



# FlashStack Converged Infrastructure Solution Design Guide for Citrix XenDesktop 7.7

February, 2016



© 2016 Pure Storage, Inc. All rights reserved. Pure Storage, the "P" Logo, and Pure1 are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. Cisco UCS Manager, Login VSI, Microsoft Windows 7, Microsoft Windows Server 2012 R2, Citrix XenDesktop, Machine Creation Service (MCS), Provisioning Server (PVS), vCenter, and vSphere are registered trademarks of Cisco, Microsoft, Login VSI, Citrix and VMware in the U.S. and other countries. The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc. 650 Castro Street, Mountain View, CA 94041

<http://www.purestorage.com>

# Contents

---

Executive Summary.....	4
Goals and Objectives .....	4
Audience .....	4
Design Guide Principles .....	5
FlashStack Introduction.....	6
Pure Storage Introduction.....	7
Infrastructure Components of the Design Guide.....	11
Design Guide Solutions Overview .....	19
VMware vSphere Configuration and Tuning .....	23
Citrix XenDesktop 7.7 Configuration and Tuning.....	26
Microsoft Windows 7 Desktop Configuration .....	28
Desktop Testing Tool – Login VSI 4.1.4.....	29
Pure Storage FlashArray Configuration .....	32
Solution Validation.....	32
Scalability Results .....	35
Test Summary .....	69
Design Considerations and Sizing Guidelines .....	70
Summary of Findings.....	72
Conclusions .....	73
About the Author.....	73

## Executive Summary

---

Pure Storage has introduced a converged infrastructure platform known as FlashStack that is built upon trusted hardware from Cisco and Pure Storage. FlashStack leverages Cisco's extremely flexible and expandable Unified Computing System (UCS™) to provide the compute horsepower, Cisco Nexus Switches for networking and Pure Storage FlashArray//m Series as the storage foundation. FlashStack partners can provide the guidance, services, and experience necessary to help you deploy FlashStack quickly and with confidence. FlashStack Authorized Support Partners provide a single point of contact for services and end-to-end support on the entire FlashStack converged infrastructure solution.

This document describes a reference architecture for deploying a Citrix XenDesktop 7.7 VDI environment on a FlashStack Converged Infrastructure using VMware vSphere 6.0, Login VSI 4.1.4 (the industry standard load testing solution for virtual desktop environments), and Microsoft Windows Server 2012R2 and Microsoft Windows 7 (64-bit) as the guest operating system. Pure Storage has validated the reference architecture within its lab – this document presents the hardware and software configuration, the test workload configuration, testing results and further offers implementation and sizing guidance for a mixed workload VDI population from hundreds to thousands of concurrent sessions.

## Goals and Objectives

---

The goal of this document is to showcase the ease of deploying a large number of virtual desktops on the FlashStack Converged Infrastructure with Cisco UCS, Cisco Nexus and Pure Storage FlashArray. We will demonstrate the scalability of both Citrix Machine Creation Service (MCS) as well as Provisioning Services (PVS)-based virtual desktop building blocks with Login VSI as a performance benchmark tool running mixed workloads (Office, Knowledge and Power users) in ratios that closely align with production customer deployments. Desktops delivered via MCS and PVS using Windows 7 64-bit desktops will be deployed incrementally and compared in number to find the high-end scalability characterization of the FlashStack. We will demonstrate linear scalability via testing with increasing cluster sizes of one, four, eight and later two entire UCS 5108 chassis with sixteen B-200M4 servers with a particular focus and emphasis on storage. In addition, we highlight the benefits of the Pure Storage FlashArray including data reduction, low latency and resiliency testing that directly impacts the user experience and provides customers an ideal solution for any XenDesktop MCS or XenDesktop PVS deployment project.

## Audience

---

The target audience for this document includes storage and virtualization administrators, consulting data center architects, field engineers, and desktop specialists who want to implement Citrix XenDesktop virtual desktops on a Pure Storage FlashStack Converged Infrastructure with VMware vSphere for virtualization. A working knowledge of VMware vSphere, XenDesktop, Login VSI, server, storage, networks and data center design is assumed but is not a prerequisite to read this document.

