



THE VIRTUAL COMPUTING
ENVIRONMENT COMPANY

www.vce.com

Vblock[®] Specialized System for Extreme Applications with Citrix XenDesktop 7.1

Version 1.1

August 2014



THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." VCE MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OR MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright 2014 VCE Company, LLC. All Rights Reserved.

VCE believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Contents

Introduction	5
Key highlights	6
Vblock® Specialized System for Extreme Applications	7
Citrix XenDesktop	7
Environment profile	8
Performance validation	10
Test methodology	10
Test procedure.....	11
VSImax and Citrix XenDesktop session concurrency	11
Timing	18
Storage capacity	20
Design guidelines	21
Virtual infrastructure design	21
vSphere configuration	22
Infrastructure and management servers	23
Cluster and storage design for PVS server and desktops	23
Citrix PVS server volume design	24
Network design	24
IP network components.....	25
PVS server networking consideration	26
Trivial File Transfer Protocol (TFTP) design	26
Technology components.....	28
Vblock Specialized System for Extreme Applications	28
Storage components	29
Compute and networking components	29
Citrix XenDesktop.....	30
VMware components	31
VMware vSphere 5.....	31
VMware ESXI 5.1	31
VMware vCenter Server.....	31
Hardware and software components.....	31
Additional guidelines	33
Server sizing guidelines.....	33
vSphere settings	34
vStorage API for Array Integration (VAAI) settings	35
HBA queue depth adjustment	36

Native vSphere storage multipathing 37
Next steps 37

Introduction

Companies are under increasing pressure to deliver enterprise applications on hosted virtual desktops, with the user experience and reliability of a conventional physical desktop. Virtual desktop infrastructure (VDI) holds the promise of not only consolidating desktop computing in the data center for improved manageability, efficiency, and security, but also for supporting new services such as bring-your-own-device and desktop mobility. The success of end-user computing deployments is dependent on user experience and how responsively virtual desktops perform. User experience expectations are increasingly set based on devices that use flash memory, such as tablets and ultrabooks. Users accustomed to the performance of an ultrabook may have performance issues using a slower virtual desktop. Successful end-user computing deployments must deliver a better-than-local desktop experience and a lower cost per desktop, compared to a physical machine, and enable IT to continue using existing desktop management applications and tools.

Flash-based storage improves I/O performance and efficiency for many applications, including server and desktop virtualization, database acceleration, and cloud infrastructure. The Vblock[®] Specialized System for Extreme Applications from VCE is an enterprise and service provider-ready system that brings the benefits of the world's most advanced converged infrastructure to meet the demands of the most extreme applications. A flash-based array delivers the extraordinary performance required for scalable VDI deployments, while dramatically lowering costs and enhancing the user experience.

The Vblock Specialized System for Extreme Applications is ideal for demanding implementations, such as VDI. By leveraging the XtremIO storage system in the Vblock System, you can take advantage of several benefits for VDI deployments, such as Citrix XenDesktop. VDI offers users ready access to virtual desktops in the data center. Administrative tasks are simplified, security and data protection are enhanced, and remote access is improved.

The purpose of this paper is to demonstrate that the Vblock System meets the functional and technical requirements of a 2500-user deployment in a Citrix XenDesktop version 7.1 VDI environment with Citrix Provisioning Services (PVS) and VMware vSphere.

Note: The paper contains a solution architecture that is intended for real-world deployment and mix of desktop applications. As a result, the test metrics reflect a lower density of virtual desktops to blades than is possible in tests designed to show maximum potential VDI densities. The results outlined in this solution architecture are specific to the Vblock Specialized System for Extreme Applications. You may experience different results in your environment.

Key highlights

Extensive user-experience and operations testing, including Login VSI desktop performance benchmark testing, revealed world-class operational performance. A simple design architecture and efficient use of storage preserved ease of use at an attractive price point. The results are summarized here and further described later in this paper.

 <p>2,500 streaming desktops - deployed</p>	<p>4 hours, 25 minutes</p>	 <p>LoginVSI</p> <ul style="list-style-type: none"> 2,500 active sessions 100% session concurrency 	<p>Maximum</p> <p>Not reached</p>
 <p>Average DHCP lease ACK per desktop</p>	<p>4 seconds</p>	 <p>Front-end data store latency</p>	<p>0.6 ms</p>
 <p>2,500 streaming desktops - restart and ready</p>	<p>14 minutes</p>	 <p>Total storage footprint</p>	<p>4.6 TB</p>

Key results demonstrate that Vblock Specialized System for Extreme Applications easily supports a deployment of 2,500 Citrix XenDesktop PVS streaming hosted virtual desktops with 100% concurrency, acceptable CPU, memory, and storage use, along with acceptable application response times.

Our test results demonstrate that it is possible to deliver quality user experience at scale for every desktop, with headroom for any desktop to burst to thousands of IOPS as required to drive user productivity. This is due to the EMC XtremIO storage platform, which provides considerably higher levels of application performance and lower virtual desktop costs than alternative platforms. The high performance and simplicity of the Vblock Specialized System for Extreme Applications with the EMC XtremIO array contributed significantly to the overall success.

Vblock[®] Specialized System for Extreme Applications

The Vblock Specialized System for Extreme Applications is specifically designed to unlock the full potential of the XtremIO all-flash array. Its powerful architecture is designed to deliver the performance, reliability, and IT flexibility needed in today's enterprise data centers. VDI deployments with more than 1000 desktops are ideally suited to the power and flexibility of this Vblock System. Benefits include:

- A superior end-user experience
- Lower cost per desktop
- Complete flexibility in view deployments, allowing the use of full- or linked-clone technology interchangeably without drawbacks
- Assured project success from pilot to large-scale production
- A fast, simple method of performing high-volume desktop cloning, even during production hours

The Vblock Specialized System for Extreme Applications relies on the EMC XtremIO flash array to deliver performance in the millions of IOPS with sub-millisecond response times. This enables VDI users to boot up and log on quickly from anywhere and enjoy better access to the applications and information they need to get work done.

Citrix XenDesktop

Citrix XenDesktop delivers Windows desktops and apps as secure mobile services. Users can self-select apps from an easy-to-use "store" that is securely accessible from tablets, smartphones, PCs, Macs, and thin clients.

Provisioning Services (PVS) allows computers to be provisioned in real-time from a single shared-disk image. This frees administrators from managing and patching individual systems. Instead, all image management is done on the master image. By streaming a single shared disk image instead of copying images to individual machines, PVS enables organizations to reduce the number of disk images they manage, even as the number of machines grows, providing the efficiencies of a centralized management with the benefits of distributed processing.

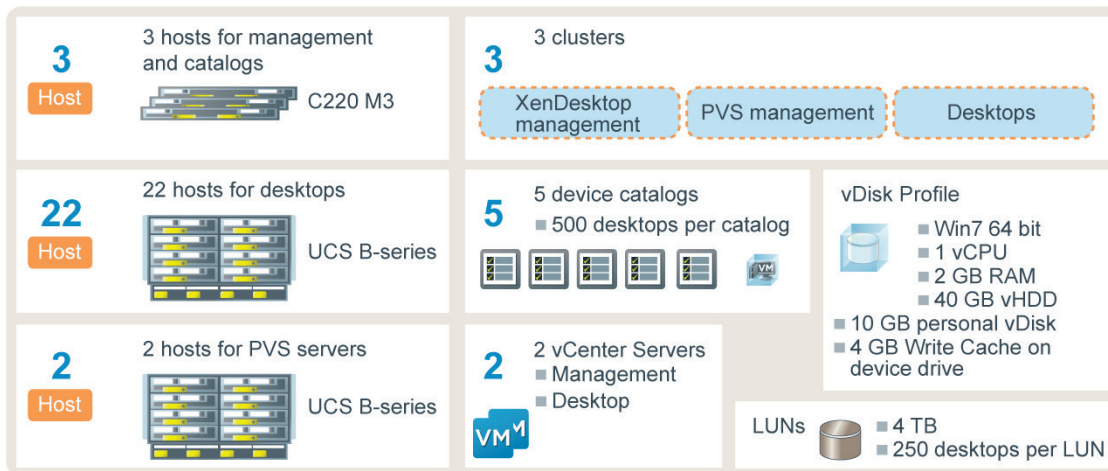
All Citrix XenDesktops components were hosted on the VCE Vblock Specialized System for Extreme Applications, which incorporates Cisco networks, Cisco UCS server blades, and EMC XtremIO Flash-based arrays to ensure the performance and responsiveness needed to support massively scalable solutions.

