VDI workload.
Login VSI Launchers	A Login VSI launcher is a Windows system that launches desktop sessions on target virtual desktop machines.
Task worker (Login VSI Light workload)	A Task Worker is a user who performs repetitive functions with a small set of applications. Examples of a Task Worker include call center employees, data entry workers, and administrative task performers.

Refer to Appendix A for more detailed information about the hardware and software components used for evaluating this solution.

3.2 XenDesktop VDI delivery using provisioning servers

Citrix XenDesktop 5 supports two different features to host virtual desktops, Machine Creation Services and Provisioning Services. The hosting methodology should be selected based on the overall goals of the organization, resource availability, and management activities.

For more detailed information on these two approaches, please refer to:

<http://support.citrix.com/servlet/KbServlet/download/26449-102-650176/XD%20-%20Planning%20Guide%20-%20Desktop%20Image%20Delivery.pdf>

In our tests, we used provisioning services for streaming the OS image to all virtual desktops. The key factors which influenced this decision were:

- Provisioning Services creates and de-provisions virtual desktops from a single desktop image on demand, optimizes storage utilization, and provides a new virtual desktop to each user each time they log on.
- Portions of the disk image are cached within the RAM on the provisioning servers (PVS). If sufficient RAM is configured on the provisioning servers, most of the read I/O is handled by memory rather than from the SAN. Memory access is faster than the disk. This significantly reduces storage requirements which results in reduced costs and also better user experience.
- Adding new desktops has a minimal impact on storage as the target desktops do not require any storage infrastructure operations.
- Management is very simple because only a single disk image needs to be maintained. If any software updates or patches need to be installed, only a single image needs to be updated and the same updated image is streamed to all virtual desktops.

Figure 2 is a block diagram that shows a high level overview of the Citrix XenDesktop VDI delivery model using provisioning services. This process and figure show the VDI delivery model:

1. The end user logs in using the Web interface (A).
2. The Web interface (A) passes the credentials of the user to the Desktop delivery controller (E).

3. The Desktop delivery controller authenticates the user using Active Directory (D) and starts the VM (C) on the Hypervisor.
4. The VM (C) contacts the DHCP server (D) to find an IP address and the location of the boot image.
5. The VM (C) starts using the boot image received from the Provisioning server (G) through the network.
6. The Desktop delivery controller (E) assigns the user a VM after verifying the license through the license server (B) and connects the user to the VM through ICA.
7. The end user can login to the virtual desktop using the Client device (F).




Figure 2 Citrix XenDesktop VDI delivery model

3.3 Provisioning server components

The key components of the provisioning server are:

- **PVS vDisk**: vDisk exists as disk image files on a PVS stored on a shared storage device. The vDisk images are configured to be in Private, Standard, or Different Disk mode. Private mode gives exclusive read-write access to a single desktop while vDisk in Standard or Different Disk mode is shared with read-only permission among multiple desktops.
- **PVS write cache**: Any writes made to the desktop operating system are redirected to a temporary area called the write cache. The write cache can exist as a temporary file on a PVS, in the memory of the virtual desktop, or on the local hard drive of the virtual desktop.




Figure 3 Virtual desktop disk composition in Citrix XenDesktop VDI environment

As shown in Figure 3, the shared read-only vDisk and temporary write cache are served from EqualLogic PS6010XVS array(s). The user data is kept separate by using a separate CIFS share.

3.4 EqualLogic storage array configuration

We used a single EqualLogic PS6010XVS array to store the vDisk of the PVS and also the write cache for all virtual desktops. Only 2.1 TB on the array was used to configure virtual desktops which resulted in more than 15% head room on the storage array.

The volume layout used for configuring vDisk and write cache on a single PS6010XVS array is shown in Table 2.

Table 2 EqualLogic volume layout on a single PS6010XVS array

Volume name	Reported size	Capacity
vDisk-vol	150 GB	Total member capacity: 2.48 TB Capacity used by volumes: 2.1 TB (84.8%) Free member space: 384.87 GB (15.2%)
VDIVOL1	500 GB	
VDIVOL2	500 GB	
VDIVOL3	500 GB	
VDIVOL4	500 GB	

- The volume named “vDisk-vol” was used to host the master vDisk image. This volume was accessed by two provisioning servers using Windows Guest iSCSI initiators.
- Four volumes, each with 500 GB of capacity, were created for allocating write cache for virtual desktops.
- Each of the volumes was presented to the ESXi server as an individual datastore.
- In the two array configuration, four more volumes, each with 500GB capacity, were used to allocate write cache for the virtual desktops.

In addition to the PS6010XVS array, we used a PS6510E array to provide a CIFS share for all virtual desktops through a Windows OS-based file server VM. This CIFS share was used to redirect the user data and roaming profiles of virtual desktops to a centralized location. A single volume with 2 TB of capacity was created on the PS6510E and this volume hosted the user profiles and user data folders for all virtual desktops via the file server.

In one test, we also evaluated how the number of virtual desktops can be scaled by adding an additional EqualLogic array. The volume layout remained the same in the two array configuration. Only around 4 TB of capacity on the arrays was used with more than 18% head room on the storage pool.

3.5 Citrix XenDesktop networking

We used 16 Dell PowerEdge M610 blade servers within a Dell PowerEdge M1000e Modular Blade Enclosure as the platform for hosting virtual desktops. All 16 M610 blade servers were completely dedicated to hosting virtual desktops. Citrix infrastructure components such as PVS, DDC, and additional components such as Active Directory and vCenter were hosted on separate rack-based servers.

3.5.1 Network architecture

Figure 4 illustrates the network connectivity between the blade server chassis and the storage array. The connectivity from a single M610 blade server is shown. The connectivity from the ESX servers hosting PVS, DDC, user data CIFS shares, and Login VSI workload launchers are not shown in this diagram.




Figure 4 Server LAN and iSCSI SAN connectivity

- Each PowerEdge M610 server has one on-board Broadcom NetXtreme 5709 dual port NIC mezzanine card and this was assigned to Fabric A. The M610 servers had two additional dual-port Broadcom NetXtreme 57711 10 GbE NIC mezzanine cards and they were assigned to Fabric B and Fabric C.
- Port B1 (Fabric B) and Port C1 (Fabric C) on each server was used for iSCSI SAN connectivity. Similarly, ports B2 and C2 were dedicated for VDI network. This type of network connectivity ensures that the configuration is resilient even if one of the mezzanine cards fails.
- Dell PowerConnect M6220 switches were installed in Fabric A1 and A2. Dell PowerEdge M8024 switches were installed on Fabric B1 and Fabric C1. Dell PowerEdge M6220 switches were installed on Fabric B2 and Fabric C2.

- Two PowerConnect 8024F switches were used as external SAN switches. We used four 10 GbE SFP+ uplink modules from M8024 switches on Fabric B1 and Fabric C1 to the two PC8024F switches which are connected via a LAG of 4x10 GbE links.
- The other two ports on Fabric B and Fabric C (Ports B2 and C2) were dedicated for VDI traffic. The two M6220 switches on Fabric B2 and C2 were uplinked to two external PowerConnect 6248 switches using two 10 GbE SFP+ uplink modules. The two external PowerConnect switches were used exclusively for VDI connectivity and they were stacked together for high availability and better performance. This network connectivity is not shown in the above diagram. Please refer to Appendix C for detailed network connectivity diagrams for each of the different networks.

