



White Paper

50,000-Seat VMware View Deployment

Jack McLeod, Chris Gebhardt, Chad Morgenstern, Brian Collins, and Rachel Zhu, NetApp
Ravi Venkat, Cisco
Mac Binesh and Fred Schimscheimer, VMware
Param Desai, Wyse
Arthur Perkins, Fujitsu America
August 2011 | WP-7108

TABLE OF CONTENTS

1	AUGUST 2011 UPDATE	4
2	OVERVIEW	11
	GOAL	11
	DESIGN APPROACH	11
2	PARTNERS	14
	NETAPP	14
	VMWARE	14
	CISCO	16
	FUJITSU	16
	WYSE TECHNOLOGY	16
3	INFRASTRUCTURE COMPONENTS	17
	NETAPP KILO-CLIENT	17
	ACTIVE DIRECTORY	18
	NETWORK	19
	SERVERS	20
	STORAGE	31
	VMWARE VSPHERE 4.1	32
	VMWARE VIEW 4.5	34
	WINDOWS 7	35
	WYSE TECHNOLOGY	38
4	REFERENCES	39
5	ACKNOWLEDGEMENTS	39

LIST OF TABLES

Table 1)	The creation scenario.	5
Table 2)	The initial login scenario.	5
Table 3)	The “Tuesday morning” or typical login scenario.	6
Table 4)	The power-on scenario.	6
Table 5)	The “Monday morning” or profile load login scenario.	6
Table 6)	The steady-state scenario.	6
Table 7)	Server configuration per AD environment.	12
Table 8)	PRIMERGY RX200 S5: If efficiency is the key decision factor.	21
Table 9)	FAS controller A.....	31
Table 10)	FAS controller B.....	32

Table 11) NetApp controller physical configuration.	32
Table 12) View 4.5 setup.	34

LIST OF FIGURES

Figure 1) Read and write operations per second.	8
Figure 2) Read and write throughput.	8
Figure 3) Read operation breakdown for power-on.	9
Figure 4) Read operation breakdown for initial login.	9
Figure 5) Read operation breakdown for steady state.	10
Figure 6) Desktop virtualization POD design (graphic provided by VMware).	12
Figure 7) Active Directory overview (graphic provided by VMware).....	13
Figure 8) Active Directory.	18
Figure 9) Front view of I/O module (graphic provided by Fujitsu).....	23
Figure 10) Rear view of I/O module (graphic provided by Fujitsu).	23
Figure 11) Cisco UCS (graphic provided by Cisco).	27
Figure 12) Storage VMotion.	33
Figure 13) Create Rapid Clones Wizard.	35

1 AUGUST 2011 UPDATE

In August of 2010, NetApp and a number of partners published this white paper describing the deployment of a 50,000-seat virtual desktop infrastructure (VDI) environment using NetApp storage, Cisco Unified Computing System™ (Cisco UCS™) and Cisco Nexus®, VMware software, and Fujitsu servers. This white paper focused solely on the high-level architectural design and the specifics of the technology that each of the partners brought to the table to allow for this deployment. This initial effort called for using NetApp FAS3170 storage controllers and defined modular units of storage and servers called “pools of desktops” (PODs), based on the hardware and software needed for deploying 5,000 seats into a virtualization infrastructure. This initial white paper defined a POD as follows:

- 60 ESX® 4.1 hosts (Cisco UCS or Fujitsu PRIMERGY)
- 1 FAS3170A high-availability (HA) cluster
- 96 15K RPM Fibre Channel drives
- 2 512GB Flash Cache cards
- 2 VMware vCenter™ Servers
- 3 VMware View™ Connection Servers running PC-over-IP (PCoIP) connection protocol
- 5,000 Microsoft® Windows® 7 virtual desktop virtual machines (VMs)

Shortly after the publication of this paper in August 2010, NetApp refreshed its midrange storage offerings to include the FAS3270 storage systems, which improved on the capacity and performance of the FAS3170 storage systems used in the initial white paper. As a result, we used a FAS3270 storage system instead of the FAS3170 used in the initial white paper because it significantly improves the performance and scalability of the solution. In addition, subsequent early testing showed that being able to effectively support 5,000 VDI desktops required that we add 30 more servers so that adequate memory resources were available during the testing. As a result, we now define a POD as follows:

- 90 ESX 4.1 Servers
 - Fujitsu PRIMERGY RX200-S5 with 2 Quad Core Nehalem CPUs with hyperthreading
 - 48GB main memory
- 1 FAS3270A HA cluster
- 96 15K RPM Fibre Channel drives
- 2 512GB Flash Cache modules
- 2 VMware vCenter Servers
- 3 VMware View Connection Servers running PC-over-IP (PCoIP) connection protocol
- 5,000 Microsoft Windows 7 virtual desktop VMs

Under this new definition, each NetApp FAS3270 controller supported 45 ESX servers and 2,500 Windows 7 persistent virtual desktops. Because of the large hardware requirements to create a full POD, we chose to limit these subsequent tests to using what we describe as half a POD or 2,500 virtual desktops being served by one of the two FAS3270 storage controllers. Because each FAS3270 storage controller is actually independently serving 2,500 virtual desktops, the performance measured for the 2,500 virtual desktops can simply be doubled to account for the full POD supporting 5,000 virtual desktops.

For these tests we used VMware View 4.5 and VMware vSphere 4.1 to deploy a scenario consisting of 2,500 virtual desktops. In addition, we used the VMware Reference Architecture Workload Code (RAWC) tool to generate a workload typical of what might be found in a VDI environment.

We tested with both the current General Availability (GA) release (Data ONTAP® 8.0.1) and Data ONTAP 8.1. We selected Data ONTAP 8.1 to use because it provides additional performance gains and spindle reduction benefits to the NetApp virtual storage tiering (VST) capabilities. VST allows customers to benefit from NetApp storage efficiency and at the same time significantly increase I/O performance. VST

is natively built into the Data ONTAP operating system and works by leveraging block-sharing technologies such as NetApp primary storage deduplication and file/volume FlexClone to reduce the amount of cache required and eliminate duplicate disk reads. Only one instance of any duplicate block is read into cache, thus requiring less cache than traditional storage solutions. Because VMware View implementations can see as great as 99% initial space savings using NetApp space-efficient cloning technologies, this translates into higher cache deduplication and high cache hit rates. VST is especially effective in addressing the simultaneous system boot or “boot storm” and login of hundreds to thousands of virtual desktop systems that can overload a traditional legacy storage system. With Data ONTAP in 8.1, deduplicated data blocks in main memory can be shared more efficiently than in previous releases. This allows for a larger working set to fit in the storage controller’s main memory, resulting in faster access times and reduced CPU.

During our testing, we confirmed that there is much more to a VDI workload than simply what goes on during steady state when users are accessing their desktops during the course of normal business. Although it is important to understand the characteristics of this phase, there are also situations in which reboots, logins, profile creation, and manipulation of large numbers of virtual desktops place heavy burdens on the storage supporting the overall VDI environment. Failure to understand and plan for these workloads can have a serious negative impact on end-user experience and project success.

In the new NetApp technical report, “TR-3949: NetApp and VMware View 5,000-Seat Performance Report,” we examined a number of different common VDI scenarios from the perspective of the fictitious Acme Corporation. Acme has deployed a 2,500-seat VDI environment using NetApp storage and wants to understand the different types of workloads that might be encountered during a typical workweek in which users typically log in in larger numbers, start and stop applications, and conduct routine tasks using the applications that have been provided to them. Additionally, there are times when routine maintenance requires that all of the virtual desktops be powered off, which then requires subsequent power-on and booting of large numbers of VMs simultaneously.

These scenarios and their outcomes as measured by user experience are defined at a high level in Table 1 through Table 6. The user experiences in “TR-3949: NetApp and VMware View 5,000-Seat Performance Report,” are measured by Liquidware Labs Stratusphere UX, which uses a weighting algorithm to determine “goodness” of the user experience. The power-on and creation scenarios are reported in the elapsed time. Note that the application mixture selected is one that VMware advised would generate 12 input/output operations per second (IOPS)/desktop, the expected workload of “knowledge users.”

Table 1) The creation scenario.

Test	Primary Measure of Success	Creation Time
Create the 2,500 VMs and measure how long it takes.	Length of time the creation takes to be ready for first power-on	3 hours from beginning of the clone process to 2,500 VMs ready for power-on

Table 2) The initial login scenario.

Test	Primary Measure of Success	User Experience
2,500 users log in for the first time over a half-hour period and begin working. This login triggers 2,500 profile creations.	User experience as measured by Liquidware Labs UX	<ul style="list-style-type: none"> • Data ONTAP 8.0.1: 62% of users had a good user experience, 38% fair. • Data ONTAP 8.1: 97% of users had a good user experience, 3% fair.

Table 3) The “Tuesday morning” or typical login scenario.

Test	Primary Measure of Success	User Experience
2,500 users log into VMs previously accessed and not rebooted since. At login, the users begin working.	User experience as measured by Liquidware Labs UX	<ul style="list-style-type: none"> • Data ONTAP 8.0.1: 100% of the users had a good user experience. • Data ONTAP 8.1: 100% of the users had a good user experience.

Table 4) The power-on scenario.

Test	Primary Measure of Success	Power-On Time
VMware vCenter controls the mass power-up of 2,500 virtual desktops. Neither profiles nor application DLLs must be loaded from disk.	Total time to power-on	<ul style="list-style-type: none"> • Data ONTAP 8.0.1: 36 minutes • Data ONTAP 8.1: 21 minutes

Table 5) The “Monday morning” or profile load login scenario.

Test	Primary Measure of Success	User Experience
2,500 users log into VMs previously accessed but since rebooted. At login, the users begin working. Profiles and application DLLs must be loaded from disk.	User experience as measured by Liquidware Labs UX	<ul style="list-style-type: none"> • Data ONTAP 8.0.1: 100% of the users had a good user experience. • Data ONTAP 8.1: 100% of the users had a good user experience.

Table 6) The steady-state scenario.

Test	Primary Measure of Success	User Experience
2,500 users have completed their logins and have opened and closed all applications at least once but continue their day’s work.	User experience as measured by Liquidware Labs UX	<ul style="list-style-type: none"> • Data ONTAP 8.0.1: 100% of the users had a good user experience. • Data ONTAP 8.1: 100% of the users had a good user experience.

At a high level, through our testing we determined that each of the scenarios generated a workload unique to itself. The power-on workload was characterized as being read heavy with a high IOPS count. The “Monday morning” scenario as well as the initial login scenario generated workloads fairly balanced between reads and writes, although the intensity of the initial login workload was the greater of the two. The “Tuesday morning” login scenario as well as the steady-state workload scenario generated mostly write operations. The workload generated by these last two scenarios was far less than the other scenario; of the two, the login generated more than the steady state.

The characterization of each of these scenarios has an impact on end users and their overall experience. Power-on of VMs is not normally associated with end-user experience, except when events such as HA, disaster recovery (DR), or any unplanned maintenance occur. In this case, booting quickly is extremely important to the end users. For login, regardless of the day of the week, no user wants to wait 15 minutes until their desktop is ready for use. And while users are interacting with their desktops, they want performance as good as or better than that provided by a physical desktop. Failure to characterize these

different workloads and their impact on the infrastructure and the end user will ultimately affect the success of the project.

In addition, we found that these different categories of workloads exhibited characteristics that were sometimes vastly different from the anecdotal classification of a VDI workload as being small-block random reads. For example, we found the following:

- The disk writes are primarily small, in general 4kB, but may be much larger; 1MB write operations have been recorded during large file-save operations.
- Read operations ranged in size from sub-4kB to 1024kB and were fairly balanced in quantity between random and sequential.

In other words, although virtual desktop workloads do generate small random workloads, they also generate large random operations as well as small and large sequential operations. Our tests show that the industry-standard definition of a VDI workload (meaning just looking at IOPS) is much too simplistic and that correctly sizing storage for a VDI environment requires looking deeper.

The graphs in Figure 1 and Figure 2 demonstrate the differences in workload between the scenarios described as they might play out over a typical day or week in the life of Acme Corporation's VDI environment as displayed in the format "a day in the life of." In the August 2010 version of this white paper, the environment was architected using state-of-the-art technologies. Although we did test the original configuration, many revisions of software and hardware have become available, so the newer versions were tested as well. Although the results of the original test validated the architecture, all results detailed in "TR-3949: NetApp and VMware View 5,000-Seat Performance Report," were achieved with a FAS3270A, 15K drives, and Data ONTAP 8.0.1 and 8.1. For the purposes of "TR-3949: NetApp and VMware View 5,000-Seat Performance Report," this architecture represents our "golden" configuration: "golden" because this particular configuration produced the best user experiences, the highest throughput, and the lowest latencies, and it did so with the lowest cost/IO operation. These graphs showing throughput and IOPS are from the perspective of the storage controller.

Observe the clear delineation between the workload phases as measured in terms of both operations per second and throughput. Notice that the power-on scenario is dominated primarily by a relatively large number of read operations that drive the throughput past peaks of 700MB per second. In addition, the initial login scenario is fairly balanced in operations between read/write operations, with the read operations driving peak throughput of over 400MB per second. Finally, after users have powered on their VMs and logged into their desktops, the steady-state workload shows a significant drop in both IOPS and throughput compared to the power-on and login scenarios. Notice further that, as shown in Figure 1 and Figure 2, once steady state is achieved, the workload shifts toward an 80:20 distribution of write-to-read operations but toward a fairly balanced distribution in terms of throughput as more of the work associated with the initial read activity (for example, loading application libraries and opening/reading files) becomes cached in the VM guest operating system (OS).

Figure 1) Read and write operations per second.