Environment profile

The Vblock Specialized System for Extreme Applications one-cabinet system consists of a dedicated three-chassis, 24-blade Cisco UCS blade environment. Infrastructure and management servers are isolated on three Cisco UCS C220 rack-mount servers, with storage for all virtual desktops and infrastructure servers hosted on the same EMC XtremIO storage array (single X-Brick) environment.

Note: Testing was performed on a pre-release version of the Vblock Specialized System for Extreme Applications.

The validation environment was set up as follows:

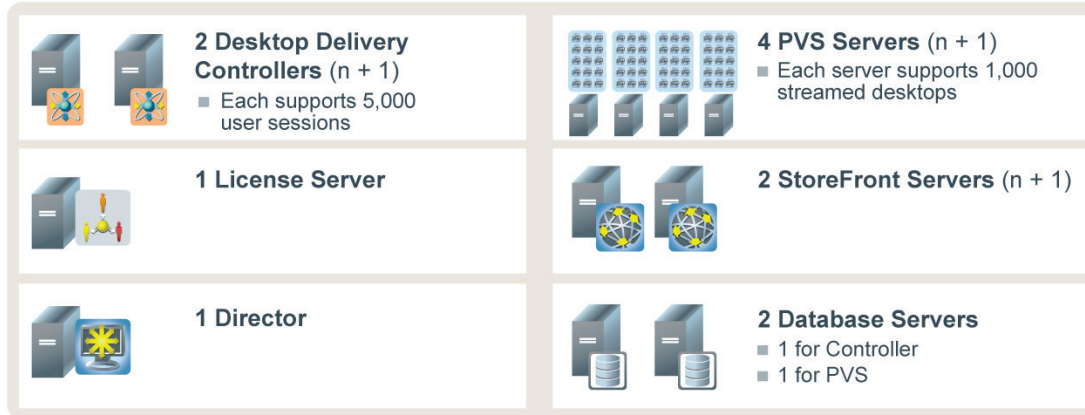


Environment design

- Three hosts for management and monitoring on the AMP for Vblock Specialized System for Extreme Applications (Cisco UCS C220 M3 servers)
- One 22-node vSphere cluster for streaming persistent desktops on three Cisco UCS Blade Chassis, 24 Cisco UCS B200 M3 Blade Servers (16-core, 256GB or RAM, 16-core Intel E2680, 2.7 GHZ, 256GB of RAM)
- Two hosts for Citrix Provisioning Servers
- Three clusters: one for XenDesktop management, one for PVS management, and one for desktop management
- 500 desktops per device catalog
- Two vCenter Servers: one for desktops and one for management servers
- Desktop Image: Windows 7 64-bit, 1 vCPU, 2GB of RAM, 40GB vHDD
- Standard LUN size: 4TB

Note: We used a separate Vblock System LoginVSI 4.0 Launcher: 8-node vSphere cluster on Cisco UCS B200 M2 Blade Server.

For validation of Citrix PVS, we used the following configuration.



Validation configuration

Performance validation

The main tool used during testing was Login VSI 4.8. We used it to generate desktop workloads and gather in-session VDI performance data. Login VSI is an industry-standard tool designed to measure the maximum capacity of VDI infrastructures by simulating unique user workloads. It provides an in-session workload representative of a typical user. We used a predefined medium workload that simulates a knowledge workload.

The simulated users work with the same applications as typical employees, such as Microsoft Word, Excel, Outlook, and Internet Explorer. The results of several testing measurements are compiled in a metric known as VSI_{max}, which quantifies the maximum capacity of VDI workloads running on a given infrastructure while delivering an acceptable user experience.

Test methodology

We configured Login VSI to run a medium workload against a Citrix XenDesktop catalog of virtual desktops, with the tests set up to log users on to virtual desktops incrementally every 30 seconds between sessions per physical host (blade).

We allocated 2 GB RAM to each hosted virtual desktop and used the Login VSI parameters in addition to the medium workload. We generated the application workloads described in the following table.

Application	Workload
Microsoft Outlook	Browse 10 messages.
Internet Explorer	Leave one instance open and continuously browse a second instance.
Microsoft Word	Measure response time for one instance. Review and edit a document in a second instance.
Microsoft Excel	Open a very large randomized spreadsheet.
Microsoft PowerPoint	Review and edit a presentation.
7-Zip	Use the CLI to zip all session output.

We measured response times for the following transactions:

Transaction	Description
Refresh	Copy a random document, a PowerPoint file, and an Outlook file from the server to a local drive.
Load	Launch Microsoft Word.
Open	Open a Microsoft Word document.
Zip	Compress a random PowerPoint file.

Test procedure

We used the following test procedure to ensure consistent results:

- 1 Cleanly started all desktop virtual machines and clients before each test.
- 2 Idled all desktop virtual machines and client launchers until start up services on the operating system settled down and memory and CPU on the launchers showed no usage.
- 3 Executed at least two Login VSI loops in each active test phase session.
- 4 Logged off all users after VSI completion.
- 5 Generated test run reports and data.
- 6 Shut down all desktop virtual machines and clients.

VSI_{max} and Citrix XenDesktop session concurrency

We performed three tests to demonstrate session concurrency.


Test 1: Single cluster, 2500 desktops, 100% session concurrency

First we tested 2,500 desktops with 100% concurrency. Test results showed **VSI_{max} not reached**.

At this level of session concurrency, host CPU resources are pushed to the point of CPU saturation and user experience starts to break down. This is not a viable, sustained level of CPU usage, because little CPU headroom is available for burst capability or failover. The sustained upper CPU usage threshold for most production implementations is 85%. Sustained CPU usage above 85% ordinarily causes a high CPU usage alarm in VMware vSphere.

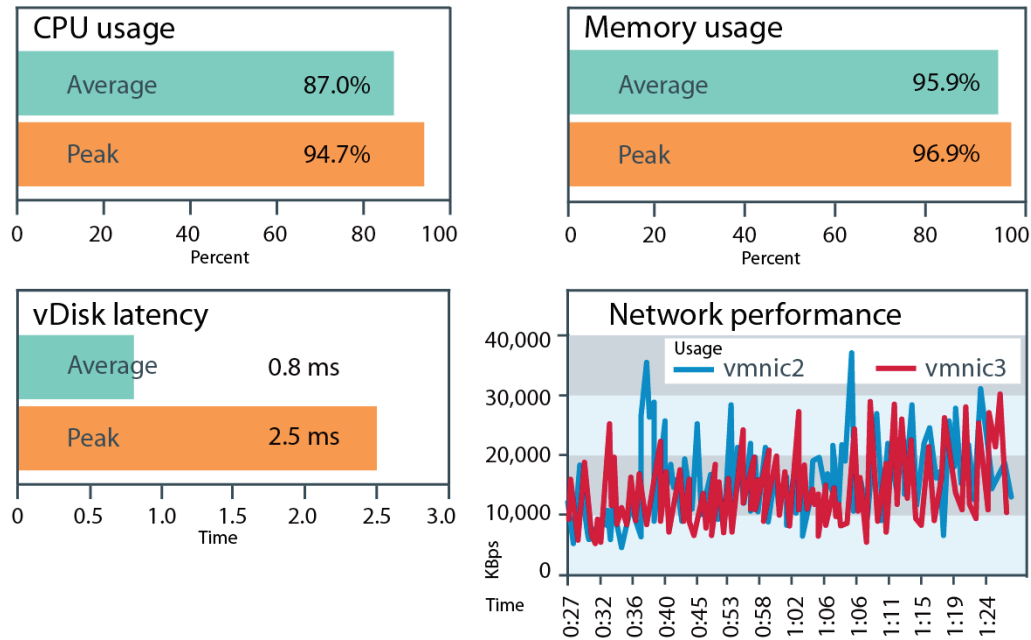
Highlights:

- Desktop access using the ICA protocol and Medium Workload (with Flash enabled by default)
- Mixed host performance (CPU saturated with 94% usage and memory usage at 95%)
- Excellent desktop performance (all vDisk latency under 0.8 ms for reads and writes)
- Peak of 9,773 IOPS on XtremIO storage
- Storage performance: 302 MB/s bandwidth and 10,198 IOPS latency



LoginVSI user experience: VSImax not reached

Single Cluster, 2500 Desktops



Host and vDisk performance



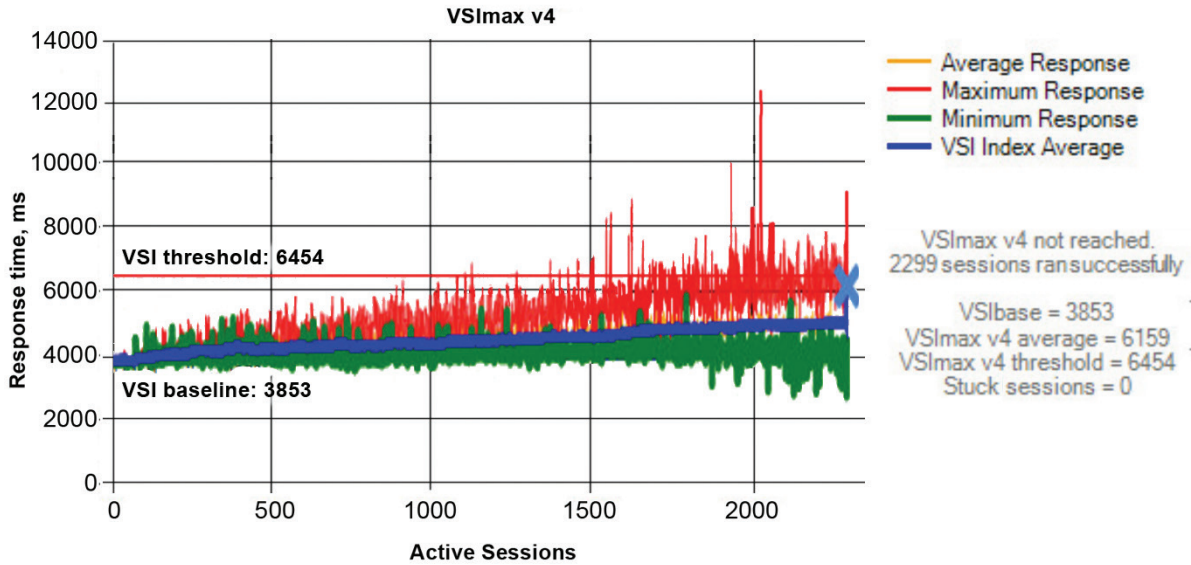
Test 1: Storage performance

Test 2: Single cluster, 2300 desktops, 88% session concurrency

Next, we tested 2,300 desktops with 88% concurrency. This showed much better host CPU resource usage. All 2,300 sessions demonstrated excellent performance, and host CPU resources remained at 88%. All 2,300 desktops were powered on and available while the Login VSI workload ran on 88% of the available desktops. Typical production VDI environments exhibit concurrent desktop usage of 80% of total available capacity.





Highlights:

- Login VSI VSImax not reached
- Better host performance (CPU usage at 88% and memory usage at 86%)
- Excellent desktop performance (all vDisk latency under 0.6 ms)
- Peak of 8,321 IOPS on XtremIO storage
- Storage performance: 280 MB/s bandwidth and 8,103 IOPS latency

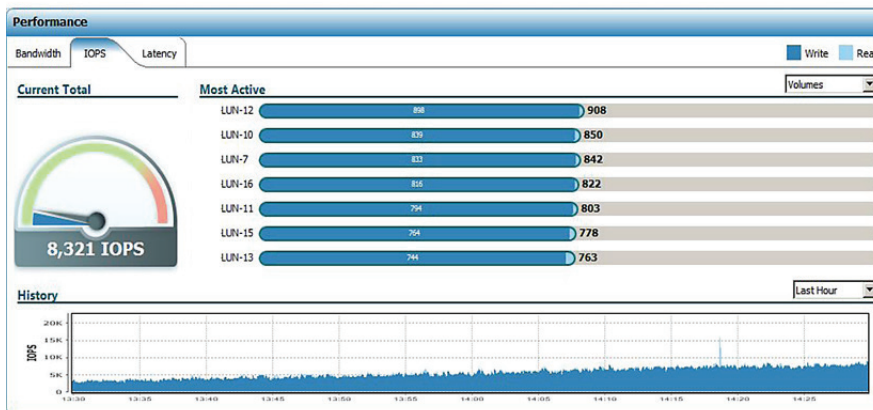
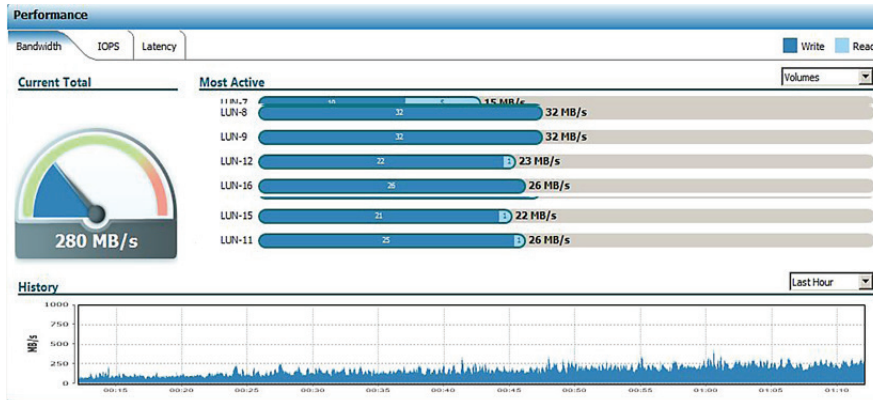


User experience: VSImax not reached

Single Cluster, 2300 Desktops



Host and vDisk performance




Test 2: Storage performance

Test 3: Single cluster, 2000 desktops, 80% session concurrency

Next, we tested 2,000 desktops with 80% concurrency. This showed much better host CPU resource usage. All 2,000 sessions demonstrated excellent performance. Host CPU resources remained average at 68% and peak at 81%. All 2,000 desktops were powered on and available while the Login VSI workload ran on 80% of the available desktops. Typical production VDI environments exhibit concurrent desktop usage of 80% of total available capacity.





Highlights:

- Login VSImax not reached
- Excellent host performance (CPU peak usage at 81% and memory peak usage at 80%)
- Excellent desktop performance (all vDisk latency under 0.5 ms)
- Peak of 6,823 IOPS on XtremIO storage
- Storage performance: 175 MB/s bandwidth and 7,092 IOPS latency

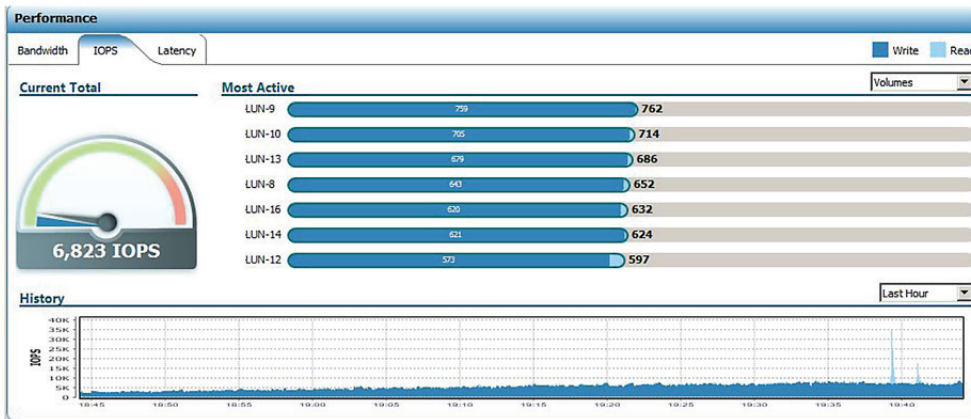
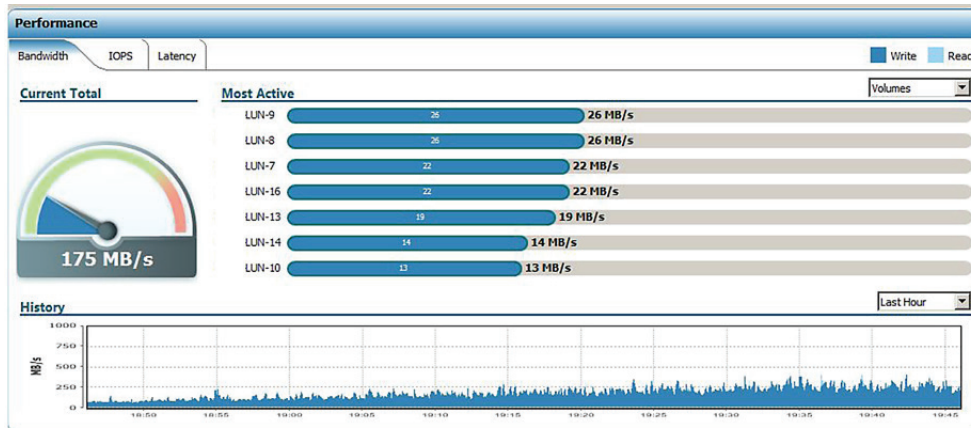