3.5.2 Network characteristics

Provisioning servers stream the master disk image to all virtual desktops over the network. This means that designing the network is very critical for achieving optimal VDI performance. While designing a XenDesktop VDI solution, the most important network characteristics to be considered are summarized in the table below.

Table 3 Network characteristics

Type of network	Description
VDI network	Represents the network traffic which happens when provisioning servers stream the OS disk image to all virtual desktops over the network. The virtual desktops contact the Active Directory server and DHCP server using the same network.
iSCSI SAN	Dedicated iSCSI SAN network through which all virtual desktop and other infrastructure components access the EqualLogic storage.
Infrastructure network	Represents the network traffic due to the communication between different infrastructure components such as the Web interface server, Active Directory, Microsoft SQL Server, and VMware vCenter. The management network of all ESXi servers and storage arrays was also configured on this network.

3.5.3 vSwitch and VLAN configuration

The EqualLogic arrays were connected to the servers using a dedicated iSCSI SAN. A dedicated vSwitch was used for each type of network.

The high-level diagram showing the different vSwitches and VLANs is shown in Figure 5.




Figure 5 vSwitches and VLANs

The infrastructure network was further divided into three VLANs to segregate network traffic into different classes. The three VLANs used in the test configuration were:

- Management VLAN
- Infrastructure VLAN
- vMotion@ VLAN

For more detailed information on how each of these different networks was set up and how the VLANs were configured, please refer to Appendix C .

4 Citrix XenDesktop test methodology

4.1 Test objectives

The primary objectives of our testing were:

- Develop best practices and sizing guidelines for a Citrix XenDesktop Provisioning Services based VDI solution deployed on Dell EqualLogic PS6010XVS series storage, Dell PowerEdge blade servers, and Dell PowerConnect switches with VMware vSphere 4.1 as the server virtualization platform.
- Determine how many virtual desktops can be deployed using a single PS6010XVS array with acceptable user experience indicators for task worker profiles.
- Determine the performance impact on the storage array during peak I/O activity such as boot storms and login storms.
- Determine how to scale the number of virtual desktops by adding additional storage array members.

The testing focused on the data center aspects of VDI and detailed statistics and metrics were captured and analyzed for use cases including desktop boot-up, user logon, user workload execution (steady state), and user logoff.

Note that the applications were delivered directly through desktop images instead of using Citrix XenApp platform for application virtualization. The objective for this project was storage characterization for the Citrix XenDesktop based VDI solution and the application delivery mechanism does not impact this objective.

4.2 Test approach

The key VDI use cases which were validated are listed below:

- Task worker workload characterization
- Boot storm
- Login storm
- Scaling virtual desktops by adding an EqualLogic array

Before each test, the virtual desktops were restarted so that the results were not influenced by the previous tests. When the virtual desktops were pre-booted, the tests were executed after waiting for at least 20 minutes after the virtual desktops registered in the DDC. Only non-persistent desktops were used where any changes made to the desktop image are lost as soon as the user logs off. However, changes to user profiles were preserved by storing them on the CIFS file share.

4.3 Test tools

We used several test and monitoring tools to perform tests and to capture the key performance metrics.

4.3.1 Load generation

Login VSI is a VDI benchmarking tool which can be used to determine the maximum number of desktops that can be run on a physical or virtual server. Login VSI simulates a realistic VDI workload using the Autolt script within each desktop session to automate the execution of generic applications.

The tool's "Light" workload was used to simulate the task worker workload. The characteristics of "Light" workload are as follows:

- Applications used include Microsoft Internet Explorer®, Adobe® Acrobat® Reader®, Adobe Flash® player, Adobe Shockwave® player, Freemind, Kid-Key-Lock, and Bullzip PDF printer
- Only two applications are open simultaneously.
- Workload is executed in 12 minute loops and the total idle time between each loop is about 1 minute and 45 seconds.

Find more details about the Login VSI tool at <http://www.loginvsi.com/en/login-vsi/product-overview>. More details regarding the different types of workloads supported by Login VSI and their characteristics are described in <http://www.loginvsi.com/en/admin-guide/workloads>.

4.3.2 Liquidware Labs Stratusphere UX

The complex setup and different configuration layers involved in VDI makes it very challenging to measure the user experience. VDI performance and functionality depends on various factors such as resource requirements like CPU, memory, network, and storage. The other key consideration is the user experience with respect to time spent while trying to login and response times of applications launched within the desktop.

Stratusphere UX by Liquidware Labs is designed for desktop administrators and support personnel to measure and analyze end user experience and application response times. We used Stratusphere UX for each test to gather metrics related to user experience and desktop performance.

4.3.3 Monitoring tools

We used Dell EqualLogic SAN HeadQuarters (SAN HQ) for monitoring the storage array performance. The system resource utilization and performance metrics at the Citrix infrastructure layer and ESXi hypervisor layer were monitored using VMware vCenter.

Detailed performance metrics were captured from the hypervisor, virtual desktop, storage array, Citrix Infrastructure layer, and load generator.

4.4 Test criteria

The primary focus of our tests was to determine the maximum number of desktops which can be deployed using a single PS6010XVS array and then study how the number of VMs can be scaled by adding additional storage arrays. We used specific test criteria as described in the following sections while performing the tests.

4.4.1 Storage capacity and I/O latency

We used the following performance criteria at the storage layer.

- Maintain less than 20 ms disk latency
- Maintain at least 15% spare capacity on the storage array

The typical industry standard latency limit for storage disk I/O is around 20 ms. Maintaining this limit will ensure good user application response times when there are no other bottlenecks at the infrastructure layer

4.4.2 System resource utilization on the hypervisor infrastructure

The primary focus of our testing was storage, and we ensured that no other component in the VDI stack became a bottleneck while conducting these storage characterization tests. We used the following metrics at the hypervisor infrastructure layer to ensure solution consistency:

- CPU utilization on any ESXi servers not reaching 95% at any point in time
- Minimal or no memory ballooning on ESXi servers
- Total network bandwidth utilization not to exceed 90%
- Storage network retransmission less than 0.5%

4.4.3 Virtual desktop user experience

VDI configurations involve many components at different layers – application, hypervisor, server, storage and network. As a result, multiple metrics need to be monitored at different layers to ensure that the environment is healthy and performing appropriately for users. Liquidware Labs Stratusphere UX was used to assess the performance of each desktop (application level).

We used the following user experience criteria:

- At least 90% of the virtual desktops are in the 'Good' performance category.
- No virtual desktop should be in the 'Poor' performance category.

The various criteria used to determine user experience and also other detailed information about how this metric is calculated using the Stratusphere tool is available in the following document:
<http://www.liquidwarelabs.com/docs/VDI%20Fit%20and%20VDI%20UX%20White%20Paper.pdf>

5 Citrix XenDesktop test results and analysis

This section describes the different XenDesktop VDI characterization tests conducted and also the key findings from each test. The Task worker user type represents the majority of VDI users in the industry today and we focused our testing on this workload profile.

For all single array tests, we used eight M610 blade servers to host 630 virtual desktops. We uniformly distributed around 80 desktops on each blade M610 server so that server performance limits were not reached at any point. For two array tests, 16 M610 blade servers were used to host 1270 desktops.