## Design Guide Principles


---

The guiding principles for implementing this reference architecture are:

- **Repeatable:** Create a scalable building block that can be easily replicated at any customer site. Publish the version of various firmware under test and weed out any issues in the lab before customers deploy this solution.
- **Virtualized:** Implement every infrastructure component as a virtual machine.
- **Available:** Create a design that is resilient and not prone to failure of a single component. For example, we include best practices to enforce multiple paths to storage, multiple NICs for connectivity, and high availability (HA) clustering including dynamic resource scheduling (DRS) on vSphere. Additionally, we will simulate scenarios such as a VDI bootstorm to highlight storage array resiliency under a common resource-intensive event.
- **Efficient:** Take advantage of inline data reduction and low latency of the Pure Storage FlashArray by pushing the envelope of VMs per server density.
- **Simple:** Avoid unnecessary and/or complex tweaks to make the results look better than a normal out-of-box environment.
- **Scalable:** By reporting the linear scaling of XenDesktop environments within the FlashStack architecture and by incrementing the number of UCS hosts, we will show non-disruptive operations, exceptional end-user experience, outstanding VM per host density and best-in-class flash storage performance.

# FlashStack Introduction

---



FlashStack is a converged infrastructure solution that brings the benefits of an all-flash storage platform to your converged infrastructure deployments. Built on best of breed components from Cisco and Pure Storage, FlashStack provides a converged infrastructure solution with high performance all-flash storage that is simple, flexible, efficient, and costs less than legacy converged infrastructure solution based on traditional disk.

FlashStack CI is available from accredited FlashStack Partners who help provide an excellent converged infrastructure ownership experience. FlashStack Partners have the knowledge and experience necessary to help streamline the sizing, procurement, and delivery of your entire system.

Key Benefits of the FlashStack solution are:

## 1. Consistent Performance and Scalability

- Consistent sub-millisecond latency with 100% flash storage.
- Consolidate 100's of enterprise-class applications in a single rack.
- Scale easily, without disruption.
- Repeatable growth through multiple FlashStack CI deployments.

## 2. Operational Simplicity

- Fully tested, validated, and documented for rapid deployment
- Reduced management complexity
- Auto-aligned 512-byte architecture eliminates storage alignment headaches
- No storage tuning or tiers necessary

## 3. Lowest TCO

- Dramatic savings in power, cooling and space with 100% Flash.
- Industry leading data reduction that is typically 2x better than competitors
- Free FlashArray controller upgrades every three years with Forever Flash™

## 4. Enterprise Grade Resiliency

- Highly available architecture and redundant components and no single point of failure
- Non-disruptive operations
- Upgrade and expand without downtime or performance loss
- Native data protection: snapshots and replication

# Pure Storage Introduction

---

Who knew that moving to all-flash storage could help reduce the cost of IT? FlashArray//m makes server and workload investments more productive, while also lowering storage spend. With FlashArray//m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient, accelerating your journey to the cloud.



FlashArray//m's performance can also make your business smarter by unleashing the power of real-time analytics, driving customer loyalty, and creating new, innovative customer experiences that simply weren't possible with disk. All by Transforming Your Storage with FlashArray//m.

FlashArray//m enables you to transform your data center, cloud, or entire business with an affordable all-flash array capable of consolidating and accelerating all your key applications.

Mini Size—Reduce power, space and complexity by 90%

- 3U base chassis with 15-120+ TBs usable
- ~1kW of power
- 6 cables

Mighty Performance—Transform your datacenter, cloud, or entire business

- Up to 300,000 32K IOPS
- Up to 9 GB/s bandwidth
- <1ms average latency

Modular Scale—Scale FlashArray//m inside and outside of the chassis for generations

- Expandable to ~½ PB usable via expansion shelves
- Upgrade controllers and drives to expand performance and/or capacity

Meaningful Simplicity—Appliance-like deployment with worry-free operations

- Plug-and-go deployment that takes minutes, not days
- Non-disruptive upgrades and hot-swap everything
- Less parts = more reliability

The FlashArray//m expands upon the FlashArray's modular, stateless architecture, designed to enable expandability and upgradability for generations. The FlashArray//m leverages a chassis-based design with customizable modules, enabling both capacity and performance to be independently improved over time with advances in compute and flash, to meet your business' needs today and tomorrow.

The Pure Storage FlashArray is ideal for:

**Accelerating Databases and Applications** Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.

**Virtualizing and Consolidating Workloads** Easily accommodate the most IO-hungry Tier 1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.