Login VSI user experience: VSImax not reached

Single Cluster, 2000 Desktops



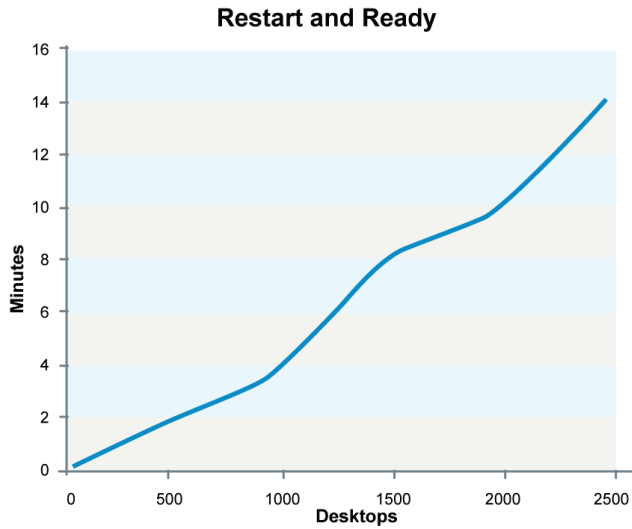
Host and vDisk performance



Test 3: Storage performance

Timing

The Citrix PVS console measured the 2,500 desktops powering on in 14 minutes.

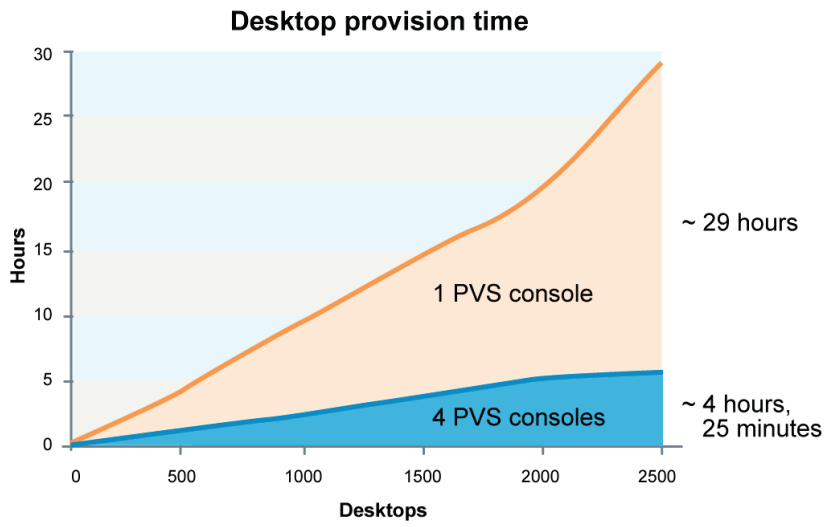


Restart and ready time for 2,500 desktops

Streaming desktop provisioning time

We observed fast, consistent, and reliable desktop provisioning, even with background workload on the compute and shared storage environment. We attempted to provision 2,500 streaming desktops using two methods to compare the PVS provisioning time.

Test	Time
2,500 desktops provisioned with one PVS Server Management Console	29 hours
2,500 desktops provisioned with four PVS Server Management Consoles	4 hours 25 minutes

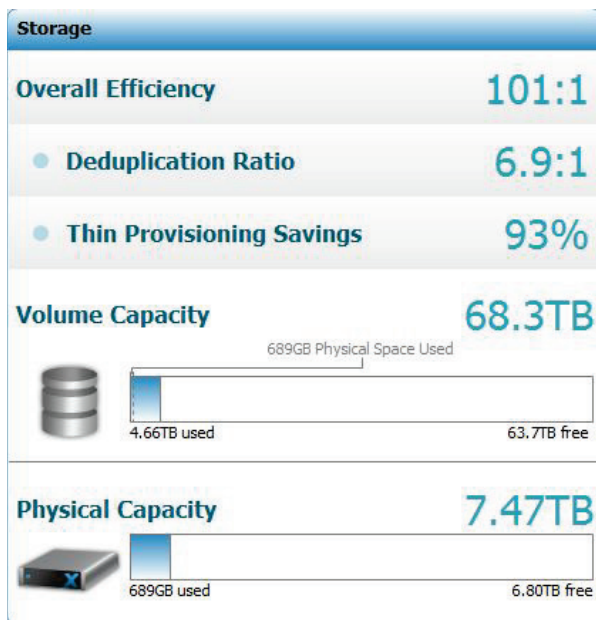


Time to provision 2,500 desktops

Storage capacity

The volume capacity configured and presented to VMware vSphere totaled 68.3 TB. The actual storage footprint for 2,500 desktops and infrastructure servers was approximately 4.6 TB. Due to the efficiency of the XtremIO inline data reduction capability, desktop and server virtual machines occupied a physical storage footprint of only 689 GB. This represents:

- Overall efficiency ratio: 101:1
- Deduplication ratio: 6.9:1
- vSphere Thin Provisioning savings: 93%

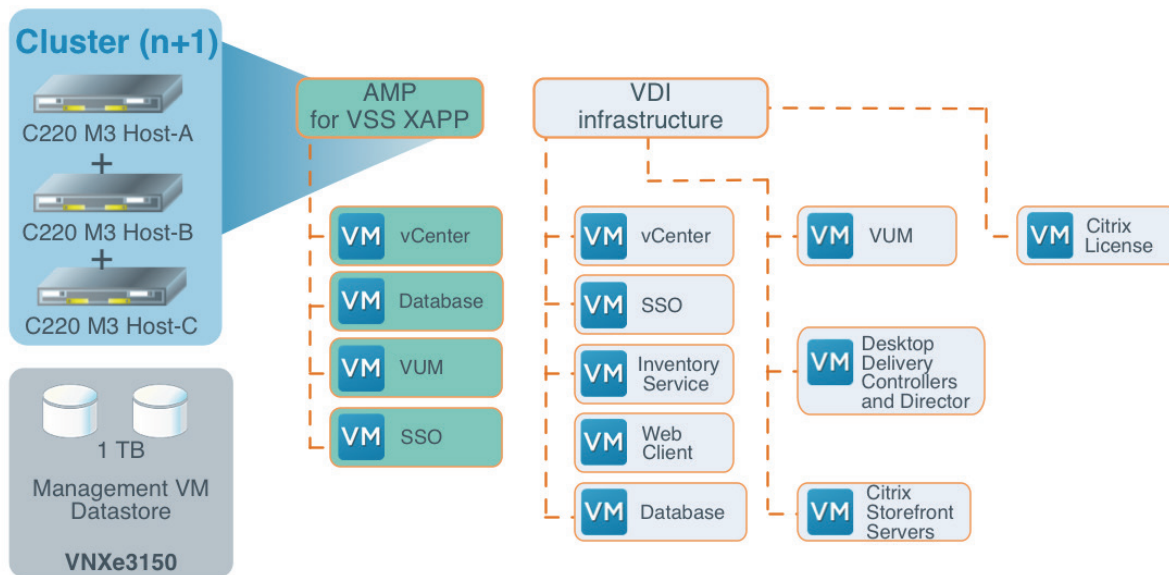


Storage footprint

Design guidelines

Virtual infrastructure design

The AMP for Vblock Specialized System for Extreme Applications manages the three management hosts and the datastore. The datastore with all virtual machines is in the EMC VNXe. VMware vCenter manages the virtual machines at the right.