5.1 Test scenarios

We tested the following scenarios:

- **Task worker – 100% Pre-Booted**

In this test we pre-booted 630 virtual desktops and allowed the desktops to be in idle state for more than 20 minutes before starting the Login VSI workload to simulate task worker profile.

- **Task worker – 20% Pre-Booted**

There are many instances in the industry where pre-booting all virtual desktops might not be possible due to infrastructure resource constraints. So in these scenarios, a minimum number of desktops are pre-booted and the remaining desktops will boot as the user logs in. In this test, only 126 virtual desktops (20% of the 630) were pre-booted and the remaining desktops were booted on demand while starting the desktop sessions.

- **Virtual desktop scaling with two arrays**

In this test, we wanted to validate the scalability of the VDI solution by adding an additional storage array. To support 630 virtual desktops, we used eight blade M610 servers and one PS6010XVS array as explained in the previous section. We added one more PS6010XVS array to the same pool and also eight more blade M610 servers to the ESXi cluster hosting virtual desktops and tested 1270 virtual desktops.

- **Boot storm**

The boot storm represents the worst-case scenario where all the virtual desktops are started at the same time and they all contend for the system resources simultaneously. We wanted to evaluate if the storage array hosting the desktops is capable of handling this kind of huge spike in I/O within acceptable latency limits.

- **Login storm**

Login storms are caused when a large number of users login to the virtual desktops almost at the same time. Login storm scenarios are much more critical compared to boot storms in a VDI environment. The boot storms can be slightly reduced by pre-booting the virtual desktops. However, login storms need to be handled by the underlying storage efficiently, because the login storm is difficult to prevent in a user environment.

All the virtual desktops were pre-booted and then a login storm was simulated by simultaneously logging into all desktops using Login VSI.

As explained in section 3.3, the key components of a VDI hosted using PVS are vDisk, write cache, and the CIFS share. The vDisk and write cache components are the most critical components and have a huge impact on the underlying storage. CIFS share is a recommended way to separate user data from system data because it helps simplify the user data management.

We analyzed the I/O activity on vDisk, write cache, and CIFS share for all of the above test scenarios. The I/O activity on vDisk and CIFS share was similar for all of the above test scenarios and they are discussed in detail in sections 5.2 and 5.3. In later sections only write cache I/O behavior will be discussed because the vDisk and user CIFS share I/O behavior were similar across all tests.

5.2 vDisk I/O Characteristics

vDisk is the most critical component because the master OS disk image is stored on this volume. This master disk image is streamed to all virtual desktops over the network (designed as VDI network in our test) by the provisioning servers. There is a significant amount of read I/O activity on this volume specifically when the virtual desktops are being booted.

In our tests, we used two PVS servers which were virtualized. Each of them was allocated with 16 GB RAM and as a result we had 32 GB of cumulative RAM available on the PVS servers. Because our vDisk was 25 GB, most of the vDisk contents were cached and PVS streamed the data blocks directly from the server cache resulting in very low IOPS on the underlying storage. As a result, we observed low read IOPS on the storage array after initial login during steady state. A single volume named "vDisk-vol" was used to host the master OS disk image. This volume was made read-only and accessed by two provisioning servers using Windows guest iSCSI initiators.

The SAN HQ chart below shows the I/O characteristics observed on the vDisk volume during login storm and steady state.




Figure 6 Total IOPS on vDisk volume

As shown in Figure 6, we observed very minimal READ I/O activity (122 IOPS) on the vDisk volume during the login storm. Each PVS server had 16 GB of RAM and as a result, most of the data which needed to be streamed was cached in the server RAM. Once the disk image blocks got cached, further read requests from virtual desktops were served directly from PVS RAMs which significantly reduced the storage IOPS later on.

The I/O pattern observed on vDisk was same for all the different test scenarios (Refer to section 5.1).

5.3 CIFS share I/O characteristics

The EqualLogic PS6510E array was used to host the user shares and user profiles of all virtual desktops through a Windows-based file server. This CIFS share was used to redirect the user data and roaming profiles of virtual desktops to a centralized location. A single volume with 2 TB of capacity was created on the PS6510E and this volume was mounted on the file server to host the user profiles and user data folders for all virtual desktops.

When simulating a steady state workload for the task worker profile using Login VSI with all 630 virtual desktops, we observed a peak of 300 IOPS and when 1270 desktops were hosted a peak of 600 IOPS were observed on the PS6510E used for hosting the CIFS share.



Figure 7 Total IOPS on PS6510E array hosting CIFS Share – 1270 virtual desktops

As shown in Figure 7, the total IOPS requirement on the CIFS share was not significant. Redirecting user profiles and shares helped to reduce the IOPS requirement on the main PS6010XVS array used for storing and streaming desktop images along with the desktop write caches. It also allows better manageability of the virtual desktops in a VDI environment.

The IOPS requirement did not vary much for vDisk and CIFS share for the different test scenarios. But the IOPS on write cache varied significantly for each test scenario and detailed analysis of write cache IOPS is explained in the following sections.

5.4 Task worker – 100% Pre-boot and login storm

The key observations and findings from the 100% Pre-boot test scenario are analyzed in this section.

5.4.1 Storage layer: EqualLogic array performance

To support 630 virtual desktops, we used eight blade M610 servers and hosted 80 VMs on each blade server. A single PS6010XVS array was used to host the vDisk and write cache of all the virtual desktops. Four volumes, each with 500 GB of capacity, were dedicated for allocating write cache for virtual desktops. Each of the volumes was presented to the ESXi server as an individual datastore. Each of these volumes hosted the write cache for around 160 virtual desktops. Each virtual desktop was configured with 1 GB of RAM and 2 GB of write cache. Within a desktop VM, the write cache was mounted as a 2 GB virtual hard drive (vmdk).

Login VSI workload was used to simulate the task worker workload profile and detailed storage and infrastructure performance metrics were captured. All virtual desktops were pre-booted before the Login VSI workload started.

Some of the key performance metrics are analyzed below.

Volume level IOPS

We observed around 1100 to 1200 IOPS on each volume and both average read and write latencies were well below 2 ms. The total IOPS and the disk latency observed on one of the write cache volumes is shown in Figure 8.




Figure 8 Total IOPS and latency per volume

SAN HQ charts for the other three volumes are not included here, but similar performance was observed on the other three volumes.

Aggregate member level IOPS

The total IOPS on the PS6010XVS array during steady state when all the VMs logged in was around 4500 which is around 7.3 IOPS per VM as shown in Figure 9. There was some read I/O activity during the beginning of the test, but it was not more than 150 to 200 IOPS. Once steady state was reached, it was only the write IOPS which were generated because of write cache activity on each desktop. The

write cache activity on each desktop included all temporary OS writes such as paging. The OS disk image reads were satisfied via network streaming through the provisioning servers and the vDisk image was stored in a separate volume.




Figure 9 Total IOPS at member level

As you can see in Table 4, most of the IOPS were handled by high performance SSD drives and there was very little I/O on the SAS drives. The SSD drives acted as the cache on the storage array and were able to handle the write cache I/O load with an average disk latency of less than 2 ms.