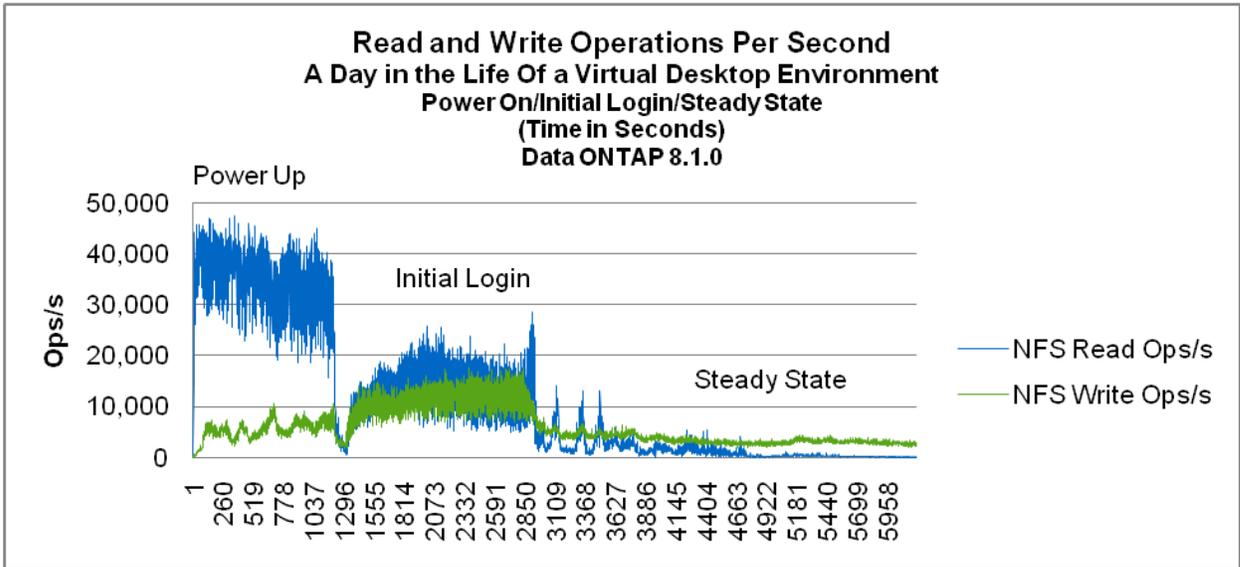
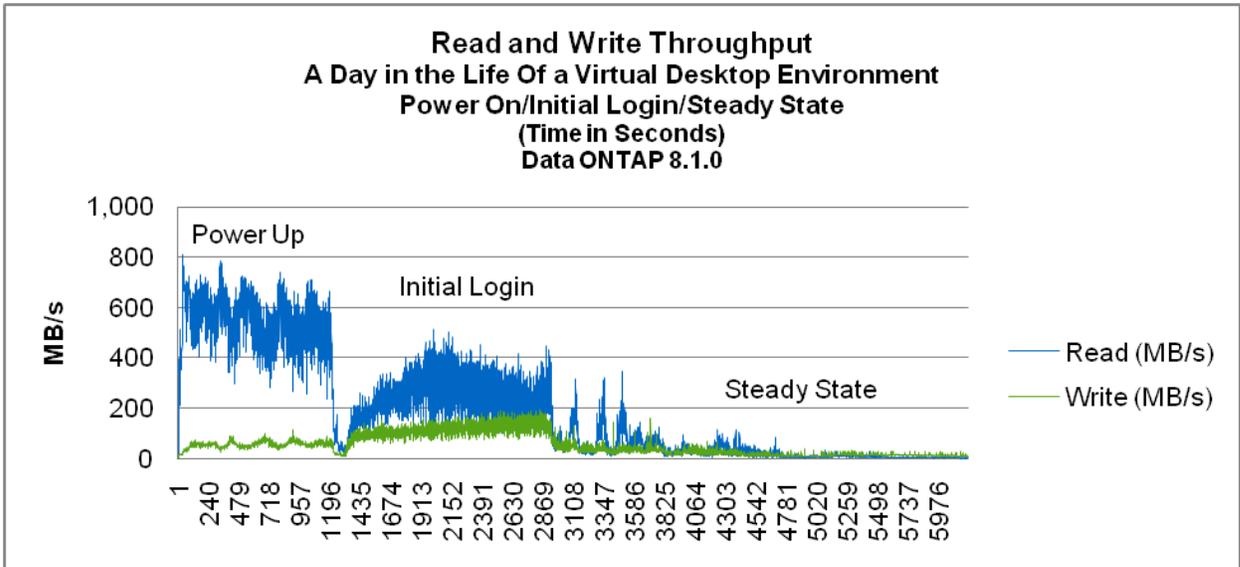


Figure 2) Read and write throughput.



The graphs in Figure 3 through Figure 5 call out the operations sizes as reported by the storage controller. Notice that as the workload distribution and its quantity vary across the workload scenarios, the operations sizes and sequential mixtures vary as well. The constant across this set of workloads is that the percent of random/sequential read operations remains fairly balanced.

Figure 3) Read operation breakdown for power-on.

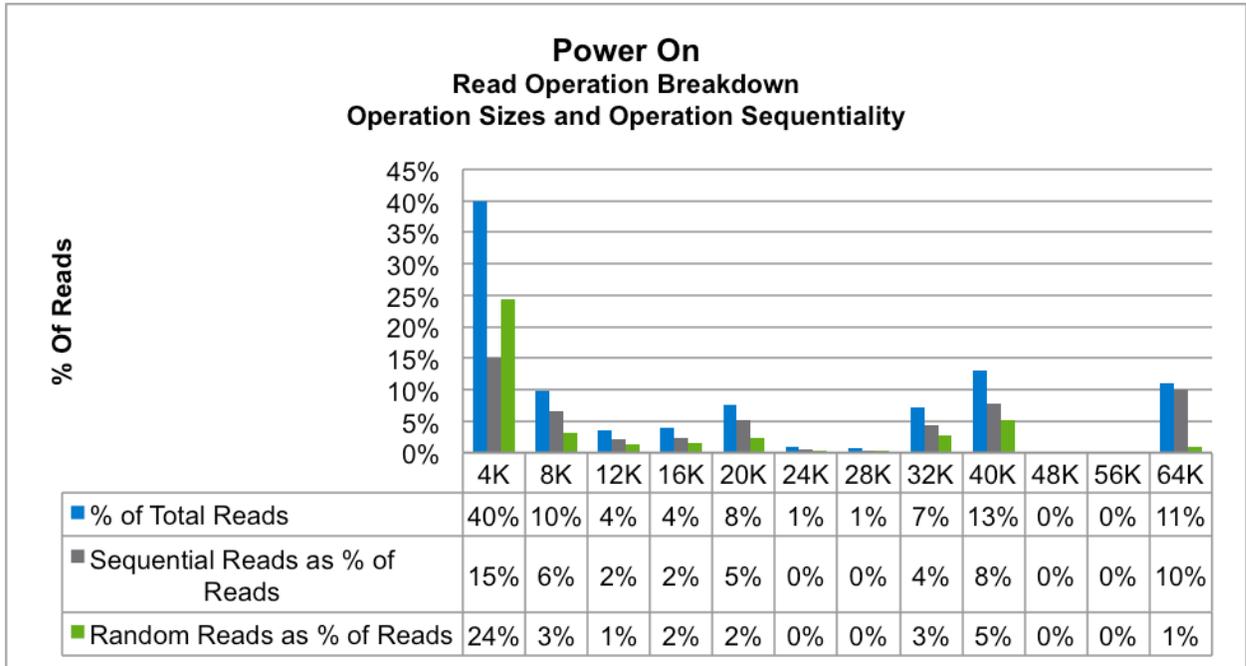


Figure 4) Read operation breakdown for initial login.

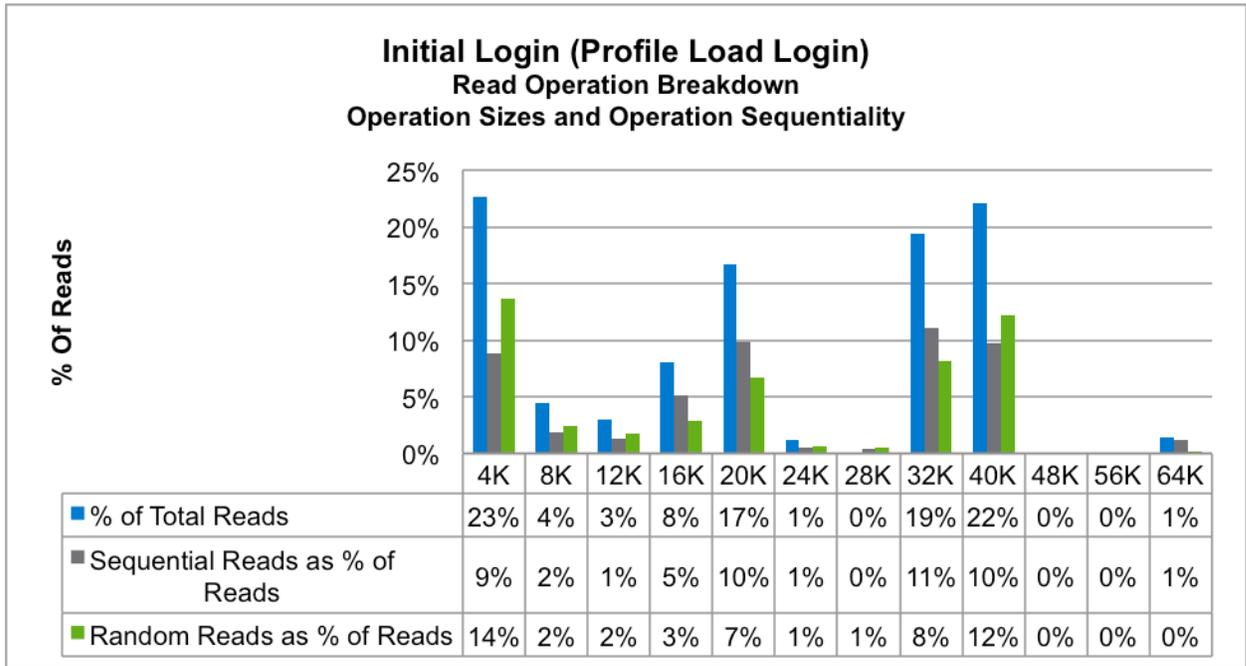
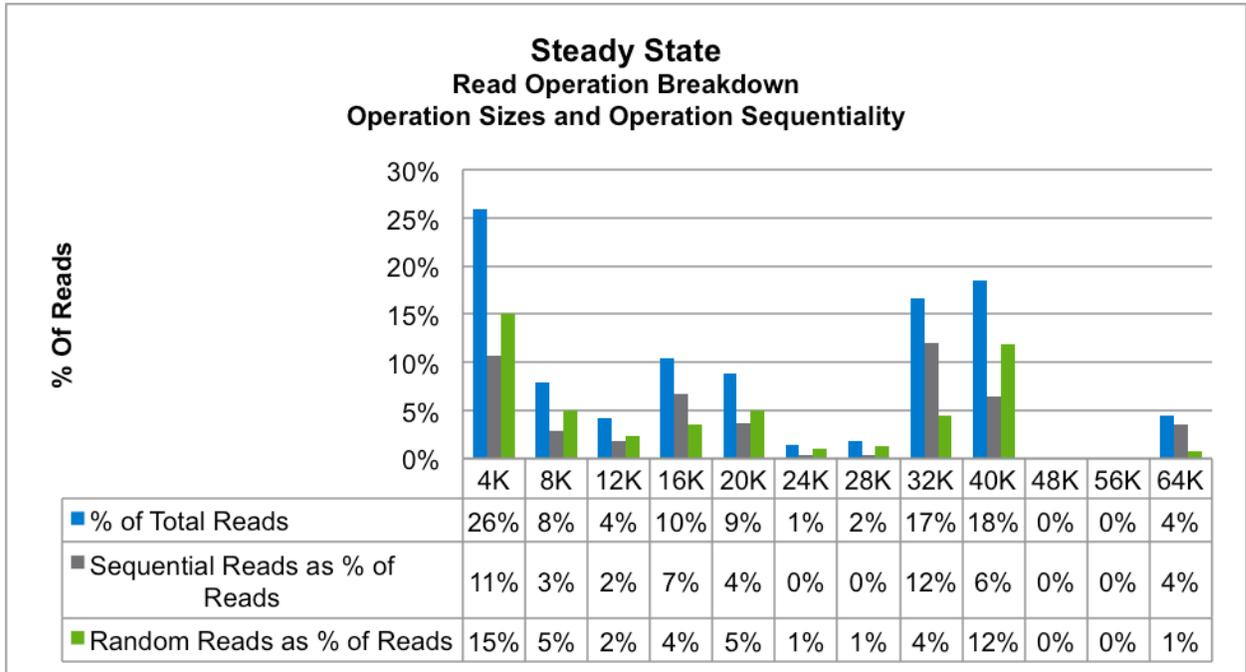


Figure 5) Read operation breakdown for steady state.



In the end, our testing showed that delivering great performance in a VDI environment requires more than just the capability to support the steady-state workload (generally thought of to be 12 IOPS per virtual desktop). We found a number of other areas that must also be considered in order to create an overall good user experience.

The remainder of this white paper is the original version published in August 2010.

2 OVERVIEW

Desktop virtualization is a hot topic throughout the virtualization industry. Organizations view desktop virtualization as a way to control costs and use limited resources to manage large-scale desktop infrastructures while increasing security and deployment efficiency. NetApp, VMware, Cisco, Fujitsu, and Wyse joined forces to create an architectural design for a 50,000-seat VMware® View™ architecture. This project demonstrates the most reliable methods for deploying a scalable and easily expandable environment from 5,000 to 50,000 seats in an efficient manner.