**Delivering the Ultimate Virtual Desktop Experience** Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance for under \$50/desktop.

**Protecting and Recovering Vital Data Assets** Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with FlashRecover.

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:

**Consistent Performance** FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.

**Less Cost than Disk** Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure. With VDI workloads data reduction is typically > 10:1.

**Mission-Critical Resiliency** FlashArray delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

**Disaster Recovery Built-In** FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

**Simplicity Built-In** FlashArray offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray the way you like it: Web-based GUI, CLI, VMware vCenter, Windows PowerShell, Python, REST API, or OpenStack.




Table 1: Pure Storage FlashArray//m Series

## FlashArray//m Specifications

	//m20	//m50	//m70
<b>Capacity</b>	<ul style="list-style-type: none"> <li>Up to 120+ TBs effective capacity*</li> <li>5 – 40TBs raw capacity (base chassis)</li> </ul>	<ul style="list-style-type: none"> <li>Up to 250+ TBs effective capacity*</li> <li>30 – 88TBs raw capacity (w/shelves)</li> </ul>	<ul style="list-style-type: none"> <li>Up to 400+ TBs effective capacity*</li> <li>44 – 136TBs raw capacity (w/shelves)</li> </ul>
<b>Performance</b>	<ul style="list-style-type: none"> <li>Up to 150,000 32K IOPS**</li> <li>&lt;1ms average latency</li> <li>Up to 5 GB/s bandwidth</li> </ul>	<ul style="list-style-type: none"> <li>Up to 220,000 32K IOPS**</li> <li>&lt;1ms average latency</li> <li>Up to 7 GB/s bandwidth</li> </ul>	<ul style="list-style-type: none"> <li>Up to 300,000 32K IOPS**</li> <li>&lt;1ms average latency</li> <li>Up to 9 GB/s bandwidth</li> </ul>
<b>Connectivity</b>	<ul style="list-style-type: none"> <li>8 Gb/s Fibre Channel</li> <li>10 Gb/s Ethernet iSCSI</li> <li>Management and Replication ports</li> </ul>	<ul style="list-style-type: none"> <li>16 Gb/s Fibre Channel</li> <li>10 Gb/s Ethernet iSCSI</li> <li>Management and Replication ports</li> </ul>	<ul style="list-style-type: none"> <li>16 Gb/s Fibre Channel</li> <li>10 Gb/s Ethernet iSCSI</li> <li>Management and Replication ports</li> </ul>
<b>Physical</b>	<ul style="list-style-type: none"> <li>3U</li> <li>742 Watts (nominal draw)</li> <li>110 lbs. (49.9 kg) fully loaded</li> <li>5.12” x 18.94” x 29.72” FlashArray//m chassis</li> </ul>	<ul style="list-style-type: none"> <li>3U – 7U</li> <li>1007 - 1447 Watts (nominal draw)</li> <li>110 lbs. (49.9 kg) fully loaded + 44 lbs. per expansion shelf</li> <li>5.12” x 18.94” x 29.72” FlashArray//m chassis</li> </ul>	<ul style="list-style-type: none"> <li>5U – 11U</li> <li>1439 – 2099 Watts (nominal draw)</li> <li>110 lbs. (49.9 kg) fully loaded + 44 lbs. per expansion shelf</li> <li>5.12” x 18.94” x 29.72” FlashArray//m chassis</li> </ul>

\* Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1, below the global average of the FlashArray user base.

\*\* Why does Pure Storage quote 32K, not 4K IOPS? The industry commonly markets 4K IOPS, but real-world environments are dominated by IO sizes of 32K or larger. FlashArray//m adapts automatically to 512B-32KB IO for superior performance, scalability, and data reduction.

## Purity Operating Environment

Purity implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashArray//m.

**Storage Software Built for Flash**—The FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity to both extend the life and ensure the maximum performance of consumer-grade MLC flash.

**Granular and Adaptive**—Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity’s data and flash management services to operate at the highest efficiency.

**Best Data Reduction Available**—FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.

**Highly Available and Resilient**—FlashProtect implements high availability, dual-parity RAID-3D, non-disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.

**Backup and Disaster Recovery Built In**—FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully-integrated in the FlashArray and leverage the native data reduction capabilities.