Table 4 Detailed view of total IOPS at disk level

Member	Pool	Disk	Description	Average IOPS	Read I/O rate	Write I/O rate
XenVDI-XVS	VDI-Pool	0	SSD 100 GB SATA-II	1335.4	9.5 MB/sec	13.8 MB/sec
XenVDI-XVS	VDI-Pool	1	SSD 100 GB SATA-II	1334.1	11.1 MB/sec	15.4 MB/sec
XenVDI-XVS	VDI-Pool	2	SSD 100 GB SATA-II	1391.6	9.7 MB/sec	14.0 MB/sec
XenVDI-XVS	VDI-Pool	3	SSD 100 GB SATA-II	1403.7	11.8 MB/sec	16.0 MB/sec
XenVDI-XVS	VDI-Pool	4	SSD 100 GB SATA-II	1726.9	10.3 MB/sec	14.6 MB/sec
XenVDI-XVS	VDI-Pool	5	SSD 100 GB SATA-II	1366.7	10.4 MB/sec	14.8 MB/sec
XenVDI-XVS	VDI-Pool	6	SSD 100 GB SATA-II	1381.5	9.9 MB/sec	14.2 MB/sec
XenVDI-XVS	VDI-Pool	7	SSD 100 GB SATA-II	1380.6	11.4 MB/sec	15.7 MB/sec
XenVDI-XVS	VDI-Pool	8	15K 450 GB SAS	1.1	17.0 KB/sec	8.5 KB/sec
XenVDI-XVS	VDI-Pool	9	15K 450 GB SAS	1.2	25.6 KB/sec	8.5 KB/sec
XenVDI-XVS	VDI-Pool	10	15K 450 GB SAS	1.1	17.0 KB/sec	8.5 KB/sec
XenVDI-XVS	VDI-Pool	11	15K 450 GB SAS	1.1	26.6 KB/sec	8.5 KB/sec
XenVDI-XVS	VDI-Pool	12	15K 450 GB SAS	1.3	17.0 KB/sec	8.5 KB/sec
XenVDI-XVS	VDI-Pool	13	15K 450 GB SAS	<1.0	17.0 KB/sec	0 KB/sec
XenVDI-XVS	VDI-Pool	14	15K 450 GB SAS	1.2	25.6 KB/sec	8.5 KB/sec
XenVDI-XVS	VDI-Pool	15	15K 450 GB SAS	0	0 KB/sec	0 KB/sec

Note: The IOPS on disk 15 were zero because it was a hot spare.

5.4.2 Login storm I/O

All 630 users were able to login to the virtual desktops within 20 minutes. A significant increase in IOPS was observed as all the users logged in to the virtual desktop simultaneously. When users login, a significant amount of read and write I/O happens on the underlying storage from the write cache and also because of read activity from the vDisk image. We show only the write IOPS in Figure 10 because the vDisk read I/O was the same in all tests as described in section 5.2.

We used the Live View feature of SAN HQ to monitor the real-time IOPS on the PS6010XVS array. As shown in Figure 10, a peak of 6846 IOPS was observed on the array when all 630 users logged into the virtual desktops. The write IOPS during the peak were higher (6846 IOPS) than during steady state (4500 IOPS) as described earlier. This is due the write cache initialization activities caused by the OS on each desktop during login. After login most of the write activity is due the application activities to serve the simulated workload.

The login storm typically results in higher IOPS due to factors such as user profile activity, software initialization, and different services which are started upon login. Also the initial launch of applications and software services takes a large amount of IOPS and as applications are re-launched, the IOPS are reduced because the Windows OS has cached the applications in memory.




Figure 10 IOPS monitored using 'Live View' feature of SAN HQ during login storm – 630 VMs

The storage array disk latency never went beyond 5 to 6 ms even at the peak I/O during the login storm. We observed an increase in CPU and network resource utilization at peak login, but it was well within the acceptable limits.

5.4.3 Hypervisor layer: ESXi host performance

During the test, we measured CPU, memory, network, and disk performance on all ESXi servers hosting the virtual desktops.

The key observations on the ESXi servers were:

- CPU utilization did not reach more than 95% at any point in time, ensuring sufficient resources for all of the virtual desktops.
- The average disk latency of all disks on the ESXi server was less than 2 ms.
- There was no memory ballooning and the peak memory utilization was less than 50%.
- Network design is very critical because provisioning servers will stream the master VM image over the network to all virtual desktops. The peak network utilization on any ESXi servers was not more than 25 to 30%.

5.4.4 Citrix infrastructure layer: PVS and DDC performance

We monitored the resource utilization on the ESXi servers hosting the XenDesktop infrastructure VMs while running the test. Since the two provisioning servers and DDCs were load balanced, the performance on each ESXi server was similar.

CPU utilization:

PVS and DDC servers had two virtual CPUs and 16 GB of memory allocated. We observed slightly higher CPU utilization (around 5%) during the login storm due to streaming activity. After that, the utilization remained consistently below 2% to 3%.

The slightly higher CPU utilization during the initial stages was caused by various factors such as:

- During login, there is usually extra processing at the DDC to authenticate and allocate a specific user to one of the available pre-booted desktops.
- There is also increased processing on the PVS layer because there is an increased amount of activity related to reading the master disk image and streaming it to all virtual desktops.

Memory utilization:

The ESXi servers hosting provisioning servers and DDCs had sufficient memory resources available and we did not observe any memory ballooning or swapping at any point during the test.

VDI Network utilization:

We dedicated two virtual NICs (vmnic3 and vmnic5) for handling VDI streaming traffic. These two virtual NICs were teamed together to create a vSwitch and presented to DDC as gigabit adapter using the VMXNET3 driver. The peak network utilization did not go beyond 40 MB at any point of time during the test.

5.4.5 User experience: Application performance

The virtual desktop user experience and desktop performance was measured using Stratusphere UX. Stratusphere UX generates a magic quadrant style scatter plot which shows the performance of virtual desktops based on machine and I/O experience. All 630 VMs were represented as GOOD performing desktops in the Stratusphere UX scatter plot as shown in Appendix D .

5.4.6 Write cache utilization on virtual desktops

In our tests, we configured 2 GB as the write cache on each virtual desktop. The "Cache on Device Hard Drive" option was used to mount the write cache. We measured the write cache space utilization individually on specific virtual desktops and it never reached 75% even after running the workload on 630 virtual desktops for more than ten hours.

5.5 Task worker – 20% Pre-boot

This test was designed in the same way as the "100% Pre-boot" scenario except that only 20% of the total desktops were pre-booted. In this test, only 126 virtual desktops (20% of the 630) were pre-booted and the remaining desktops got booted when additional users started logging into the desktops.

5.5.1 Storage layer: EqualLogic array performance

The maximum IOPS on the storage array during login storm and steady state was almost the same as the 100% Pre-boot task worker test, but the login time increased significantly.




Figure 11 Total IOPS at member level – 20% Pre-boot scenario

In this scenario, the time required to login to 630 virtual desktops was almost 100 minutes. In the 100% Pre-Boot scenario it was just 20 to 25 minutes. This is because 80% of the desktops had to be booted before logon and this resulted in increased logon duration. In this scenario, a user will experience considerably longer login duration. The read IOPS on storage was low due to streaming of the disk image from provisioning servers. As seen in Figure 11, the write IOPS constantly varied between 4000 and 5500 IOPS for more than 40 minutes during boot and login time. This variation in IOPS was due to the fact that both boot and login operations occurred simultaneously during this period. The maximum write IOPS observed was similar to the 100% Pre-boot scenario, but it was sustained for a longer time due to the boot and login happening simultaneously.