GOAL

The purpose of this project is to design and deploy a 50,000-seat VMware View 4.5 virtual desktop environment using VMware vSphere™ 4.1 on NetApp® shared storage. Cisco, VMware, NetApp, Fujitsu, and Wyse are combining their respective technologies to create a dynamic and easily expandable virtualization infrastructure. The partners want to make this environment scalable and easy to replicate to give customers the ability to use the architecture detailed in this white paper as an example for their own deployments. This deployment uses a Pool of Desktops (POD) concept that is based on the hardware and software needed for deploying 5,000 seats into a virtualization infrastructure.

This document provides an initial overview of the ongoing 50,000-seat desktop virtualization project. Because this is an ongoing project, this white paper only discusses the architectural design and the partner involvement.

DESIGN APPROACH

THE DEPLOYMENT SCENARIO

The scenario deploys a 50,000-seat environment on 10 NetApp FAS3170 high-availability (HA) pairs using the Network File System (NFS) protocol. A 20,000-seat deployment and six 5,000-seat regular deployments are used in this scenario.

The deployment uses a 50% read/write mix and allows a minimum of 10% CPU availability on each controller. In addition, each virtual machine (VM) is assumed to use 12 input/output operations per second (IOPS) in the configuration. Based on these assumptions, a maximum of 5,000 seats can be run on a FAS3170 HA pair.

SOFTWARE NEEDED FOR DEPLOYMENT

- NetApp System Manager 1.01
- NetApp VSC 2.0 (RCU 3.1 Plug-in)
- VMware vSphere (ESX 4.1 and vCenter™ Server 4.1)
- VMware View Manager and Composer 4.5
- Microsoft® Windows® 2008 Enterprise
 - DHCP
 - DNS
 - DHCP Relay Agents
 - Active Directory® (AD)

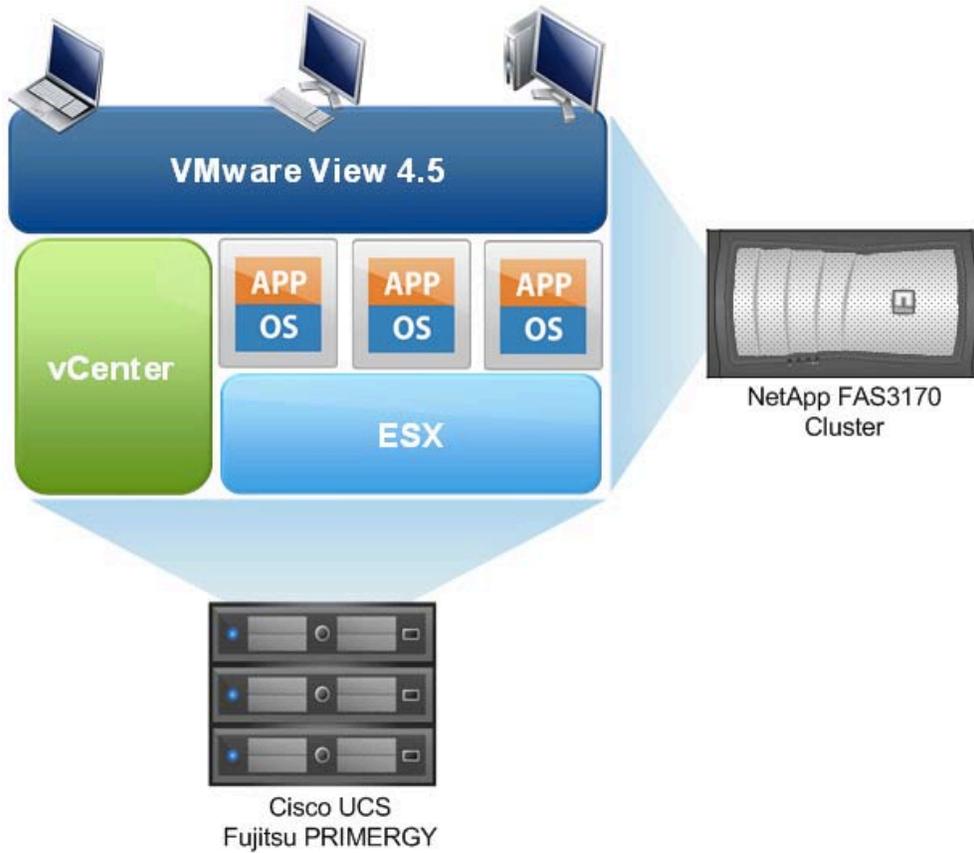
POD DEFINITION

This deployment uses a POD design because it is modular and has an easy-to-replicate process. As shown in Figure 6, the POD consists of the following components:

- 60 ESX 4.1 Hosts (Cisco UCS or Fujitsu PRIMERGY)

- 1 FAS3170 HA Cluster
- 2 vCenter Servers
- 3 Connection Servers (running PCoIP®)
- 5,000 Microsoft Windows 7 VMs

Figure 6) Desktop virtualization POD design (graphic provided by VMware).



PHYSICAL SERVER CONFIGURATION

The server specifications for this configuration are shown in Table 7. The table includes the number of servers per Active Directory environment.

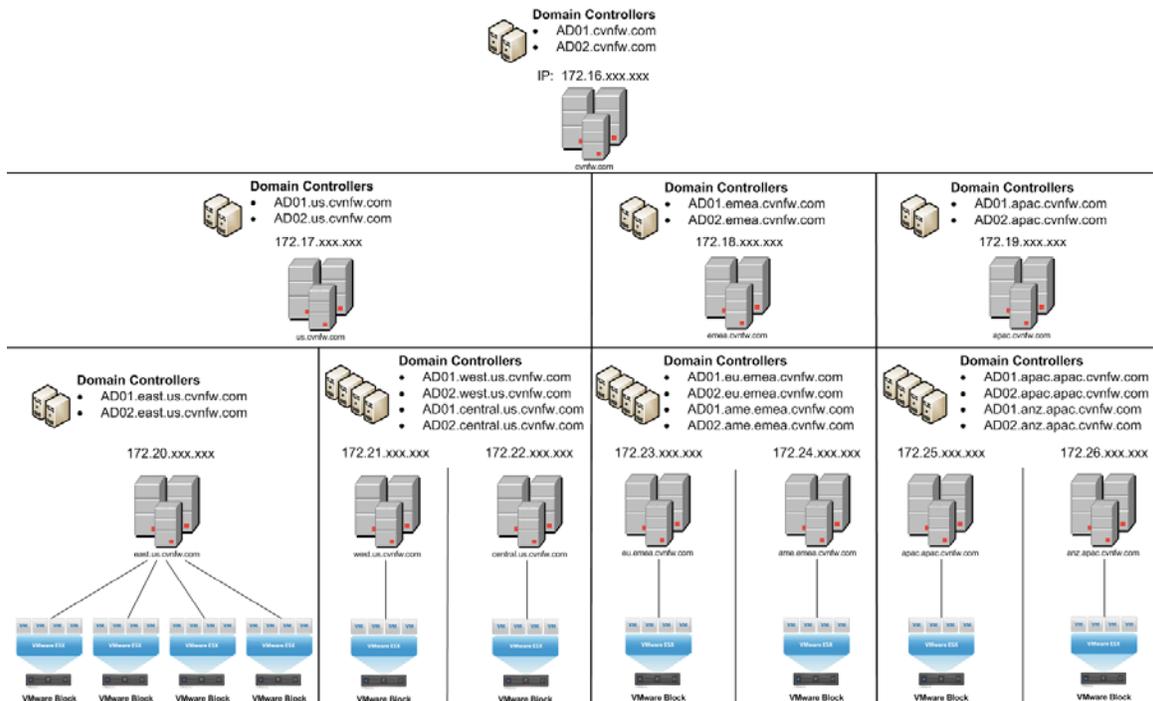
Table 7) Server configuration per AD environment.

vCenter Location	AD Domain	# Servers	Manufacturer	RAM (GB)
New York 1	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY	48
New York 2	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY	48
New York 3	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY	48
New York 4	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY	48
New York 5	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY	48
New York 6	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY	48

vCenter Location	AD Domain	# Servers	Manufacturer	RAM (GB)
New York 7	EAST.us.cvnfw.com	30	Fujitsu PRIMERGY/Cisco UCS	48/48
New York 8	EAST.us.cvnfw.com	7/15/16	Cisco UCS/Cisco UCS/Fujitsu PRIMERGY	48/24/48
Chicago	CENTRAL.us.cvnfw.com	50	Cisco UCS	24
Dallas	CENTRAL.us.cvnfw.com	52	Cisco UCS	24
Phoenix	WEST.us.cvnfw.com	46	Cisco UCS	24
San Francisco	WEST.us.cvnfw.com	54	Cisco UCS	24
Tokyo	APAC.apac.cvnfw.com	30	Fujitsu PRIMERGY	48
Singapore	APAC.apac.cvnfw.com	30	Fujitsu PRIMERGY	48
Sydney	ANZ.apac.cvnfw.com	30	Fujitsu PRIMERGY	48
Auckland	ANZ.apac.cvnfw.com	30	Fujitsu PRIMERGY	48
London	EU.emea.cvnfw.com	30	Fujitsu PRIMERGY	48
Munich	EU.emea.cvnfw.com	30	Fujitsu PRIMERGY	48
Dubai	AME.emea.cvnfw.com	30	Fujitsu PRIMERGY	48
Johannesburg	AME.emea.cvnfw.com	30	Fujitsu PRIMERGY	48

The AD specifications for this configuration are also shown in Figure 7.

Figure 7) Active Directory overview (graphic provided by VMware).



2 PARTNERS

NETAPP

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Our dedication to the principles of simplicity, innovation, and customer success has made us one of the fastest-growing storage and data management providers today.

Customers around the world choose us for our “go beyond” approach and broad portfolio of solutions for server-to-storage virtualization, business applications, data protection, and more. Our solutions and virtualized storage technologies can provide nonstop availability of critical business data and speed product development, so you can deploy new capabilities with confidence and get to revenue faster than ever before.

NetApp storage solutions for VMware virtual desktop environments simplify data management and save time while leveraging existing investments in heterogeneous storage arrays across FC, iSCSI, NFS, and CIFS. With NetApp storage, you can:

- Reduce virtual desktop storage use by 50% or more by deduplicating redundant data stored across virtual desktops, user directories, and backup and disaster recovery (DR) copies.
- Provision thousands of virtual desktops in minutes with nearly instant, low-overhead storage cloning.
- Provide users continuous access to their virtual desktops with 99.999% storage availability, automated DR, and performance acceleration that addresses desktop boot or log-in storms.
- Back up all virtual desktops and user data directly on the storage system while cost effectively keeping a daily history of every desktop for months or years.

Discover our passion for helping companies around the world go further, faster at www.netapp.com.

VMWARE

VSPHERE 4.1

Leveraging more than a decade of industry-leading technology, VMware vSphere sets a new standard as the most robust and reliable virtualization platform that forms a good foundation for building cloud infrastructures. VMware vSphere is used in over 170,000 deployments in companies of all sizes throughout the world. The latest release of VMware vSphere (version 4.1) pushes the platform even further by:

- Dramatically expanding the scale of the platform to an unmatched number of VMs and virtualized hosts, enabling the creation of private and public clouds at even lower operational costs than before
- Introducing new capabilities for the most efficient aggregation of IT resources into an elastic pool of computing power that applications can consume on demand as a utility-like service

The new and enhanced features that VMware vSphere 4.1 delivers, grouped by the major benefits of the platform, are as follows:

- Efficiency through utilization and automation
- Agility with control
- Freedom of choice

VMware vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the data center. It transforms data centers into simplified cloud computing infrastructures and enables IT organizations to deliver flexible and reliable IT services.

VMware vSphere is the industry's most complete virtualization platform, with infrastructure services that transform IT hardware into a high-performance shared computing platform, and application services that help IT organizations deliver the highest levels of availability, security, and scalability.

Note: For a complete and detailed list of the new features and capabilities in VMware vSphere 4.1, visit www.vmware.com/go/vshpere.

VMWARE VIEW

Purpose-built for delivering desktops as a managed service, VMware View™ provides a positive end-user experience and transforms IT by simplifying and automating desktop management. Centrally maintaining desktops, applications, and data reduces costs, improves security, and, at the same time, increases availability and flexibility for end users. Unlike other desktop virtualization products, VMware View is a tightly integrated end-to-end solution built on the industry-leading virtualization platform. VMware View allows customers to extend powerful business continuity and disaster recovery features to their desktops and standardize on a common platform from the desktop through the data center to the cloud.

The VMware View solution provides a wide range of benefits.

- **Simplify and automate desktop management.** VMware View lets you manage all desktops centrally in the data center and provision desktops instantly to new users, departments, or offices. Create instant clones from a standard image and create dynamic pools or groups of desktops.
- **Optimize end-user experience.** The VMware View PCoIP display protocol provides a superior end-user experience over any network. Adaptive technology provides an optimized virtual desktop delivery on both the LAN and the WAN. Address the broadest list of use cases and deployment options with a single protocol. Access personalized virtual desktops complete with applications and end-user data and settings anywhere and anytime with VMware View.
- **Lower costs.** VMware View reduces the overall costs of desktop computing by up to 50% by centralizing management, administration, and resources and by removing IT infrastructure from remote offices.
- **Enhance security.** Because all data is maintained within the corporate firewall, VMware View minimizes risk and data loss. Built-in Secure Socket Layer (SSL) encryption provides secure tunneling to virtual desktops from unmanaged devices or untrusted networks.
- **Increase business agility and user flexibility.** VMware View accommodates changing business needs, such as adding new desktop users or groups of users, while providing a consistent experience to every user from any network point.
- **Provide built-in business continuity and disaster recovery.** VMware View is built on industry-leading VMware vSphere, allowing you to easily extend features such as HA and Fault Tolerance to your desktops without having to purchase expensive clustering solutions. Automate desktop backup and recovery as a business process in the data center.
- **Standardize on a common platform.** VMware View includes VMware vSphere and brings all the benefits and enterprise features of the data center to the desktop. Extending features such as VMotion™, HA, Distributed Resources Scheduler, and Fault Tolerance to your desktops provides a built-in disaster recovery and business continuity solution. Optimized specifically for desktop workloads, VMware vSphere is able to handle the high loads associated with desktop operations such as boot-up and suspend operations. Standardize your virtualization platform, and use a single solution to manage both servers and desktops from the data center through to the cloud.