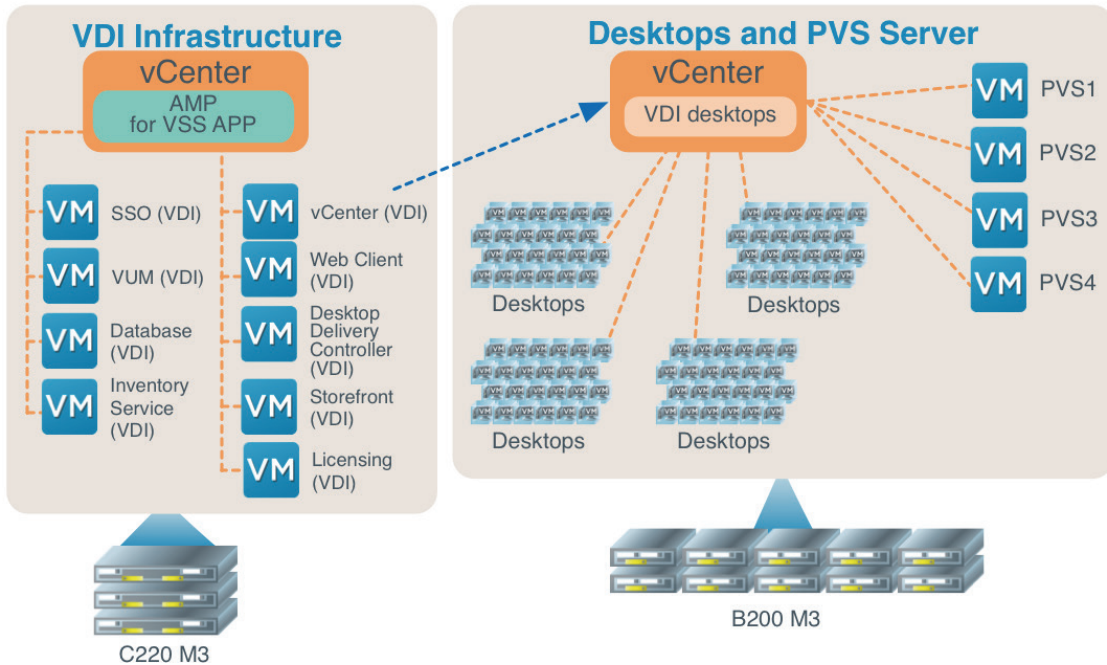


- Two vCenters reside in the AMP for VSS XAPP C220s:
- AMP for VSS XAPP – Manages all C220 ESXi hosts and AMP for VSS XAPP VMs.
 - VDI infrastructure – Manages all B200 ESXi hosts, PVS server VMs and VDI desktop VMs.

Logical design: Management

Note: A dual cabinet system has four C220 servers in a cluster.

As shown in the following diagram the Cisco UCS C220 M3 servers manage the virtual machines, and the Cisco UCS B200 M3 servers host the desktops and the four Citrix PVS servers. Citrix PVS created 2,500 virtual desktops.



Virtual infrastructure

vSphere configuration

We deployed separate virtual datacenters, one for managing and monitoring servers and one for hosted desktops. This is the standard, recommended configuration for most production Citrix XenDesktop deployments.

Infrastructure and management servers

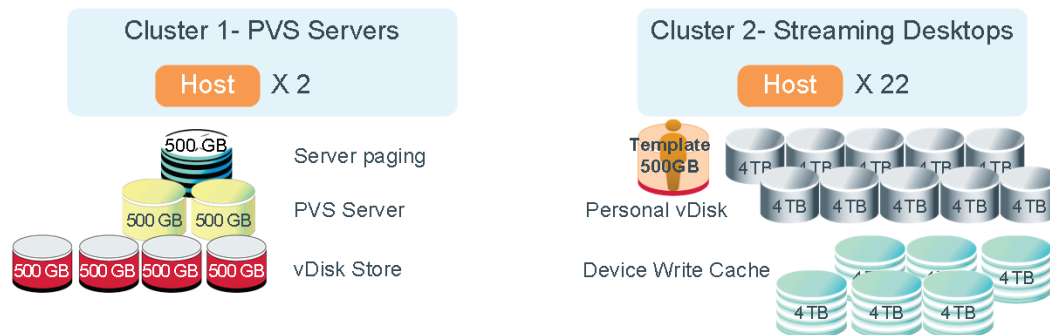
We deployed two VMware vCenter Server instances, one for desktops and one for infrastructure and management servers, in conformance with Citrix XenDesktop 7.1 architecture planning guidelines.

Note: As of Citrix XenDesktop 7.1, VMware supports managing up to 10,000 desktops with a single vCenter 5.2 server instance. This is a significant revision from prior versions.

All additional vCenter roles (inventory, SSO, vCenter) for the desktop vCenter were divested to separate servers to avoid any resource contention that might have resulted from combining roles on a busy vCenter Server.

Cluster and storage design for PVS server and desktops

For desktop virtual machines, we deployed two vSphere clusters with 24 nodes (with N + 1 HA readiness). The PVS servers used two hosts and the streaming desktops used 22 hosts.



Cluster and storage design

The first cluster has two hosts with the following LUNs:

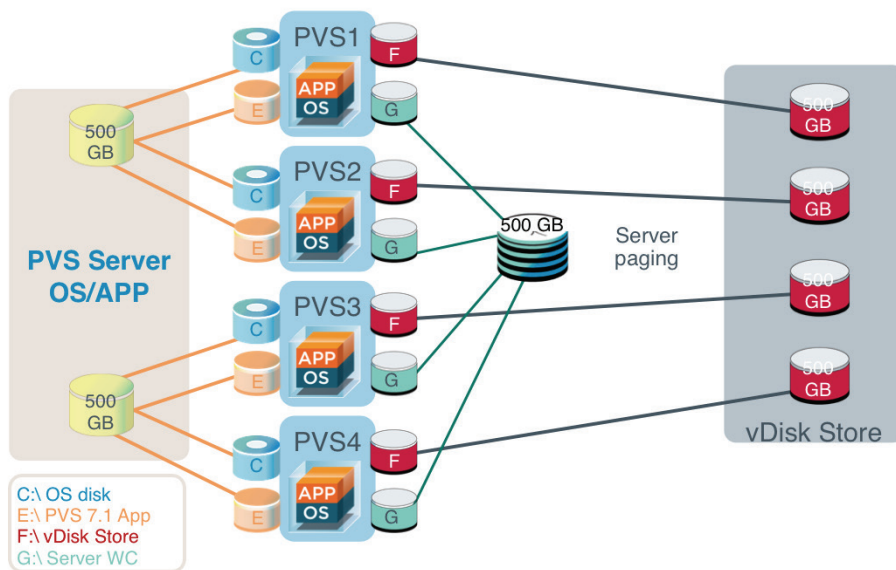
- Four 500 GB LUNs mapped for the PVS vDisk store
- Two 500 GB LUNs mapped for the PVS application
- One 500 GB LUN mapped for Windows 2012 operating system deployment for PVS server

The second cluster has 22 hosts with the following LUNs:

- Ten 4 TB LUNs mapped for personal vDisk
- Six 4 TB LUNs mapped for PVS device write cache
- One 500 GB LUN mapped to store templates and other housekeeping tasks

Citrix PVS server volume design

Each of the four virtual PVS servers has specific volumes mapped so that each can support specific workload to provide optimum performance. We used C: drive for the OS disk, E: drive for the PVS application, F: drive for the vDisk Store (mapped to each individual LUN in X-Brick), and G: drive for server paging.