The other performance characteristics at the hypervisor layer, Citrix Infrastructure layer, user experience measures on virtual desktops, and write cache utilization were the same as the observations captured during 100% Pre-boot test scenario.

5.6 Virtual desktop scaling with two arrays

This test was the same as the 100% Pre-Boot test described in section 5.4. Here we wanted to validate the scalability of the VDI solution by adding an additional storage array. We added one more PS6010XVS array to the same pool and also added eight more blade M610 servers to the VMHA ESXi cluster hosting virtual desktops. In this test, two PS6010XVS arrays with a total of eight volumes each with 500 GB capacity were used. They were individually mapped to their respective datastores and presented to the 16 ESXi servers. A single vDisk volume was used to store the base OS disk image which got streamed to all virtual desktops.

5.6.1 Storage layer: EqualLogic array performance

Figure 12 shows the total IOPS on the EqualLogic pool which consisted of two PS6010XVS arrays in the same pool. All the users logged into the virtual desktops within 25 minutes. The peak IOPS during login storm reached to almost 12400, which we monitored with SAN HQ Live View.




Figure 12 Total IOPS on the EqualLogic pool

We verified that system resource utilization at the ESXi layer and Citrix infrastructure layer were within the acceptable limits as defined in the section 4.4. More than 94% of the desktops were represented as 'GOOD' performing desktops by Stratusphere UX. The Stratusphere UX scatter plot is shown in Appendix D .

The test results for both one array and two array configurations are summarized in Table 5.

Table 5 IOPS in two array test

Number of PS6010XVS arrays	Number of Blade M610 servers	Number of virtual desktops	IOPS (Steady state)
1	8	630	4500
2	16	1270	9000

These results show that almost linear scaling could be achieved by adding additional EqualLogic array members when there is no bottleneck on other infrastructure components.

5.6.2 Login storm I/O

All 1270 users were pre-booted and the users were able to login to the virtual desktops within first 25 minutes. Around 12400 IOPS were observed cumulatively on both EqualLogic arrays in the pool.



Figure 13 IOPS on one of the EqualLogic array during login storm – 1270 VMs

Note: The above SAN HQ chart shows the peak IOPS of 6200 during the login storm on one of the array members. Almost the same number of IOPS was observed on the other EqualLogic member.

The storage array disk latency never went beyond 5 to 6 ms even at the peak I/O during the login storm. We observed an increase in CPU and network resource utilization at peak login, but it was well within the acceptable limits.

5.7 Boot storm

We restarted all virtual desktops from VMware vCenter to simulate the boot storm. The peak IOPS when 1270 virtual desktops were booted was around 15000. The average latency on both storage arrays was less than 10 ms and the storage arrays were able to handle this I/O burst with no significant increase to the latency.

5.7.1 Storage layer: EqualLogic array performance




Figure 14 Total IOPS and average latency on VDI Pool

All 1270 desktops were booted within 10 to 12 minutes. The storage arrays were able to handle this spike in IOPS and the average latency remained below 10 ms through the test. As shown in Figure 14, the boot storm consisted of both read and write I/O. Read I/O is a smaller proportion compared to the write I/O due to the disk image streaming from provisioning servers.

Read I/O is caused by vDisk read by provisioning servers and reads to the desktop write cache during initialization. Write I/O is primarily caused by write cache initialization during login and steady state application activity. The following SAN HQ chart shows the amount of read and write IOPS on an individual write cache volume.




Figure 15 Read and write IOPS on an individual volume

We did not observe any bottlenecks on the ESXi servers with respect to CPU or memory resources during the boot operation. The peak CPU utilization observed was around 60% during the boot process. The system resources on the Citrix Infrastructure ESXi servers hosting provisioning servers and DDCs were hardly used during the boot storm. The CPU resource utilization on each of the Citrix Infrastructure ESXi servers while booting VMs was less than 10%.

The network utilization reached a peak of 60 MB during the boot storm and gradually declined once the master disk image was streamed to all virtual desktops.

This test clearly demonstrated that the storage array was never the bottleneck in our test configuration and could easily handle the VDI boot storm scenarios.

5.8 Results summary

The key observations from the test results are listed below.

- A single PS6010XVS array was able to host 630 virtual desktops and support the task worker type of I/O activity. Two arrays were able to host 1270 desktops.
- The VDI I/O represented write-intensive I/O on write cache disks (almost 90% write) and less read I/O (10% read) on the vDisk. The I/O on vDisk was further reduced because most of the vDisk data blocks were cached on the PVS RAM.

- None of the system resources on the ESXi servers hosting virtual desktops reached maximum utilization levels at any point in time.
- In the 100% Pre-boot scenario, the login completed within 20 to 25 minutes.
- In the 20% Pre-boot scenario, the login duration increased to 100 minutes.
- During boot storm simulation, we observed nearly 15000 IOPS on the two array configuration.
- The Stratusphere UX tool represented 100% of the virtual desktops in 'GOOD' category on the one array configuration. On the two array configuration, the virtual desktops represented as "GOOD" were more than 98%.

6 Sizing guidelines for EqualLogic SANs

Virtual desktop usage in enterprise environments follows predictable I/O patterns. For example, at the beginning of the workday most employees login into their desktops within a relatively short timeframe. 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 and afternoon hours and low activity during break hours would be expected. These phases of activity cause different I/O patterns and loads in the VDI environment SAN. There may also be some rare scenarios such as recovery after an unexpected power shutdown which can cause a boot storm. The storage array should be able to handle this type of workload within acceptable latency limits.

The first step to design and deploy a storage infrastructure for VDI is to understand the storage characteristics and performance requirements of the application. The most important storage characteristics to be considered while designing the storage infrastructure are:

- Capacity
- Performance
 - IOPS
 - Average disk latency
- Type of I/O (Sequential or Random)
 - Read and write ratio

In this chapter, we will summarize the key performance and capacity guidelines specifically focused on Citrix XenDesktop VDI environments deployed using EqualLogic storage. These guidelines are based on the results and key observations captured while completing the tests explained in section 5.

6.1 Storage capacity considerations

6.1.1 vDisk capacity requirements

The vDisk size was determined based on these key considerations:

- The size of the Windows 7 OS image.
- Generally it is recommended to keep at least one or two clean backups of the base OS images.
- Also there may be some requirements such as keeping multiple versions of vDisks in case of major software updates or new patches.

The size of the OS image depends on the specific use case and also the number of applications and software installed. These estimates were based on an OS installation with only the applications to run the Login VSI workload. These estimates may not accurately reflect the required sizes of vDisks in every configuration, but they provide a guideline to determine space requirements for vDisk images for task worker I/O profiles.

We recommend using standard mode vDisks where the single master OS image is streamed to all virtual desktops over the network. This type of implementation is very efficient because administrators need to manage and maintain a single image instead of each individual virtual desktop and it reduces storage requirements significantly because only one OS image needs to be stored.

For more information regarding the different Provisioning Services vDisk modes, please refer to the following link: <http://support.citrix.com/proddocs/topic/provisioning-56/pvs-technology-overview-vDisk-modes.html>

6.1.2 Write cache capacity requirements

The write cache is a temporary disk on each virtual desktop which contains the modified blocks from each user session. The write cache file is deleted during each reboot cycle.