VMware View provides unified access to virtual desktops and applications that run in a central secure data center and are accessible from a wide variety of devices. VMware View Composer streamlines image management while reducing storage needs through the use of VMware Linked Clone technology.

CISCO

Cisco Systems, Inc., is the worldwide leader in networking for the Internet. Today, networks are an essential part of business, education, government, and home communications, and Cisco Internet Protocol-based (IP) networking solutions are the foundation of these networks.

Cisco hardware, software, and service offerings are used to create Internet solutions that allow individuals, companies, and countries to increase productivity, improve customer satisfaction, and strengthen their competitive advantage. The Cisco name has become synonymous with the Internet, as well as with the productivity improvements that Internet business solutions provide. The Cisco vision is to change the way people work, live, play, and learn.

- **Founded.** Cisco was founded in 1984 by a group of computer scientists from Stanford University.
- **Incorporation.** Cisco was incorporated on December 10, 1984, in California.
- **Stock Symbol.** NASDAQ: CSCO
- **Employees.** As of the end of Q2 FY'2010 (ending January 23, 2010), Cisco had 68,574 employees worldwide.
- **Corporate Headquarters.** San Jose, California, U.S.

FUJITSU

Fujitsu is a leading provider of information and communication technologies (ICT)-based business solutions for the global marketplace. With approximately 170,000 employees supporting customers in 70 countries, Fujitsu combines a worldwide corps of systems and services experts with highly reliable computing and communications products and advanced microelectronics to deliver added value to customers. Headquartered in Tokyo, Fujitsu Limited (TSE:6702) reported consolidated revenues of 4.6 trillion yen (US\$50 billion) for the fiscal year ending on March 31, 2010. For more information, refer to www.fujitsu.com.

Fujitsu America Inc. is a leading ICT solutions provider for organizations in the United States, Canada, and the Caribbean. Fujitsu enables clients to meet their business objectives through integrated offerings including consulting, systems integration, managed services, outsourcing, and for enterprise applications, data center, and field services operations, based on server, software, storage, and mobile technologies. Fujitsu provides industry-oriented solutions for the manufacturing, retail, healthcare, government, education, financial services, and communications sectors. For more information, refer to <http://solutions.us.fujitsu.com/>.

WYSE TECHNOLOGY

Wyse Technology is the inventor and global leader in cloud client computing, leveraging its industry-leading thin and zero client computing-based desktop virtualization software, hardware, and services. Cloud Client Computing is the ultimate client computing solution for our time. It replaces the outdated computing model of the unsecure, unreliable, ungreen, and expensive PC. It delivers the most optimized security, reliability, total cost of ownership, and user experience with the lowest energy usage. It simply connects all the dots: thin and zero client computing, unified communications, desktop virtualization, and the Web for users to reach the clouds—in a private, public, government, or hybrid cloud. It is software. It is hardware. It is services. It is in business. It is at home. It is on the go. It is freedom—so users can focus on what is important. Wyse partners with industry-leading IT vendors including, Cisco, Citrix, CSC, IBM, Microsoft, NetApp, and VMware. Wyse also partners with globally recognized distribution and service partners along with its award-winning partner programs to service any customer, anywhere in the world, at anytime. Wyse is headquartered in San Jose, California, U.S., with offices worldwide.

Wyse is the worldwide leader in Thin and Zero Client Computing, offering thin and zero clients before anyone else. Wyse software makes it easy to provision, manage, accelerate, or optimize the user experience. A thin or zero client can even be serviced from one central location. After all, it's much easier

and more cost effective to manage several servers rather than thousands of individual desktop PCs. And with no moving parts, Wyse thin and zero clients deliver greater reliability, availability, and a lower cost of ownership than other solutions.

3 INFRASTRUCTURE COMPONENTS

NETAPP KILO-CLIENT

The NetApp Kilo-Client is a test environment that allows NetApp to quickly configure and boot a large number of physical and/or virtual clients to run tests against NetApp storage hardware and software. The deployment is performed on the third generation (or Kilo 3G) of this environment, which takes advantage of Fujitsu PRIMERGY RX200 S5 rackable servers and Cisco UCS B200 M1 blades.

The real key to the Kilo-Client is its ability to perform fast, flexible, and space-efficient booting. As in any cloud infrastructure, the Kilo-Client has the capability to quickly repurpose any number of clients for any task—physical or virtual. The Kilo-Client uses a combination of Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) boot to boot each physical server and an NFS boot to support VMs booting on servers configured to run virtualization. The FC boot in combination with our dynamic logical unit number (LUN) cloning capability is used to rapidly and efficiently boot physical and virtual servers.

The Kilo-Client maintains a set of "golden" boot images (as FC LUNs) for each operating system and application stack. Through the use of NetApp SnapMirror® and NetApp FlexClone®, hundreds of clones can quickly be reproduced for each physical server being configured for a test. This approach allows near-instantaneous image provisioning with a near-zero footprint.

The process of booting VMs builds on the same steps:

1. Boot VMware ESX on each host for the test.
2. Register those hosts dynamically in the VMware vCenter.
3. Prepare the correct network settings and datastores for the VMs.
4. Use VSC 2.0 to clone and register the appropriate number and types of VMs.
5. Dynamically register the servers in DNS and DHCP, and boot the VMs.
6. Check to make sure that everything is correct.

COMPLETE AUTOMATION

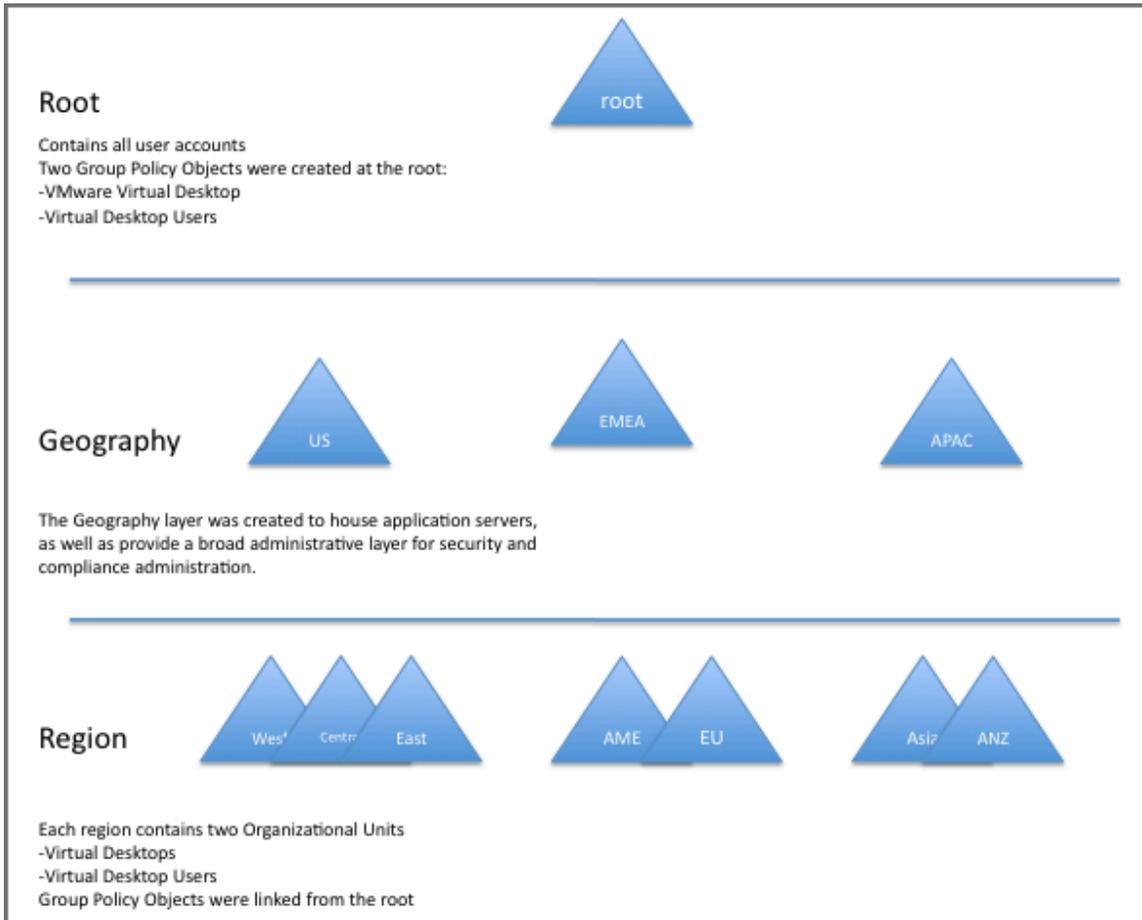
Over the past several years, Perl scripts have been created that work in conjunction with NetApp and VMware tools to automate the steps to boot VMs so that 500 to 1,000 VMs can be routinely deployed in 2 to 3 hours.

PHYSICAL LAYOUT

The design concentrates all boot infrastructures in one location. Servers and storage systems are grouped into pods that include just the necessary switching (IP and FC) to meet the needs of the pod. This makes the pods easy to replicate. In addition, it is easy to grow and scale the Kilo-Client in any dimension by adding another pod of that type. Manual setup and teardown are no longer required. Therefore, new pods can be deployed anywhere in the data center as more space is needed, allowing the data center to operate with maximum efficiency.

ACTIVE DIRECTORY

Figure 8) Active Directory.



DOMAIN STRUCTURE

This design uses the following business criteria:

- All user accounts must be in the root domain.
- Application servers must reside in as few domains as possible for simplified administration, reduced security exposure, and manageability.
- Regional administrators must have full and autonomous control over the administration of the machines, security, and policies in their sphere of responsibility.

This leads to a three-tier design containing a single forest root domain to contain the user accounts, enterprise-wide Group Policy Objects (GPOs), and the virtual desktop users' security group.

The application servers are contained in three geographically oriented sub-domains. Each geographic sub-domain contains two or three regional domains at the third and bottom tier.

TOP LEVEL

The top-level domain contains four Windows Server 2008R2 servers providing the Directory Service and DNS roles. All user accounts are contained in this domain. The exceptions are the domain and local administrator accounts in the sub-domains, and the service accounts as needed at the second-level (application server) and third-level (regional) domains.

Two GPOs are created as follows:

- **Virtual Desktops.** The Virtual Desktops GPO is created to contain the various VM settings by using the template provided from VMware, which is a recommended best practice by VMware.
- **Virtual Desktop Users.** The Virtual Desktop Users GPO is created to contain the virtual desktop user settings by using the template provided from VMware, which is a recommended best practice by VMware.

All user accounts that use the virtual desktop machines are placed in an Organizational Unit (OU) named Virtual Desktop Users. This OU contains a security group for the virtual desktop users as well as a link to the Virtual Desktop Users GPO.

GEOGRAPHICAL TIER

These domains are largely vacant because the focus of this exercise is the virtual desktops at the bottom tier exclusive of any enterprise application servers. Each of the three geographical layer domains contains two Windows Server 2008R2 servers providing Directory Service and DNS roles.

REGIONAL TIER

At the lowest level of the forest, the regional tier, each domain contains three Windows Server 2008R2 servers. Two servers contain the Directory Service and DNS roles. The third server provides the DHCP service for the region.

Each domain contains an OU named Virtual Desktops. All of the deployed virtual desktops are placed in this OU. A link is created to the Virtual Desktop Machines GPO contained in the root domain.

NETWORK

This deployment uses a network design with Cisco® Nexus 7000 switches and Cisco Nexus 5020 switches. All of Cisco's best practices are followed in the setup of the Nexus environment. For more information on configuring a Cisco Nexus environment, visit www.cisco.com.

The goal in using a Cisco Nexus environment for networking is to integrate its capabilities to logically separate public IP traffic from storage IP traffic. By doing this, the chance of issues developing from changes made to a portion of the network is mitigated.

This configuration uses Cisco Nexus 5020 switches, which support vPCs. Therefore, a logical separation of the storage network from the rest of the network is achieved while providing a high level of redundancy, fault tolerance, and security. The vPC also provides multipathing, which allows you to create redundancy by enabling multiple parallel paths between nodes and load balancing traffic where alternate paths exist.

SERVERS

FUJITSU

A portion of this deployment uses 468 Fujitsu® PRIMERGY RX200 S5s. Each of these Fujitsu PRIMERGY RX200 S5 is configured with 48GB of RAM and two 2.26 GHZ Intel® Xeon® Nehalem E5520 processors. The processors are all quad-core, 8-threaded processors delivering a total of 8 cores and 16 threads per system.