**Pure1 Manage**—By combining local web-based management with cloud-based monitoring, Pure1 Manage allows you to manage your FlashArray wherever you are – with just a web browser.

**Pure1 Connect**—A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashArray//m to all your data center and cloud monitoring, management, and orchestration tools.

**Pure1 Support**—FlashArray//m is constantly cloud-connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.

**Pure1 Collaborate**—Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.

## Experience Evergreen Storage



Tired of the 3-5 year array replacement merry-go-round? The move to FlashArray//m can be your last data migration. Purchase and deploy storage once and once only – then expand capacity and performance incrementally in conjunction with your business needs and without downtime. Pure Storage’s vision for Evergreen Storage is delivered by a combination of the FlashArray’s stateless, modular architecture and the ForeverFlash business model, enabling you to extend the lifecycle of storage from 3-5 years to a decade or more.

# Infrastructure Components of the Design Guide

## Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS) is a next-generation data center platform that unites compute, network, storage access, and virtualization into an organized structure aimed to reduce total cost of ownership and introduce vastly improved infrastructure deployment mechanisms at scale. UCS incorporates a unified network fabric with scalable, modular and powerful x86-architecture servers. With an innovative and proven design, Cisco UCS delivers an architecture that increases cost efficiency, agility, and flexibility beyond what traditional blade and rack-mount servers provide. Cisco makes organizations more effective by addressing the real problems that IT managers and executives face and solves them on a systemic level.



Figure 1: Cisco Unified Computing System

### *Greater Time-on-Task Efficiency*

Automated configuration can change an IT organization's approach from reactive to proactive. The result is more time for innovation, less time spent on maintenance, and faster response times. These efficiencies allow IT staff more time to address strategic business initiatives. They also enable better quality of life for IT staff, which means higher morale and better staff retention—both critical elements for long-term efficiency.

Cisco UCS Manager is an embedded, model-based management system that allows IT administrators to set a vast range of server configuration policies, from firmware and BIOS settings to network and storage connectivity. Individual servers can be deployed in less time and with fewer steps than in traditional environments. Automation frees staff from tedious, repetitive, time-consuming chores that are often the source of errors that cause downtime, making the entire data center more cost effective.

### *Easier Scaling*

Automation means rapid deployment, reduced opportunity cost, and better capital resource utilization. With Cisco UCS, rack-mount and blade servers can move from the loading dock and into production in a “plug-and-play” operation. Automatically configure blade servers using predefined policies simply by inserting the devices into an open blade chassis slot. Integrate rack-mount servers by connecting them to top-of-rack Cisco Nexus® fabric extenders. Since policies make configuration automated and repeatable, configuring 100 new servers is as straightforward as configuring one server, delivering agile, cost-effective scaling.

### *Virtual Blade Chassis*

With a separate network and separate management for each chassis, traditional blade systems are functionally an accidental architecture based on an approach that compresses all the components of a rack into each and every chassis. Such traditional blade systems are managed with multiple management tools that are combined to give the illusion of convergence for what is ultimately a more labor-intensive,

error-prone and costly delivery methodology. Rack-mount servers are not integrated and must be managed separately or through additional tool sets, adding complexity, overhead, and the burden of more time.

Architecturally, Cisco UCS blade and rack-mount servers are joined into a single virtual blade chassis that is centrally managed yet physically distributed across multiple blade chassis, rack-mount servers, and even racks and rows. This capability is delivered through

Cisco® fabric interconnects that provide redundant connectivity, a common management and networking interface, and enhanced flexibility. This larger virtual chassis, with a single redundant point of management, results in lower infrastructure cost per server, with fewer management touch points, and lower administration, capital, and operational costs.

### Cisco Nexus 9396PX Switch

The Cisco Nexus 9396PX delivers proven high performance and density, low latency, and exceptional power efficiency in a broad range of compact form factors.



Figure 2: Cisco Nexus 9396 Switch

Operating in Cisco NX-OS Software mode or in Application Centric Infrastructure (ACI) mode, these switches are ideal for traditional or fully automated data center deployments. In our setup, Cisco UCS 6248UP FI is connected through 10GbE to Nexus 9K for uplinks in bow-tie setup and the Nexus 9K is connected to the external network.