The options for storing write cache are:

- Cache on PVS
- Cache on target device (virtual desktop) local RAM
- Cache on target device (virtual desktop) local Hard Drive

For more information regarding the different write cache destinations for standard vDisk images, please refer to the following link: <http://support.citrix.com/proddocs/topic/provisioning-56/pvs-technology-overview-write-cache-intro.html>

The size of the write cache depends on several factors such as type of applications used, user workload, and frequency of reboots. It is highly recommended to run a pre-deployment environment to determine the exact write cache usage. For the applications we tested, we found that 2 GB of write cache was sufficient for the workload under test.

6.2 Storage performance considerations

The key results from our tests from section 5 are summarized in the table below.

Table 6 Test results for one array and two array configurations

Test configuration	1 x EqualLogic PS6010XVS	2 x EqualLogic PS6010XVS
Total number of desktops	630	1270
Number of datastores used to store virtual desktop write cache	4	8
Number of virtual desktops per datastore for write cache	~160	~160
Number of blade PowerEdge M610 servers used to host virtual desktops	8	16
Number of virtual desktops per ESXi server	80	80
Average latency	< ~10ms	< ~10ms
Peak IOPS during login storm	6846	12400

Peak IOPS during boot storm	10000	15000
Peak IOPS during steady state	4500	9000
Capacity buffer	15%	18%

Based on our test results, I/O characteristics observed for the task worker type of workload profile during the lifecycle of a virtual desktop can be summarized as shown in Table 7.

Table 7 I/O characteristics during the lifecycle of a VM

Type of I/O	IOPS per desktop	Read/Write ratio	I/O Block size
Steady state	7 - 8	0.2% / 99.8%	Read: 1KB / Write: 9KB
Login storm	11 - 13	3% / 97%	Read: 23KB / Write: 5KB
Boot storm	16 - 18	13% / 87%	Read: 32KB / Write: 5KB

Note: This IOPS estimation is based on the workload simulated by using Login VSI "Light" workload. The number of IOPS during login storm and steady state may vary depending on the actual virtual desktop environment. It is recommended to run a pre-deployment test to determine the exact IOPS requirement.

Boot storm was simulated by restarting all Windows 7 virtual desktops from VMware vCenter. The boot storm IOPS number can be used as a guideline for determining boot storm I/O for Windows 7 VDI environments only when the Windows 7 image is created and tuned as described in section 7.4.3.

6.3 Sizing Calculations

IOPS is not the only key criteria while designing the storage solution. Calculating the storage capacity is also an important aspect of designing storage for VDI environments.

Storage capacity requirements for pooled desktops using a PVS primarily depend on the following factors:

- Number of vDisks to be created and maintained
- vDisk mode to be used
- OS installed
- Additional applications to be installed
- Storage location of the write cache file and size
- The type of I/O activity for each virtual desktop

The following formula may be used for determining how many virtual desktops can be deployed.

$$\text{NumVM} = (\text{ArrayCapacity} - (15\% * \text{ArrayCapacity} + \text{vDiskSize} * \text{NumvDisk})) / (125\% * (\text{VMem} + \text{WCacheSize}))$$

In this formula, we use 15% headroom for the array capacity and 25% headroom for VM RAM and the write cache allocated on the virtual desktop.

The following table describes the formula components and provides a sample calculation.

Table 8 Formula components

Variable	Description	Value in our test configuration
NumVM	Number of virtual desktops to be deployed.	
VMMem	Amount of RAM allocated on each virtual desktop.	1 GB
vDiskSize	Size of the master OS disk image.	25 GB
NumvDisk	Number of vDisks	6
WCacheSize	Size of the Write cache.	1.5 GB
ArrayCapacity	Total capacity of the PS6010XVS array in GB.	2539.5 GB

The size of the write cache and RAM on the virtual desktop varies depending on the applications used in different customer environments. The actual usage needs to be monitored and the formula should be used as guideline for sizing storage capacity in deploying virtual desktops.

Sample capacity calculation from our test configuration:

$$\text{NumVM} = (2539 - (0.15 \times 2539 + 25 \times 6)) / (1.25 \times 1 + 1.25 \times 1.5)$$

This calculation shows that we can deploy around 640 desktops using a single PS6010XVS array to meet our test capacity requirements. In our test configuration, we deployed 630 virtual desktops and, from a performance standpoint, the storage disk latency was well within 10 ms and there was no bottleneck with storage array performance. The observed IOPS for our workload is described above in Table 6.

If the number of virtual desktops needs to be scaled, then we recommend adding additional EqualLogic arrays. Based on our test results, linear scaling can be achieved by adding additional storage arrays as long as there is no bottleneck on other system resources. It is important to note that when sizing for desktop count, both the capacity and IOPS needs should be considered. Either capacity or IOPS requirement may warrant more storage arrays than required by the other. It is a best practice to determine the number of arrays needed for capacity and IOPS separately. The maximum of these two values can be considered as the final requirement. This approach ensures that the underlying storage can handle both performance and capacity requirements needed by the VDI environment.

7 Best practices

7.1 Desktop profiles and I/O storms

7.1.1 Implement roaming profiles and folder redirection

Using a separate storage array and redirecting user profile/folders to a file server using that array is highly recommended for better management and increased performance of VDI environments. Implementing roaming profiles and folder redirection helped in reducing the performance impact during user logon and also allowed user data to be persistent across boot. More information related folder redirection and roaming profiles can be found in Appendix B

7.1.2 Use XenDesktop Power management feature

Citrix XenDesktop has a feature called Power Management that can be used to maintain the required number of idle virtual desktops based on user logins.

By default, the Power management feature will ensure that 10% of the total number of desktops are available. It is recommended to set the Power Management feature to 100% so that all the virtual desktops are pre-booted. This will result in reduced login time and better user experience. However it may not be always possible to pre-boot 100% of the virtual desktops due to resource or power constraints in a data center. In these scenarios, we recommend pre-booting as many desktops as possible just before users start logging in to the virtual desktops.

7.1.3 Boot and login storm considerations

To avoid I/O bursts due to boot storms, we recommend pre-booting all virtual desktops before the first set of users start logging on at the beginning of a workday. It may not be possible to pre-boot all virtual desktops due to infrastructure limitations. In this case, we recommend pre-booting as many desktops as possible to minimize desktop boot time and improve user experience.

If the storage layer is designed based only on the steady state IOPS, the user will experience degraded performance during the duration of the login storm. We recommend designing the storage after considering the boot and login storms.

Based on our test results, VDI workload using provisioning servers is highly write-intensive (90% write and 10% read I/O). Because of this, using SSD drives will greatly improve the performance. You also need sufficient capacity on the array to host write cache for the virtual desktops. We recommend PS6010XVS hybrid arrays because they offer the right combination of performance and capacity suitable for VDI environments.

7.2 ESXi host configuration

We recommend implementing the following recommendations on all ESXi servers hosting virtual desktops and Citrix infrastructure:

- Install and configure EqualLogic Multipathing Extension Module (MEM) for vSphere 4.1
- Separate the virtual switch to segregate iSCSI SAN traffic, VDI traffic, and Management/Infrastructure network I/O
- Every network path on the servers must be assigned at least two NICs for high availability.

7.3 Network configuration

The key recommendations related to network design are listed below.