Introduction and Highlights of the S5 Generation PRIMERGY Dual-Socket Rack Servers

The PRIMERGY RX200 S5, offered in a flat rack housing of only one height unit (1U), provides a significantly improved performance, higher extendibility, greater reliability, and decreased power consumption compared with previous generation servers. This means that it can offer more for less—by delivering increased processing power and consuming less power compared with the older systems. A key part of this energy efficiency is our innovative PRIMERGY Cool-safe™ system design of the new dual S5 rack server generation. PRIMERGY Cool-safe™ is a design methodology that improves system cooling, decreases power consumption, and improves the reliability of PRIMERGY servers through the optimization of power subsystems, airflow, fan design, chassis design, and power management features. In the PRIMERGY RX200 S5, the efficiency of the hot-pluggable, redundant power supply units has been increased to 89% to improve energy use. Depending on the load, sensor-controlled, hot-pluggable, and redundant fans are used to maintain an optimal temperature in the server at all times. These fans are supported by an innovative, extremely air-permeable front panel with a honeycomb design and barrier-free, internal airflow channels inside the server.

The integrated ServerView Power Management enables the monitoring and active planning of maximum consumption loads, thus enabling cost efficiency. The result is a high-performance server, increased reliability for major components, and reduced data center cooling requirements.

The careful design of this server and the strict adherence to performance and reliability resulted in the RX200 S5 achieving a first-place score in VMware's VMmark Benchmark for dual-socket systems. On August 11, 2009, Fujitsu achieved a VMmark score of 24.20@17 tiles with a PRIMERGY RX200 S5 and VMware ESX 4.0. This score as well as the detailed results and configuration data can be seen at www.vmware.com/products/vmmark/results.html. With a score of 24.20@17 tiles, the PRIMERGY RX200 S5 was (at the time the benchmark result was published) the highest scoring 1U rack server with 8 processing cores and was rated third in the VMmark ranking for servers in the 8-core.

PRIMERGY servers with Cool-safe™ allow you to get the most performance out of today's multicore processor designs by providing a well-cooled, energy efficient, reliable, and expandable platform for server deployments and for virtualization.

TECHNICAL HIGHLIGHTS

The common design principles present across the S5 generation of PRIMERGY Dual-Socket servers are:

- Increased power efficiency due to the energy-efficient Cool-safe™ design, EPA-compliant PSU (≥ 89% energy efficiency), cable-free design of the motherboard wherever possible, and new power management features such as Power Consumption Limiting and Budgeting
- Dual-Core, Quad-Core, or Turbo Quad-Core Intel Xeon processor 5500 series with up to a maximum of 1333 MHz bus speed for DDR3 RAM and 8MB of shared TLC. These processors support the new serial QuickPath Interconnect (QPI) links and offer outstanding performance and a balanced architecture that incorporates next-generation memory and I/O technologies

- On-board dual Gb Ethernet controller with TCP/IP acceleration and iSCSI support
- Integrated Remote Management Controller (iRMC S2, IPMI 2.0 compliant) with a dedicated service LAN port and enhanced graphics, an optional iRMC S2 advanced pack offering full graphical console redirection, and a virtual remote storage feature
- Diagnostic LEDs inside the server with additional blue chassis ID LEDs at the front and rear to allow support teams to locate the server in case of failure
- Redundant and hot-plug system fans and power supply options
- Optional Trusted Platform Module (TPM) for enhanced security when storing cryptographic keys
- A choice of Integrated Customer Self-Service (CSS) modules (system LEDs only, Local Service Panel LSP with additional diagnostic LEDs, or a diagnostic LCD Display LSD). These modules show problem areas of the server and allow for the self-service of failing components locally at the server

PRIMERGY RX200 S5 SPECIFIC FEATURES

- Up to 96GB (48GB per CPU), DDR3-1066, and 1333 MHz memory with a high degree of availability provided by advanced memory protection mechanisms including SDDC hot-spare memory and memory mirroring support
- Four-port on-board SATA controller, ICH10R, for SATA variant with up to four HDDs/SSDs and an optical drive
- Optional 4/8-port modular RAID controller with SAS/RAID 0, 1, 1E (LSI1064 or 1068) or RAID 5, 6 option (LSI1078) with a 256 or 512MB RAID cache and an optional BBU in the internal PCIe Gen2 x4 slot, if at least one internal HD is configured
- Two models with support for either 8 x 2.5-inch SAS (36, 73, 146, or 300GB), or SSD drives (32 or 64GB), or 6 x 2.5-inch SAS (36, 73, 146, or 300GB), or SSD drives (32 or 64GB) plus 1 x optical drive (DVD writer or Blue-ray reader)
- Three PCIe Gen2 slots (1x x8 low profile, 1x x8 full height or low profile, 1x x4 low profile); the x4 slot is reserved for the optional modular RAID controller
- Front-VGA, 3x front USB, 3x rear USB, and 1x internal USB (a total of 7)

CUSTOMER BENEFITS

Table 8 shows the most important features of the PRIMERGY Dual S5 servers and the potential benefits for customers.

Table 8) PRIMERGY RX200 S5: If efficiency is the key decision factor.

Features	Benefits for the Customer
<ul style="list-style-type: none"> • Highly efficient power supply units (89% or more), EPA (Environmental Protection Agency) compliant • Sensor-controlled fan management • Power limiting and power budgeting • 2.5-inch hard disks with low consumption 	<ul style="list-style-type: none"> • High performance especially for efficient energy use • Fan management as required for saving energy and reducing noise levels • Individually defined limits for maximum power consumption • 2.5-inch hard disks save up to 20% energy compared with regular 3.5-inch disks

Features	Benefits for the Customer
<ul style="list-style-type: none"> • Memory mirroring option • Hot-plug redundant power supplies and fans, hot-plug hard disks • Cool-safe system design with high air throughput • Integrated iRMC S2 plus optional Advanced Pack • Modular RAID for levels 0, 1, 5, 6, 10, 50, and 60 with 256MB or 512MB RAID cache and BBU option • Individual service packages 	<ul style="list-style-type: none"> • High availability and reliability • Security level for each application scenario • Increased performance, longer lifespan for components used, optimal performance per watt ratio • Easy, fast remote management access from anywhere maximizing uptime • Low-priced, powerful data security • Tailor-made service for the respective requirements
<ul style="list-style-type: none"> • Dual-Core, Quad-Core, and Turbo Quad-Core Intel Xeon processor 5500 series • Up to 96GB state-of-the-art DDR3 main memory • 2 free PCIe Gen2 Slots, doubled I/O throughput • 2 x GbE LAN with TCP/IP acceleration • Up to 8 x 2.5 hot-plug SAS and SATA hard disks in 1U • Certified with VMware • VMware ESXi embedded via UFM 	<ul style="list-style-type: none"> • More virtual machines and applications can be used on one server • Higher application loads and extended usage options • Double I/O bandwidth compared with previous gen slots; SAN and network loads achieve optimal throughput • More than 2TB of low-priced internal hard disk storage • Full compatibility with virtualization platforms
<ul style="list-style-type: none"> • Customer self-service module (LEDs) • Switchable service LAN (shared or dedicated) • Illuminated green controls for hot-plug components • Fully extendable telescopic rails 	<ul style="list-style-type: none"> • Cost-reducing and proactive customer self-service • Physically separated service access • Easy to use with standardized labeling • Comfortable rack installation and server operation
<ul style="list-style-type: none"> • ServerView Suite is a set of proven management tools for the efficient management of physical and virtual resources throughout the entire lifecycle that provides excellent installation: <ul style="list-style-type: none"> – Stable operations – Secure updates – Exact (remote) maintenance – Easy integration in specific corporate management solutions 	<ul style="list-style-type: none"> • The key to high-level IT benefits and reduced operational and service costs: <ul style="list-style-type: none"> – Greater reliability – Lower downtimes – Improved service quality

PRIMERGY RX200 S5 PRODUCT DESCRIPTION

Chassis

The housing is designed according to the new PRIMERGY family design. User serviceable touch points are identified in green, a handy System ID card, an improved ventilation system with a honeycomb design front bezel, and unified operation panels. The improved, efficient cooling (air ventilation ratio S4 to S5 system 38.5% to 56%) avoids high temperatures, results in a lower fan speed and less noise, and is optimized to support the new Intel Xeon processors. The six redundant double fans blow directly at the hot spots such as the CPUs and the memory.

The chassis is optimized for 19-inch rack integration and is 1 height unit. Telescopic rails and cable management are offered as an option with either full or partial extraction rail kits. Nineteen-inch rack cabinets with a depth of at least 900 mm are recommended for integration.

Eight LEDs on the chassis allow the status to be checked at a glance and provide excellent serviceability. Three LEDs at the front of the operating panel indicate power/standby, hard disk activity, and Global Error (Field Replaceable Unit [FRU] and Customer Replaceable Unit [CRU]), and a fourth is used for identifying the system. This is enhanced by additional LEDs (system status, identification, LAN activity, and LAN mode) at the rear of the system.

The Customer Self Service (CSS) module for PRIMERGY RX200 S5 is integral. For certain components, the CSS module provides the opportunity for customers to do a component replacement. In the event of an error, an additional LED indicates this self-service option. The CSS module shows the category of the self-service component by using a specific LED.

The new PRIMERGY servers introduce two classes for the components. Class one is made up of components that can only be replaced by the field service. These are FRU components. The new second class consists of components that can be replaced by the customer without having to call the service. These components are called CRU components. Both components are displayed by the Global Error LED (amber FRU, yellow CRU). A second, yellow LED was introduced to give the customer a visual indicator. The yellow LED indicates that the customer can solve the problem on his/her own, and the amber LED requires a call for field service. A shining LED indicates that the component will fail soon. A blinking LED indicates that the component is defective.

Additional PRIMERGY diagnostic LEDs are integrated for the main memory, the CPU, the redundant hot-plug fans, and the power supply. The power is supplied from a hot-plug power supply unit with an output of 770 W and a redundancy option (1 + 1 x 770 W each). Six redundant hot-plug double fans are standard (5 + 1 redundancy). In addition to the LEDs, the operating panel at the front contains an on/off switch and a hidden reset and NMI switch.

Figure 9) Front view of I/O module (graphic provided by Fujitsu).

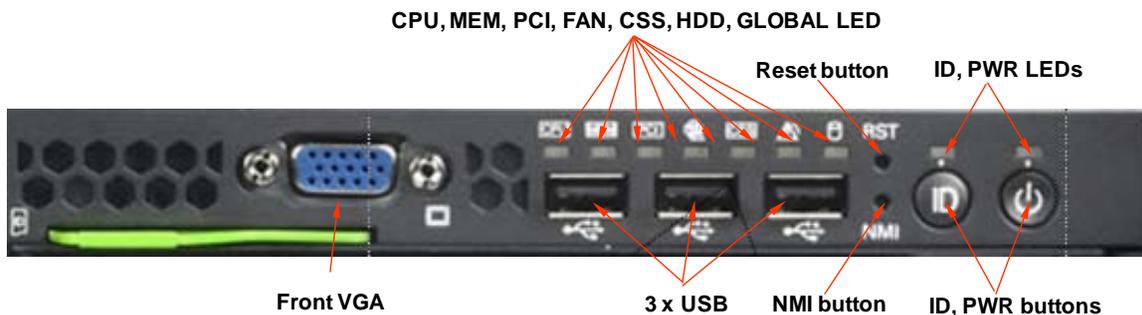
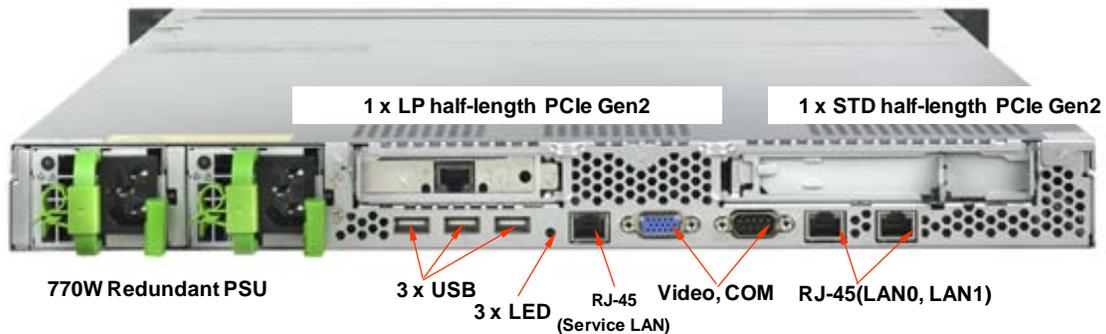


Figure 10) Rear view of I/O module (graphic provided by Fujitsu).



Motherboard

Highly efficient voltage regulators (maximum of 87% efficiency) for CPUs and memory lead to further energy savings. The measurement of voltage and current takes place directly at different board positions and is monitored and partly controlled by the Power Consumption functionality of Server Management. New Server Management features such as Power Consumption Limiting (after reaching a predefined threshold an action will be initiated: alert, graceful shutdown, or power off) and Budgeting (based on predefined policies, the system controls the power consumption by itself, preventing a given threshold from being reached) are available. Supporting energy-efficient BIOS features in conjunction with the OS and HW deliver further saving functions for CPU, PCIe devices, QuickPath Interconnect, and memory. (For more information on energy efficiency, refer to the Cool-safe™ documents on the Fujitsu Web site.)

Processors

Two processor sockets are available for up to two leading performance Dual-Core, Quad-Core, or Turbo Quad-Core Intel Xeon processor 5500 series.

All processors integrate Intel® QuickPath Technology, Intel® Intelligent Power Technology, and Intel® Virtualization Technology. Intel® VT FlexMigration, Intel® VT FlexPriority, and Intel® 64 architecture are standard on all SKUs. Higher-frequency versions of the processors support Demand-Based Switching (DBS). Quad-Core CPUs with 8MB TLC also support the new Turbo Boost mode and Hyper-Threading.

It is possible to upgrade from one CPU to a second CPU at a later time. The second CPU must be an identical type and clock rate.

A dual-processor system requires a multiprocessor operating system.

Main Memory

The main memory is designed for a maximum of 96GB registered DDR3 RDIMM RAM with ECC. In addition, Memory Scrubbing and SDDC (Chipkill) functionality is integrated. Optionally, Memory Mirroring and hot-spare memory are supported.