### Cisco MDS 9148S Switch

The Cisco MDS 9148S 16G Multilayer Fabric Switch is the next generation of the highly reliable Cisco MDS 9100 Series Switches. It includes up to 48 auto-sensing line-rate 16-Gbps Fibre Channel ports in a compact easy to deploy and manage 1-rack-unit (1RU) form factor.



Figure 3: Cisco MDS 9148

In all, the Cisco MDS 9148S is a powerful and flexible switch that delivers high performance and comprehensive Enterprise-class features.

## VMware vSphere 6.0

VMware vSphere is the industry-leading virtualization platform for building cloud infrastructures. It enables IT to meet SLAs (service-level agreements) for the most demanding business critical applications, at the lowest TCO (total cost of ownership). vSphere accelerates the shift to cloud computing for existing data centers and also underpins compatible public cloud offerings, forming the foundation for the industry's only hybrid cloud model. With the support of more than 3,000 applications from more than 2,000 ISV partners, vSphere is the trusted platform for any application.

- VMware vSphere Hypervisor Architecture provides a robust, production-proven, high-performance virtualization layer. It enables multiple virtual machines to share hardware resources with performance that can match (and in some cases exceed) native throughput.
- Each vSphere Hypervisor 6.0 instance can support as many as 480 logical CPUs, 12TB of RAM, and 1024 virtual machines. By leveraging the newest hardware advances, ESXi 6.0 enables the virtualization of applications that were once thought to be non-virtualizable.
- VMware ESXi™ 6.0 has dramatically increased the scalability of the platform. With vSphere Hypervisor 6.0, clusters can scale to as many as 64 hosts, up from 32 in previous releases. With 64 hosts in a cluster, vSphere 6.0 can support 8,000 virtual machines in a single cluster. This enables greater consolidation ratios, more efficient use of VMware vSphere Distributed Resource Scheduler™ (vSphere DRS), and fewer clusters that must be separately managed.
- VMware vSphere Virtual Machine File System (VMFS) allows virtual machines to access shared storage devices (Fibre Channel, iSCSI, etc.) and is a key enabling technology for other vSphere components such as VMware vSphere Storage vMotion®.
- VMware vSphere Storage APIs provide integration with supported third-party data protection, multipathing and storage array solutions.

## Citrix XenDesktop 7.7

Citrix XenDesktop provides both desktop and application virtualization solutions, which provide a streamlined approach to deliver, protect, and manage Windows desktops and applications to the end user so they can work anytime, anywhere, on any device.

### *Key Features*

XenDesktop 7.7 leverages desktop virtualization and builds on these capabilities, allowing IT to deliver virtualized and remote desktop and applications through a single platform and supports users with access to all their Windows and online resources through one unified workspace.

XenDesktop 7.7 supports the following key functionalities:

- Desktops and Applications Delivered through a Single Platform – Deliver virtual or remote desktops and applications through a single platform to streamline management and easily entitle end users.

- Unified Workspace – Securely delivers desktops, applications, and online services to end users through a unified workspace, providing a consistent user experience across devices, locations, media, and connections.
- Closed Loop Management and Automation – Consolidated control, delivery and protection of user compute resources with cloud analytics and automation, cloud orchestration and self-service features.
- Optimization with the Software-Defined Data Center – Allocates resources dynamically with virtual storage, compute, and networking to manage and deliver desktop services on demand.
- Central Image Management – Central image management for physical, virtual, and BYO devices.
- Hybrid-cloud flexibility – Provides an architecture built for onsite and cloud-based deployment.

### Citrix XenDesktop Architecture and Components

This section describes the components and various layers that make up a XenDesktop deployment as shown in [Figure 4](#) below.




Figure 4: Citrix XenDesktop 7.7 layers and components

### Delivery Controller

The Delivery Controller is the central management component of any XenApp or XenDesktop Site. Each Site has one or more Delivery Controllers. It is installed on at least one server in the data center. The

Controller consists of services that communicate with the hypervisor to distribute applications and desktops, authenticate and manage user access, broker connections between users and their virtual desktops and applications, optimize user connections and load-balance these connections.

The Controller manages the state of the desktops, starting and stopping them based on demand and administrative configuration. In some editions, the Controller allows you to install Profile management to manage user personalization settings in virtualized or physical Windows environments.

### *StoreFront*

StoreFront authenticates users to Sites hosting resources and manages stores of desktops and applications that clients access. It hosts your enterprise application store, which lets you give users self-service access to desktops and applications you make available to them. It also keeps track of users' application subscriptions, shortcut names, and other data to ensure that they have a consistent experience across multiple devices.