Citrix PVS server volume design

Network design

Each fabric interconnect has multiple ports reserved for 10 GbE ports. These connections are formed into a port channel to the Cisco Nexus switch and carry IP traffic destined for the desktop network 10 GbE links. In a unified storage configuration, this port channel can also carry IP traffic to the X-Blades within the storage layer.

The Cisco Nexus 5548UP switches in the network layer provide 10 GbE IP connectivity between the infrastructure and the outside world. In unified storage architecture, the switches connect the fabric interconnects in the compute layer to the XtremIO X-Bricks in the storage layer.

In a segregated architecture, the Cisco Nexus 5548UP series switches in the network layer provide Fibre Channel (FC) links between the Cisco fabric interconnects and the EMC XtremIO storage array. These Fibre Channel connections provide block-level devices to blades in the compute layer.



Networking infrastructure

IP network components

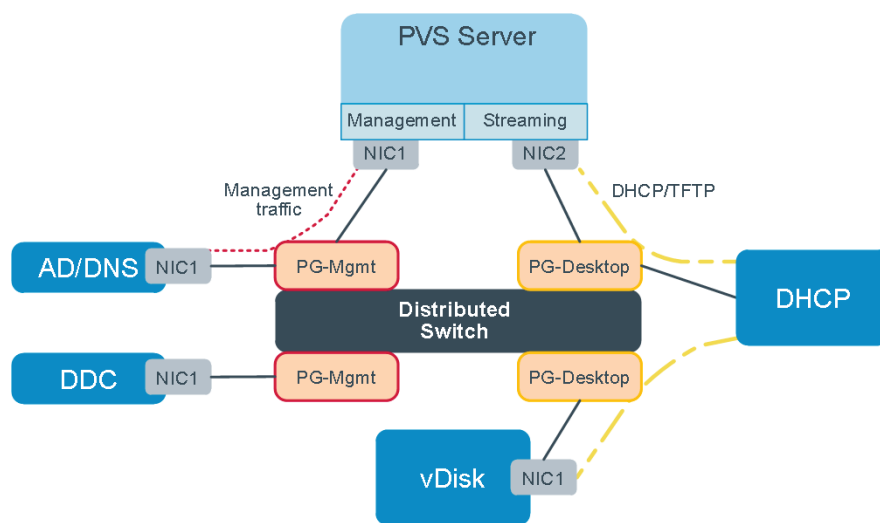
To support the Ethernet and SAN requirements in the traditional segregated network architecture, two Cisco Nexus 5548UP switches provide 10 GbE Ethernet and Fibre Channel (FC) connectivity.

The Cisco Nexus 5548UP switches have 32 integrated, low-latency unified ports, each providing line-rate 10 GB Ethernet or FC connectivity. The Cisco Nexus 5548UP switches each have one expansion slot that can be populated with a 16-port unified port expansion module.

PVS server networking consideration

Each PVS server has two NICs, each of which connects to a distributed switch inside vCenter.

- NIC1 is assigned a management VLAN network to take only management traffic through management port group (PG-Mgmt). It handles all management work, including talking to Desktop Delivery Controller (DDC) and Active Directory (AD).
- NIC2 is assigned a VLAN subnet, which is dedicated for streaming and DHCP network. It is connected with a port group on a distributed switch dedicated to desktops (PG-Desktop). This controls all DHCP/TFTP traffic between desktops and PVS servers.



Load balancing design

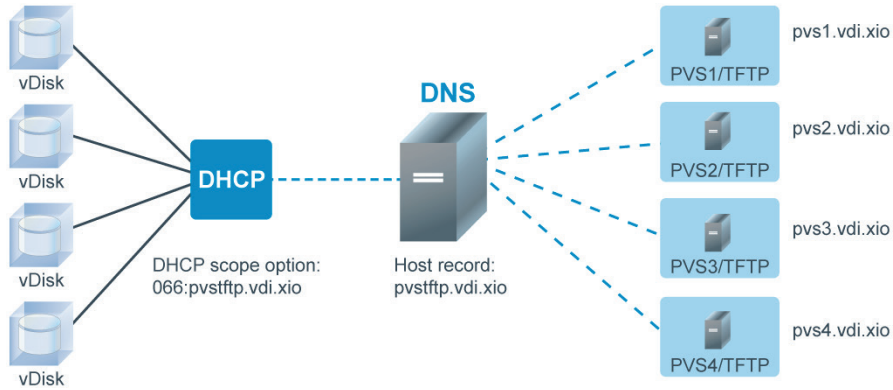
Trivial File Transfer Protocol (TFTP) design

Citrix PVS is a software streaming technology that allows servers and desktops to be provisioned and re-provisioned in real time from a single shared-disk image. The service is commonly integrated with Citrix virtualization solutions to optimize operating system delivery and management.

Often these solutions are critical to the operation of organizations and require high availability. Providing high availability for PVS requires designed each component of the service without a single point of failure across the network. The bootstrap file, named ardbp32.bin, is a key component that may be delivered to PVS target devices in order for them to communicate with PVS over the network. The bootstrap file is typically delivered via Trivial File Transfer Protocol (TFTP) services hosted on provisioning servers.

The core concept to providing high availability to TFTP is being able to deliver the bootstrap file to the PVS target device. Once the bootstrap is delivered, the target device can communicate to multiple PVS servers. There are many different ways for providing high availability for the TFTP services, but in this design, we used the DNS round robin method to do the load balancing and high availability of TFTP service.

With DNS round robin, a TFTP server DNS name is carried in DHCP option 66. The DNS server then has two or more A records defined for the domain name and cycle through the list of records in round robin fashion in response to DNS queries. This provides PVS target devices with redundancy in obtaining their bootfile.



Citrix PVS TFTP design

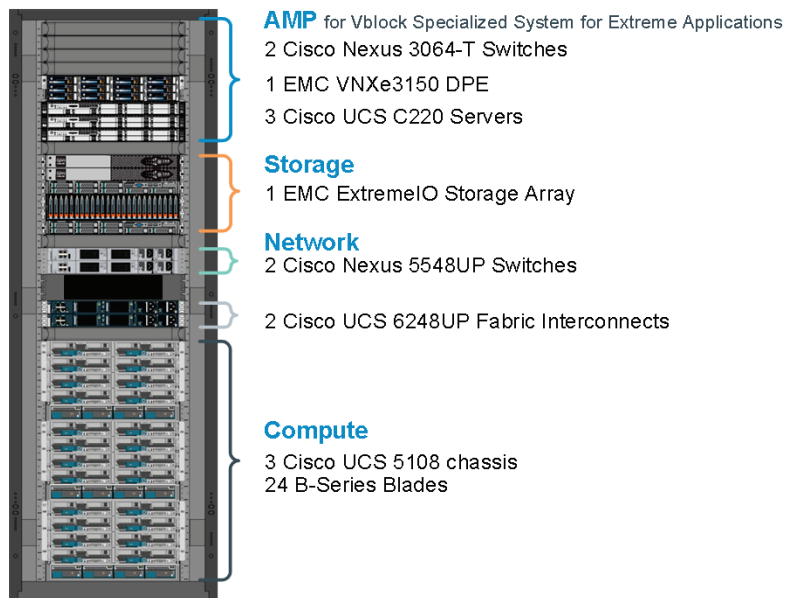
Technology components

This section summarizes key technologies used.

Vblock Specialized System for Extreme Applications

The Vblock Specialized System for Extreme Applications provides support for databases, virtualization, VDI, and extreme applications that require:

- High levels of IOPS
- Consistent < 1 ms response times
- A high volume of random access virtual machines and applications
- Virtual machines and applications with high levels of redundant data that can benefit from the powerful inline data reduction provided by XtremIO



Vblock Specialized System for Extreme Applications

Storage components

The Vblock Specialized System for Extreme Applications uses EMC XtremIO storage arrays as primary storage for Citrix XenDesktop virtual desktops and the Citrix PVS servers. The AMP for Vblock Specialized System for Extreme Applications is used for management and monitoring.

XtremIO is a highly scalable, all-flash storage array that combines multi-level cell flash with sophisticated wear leveling, data reduction, and write abatement technology to deliver extended flash endurance that makes the system both enterprise-reliable and reasonably priced.