There are six memory slots for a maximum of 48GB registered DDR3 RAM per CPU with 8GB RDIMMs or a maximum of 12GB unbuffered (ub) DDR3 RAM per CPU available with 2GB UDIMMs.

A maximum of 96GB registered or 24GB ub RAM for two CPUs is possible. The memory area is divided into three channels per CPU with two slots per channel. Slot number 1 of each channel belongs to memory bank 1, while slot number 2 belongs to memory bank 2. A mix of registered and unbuffered modules is not allowed.

DDR3 1066 and 1333 MHz modules can be mixed depending on the operation mode that is set in the system BIOS, but they run with the slower speed.

SDDC is supported only for registered memory modules.

- The following configuration is possible in the Independent Channel Mode:
 - Each slot can optionally be equipped either with registered x4/x8 organized DDR3 modules: 2GB single rank, 4GB, or 8GB dual rank, or with unbuffered x8 organized DDR3 modules: 1GB single rank and 2GB dual rank.
- The following configuration is possible in the Mirrored Channel Mode:
 - Each memory bank can optionally be equipped with 2x 2GB single rank, 2x 4GB, or 2x 8GB dual rank registered x4/x8 organized DDR3 modules.
 - In each memory bank channel, A and B of CPU 1 or channel D and E of CPU 2 must be equipped with identical modules. Channel C of CPU 1 and channel F of CPU 2 are not equipped.
 - Channel B contains the mirrored memory of channel A of CPU 1.
 - Channel E contains the mirrored memory of channel D of CPU 2.
 - No special order codes for UDIMMs are offered for this mode; operation is technically possible, but at the user's risk.
- The following configuration is possible in the Spare Channel Mode:

- Each memory bank can optionally be equipped with 3 x 2GB single rank, 3 x 4GB, or 3 x 8GB dual rank registered x4/x8 DDR3 modules.
- Each slot of one bank must be equipped with identical modules for the Spare Channel Mode.
- Channels A and B of CPU 1 as well as channels D and E of CPU 2 contain the active memory modules.
- Channel C of CPU 1 as well as channel F of CPU 2 contain the spare module.
- No special order codes for UDIMMs are offered for this mode. The operation is technically possible, but at the user's risk.

The following also apply to the configurations:

- In Independent Channel Mode, a minimum of one memory module must be configured for each CPU. An additional memory extension can still be configured up to five times per CPU, or
- In Mirrored Channel Mode, one bank must be equipped with two modules (channel A+B for CPU 1 or D+E for CPU 2).
- Additional memory extensions can still be configured up to one time per CPU, or
- In Spare Channel Mode or Performance Mode, one bank must be equipped with three modules (channel A+B+C for CPU 1 or D+E+F for CPU 2). Additional memory extensions can still be configured up to one time per CPU.

On-board or Integrated Controllers

A four-port SATA controller included in the Southbridge ICH10R is available for SATA Raid 0/1 for the 4x 2.5" hot-plug SATA HDDs or SSDs only. One additional SATA channel for 1x DVD is integrated in the ICH10R also. A modular SAS RAID 0/1/1E controller with IME (Integrated Mirroring Enhanced) for Windows and Linux® support is optional based on the LSI1064 chipset with 4 ports or LSI1068 with 8 ports. IME is a RAID 1 implementation, which can also operate with an odd number of hard disks. As an alternative, the optional 8-port modular SAS RAID 5, 6 controller based on the LSI 1078 chipset can be chosen with a 256 or 512MB RAID cache and an optional BBU.

The system has a Dual Gb LAN controller with TCP/IP acceleration and iSCSI support (Intel Zoar) that can be configured in failover mode (AFT support under Microsoft Windows).

The integrated Remote Management Controller S2 on-board server management controller comes with, among other things, enhanced graphics and an additional 10/100Mb Service LAN that is also switchable to one of the Gb LAN ports.

The graphics controller is built into the iRMC S2. Out-of-band management is possible in the case of defective hardware or software with the iRMC S2 (IPMI 2.0 compliant). The optional iRMC Advanced Pack as a graphical console redirection and the connection of local storage to the remote server (license option) are also available.

- Interfaces at the rear are:
 - 1 x Serial RS-232-C (9-pin), usable for iRMC or system or shared, 1 x VGA (15-pin), 3 x USB 2.0, 2 x LAN RJ45, 1x Service LAN
- Interfaces at the front are:
 - 3 x USB 2.0, 1 x VGA (15-pin), usable as alternative to the rear side VGA connector
- Internal interfaces are:
 - 1 x USB 2.0 (UHCI) with 480Mb and a UFM connector (for use of a 2GB UFM with boot-enabled Hypervisor [VMware ESXi]), no USB wakeup, 4x SATA for 4x SATA HDD/SSD, 1x SATA for DVD (only usable with 6 x 2.5-inch HDD + DVD option)

A new Trusted Platform Module for the safer storage of keys is available as an option. Export restrictions eventually must be considered.

One USB Flash Module (UFM with 2GB) can be inserted by using an internal USB connector.

I/O Slots

The PRIMERGY RX200 S5 rack server can be extended flexibly by means of two external PCIe Gen2 slots and one internal:

- 1x PCIe x8 Gen2 (standard or low-profile cards, 170-mm length)
- 1x PCIe x8 Gen2 (low-profile cards, 170-mm length)
- 1x PCIe x4 Gen2 (internal for modular RAID controller only)

CISCO UNIFIED COMPUTING SYSTEM

Two hundred twenty-three Cisco Unified Computing System™ (UCS) B200 M1 blades are used for a portion of this deployment. Each Cisco UCS B200 M1 blade has two 2.26 GHZ Intel Xeon Nehalem E5520 processors. The processors are all quad-core, 8-threaded processors delivering a total of 8 cores and 16 threads per system. Forty-eight of the Cisco UCS B200 M1 blades have 48GB of RAM, and 208 have 24GB of RAM.

The Cisco UCS is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10-Gigabit Ethernet (GbE) unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

The main system components include:

- **Compute.** The system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon 5500 Series processors. The blade servers offer patented Cisco Extended Memory Technology to support applications with large datasets and to allow more VMs per server.
- **Network.** The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates what are currently three separate networks: LANs, SANs, and high-performance computing networks. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing power and cooling requirements.
- **Virtualization.** The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access.** The system provides consolidated access to both SAN storage and network attached storage (NAS) over the unified fabric. Unifying storage access means that the Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet, and iSCSI, providing customers with choice and investment protection. In addition, administrators can preassign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management while helping to increase productivity.
- **Management.** The system integrates all the system components, enabling the entire solution to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager provides an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application-programming interface (API) to manage all system configurations and operations. Cisco UCS Manager helps increase IT staff productivity, enabling storage, network, and server administrators to collaborate on defining service profiles for applications. Service profiles are logical representations of desired physical configurations and infrastructure policies. They help automate provisioning and increase business agility, allowing data center managers to provision resources in minutes instead of days.

Working as a single, cohesive system, these components unify technology in the data center. They represent a radical simplification in comparison to traditional systems, helping simplify data center operations while reducing power and cooling requirements. The system amplifies IT agility for improved

business outcomes. The Cisco UCS components, shown in Figure 11, include, from left to right, fabric interconnects, blade server chassis, blade servers, and, in the foreground, fabric extenders and network adapters.

Figure 11) Cisco UCS (graphic provided by Cisco).

Cisco UCS Unified Computing System – Blade Servers



SYSTEM COMPONENTS

Fabric Interconnect

The Cisco UCS 6100 Series Fabric Interconnects are a core part of the Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6100 Series offers line-rate, low-latency, lossless 10-GbE, and Fibre Channel over Ethernet functions. The Cisco UCS 6100 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6100 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting the unified fabric, the Cisco UCS 6100 Series provides both LAN and SAN connectivity for all blades within its domain.

From a networking perspective, the Cisco UCS 6100 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10GbE on all ports, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10GbE unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the interconnect. Significant TCO savings come from an FCoE-optimized server design that can consolidate network interface cards (NICs), host bus adapters (HBAs), cables, and switches.

The Cisco UCS 6100 Series is also built to consolidate LAN and SAN traffic onto a single, unified fabric, saving the capital and operating expenses associated with multiple parallel networks, different types of adapter cards, switching infrastructure, and cabling within racks. Fibre Channel expansion modules in the interconnect support direct connections from the Cisco UCS to existing native Fibre Channel SANs. The capability to connect FCoE to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling.

The Cisco UCS 6100 Series is equipped to support the following module options:

- Ethernet module that provides 6 ports of 10GbE using the SFP+ interface
- Fibre Channel plus Ethernet module that provides 4 ports of 10GbE using the SFP+ interface, and 4 ports of 1/2/4-Gbps native Fibre Channel connectivity using the SFP interface
- Fibre Channel module that provides 8 ports of 1/2/4-Gbps native Fibre Channel using the SFP interface for transparent connectivity with existing Fibre Channel networks
- Fibre Channel module that provides 6 ports of 1/2/4/8-Gbps native Fibre Channel using the SFP or SFP+ interface for transparent connectivity with existing Fibre Channel networks

Fabric Extenders

Cisco UCS 2100 Series Fabric Extenders bring the unified fabric into the blade server enclosure, providing 10GbE connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

The Cisco UCS 2100 Series extends the I/O fabric between the Cisco UCS 6100 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic Fibre Channel over Ethernet fabric to connect all blades and chassis together.

The fabric extender is similar to a distributed line card; therefore, it does not do any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity and enabling the Cisco UCS to scale to many chassis without multiplying the number of switches needed. This reduces TCO and allows all chassis to be managed as a single, highly available management domain.

The Cisco 2100 Series also manages the chassis environment (the power supply and fans as well as the blades) in conjunction with the fabric interconnect. Therefore, separate chassis management modules are not required.

The Cisco UCS 2100 Series Fabric Extenders fit into the back of the Cisco UCS 5100 Series chassis. Each Cisco UCS 5100 Series chassis can support up to two fabric extenders, enabling increased capacity as well as redundancy.

The UCS 2104XP Fabric Extender has four 10GbE, FCoE-capable, Small Form-Factor Pluggable Plus (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2104XP has eight 10GbE ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 80Gbps of I/O to the chassis.

Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco UCS, delivering a scalable and flexible blade server chassis for today and tomorrow's data center while helping to reduce TCO.

Cisco's first blade server chassis offering, the Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half- and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92% efficient and can be configured to support nonredundant, N+1 redundant, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for the Cisco UCS 2104XP Fabric Extenders.

A passive midplane provides up to 20 Gbps of I/O bandwidth per server slot and up to 40 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40GbE standards.

B200 M1 Blade Server

The Cisco UCS B200 M1 Blade Server is a half-width, two-socket blade server. The system uses two Intel Xeon 5500 Series processors, up to 96GB of DDR3 memory, two optional hot-swappable Small

Form Factor serial attached SCSI (SAS) disk drives, and a single mezzanine connector for up to 20 Gbps of I/O throughput. The server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

B250 M1 Blade Server

The Cisco UCS B250 M1 Extended Memory Blade Server is a full-width, two-socket blade server featuring Cisco Extended Memory Technology. The system supports two Intel Xeon 5500 Series processors, up to 384GB of DDR3 memory, two optional SFF SAS disk drives, and two mezzanine connections for up to 40 Gbps of I/O throughput. The server increases performance and capacity for demanding virtualization and large-dataset workloads with greater memory capacity and throughput.

Converged Network Adapter

The Cisco UCS M71KR-Q QLogic Converged Network Adapter (CNA) is a QLogic-based Fibre Channel over Ethernet mezzanine card that provides connectivity for Cisco UCS B-Series Blade Servers in the Cisco UCS. Designed specifically for the Cisco UCS blades, the adapter provides a dual-port connection to the midplane of the blade server chassis.

The Cisco UCS M71KR-Q uses an Intel 82598 10GbE controller for network traffic and a QLogic 4-Gbps Fibre Channel controller for Fibre Channel traffic, all on the same mezzanine card. The Cisco UCS M71KR-Q presents two discrete Fibre Channel HBA ports and two Ethernet network ports to the operating system.

The Cisco UCS M71KR-Q provides both 10GbE and 4-Gbps Fibre Channel functions using drivers from QLogic, providing:

- Risk mitigation through compatibility with current QLogic adapter-based SAN environments and drivers
- Reduced TCO through consolidation of LAN and SAN traffic over the same mezzanine card and fabric, reducing the overall number of NICs, HBAs, cables, and switches
- Integrated management with Cisco UCS Manager

EXTENDED MEMORY ARCHITECTURE

Modern CPUs with built-in memory controllers support a limited number of memory channels and slots per CPU. The need for virtualization software to run multiple OS instances demands large amounts of memory, and that, combined with the fact that CPU performance is outstripping memory performance, can lead to memory bottlenecks. Even some traditional nonvirtualized applications demand large amounts of main memory: Database management system performance can be improved dramatically by caching database tables in memory, and modeling and simulation software can benefit from caching more of the problem state in memory.

To obtain a larger memory footprint, most IT organizations are forced to upgrade to larger, more expensive four-socket servers. CPUs that can support four-socket configurations are typically more expensive, require more power, and entail higher licensing costs. Cisco Extended Memory Technology expands the capabilities of CPU-based memory controllers by logically changing the geometry of main memory while still using standard DDR3 memory. The technology makes every four DIMM slots in the expanded memory blade server appear to the CPU's memory controller as a single DIMM that is four times the size. For example, using standard DDR3 DIMMs, the technology makes four 8GB DIMMS appear as a single 32GB DIMM.