### *NetScaler Gateway*

A data-access solution that provides secure access inside or outside the LAN's firewall with additional credentials.

### *Citrix Director*

A web-based tool that allows administrators access to real-time data from the Broker agent, historical data from the site database and HDX data from NetScaler for troubleshooting and support.

### *Citrix Studio*

A management console that allows administrators to configure and manage sites, entitle users and groups of users for access and provide real-time usage data. Studio provides various wizards to guide you through the process of setting up your environment, creating your workloads to host applications and desktops and assigning applications and desktops to users. You can also use Studio to allocate and track Citrix licenses for your site.

### *Provisioning Methods*

- Machine Creation Services (MCS) — A collection of services that create virtual servers and desktops from a master image on demand, optimizing storage utilization and providing a virtual machine to users every time they log on. Machine Creation Services is fully integrated and administered in Citrix Studio.
- Provisioning Services (PVS) — Enables computers to be provisioned and reprovisioned in real-time from a single shared-disk image. Provisioning Services manages target devices as a device collection. The desktop and applications are delivered from a Provisioning Services vDisk that is imaged from a master target device, which enables you to leverage the processing power of physical hardware or virtual machines. Provisioning Services is managed through its own console.
- Existing images — Applies to desktops and applications that you have already migrated to virtual machines in the data center. You must manage target devices on an individual basis or collectively using third-party electronic software distribution (ESD) tools.

## *Provisioning Server*

A Provisioning Server is any server that has Stream Services installed, which is used to stream software from vDisks, as needed, to target devices. In some implementations, vDisks reside directly on the Provisioning Server. In larger implementations, Provisioning Servers may get the vDisk from a shared-storage location on the network.

Provisioning Servers also retrieve and provide configuration information to and from the Provisioning Services Database. Provisioning Server configuration options are available to ensure high availability and load-balancing of target device connections.

## *Provisioning Services Farm*

A farm represents the top level of a Provisioning Services infrastructure. The farm is created when the Configuration Wizard is run on the first Provisioning Server that will be added to that farm. Farms provide a farm administrator with a method for managing all components within the farm, such as:

- Product licensing
- Farm properties
- Administrative roles
- Active Directory configurations
- Provisioning Servers
- vDisk images
- Target devices
- Target device collections
- Sites
- Stores
- Views




Figure 5: Graphical representation of a PVS Farm implementation

### *vDisks*

vDisks exist as disk image files on a Provisioning Server or on a shared storage device. A vDisk consists of a VHD base image file, any associated properties files (.pvp), and if applicable, a chain of referenced VHD differencing disks (.avhd).

vDisks are assigned to target devices. Target devices boot from and stream software from an assigned vDisk image.

### *Virtual Delivery Agent (VDA)*

The VDA is installed on each physical or virtual machine in your Site that you want to make available to users. It enables the machine to register with the Controller, which in turn allows the machine and the resources it is hosting to be made available to users. VDAs establish and manage the connection between the machine and the user device, verify that a Citrix license is available for the user or session, and apply whatever policies have been configured for the session. The VDA communicates session information to the Broker Service in the Controller through the broker agent included in the VDA.

XenApp and XenDesktop include VDAs for Windows server and desktop operating systems. VDAs for Windows server operating systems allow multiple users to connect to the server at one time. VDAs for Windows desktops allow only one user to connect to the desktop at a time

### *Citrix Receiver*

Citrix Receiver clients are available for Windows, Mac, Linux, iOS, Android, Chrome OS, BlackBerry 10, Windows Phone as well as various Thin and Zero client devices to provide the connection to remote

desktops from your device of choice. The client supplies the connection to the virtual machine and communicates with StoreFront via secure access.

### Login VSI 4.1.4

Login VSI ([www.loginvsi.com](http://www.loginvsi.com)), the industry standard load testing solution for virtualized desktop environments is a tool designed to simulate a large-scale deployment of virtualized desktop systems and study its effects on an entire virtualized infrastructure. The tool is scalable from a few virtual machines running on one VMware ESXi (or other supported hypervisor) host up to hundreds to even thousands of virtual machines distributed across a cluster of ESXi hosts. Moreover, in addition to performance characteristics of the virtual desktops themselves, this tool also accurately shows the absolute maximum number of virtual machines that can be deployed on a given host or cluster of hosts. This is accomplished by using 'Launcher' Windows machines that simulate one or more end-point devices connecting in to the target VDI cluster and execute pre-defined classes of workloads that closely mimic real-world users.