The smallest unit of an XtremIO system is an X-Brick. Additional uniformly sized X-Bricks can be added for scalability in an XtremIO design. A simple user interface reduces the time required to design and provision storage. Enterprise resiliency is delivered through Fibre Channel connectivity, flash-specific dual-parity data protection, and storage presentation over the iSCSI protocol.

Compute and networking components

The Vblock Specialized System for Extreme Applications uses Cisco UCS blade enclosures, interconnects, and blade servers.

The Cisco UCS data center platform combines x86-architecture blade and rack servers with networking and storage access into a single system. Innovations in the platform include a standards-based, unified network fabric, Cisco Virtual Interface Card (VIC), and Cisco UCS Extended Memory Technology. A wire-once architecture with a self-aware, self-integrating, intelligent infrastructure eliminates the need for manual assembly of components into systems.

Cisco UCS B-Series two-socket blade servers deliver record-setting performance to a wide range of workloads. Based on Intel Xeon processor E7 and E5 product families, these servers are designed for virtualized applications and reduce CapEx and OpEx through converged network fabrics and integrated systems management.

Citrix XenDesktop

Citrix XenDesktop delivers Windows desktops and apps as secure on-demand services to any user, any device, anywhere. Users can self-select apps from an easy-to-use “store” that is securely accessible from tablets, smartphones, PCs, Macs, and thin clients.

Citrix XenDesktop offers a next-generation, user-centric desktop virtualization solution that provides a complete system for desktop delivery. For IT organizations, XenDesktop greatly simplifies the desktop lifecycle management process and drives down the cost of desktop ownership by separating the delivery of the desktop operating system from applications and user settings.

We used the following components:

Component	Description
Provisioning Services	Allows computers to be provisioned and re-provisioned in real-time from a single shared-disk image. Unique user settings are preserved in a personal vDisk. The PVS farm is at the top level of the PVS infrastructure. The farm provides a method for representing, defining, and managing logical groups of PVS components into sites.
StoreFront	Authenticates users to sites hosting resources and manages stores of desktops and applications that users access.
License Server	Manages product licenses for all XenDesktop components. You must create at least one license server to store and manage your license files.
Delivery Controller	Installed on servers in the data center, the Delivery Controller consists of services that communicate with the hypervisor to distribute applications and desktops, authenticate and manage user access, and broker connections between users and their virtual desktops and applications. The Controller manages the state of the desktops, starting and stopping them based on demand and administrative configuration.
Machine Creation Services	A collection of services that work together to create virtual servers and desktops from a master image on demand, optimizing storage utilization and providing a pristine virtual machine to users every time they log on. It connects vCenter to the ESXi hosts using the vSphere API, which then directs the ESXi host to build the hosted virtual desktops from the master image virtual machine and to create, start, stop, and delete virtual machines.
Director	A web-based tool that enables IT support and help desk teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users.
Software and tools	<ul style="list-style-type: none"> ▪ PVS Agent, the software installed in the virtual disk that allows Provisioning Services to control the machine. ▪ Citrix Receiver, which provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user’s devices, including smartphones, tables, and PCs. ▪ HDX Technology, a set of capabilities that delivers a high-definition desktop virtualization user experience.

VMware components

VMware vSphere 5

When using VMware vSphere 5 as the hypervisor for XenDesktop, the vSphere 5 components are built on an infrastructure incorporating VMware ESXi 5 hosts connected to local and shared storage. Shared storage can be on a network file system (NFS), Internet Small Computer System Interface (iSCSI), or Fibre Channel SAN, as dictated by the enterprise.

VMware ESXi 5.1

VMware ESXi delivers the base hypervisor functions for the VMware vSphere environment that hosts the virtual desktops.

VMware vCenter Server

VMware vCenter Server provides the infrastructure to manage multiple ESXi hypervisors as a single infrastructure cluster. The vCenter Server allows administrators to configure resource clusters and manage storage and high-availability functions across the environment. In a XenDesktop configuration, VMware vCenter can be configured as a virtual machine running on the ESXi infrastructure or on a physical server, based on high-availability requirements.

Hardware and software components

The following table lists the hardware used to validate this solution.

Layer	Hardware	Quantity
Compute	Cisco UCS C220 M3 rack server	3
	Cisco UCS 5108 Blade Server chassis	3
	Cisco UCS B200-M3 Blade Server	24
Network	Cisco Nexus 3064-T switch	2
	Cisco UCS 6248UP Series Fabric Interconnect	2
	Cisco Nexus 5548UP Series IP switch	2
Storage	EMC VNXe 3150 Series Unified Storage System	1
	EMC X-Brick Storage	1

Note: Citrix XenDesktop was tested on a single cabinet system. The VCE Vblock Specialized System for Extreme Applications provides either single or dual cabinet solutions.

The following table lists the software used to validate this solution.

Software	Version
Citrix XenDesktop	7.1
Provisioning Server	7.1
VMware vSphere	5.1
Microsoft SQL Server	2008 R2 Standard
Microsoft Windows server operating system	2008 R2 Enterprise
Microsoft Windows desktop operating system	Windows 7 Enterprise Edition 64-bit
Microsoft Server	2012 R2

Additional guidelines

Server sizing guidelines

All server resources were sized according to the current best practices from VMware and are listed in the following table:

Server role	vCPU	RAM (GB)	Storage (GB)	Operating system	Notes
SQL database server	2	6	100	Windows Server 2008 6(E) 64-bit R2	For VDI workload
vCenter Server – VDI desktops	4	10	60	Windows Server 2008 (E) 64-bit R2	
vCenter SSO – VDI	2	4	40	Windows Server 2008 (E) 64-bit R2	
vCenter Inventory – VDI	2	6	60		
vCenter WebClient – VDI	2	4	40	Windows Server 2008 (E) 64-bit R2	
vCenter Update Manager – VDI vCenter	2	4	120	Windows Server 2008 (E) 64-bit R2	For VDI vCenter
SQL database for Citrix Desktop Delivery Controller	4	6	C: 40 E:100	Windows Server 2008 (E) 64-bit R2	
SQL database for Citrix PVS	4	6	C:40 E:100	Windows Server 2008 (E) 64-bit R2	
Desktop Delivery Controller 1	4	4	C:40 E:50	Windows Server 2012 R2	
Desktop Delivery Controller 2	4	4	C:40 E:50	Windows Server 2012 R2	
StoreFront 1	4	4	C:40 E:50	Windows Server 2012 R2	
StoreFront 2	4	4	C:40 E:50	Windows Server 2012 R2	
License Server	2	4	C:40 E:50	Windows Server 2012 R2	
Director	1	4	C:40 E:50	Windows Server 2012 R2	

Server role	vCPU	RAM (GB)	Storage (GB)	Operating system	Notes
PVS 1	4	16	*	Windows Server 2012 R2	
PVS 2	4	16	*	Windows Server 2012 R2	
PVS 3	4	16	*	Windows Server 2012 R2	
PVS 4	4	16	*	Windows Server 2012 R2	

* For more information, see the Cluster and storage design for PVS server and desktops section.