This patented technology allows the CPU to access more industry-standard memory than ever before in a two-socket server:

- For memory-intensive environments, data centers can better balance the ratio of CPU power to memory and install larger amounts of memory without having the expense and energy waste of moving to four-socket servers simply to have a larger memory capacity. With a larger main-memory footprint, CPU utilization can improve because of fewer disk waits on page-in and other I/O

operations, making a more effective use of capital investments and a more conservative use of energy.

- For environments that need significant amounts of main memory but that do not need a full 384GB, smaller-sized DIMMs can be used in place of 8GB DIMMs. The result is that two 4GB DIMMS are typically less expensive than one 8GB DIMM.

CISCO UCS MANAGER

Cisco UCS Manager provides unified, centralized, embedded management of all software and hardware components of the Cisco UCS across multiple chassis and thousands of VMs.

By enabling better automation of processes, Cisco UCS Manager allows data center managers to achieve greater agility and scale in their server operations while reducing complexity and risk. Cisco UCS Manager provides flexible role- and policy-based management using service profiles and templates.

Cisco UCS Manager provides focused integration with industry-leading systems management partners and their solutions and provides easy adoption of the Cisco UCS through the use of existing IT staff skills, tools, and processes. Cisco UCS Manager enables custom development with an extensive XML API that exposes several thousand points of integration and provides increased system visibility and control.

The Cisco UCS Manager has the following benefits and features:

- A centralized management interface that integrates the entire set of Cisco UCS components
- Role-based administration that builds on existing skills and best practices and supports collaboration across disciplines
- Policy-based management that shifts IT's focus from maintenance to strategic initiatives
- Service profiles for a fast, consistent, compliant, and accurate configuration
- Service profile templates that help to develop consistent policies within the system for a given service or application
- Physical and VM mobility through just-in-time provisioning
- High-availability configuration when two fabric interconnects are used
- Scalability to support up to 320 servers per manager instance
- XML API, which facilitates coordination with third-party provisioning tools
- Autodiscovery, which allows Cisco UCS Manager to detect, inventory, manage, and provision any system component that is added or changed

DYNAMIC DESKTOP VIRTUALIZATION PROVISIONING AND DEPLOYMENT

IT managers expect to dynamically provision virtual desktops. For example, an IT manager in a financial institution brings servers in the data center online as employees arrive in the morning. However, when employees go home in the evening, these servers are repurposed to run analytics useful for the next day. Similar patterns can be found in other industries; for example, in large retail environments, the shift is from sales transactions to analytics and back.

For this kind of scenario to work, the server infrastructure needs to be highly agile and offer dynamic provisioning capabilities. The Cisco UCS facilitates dynamic provisioning with service profiles. The Cisco UCS service profile abstracts compute and network attributes such as identity (universal user ID [UUID], MAC address, and worldwide name [WWN]), I/O configurations, firmware versions, and BIOS boot order so that they can all be quickly applied to a hardware blade when a server needs to be deployed.

For example, the IT manager can predefine a service profile template called VDI-Service-Profile. This template has all the compute and network attributes necessary for a server to run the desktop virtualization workload. When an additional server is required, the IT manager simply initiates a new service profile from the template. The server has all the LAN and SAN properties, such as the VLAN and VSAN, required to deploy the desktops.

With this approach, the newly added server is now available to the desktop virtualization manager's compute resource pool in minutes instead of weeks. Hence, service profiles improve business agility by quickly aligning compute, storage, and network resources with the changing virtual desktop workload requirements.

Furthermore, the current server infrastructure has multiple management points for managing the various elements in the system. Separate management planes exist for managing firmware, identities, chassis, networking, and so forth, each adding operational overhead and, in some cases, license costs. Cisco UCS Manager is a unified, highly available device manager that manages all the elements of the Cisco UCS from a single management point.

VALUE PROPOSITIONS OF UCS FOR DESKTOP VIRTUALIZATION

- Service profiles templates and pools
- Different service profiles for different workload or user profiles
- Concepts of pools and templates of UCSM fit nicely with VMware View's idea of Desktop pools
- Unified fabric with high I/O bandwidth helps in scaling data-intensive workloads
- Chassis configured based on bandwidth (20/40/80 Gbps), not just connectivity
- Extended memory blade helps in running a large number of memory-intensive desktops
- Current results indicate that Microsoft Windows 7 desktops (with a large memory footprint) need a large memory for scaling

STORAGE

This deployment uses ten FAS3170 clusters. The flexibility and expandability of the FAS3170 along with its ability to run a large number of virtual desktops make this storage controller a good choice for the deployment. The FAS3170 allows a customer to consolidate diverse datasets onto a unified storage platform and provides simultaneous block and file services for business and technical applications. Also, the FAS3170 has a compelling file- and transaction-based performance with high bandwidth, 64-bit architecture, and the latest I/O technologies including support for 8Gb FC connectivity. Additionally, deduplication, FlexClone, and thin provisioning help eliminate stranded storage and drastically increase storage efficiency in a virtual desktop environment. This highly efficient storage utilization dramatically reduces your consumption of raw storage, power, cooling, and space.

OVERVIEW OF THE LOGICAL STORAGE CONFIGURATION

- Controllers A and B each hosts 2,500 VMs created using NetApp VSC 2.0.
- The VM swap file (vswap) datastore on storage controllers A and B hosts the VM swap file for all 2,500 VMs on each controller.

Table 9) FAS controller A.

Desktop Virtualization Infrastructure Component	Number
Total volumes on FAS controller A	13 (including root volume)
FlexClone gold volume	1
FlexClone volumes	9
Volume for vswap datastore	1
Volume to host template VM (to be used as the source for creating all the NetApp VSC 2.0-based VMs)	1

Table 10) FAS controller B.

Desktop Virtualization Infrastructure Component	Number
Total volumes on FAS controller B	13 (including root volume)
FlexClone gold volume	1
FlexClone volumes	9
Volume for vswap datastore	1
Volume to host template VM (to be used as the source for creating all the NetApp VSC 2.0–based VMs)	1

Table 11) NetApp controller physical configuration.

NetApp System Components	Number and/or Type
Disk shelves required	6
Size and speed of hard disk in shelves	300GB or 450GB at 15K RPM
Disk shelf type	DS14
Dual-port 10GB Ethernet NIC	2 (one per controller)
Flash Cache	2 (one per controller)
NFS licenses	2 (one per controller)
CIFS licenses	2 (one per controller)
FlexClone licenses	2 (one per controller)
FlexShare [®] licenses	2 (one per controller)

VMWARE VSPHERE 4.1

VMware vSphere is the industry's most complete virtualization platform, with infrastructure services that transform IT hardware into a high-performance shared computing platform, and application services that help IT organizations deliver the highest levels of availability, security, and scalability.

VMware vSphere 4.1 Enterprise Plus is used as the virtualization and management platform for this deployment because of the robustness of the platform. The individual, specific vSphere components for this project are vSphere ESX 4.1, VMware vCenter 4.1, and VMware View Client 4.1 along with many of their respective features. This section provides a high-level description of each component and feature used in the architecture.

VMWARE ESX 4.1

Almost 700 Cisco and Fujitsu servers are used in this virtual desktop deployment for the guest VMs. Each server is provisioned with a NetApp LUN using NetApp FlexClone. This allows rapid provisioning of each server with a very small storage footprint. Each host is then configured using proprietary scripts to give it its own personality. After the provisioning and customization of each host are performed, Host Profiles are applied to standardize and manipulate the configuration of the environment. The ESX 4.1 feature that may be most beneficial for this project is VMware ESX's memory overcommitment. This feature allows us to scale much larger than the physical memory footprint of the physical servers. A second ESX feature that helps simplify the complexity of such a large deployment is the selection of NFS as the storage protocol. Using NFS as the storage protocol is beneficial because it does not require a separate FC

network. We are able to use a 10GB Nexus IP network to provide enough bandwidth and logical separation of the storage and other related traffic without separate, distinct infrastructures. NFS also increases visibility into the storage consumption and deduplication within the environment.

VMWARE VCENTER SERVER AND CLIENT 4.1

This environment uses twenty-one 64-bit VMware vSphere vCenter servers installed on 64-bit Windows 2008 R2 servers to manage the environment. Most vCenter servers are limited to 30 hosts and 2,500 VMs. Each vCenter server has one data center and two clusters within each data center with 15 hosts each. VMware HA and DRS are then configured on each cluster to provide a highly available and scalable environment.

VMWARE HA

The environment has VMware HA enabled within the VMware clusters to make sure that, if a host fails, the failing VMs would restart on a surviving node. Additional hosts are not added within the architecture to accommodate failures. One purpose is to stress the limits of the architecture. Adding additional hosts for failover capacity decreases the number of VMs per hosts unless a N+# architecture is used, and it does not truly test the limits of the hardware platforms.

VMWARE DISTRIBUTED RESOURCE SCHEDULER (DRS)

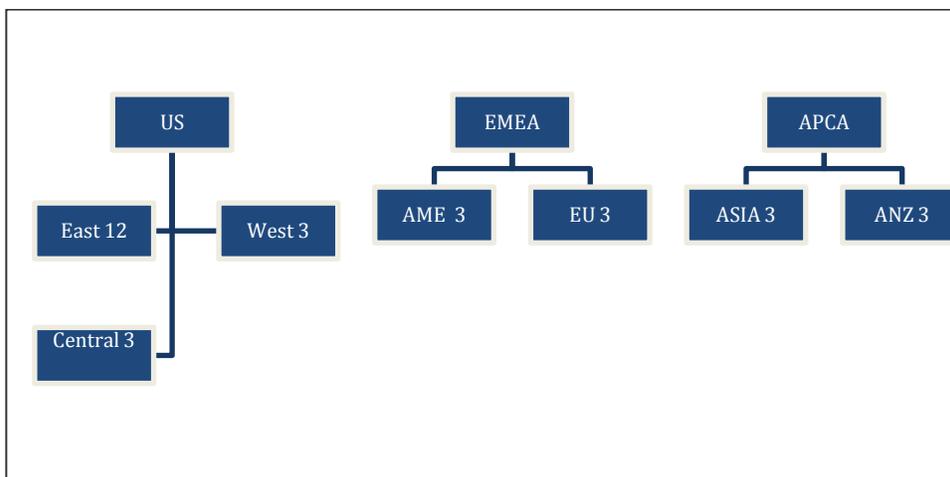
Using VMware DRS allows us to distribute the load within each cluster more evenly. This is important when deploying desktop virtualization because DRS helps provide an excellent end-user experience. Using DRS eliminates the chance of multiple people running a single host out of physical resources. VMware's DRS algorithm calculates the appropriate VMs to move to a different host to evenly distribute resources among all nodes in the cluster.

VMWARE VMOTION AND STORAGE VMOTION

VMware VMotion enables the live migration of running VMs from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity. VMotion is a core component of DRS and its ability to evenly distribute the workload across ESX servers.

Storage VMotion enables the migration of VM files from one storage datastore to another without service interruption. This component of the architecture is necessary for many reasons. First, Storage VMotion can be used to migrate VMs that require either more or fewer resources. If a user or group of users needs to move to a different storage controller within the enterprise, Storage VMotion can eliminate the desktop downtime required to move the VM.

Figure 12) Storage VMotion.



HOST PROFILES

Host profiles are used in this architecture to standardize and distribute configuration changes to all hosts within a vCenter environment. This allows us to easily mount NFS datastores to all hosts in the cluster, make and distribute firewall settings, and set and configure DNS servers and any other configuration details.

VMWARE VNETWORK DISTRIBUTED SWITCHES

Virtual Distributed Switches are used due to the large number of hosts within this environment. These switches facilitate both the management of the logical networking components and the creation and running of scripts for standard vswitches. Often, this architecture allows us to make changes only once for a group of 30 servers when new requirements for networking components (VLANS) and vNetwork Distributed Switches are implemented. Previous methods involve either writing scripts or performing each step on each host.

VMWARE VIEW 4.5

This architecture provides a standard, scalable design by using a POD strategy. Each POD consists of three VMware View connection servers to support 5,000 desktops. A connection server can support up to 2,000 virtual desktops. As an option in this design, an additional broker can be configured for the region in order to have redundancy and fault tolerance. The overall design contains 10 PODs with a total of 30 view connection servers.

In each POD, the incoming requests are split across View Connection Server instances. Each 5,000-seat region contains 3 connection servers, and each 10,000-seat region contains 5 connection servers.

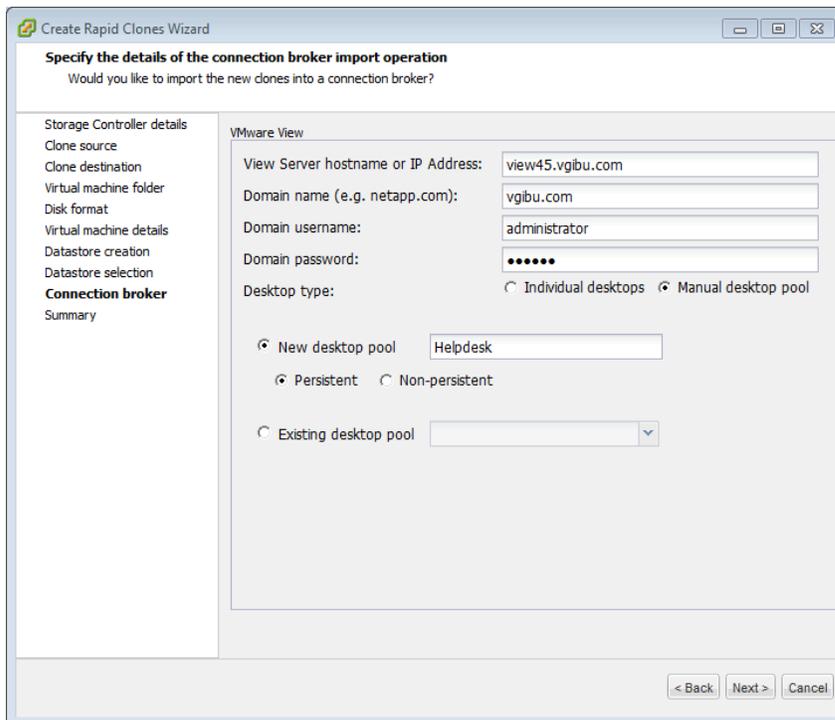
Table 12 shows the steps to set up View Connection Servers.