- We recommend using at least two dedicated physical NICs per server for each of the networks listed below:
 - VDI network
 - Management, Infrastructure, and vMotion network
 - iSCSI SAN
- Use VLANs to segregate different types of network traffic such as the Management network, Infrastructure network, and the vMotion network. This helps in improving manageability, performance, and security. See Appendix C for more detailed VLAN configuration information.
- vSwitches on ESXi 4.1 servers are configured with 120 virtual ports by default. This number of ports may not be sufficient specifically on ESXi servers hosting virtual desktops. Consider increasing this value based on the number of desktops that will be hosted on each ESXi server.
- The IP addresses of virtual desktops should be configured using DHCP server and the IP subnet should have enough IP addresses in one or more subnets for each virtual desktop.
- We recommend disabling spanning tree on switch ports connecting to end devices like server ports and storage ports. The Portfast setting should be enabled on the switch for these ports.

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

PS Series Array Network Performance Guidelines:

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

7.4 Citrix infrastructure

7.4.1 Read-only vDisk

vDisks should be presented as shared storage across multiple provisioning servers and the SAN volume should be formatted with a file system. However, Windows NTFS does not support concurrent read/write access which might result in data corruption. A clustered file system or Windows clustering services needs to be used when vDisks need to be deployed using shared SAN storage across multiple provisioning servers. This adds complexity, overhead, and cost, so we recommend using read-only LUNs for vDisk storage that can be concurrently accessed by multiple provisioning servers.

Some of the key benefits of this feature are:

- Completely eliminates the need to deploy shared/cluster file system software.
- Reduced complexity, overhead, and deployment costs.
- Maximum scalability and performance.

For more detailed information on how to configure read-only vDisk using EqualLogic arrays and its relevance, please refer to the article: <http://support.citrix.com/article/CTX119286>

7.4.2 Write cache recommendations

Storing write cache on the target devices (virtual desktops) is recommended for better performance because the file creation is local to the target and not streamed to any servers, which avoids additional network traffic.

Restarting virtual desktops clears the cache file and it is highly recommended to restart the desktops at least once per day or more frequently if possible. This clears the write cache area which is very critical for optimal performance of virtual desktops.

7.4.3 Optimizing the Windows 7 master image

Because we are using provisioning servers to stream the single vDisk to all virtual desktops, proper design and optimization of the master OS image is very critical for better user experience. There are specific system settings and configurations which enable faster logons, reduce unnecessary prompts, create quicker screen updates, and other conveniences. Also, Windows 7 includes many services by default. When using provisioning services, these services do not provide any value. Provisioning Service actually resets the virtual desktops to the original state upon reboot, so input into the services is lost after each reboot.

We recommend using the Windows 7 optimization settings described in the following Citrix article:
<http://support.citrix.com/servlet/KbServlet/download/25161-102-649933/XD%20-%20Windows%207%20Optimization%20Guide.pdf>

7.4.4 Allocate more RAM on PVS servers

It is always recommended to allocate slightly more memory to provisioning servers RAM than the base vDisk image. The vDisk attached to the PVS servers will cache to the memory, which results in vDisk being streamed from memory rather than from SAN. Memory access is always faster than the disk.

This approach will help to:

- Reduce the overall read IOPS on the underlying storage array.
- Allow very fast access to the master disk image which results in better user experience.

7.5 Storage

Detailed storage performance and capacity sizing guidelines are presented in section 5. Additional recommendations are listed below.

7.5.1 Volumes and datastores

It is recommended to create at least three or more volumes on a single EqualLogic member for better performance. Using multiple volumes helps by enabling parallel I/O from the controllers which in turn results in better performance and manageability. These volumes host the vDisk and desktop write cache.

7.5.2 Storage array characteristics

We recommend using EqualLogic PS6010XVS hybrid arrays which consist of eight SSD drives and eight 15K SAS drives for hosting the performance critical components such as vDisk and write cache. The other components such as user profile information can be stored on high capacity arrays such as PS6510E.

Based on our test results, we recommend adding additional EqualLogic array members to the same pool to scale the number of virtual desktops.

We also recommend keeping at least 15% capacity headroom on the storage array to accommodate any future requirements and for optimal array performance.

8 Conclusions

The life cycle of a virtual desktop in a VDI environment includes many stages such as boot, login, steady state, and logoff. Each of these stages results in varied amount of IOPS and can create a huge impact on the underlying storage. The test configuration used in our tests could support up to 630 desktops with one array and 1270 desktops with two arrays and there were no bottlenecks observed related to storage performance. The EqualLogic PS6010XVS arrays offers the right combination of high performance SSD drives and 15K SAS drives for ensuring that the performance and capacity requirements of a VDI environment are met. The test results also proved that linear scaling could be achieved by adding additional EqualLogic array members.

Appendix A Citrix XenDesktop solution configurations

Solution configuration - Hardware components		Description
Virtual Desktops	16 x Dell PowerEdge M610 Servers: - ESXi 4.1 - BIOS Version: 3.0.0 2 x Quad Core Intel® Xeon® X5687 3.6 Ghz Processors 96 GB RAM 2 x 146 GB 10K SAS internal disk drives 1 x Dual-port Broadcom 5709 1 GbE NIC (LAN on motherboard) 2 x Broadcom NetXtreme II 57711 10 GbE NIC, Dual-Port	ESXi 4.1 was installed on all M610 servers. 80 Virtual desktops were hosted by each ESXi server. Windows 7 (32 Bit) VMs
Provisioning Services Servers and Desktop Delivery Controllers (DDC)	2 x Dell PowerEdge R710 servers: - ESXi 4.1 - BIOS Version: 3.0.0 2 x Quad Core Intel® Xeon® X5690 3.46 Ghz Processors 96 GB RAM, 2 x 146 GB 10K SAS internal disk drives 1 x Quad-port Broadcom 5709 1 GbE NIC (LAN on motherboard) 2 x Broadcom NetXtreme II 57711 10 GbE NIC, Dual-Port	ESXi 4.1 was installed on both R710 servers. Two VMs per ESXi servers were created with Windows 2008 Server R2 OS. One VM hosted the first PVS and the other VM hosted a DDC. The other ESXi server had the exact same setup for high availability.
File server	1 x Dell PowerEdge R710 servers: - ESXi 4.1 - BIOS Version: 3.0.0 2 x Quad Core Intel® Xeon® X5690 3.46Ghz Processors 96 GB RAM 2 x 146 GB 10K SAS internal disk drives 1 x Quad-port Broadcom 5709 1 GbE NIC (LAN on motherboard) 2 x Broadcom NetXtreme II 57711 10	ESXi 4.1 was installed. Two VMs were created. The VMs were loaded with Windows 2008 R2 operating system. User data was presented to all virtual desktops using these servers as CIFS share.