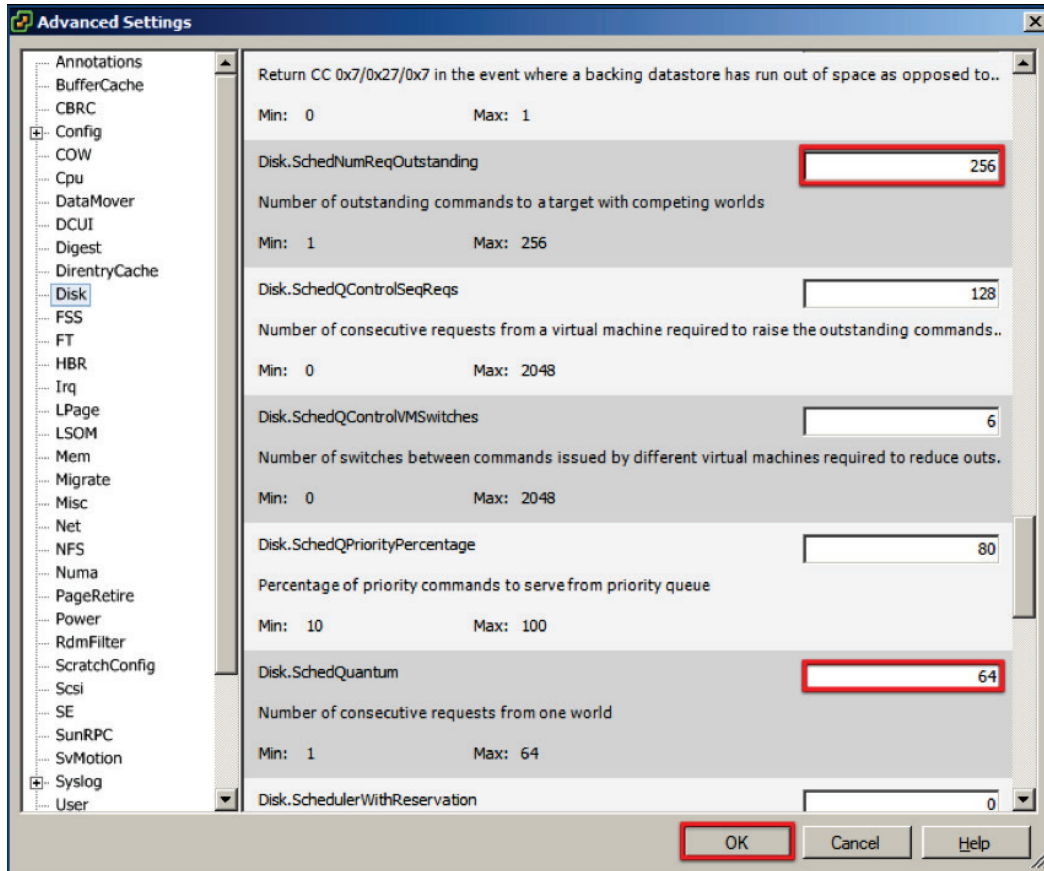
vSphere settings

Modest setting changes are required in vSphere to support EMC X-Brick storage systems:

- Adjust the number of outstanding storage target commands to 256.
- Change the number of consecutive storage commands to 64.
- Change the HBA queue depth settings to 256 for each host's HBAs.
- Change the native storage multipath policy in vSphere to Round Robin.

We configured ESXi to work with the EMC XtremIO storage array by performing the following steps.

- 1 Connect to the VMware ESX host via the vSphere Client (directly or through the vCenter Server).
- 2 Click the **Host** icon in the left pane and click **Advanced Settings**.
- 3 Click the **Disk** section.
- 4 Find the **Disk.SchedNumReqOutstanding** parameter and change it to the Max value (256).
- 5 Find the **Disk.SchedQuantum** parameter and change it to the Max value (64).
- 6 Click **OK** to apply the changes.



vStorage API for Array Integration (VAAI) settings

VAAI is a vSphere API that lowers desktop provisioning time from hours down to minutes by performing vSphere operations such as virtual machine provisioning and cloning within an array that supports VAAI. The XtremIO storage system fully supports VAAI. VAAI is enabled by default in vSphere version 5.x. No further action is required for VAAI to be used with XtremIO storage.

HBA queue depth adjustment

We adjusted the queue depth by performing the following steps:

- 1 Connect to ESX host shell as root.
- 2 Verify which HBA module is currently loaded by entering one of the following commands:
 - a For Qlogic: **esxcli system module list | grep qla**
 - b For Emulex: **esxcli system module list | grep lpfc**For example, from a host with a Qlogic HBA:
- 3 To adjust the HBA queue depth, run one of these two commands:
 - a For Qlogic:
 - b For Emulex:
- 4 Reboot the ESX host.
- 5 Connect to the ESX host shell as root.
- 6 Run the following command to confirm the queue depth adjustment:

```
# esxcli system module list | grep qla qla2xxx true true
```

```
esxcli system module parameters set -p ql2xmaxqdepth=256 -m qla2xxx
```

```
esxcli system module parameters set -p lpfc0_lun_queue_depth=256 -m lpfc820
```

```
# esxcli system module parameters list -m <driver>
```

For example, from a host with a Qlogic HBA with queue depth set to 256:

```
# esxcli system module parameters list -m qla2xxx | grep ql2maxqdepth ql2maxqdepth  
int 256 Maximum queue depth to report for target devices
```

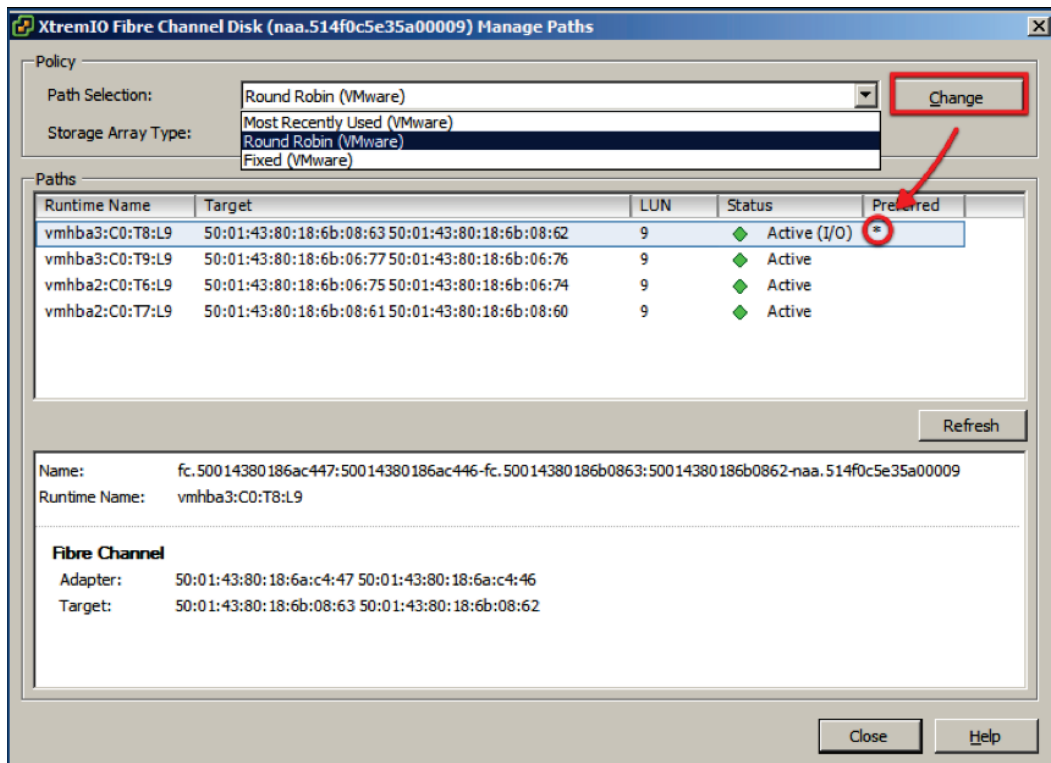
For further information about adjusting HBA queue depth with ESX, refer to [VMware Knowledge Base article 1267](#).

Native vSphere storage multipathing

XtremIO supports native multipathing technology, which is part of the VMware vSphere suite.

For the best performance, set the native Round Robin path selection policy on the XtremIO volumes presented to ESXi to ensure optimal distribution and availability of load among I/O paths to the XtremIO storage.

The following screenshot illustrates how to modify the path selection policy from the default Fixed path selection policy to the Round Robin policy:



Path selection policy

Next steps

To learn more about this and other solutions, contact a VCE representative or visit www.vce.com.



www.vce.com

About VCE

VCE, formed by Cisco and EMC with investments from VMware and Intel, accelerates the adoption of converged infrastructure and cloud-based computing models that dramatically reduce the cost of IT while improving time to market for our customers. VCE, through Vblock Systems, delivers the industry's only fully integrated and fully virtualized cloud infrastructure system. VCE solutions are available through an extensive partner network, and cover horizontal applications, vertical industry offerings, and application development environments, allowing customers to focus on business innovation instead of integrating, validating, and managing IT infrastructure.

For more information, go to www.vce.com.



© 2014 VCE Company, LLC. All rights reserved. VCE, Vblock, VCE Vision, and the VCE logo are registered trademarks or trademarks of VCE Company, LLC. and/or its affiliates in the United States or other countries. All other trademarks used herein are the property of their respective owners.