Table 12) View 4.5 setup.

Step	View 4.5 Setup
1	Use Windows 2008 R2 64-bit OS
2	Prepare Active Directory including creating an OU for View Desktops, groups for View Users, and User account for vCenter Server
3	Install View Connection Server
4	Install the View Connection Server License Key
5	Add vCenter Server instances to View Manager
6	Add a Database and Database User for View Events and configure the Event Database

The NetApp vCenter plug-in and the Virtual Storage Console 2.0 provide an automated method to provision units of the storage resources pools as datastores in a dynamic means by the VMware administrators. This provisioning model configures all storage access to follow NetApp recommended best practices in terms of availability, access control, and storage path management. VSC 2.0 also has a Provisioning and Cloning Plug-in, and can automatically create VMs and import into VMware View.

Figure 13) Create Rapid Clones Wizard.



Based on the limitation of each VMware View pool supporting 500 VMs, we are using VSC 2.0 to create 500 VMs and automatically import VMs into VMware View. We are using manual desktop pools. This pool type provides assignment of multiple users to multiple desktops, with only one active user on a desktop at a time. These types of desktops must be manually precreated using VMware full clones or tools with space-efficient VM provisioning capabilities.

The manual desktop pool supports two types of user assignments:

- **Dedicated-Assignment.** Users are assigned a desktop that can retain all of their documents, applications, and settings between sessions. The desktop is statically assigned the first time the user connects and is then used for all subsequent sessions.
- **Floating-Assignment.** Users are not assigned to particular desktops and could be connected to a different desktop from the pool each time they connect. Also, no data persistence of profile or user data occurs between sessions without using third-party software or roaming profiles.

WINDOWS 7

We are using Microsoft Windows 7 as the virtualized desktop platform because many organizations consider Windows 7 deployments in a virtual desktop environment. When used in conjunction with Windows Server 2008 R2 Group Policy Objects, Windows 7 proves to be highly customizable, flexible, and stable.

The basic image configuration consists of a Windows 7 32-bit image configured as follows:

- 2GB of RAM
- 40GB thin-provisioned hard disk assigned as drive C:
- LSI SCSI controlled
- VMXnet3 NIC
- No floppy drive

- CD/DVD set to Client Device
- Video card memory set to 128MB
- VMware Tools and VMware View Agent installed

The following applications are removed:

- Indexing
- Media Center
- Windows DVD Maker
- Tablet PC Components
- Windows Gadget Platform
- Windows Search

The following features and services are disabled because they are not applicable to the environment:

- Bit Locker
- Disk Defragmenter
- Home Group Services
- IPv6
- Security Center
- Themes
- Windows Defender
- Windows Firewall
- Windows Update

The following registry setting changes are made:

- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Terminal Server\WinStations RDP-Tcp set to 0x00000004
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk TimeoutValue=REG_DWORD:0x000000be

Other miscellaneous configuration options are set as follows:

- Disable IPv6
- Disable NetBIOS over TCP/IP
- Set background to Solid Color
- Set screensaver to None
- Set Sounds to No Sounds
- Set Windows 7 Basic Theme
- Open Windows Internet Explorer 8 and set homepage to default internal Web site
- Disable messages regarding virus protection
- Disable Check for solutions to problem reports
- Disable Automatic Computer Maintenance
- Disable Allow users to browse for troubleshooters
- Disable Allow troubleshooting to begin immediately when started
- Set Visual Effects to Best Performance
- Change Power Settings to High Performance with no sleep timer
- Turn off System Protection on drive C

- Set No GUI Boot through msconfig.exe
- Disable Windows Updates
- Disable System Notifications regarding Windows Firewall
- Disable automatic defragmentation schedule

After the image is configured, applications such as Microsoft Office are installed, and the Windows 7 VM is licensed and activated. As a last step of configuring the base image, temporary files are deleted and the machine powers off.

GROUP POLICY CONFIGURATION

Loopback processing of GPOs is enabled to support the combination of user and computer GPOs. All of the virtual desktops are placed in a Virtual Desktops organizational unit within the regional domain in which they are placed. The Virtual Desktops GPO is linked from the root domain to each regional Virtual Desktops OU.

GPO SETTINGS

VMware also provides ADMX files to set several computer group policies, which are beyond the scope of this document. These policies include settings for:

- Windows System
- View Agent configuration
- Windows Components including Internet Explorer

VMware also provides ADMX files to set several Active Directory user group policies, which are beyond the scope of this document. These policies include settings for:

- Internet Explorer
- Personalization
- Desktop

A complete list of these group policies will be included in the “VMware View 4.5 Windows 7 Optimization Guide,” which will be available from VMware at a future date.

USER WORKLOAD

VMware provides the user workload, which is called RAWC (Desktop Reference Architecture Workload Code). The RAWC workload runs on a Windows 7 or Windows XP guest operating system and is executed on each desktop VM. The RAWC workload has a set of functions that performs operations on common desktop applications including Microsoft Office, Internet Explorer, Adobe® Reader®, McAfee® Virus Scan, Windows Media® Player, and PKWARE® PKZIP®. The applications are called randomly and perform operations that mimic those of a typical desktop user, including open, save, close, minimize, and maximize windows; view an HTML page; insert text; insert random words and numbers; conduct a slide show; view a video; run a virus scan; compress files; and send and receive e-mails.

The RAWC workload has various test variables that can be configured to increase the load or adjust user behavior, including a start delay for creating “boot storms” and density (delay between application operations resulting in applications under test to run together faster, number of e-mails created and sent, and typing speed).

RAWC also supports the use of Active Directory Groups for large-scale enterprise testing. Based on the VM's Active Directory Group membership, a workload can be configured specifically for that group. This allows VMs to run varying workloads to simulate a more realistic large-scale work environment.

USER WORKLOAD PROFILES

Each Windows 7 VM is equipped to run a workload that simulates typical user behavior using an application set commonly found and used across a broad array of desktop environments. The workload has a set of randomly executed functions that perform operations on a variety of applications. Several other factors can be implemented to increase the load or adjust the user behavior, such as the number of words per minute that are typed and the delay between applications being launched. The workload configuration used for this validation includes Microsoft Outlook, Microsoft Word, Microsoft Excel, Microsoft PowerPoint, Internet Explorer, Adobe Acrobat, McAfee Virus Scan, and PKZIP. During the execution of the workload, multiple applications are opened at the same time, and windows are minimized and maximized as the workload progresses, randomly switching between each application. Individual application operations that are randomly performed include:

- **Microsoft Outlook.** Open, minimize, maximize, and close this application. Send e-mail messages.
- **Microsoft Word.** Open, minimize, and close this application. Write random words and numbers, and save the modifications.
- **Microsoft Excel.** Open, minimize, and close this application. Write random numbers, insert and delete columns and rows, copy and paste formulas, and save the modifications.
- **Microsoft PowerPoint.** Open, minimize, and close this application. Conduct a slide show presentation.
- **Adobe Acrobat Reader.** Open, minimize, and close this application. Browse through the pages in a PDF document.
- **Internet Explorer.** Open, minimize, and close this application. Browse through a page.
- **McAfee AntiVirus.** Real time scanning.
- **PKZIP.** Open and close this application. Compress a large file.

Based on the think time and words per minute used for this validation, this workload could be compared to that of a high-end task worker or a lower-end knowledge worker.

WYSE TECHNOLOGY

Wyse Technology is used as the main connection point into the environment. Each thin client is configured to access the Windows 7 VMs inside the VMware View 4.5 environment. To represent the variety of user environments found in the real world, multiple types of end-point devices from Wyse are used during this exercise. The following hardware and software are configured to the environment and are used for specific access to the virtual machines.

- **Wyse P20.** The Wyse P class delivers workstation-class performance for advanced applications including CAD, 3D solids modeling, video editing, and advanced worker level office productivity. Based on a hardware PCoIP engine, this stateless zero client does not require a local operating system. Unlike the more flexible software-based implementations, this dedicated hardware PCoIP engine delivers the highest level of display performance available. The Wyse P20 is VMware View Certified.
- **Wyse C Class.** Inside every compact Wyse C class lies a punchy VIA C7 1 GHz processor and an advanced hardware media accelerator delivering rich multimedia and video playback with up to 1080p HD resolution. Outside, you'll find advanced DVI-I display connectivity for one or two displays, four free USB 2.0 ports, PS/2 ports for the keyboard and the mouse, and Gigabit Ethernet and optional WiFi (b/g/n) support. The Wyse C Class devices running Windows Embedded, Wyse ThinOS, or Linux are VMware View Certified.
- **Wyse R Class.** Created with performance and expandability in mind, the new Wyse R class joins the uniquely integrated range of Wyse Thin Computing hardware and software solutions. The new Wyse R class features high-performance central and graphics processors coupled with full support for multiple screens, and the best PC experience available on a thin computing platform. The Wyse R Class devices running Windows Embedded and Linux are VMware View Certified.

- **Wyse PocketCloud.** Wyse PocketCloud allows IT administrators and end users to securely access their PC and virtual desktops remotely through the convenience of their iPhone®, iPod touch®, iPad™, or Google™ Android based devices. For IT administrators who manage large numbers of computers, Wyse PocketCloud allows quick and easy access to an end-user machine from anywhere on the road. For remote users who need secure access to their computers, Wyse PocketCloud is a simple and effective way of taking their computer with them in the palm of their hand. Wyse PocketCloud accesses any physical or virtualized infrastructure that supports the Microsoft RDP protocol. In addition, Wyse PocketCloud is VMware View Certified and supports the VMware View connection broker in direct and tunneling modes. This software is used in conjunction with an Apple® iPad to have quick and easy access to the virtual machines. This software uses the RDP protocol to connect into the VMware View 4.5 environment. Wyse PocketCloud is VMware View Certified.

4 REFERENCES

- VMware View Reference Architecture
www.vmware.com/resources/techresources/1084
- VMware View
www.vmware.com/products/view/
- VMware vSphere 4
www.vmware.com/products/vsphere/
- Workload Considerations for Virtual Desktop Reference Architectures
www.vmware.com/files/pdf/VMware-WP-WorkloadConsiderations-WP-EN.pdf
- TR-3705: NetApp and VMware View Solution Guide: Best Practices for Design, Architecture, Deployment, and Management
<http://media.netapp.com/documents/tr-3705.pdf>
- TR-3749: NetApp and VMware vSphere Storage Best Practices
<http://media.netapp.com/documents/tr-3749.pdf>
- TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide
<http://media.netapp.com/documents/tr-3505.pdf>
- TR-3747: NetApp Best Practices for File System Alignment in Virtual Environments
<http://media.netapp.com/documents/tr-3747.pdf>
- Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 4.1
www.cisco.com/en/US/docs/switches/datacenter/sw/4_1/nx-os/interfaces/configuration/guide/if_nxos.html
- Cisco Nexus 5000 Series Switch CLI Software Configuration Guide
www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli_rel_4_0_1a/CLIConfigurationGuide.html
- Wyse P20
www.wyse.com/products/hardware/zeroclients/P20/index.asp
- Wyse C Class
www.wyse.com/products/hardware/thinclients/C90LE/index.asp
- Wyse R Class
www.wyse.com/products/hardware/thinclients/R90L/index.asp
- Wyse PocketCloud
www.wyse.com/products/software/pocketcloud/index.asp

5 ACKNOWLEDGEMENTS

- Eric Forgette, Software Engineer, NetApp
- George Costea, Software Engineer, NetApp
- David Klem, Reference Architect, NetApp

- Wen Yu, Sr. Technical Alliance Manager, VMware
- Fred Schimscheimer, Sr. Technical Marketing Manager, VMware
- Ravi Venkat, Technical Marketing Engineer, Cisco
- Christopher Reno, Reference Architect, NetApp
- Eric Johnson, Systems Engineer, NetApp
- Jonathan Davis, Systems Engineer, NetApp
- Brandon Agee, Network Systems Administrator, NetApp
- Param Desai, Director of Product Management, Wyse
- Arthur Perkins, Sr. Staff Systems Engineer, Fujitsu America

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

© Copyright 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, FlexClone, FlexShare, and SnapMirror are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Adobe Reader is a registered trademark of Adobe, Inc. Apple, iPod, iPod touch, and iPhone are registered trademarks and iPad is a trademark of Apple Inc. Cisco is a registered trademark and the Unified Computing System is a trademark of Cisco, Inc. Citrix is a registered trademark of Citrix Systems, Inc. Fujitsu is a registered trademark and PRIMERGY Cool-safe is a trademark of Fujitsu Limited. Google is a trademark of Google, Inc. IBM is a registered trademark of International Business Machines, Inc. Intel is a registered trademark and Xeon is a trademark of Intel Corporation. McAfee Virus Scan is a registered trademark of McAfee, Inc. Microsoft, Windows, Active Directory, Internet Explorer, Microsoft Office, Windows Media Player, Windows Server, and Windows XP are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. PKZIP is a registered trademark of PKWARE, Inc. VMware is a registered trademark and vCenter, VMotion, vSphere, and VMware View are trademarks of VMware, Inc. Wyse and PCoIP are registered trademarks of Wyse, Inc. Linux is a registered trademark of Linus Torvalds. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7108-0810



www.netapp.com