	GbE NIC, Dual-Port	
Citrix XenDesktop Infrastructure Servers	1 x Dell PowerEdge R710 servers: - ESXi 4.1 - BIOS Version: 3.0.0 2 x Quad Core Intel® Xeon® X5690 3.46 Ghz Processors 96 GB RAM 2 x 146 GB 10K SAS internal disk drives 1 x Quad-port Broadcom 5709 1 GbE NIC (LAN on motherboard) 2 x Broadcom NetXtreme II 57711 10 GbE NIC, Dual-Port	The VMs were loaded with Windows 2008 R2 OS. Each VM hosted: - Citrix Web Interface - VMware vCenter - Domain controller (DNS and DHCP) - Microsoft SQL Server (DDC / PVS and vCenter Database) - Citrix License Server
XenDesktop VDI clients (Login VSI Launchers) VDI Workload Generator	3 x Dell PowerEdge R710 servers: - ESXi 4.1 - BIOS Version: 3.0.0 2 x Quad Core Intel® Xeon® X5690 3.46Ghz Processors 96 GB RAM 2 x 146 GB 10K SAS internal disk drives 1 x Quad-port Broadcom 5709 1 GbE NIC (LAN on motherboard) 2 x Broadcom NetXtreme II 57711 10GbE NIC, Dual-Port - Login VSI version 3.0	ESXi 4.1 was installed on both R710 servers. 12 VMs were created on each ESXi server. The VMs were loaded with Windows 2008 R2 operating system. 40 Login VSI launchers were initiated from each VM to simulate VDI workload.
Network	2 x Dell PowerConnect M6220 1 Gb Ethernet Switch Firmware: 4.1.0.6, A24 2 x Dell PowerConnect M8024 switch Firmware: 4.1.0.19 2 x Dell PowerConnect 8024F 10 Gb Ethernet Switch Firmware: 4.1.0.19	4 x 10 GbE uplinks to the external PC8024F switch from PCM8024 switches on the blade chassis. The PC8024 switches were used for configuring the dedicated iSCSI SAN. They were lagged using 4 x 10 GbE links.
Storage	2 x Dell EqualLogic PS6010XVS: 8 x 100GB SSD 8 x 450GB 15K SAS disks - Dual 2 port 10 GbE controllers	vDisks and write cache for each virtual desktop was stored on PS6010XVS array. User data presented as a CIFS share is stored on PS6510E array.

	○ Firmware: 5.1.1 (R189834) (H2) • 1 x Dell EqualLogic PS6510E: ○ 48 x 1TB 7.2K SATA disks ○ Dual 2 port 10 GbE controllers ○ Firmware: 5.1.1 (R189834) (H2)	
Performance Monitoring	• SAN HeadQuarters – 2.1.1 • vCenter Performance monitoring	Performance monitoring on EqualLogic arrays. Performance monitoring and capture at the ESXi host.

Solution configuration – Software components	Description/Version
Citrix XenDesktop	Version 5 Platinum Edition
Citrix Desktop Delivery Controller Server	Version 5
Citrix Provisioning Services	Version 5.6
Hypervisor	VMware vSphere 4.1 (ESXi 4.1)
Microsoft SQL Server	Version 2008 Enterprise Edition (64-Bit)
Windows 7 Enterprise	VDI Clients for characterization tests
Windows Server 2003 Enterprise R2 x64	VMs for hosting Citrix XenDesktop PVS and DDCs.
Windows Server 2008 Enterprise R2 x64	VMs for hosting Citrix XenDesktop PVS, DDCs, vCenter Server, Microsoft SQL Server, Login VSI launchers and other infrastructure VMs.

Appendix B Folder redirection and roaming profiles

User settings and user files are typically stored in the local user profile, under the Users folder. These files usually reside on the local hard drive which makes it difficult for users who use more than one computer to work with their data and synchronize settings between multiple computers. Roaming profiles and folder redirection technologies can be used to address this issue and we used both to provide seamless user experience and better manageability of the virtual desktops.

Folder Redirection lets administrators redirect the path of a folder to a new location. The location can be a folder on the local computer or a directory on a network file share. Users can work with documents on a server as if the documents were based on a local drive. The documents in the folder are available to the user from any computer on the network.

Roaming User Profiles provide a way to give users a familiar and easy-to-use working environment. Unlike a local profile, which is stored on a single computer, a roaming profile is stored on a network share, which means it can be accessed from any computer on the network.

Appendix C Network design and VLAN configuration

VDI Network

The servers listed below are part of the VDI network:

- The M610 blade servers hosting virtual desktops
- The ESXi servers hosting PVS and DDC
- The ESXi server used for hosting the file server and presenting CIFS share
- The ESXi servers used for launching Login VSI workload

The network architecture block diagram for this VLAN from the blade servers is shown in Figure 16.




Figure 16 VDI VLAN connectivity

Infrastructure network

The network architecture block diagram for this LAN from the blade servers is shown in Figure 17.




Figure 17 Infrastructure, Management, and vMotion VLAN connectivity

The infrastructure network was divided into three VLANs:

- **Infrastructure VLAN**
The infrastructure server which hosts Citrix Web Interface, Active Directory, Microsoft SQL Server, License server, and VMware vCenter was configured with a dedicated VLAN.
- **Management VLAN**
A dedicated VLAN was configured for management traffic of all the ESXi servers and virtual machines shown in Figure 17.
- **vMotion VLAN**
The vMotion traffic is routed on a dedicated VLAN on the same network. The virtual machines in the VMHA cluster might migrate from one ESXi server to other in case of a server failure, but this is a rare occurrence and the vMotion traffic is expected to be low

Appendix D Liquidware Labs Stratusphere UX

The Stratusphere UX Scatter plot for the single array test (630 Virtual Desktops) is shown in Figure 18.




Figure 18 Stratusphere UX Scatter plot for the single array test

Note: Only 626 VMs were deployed in this test as Login VSI failed to launch the workload on four VMs.

The Stratusphere UX Scatter plot for the two array test (1270 Virtual Desktops) is shown in Figure 19.




Figure 19 Stratusphere UX Scatter plot for the two array test

Note: The performance was monitored on 1200 virtual desktops. This shows more than 94% of the desktops performing as "GOOD" and the remaining desktops showed "FAIR" performance. None of the desktops showed "POOR" performance.

Related publications

The following Dell publications are referenced in this document or are recommended sources for additional information.

- PS Series Array Network Performance Guidelines:
<http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>

The following Citrix publications are referenced in this document or are recommended sources for additional information.

- XenDesktop Planning Guide: Image Delivery:
<http://support.citrix.com/servlet/KbServlet/download/26449-102-650176/XD%20-%20Planning%20Guide%20-%20Desktop%20Image%20Delivery.pdf>
- Provisioning Services vDisk Modes
<http://support.citrix.com/proddocs/topic/provisioning-56/pvs-technology-overview-vDisk-modes.html>
- Selecting the Write Cache Destination for Standard vDisk Images
<http://support.citrix.com/proddocs/topic/provisioning-56/pvs-technology-overview-write-cache-intro.html>
- Provisioning Server High Availability Considerations
<http://support.citrix.com/article/CTX119286>
- Windows 7 Optimization Guide
<http://support.citrix.com/servlet/KbServlet/download/25161-102-649933/XD%20-%20Windows%207%20Optimization%20Guide.pdf>

The following Login Consultants publications are referenced in this document or are recommended sources for additional information.

- Login VSI product overview:
<http://www.loginvsi.com/en/product-overview>
- Login VSI workloads:
<http://www.loginvsi.com/en/admin-guide/workloads>

The following Liquidware Labs publications are referenced in this document or are recommended sources for additional information.

- Measuring VDI Fitness and User Experience:
<http://www.liquidwarelabs.com/docs/VDI%20Fit%20and%20VDI%20UX%20White%20Paper.pdf>