Login VSI assists in the setup and configuration of the testing infrastructure, runs a set of application operations selected to be representative of real-world user applications, and reports data on the latencies of those operations, thereby accurately modeling the expected end-user experience provided by the environment.

Login VSI consists of the following components:

- Number of desktop virtual machines running on one or more UCS ESXi hosts to be exercised by the benchmarking tool and measured for performance against the selected workload.
- Number of client 'launcher' virtual machines running on one or more ESXi hosts on an entirely separate cluster to simulate end-users connecting into the VDI environment.
- Management Console on a Windows Server OS. [Figure 6](#) below shows a conceptual overview of a typical Login VSI layout.

# LOGIN VSI Environment Overview




Figure 6: Graphical representation of a typical Login VSI test environment

## Design Guide Solutions Overview

FlashStack consists of a combined stack of hardware (storage, network and compute) and software (Cisco UCS Manager, Citrix XenDesktop, VMware vCenter/ESXi, and Pure Storage GUI).

- **Network:** Cisco Nexus 9396 Switch, Cisco MDS 9148 and Cisco UCS Fabric Interconnect 6248UP for external and internal connectivity of IP and FC network.
- **Storage:** Pure Storage FlashArray//m20 with Fibre Channel connectivity
- **Compute:** Cisco UCS B200 M4 Blade Servers




Figure 7: Citrix XenDesktop Flashstack component and connectivity diagram

Figure 7 shows a detailed topology of the reference architecture configuration. A major goal of the architecture is to build out a highly redundant and resilient infrastructure. As such, we used powerful servers with dual Fibre Channel ports connected redundantly to two SAN switches that were connected to redundant FC target ports on the FlashArray. The servers were hosted in a vSphere HA cluster and had redundant network connectivity throughout.

### Cisco UCS Server Configuration

A pair of Cisco UCS Fabric Interconnects 6248UP, and scaling up to three chassis with twenty-four identical Intel CPU-based Cisco UCS B-series B200-M4 blade servers were deployed for hosting the virtual desktops. The UCS manager, UCS Fabric Interconnects and the components in the chassis were upgraded to 2.2.3f firmware level.

Each server had Cisco VIC 1340 cards and they were connected with four ports from each Cisco Fabric extender of the Cisco UCS chassis to the Cisco Fabric Interconnect, they were in turn connected to Cisco Nexus 5548UP Switch for upstream connectivity to access the Pure Storage FlashArray LUNs. The server configuration is described in Table 2.

Component	Description
Processor	2 X Intel Xeon E5-2670v3 @ 2.3GHz (12 Cores total, 24 Logical CPUs total)
Memory	392 GB @ 2133MHz (16 X 16GB), regular voltage
HBA	2 X 8G ports on Cisco UCS VIC 1340 (UCSB-MLOM-40G-03)
NIC	4 X 10G ports on Cisco UCS VIC 1340 (UCSB-MLOM-40G-03)
BIOS	Intel Virtualization Tech, Intel ATS Support, Intel VT-D features were enabled, Turbo Boost
vSphere	ESXi6.0.0-2494585 (Cisco Customized image)

Table 2: UCS B200-M4 server hardware configuration

### Cisco UCS Service Profile Configuration

In order to facilitate rapid deployment of UCS servers, a service profile template was created with the following characteristics:

1. We configured boot from SAN policy so that the server booted from a Pure Storage boot LUN (see [Figure 8](#) below)




Figure 8: Cisco UCS service profile template with boot from SAN policy configuration

2. We kept every other setting to the default, we didn't tweak any parameters
3. The Ethernet and FC adapter policy was set to VMware policy
4. The BIOS defaults were used for the B200-M4 blade servers
5. We configured two vHBA FC adapters and four vNIC Eth adapters on the Cisco VIC cards to avoid any single point of failure.
6. We deployed sixteen service profiles from the template and associated it with the blade servers in the two chassis. [Figure 9](#) below shows the Cisco UCS manager snapshot of service profile setup for the tests.




Figure 9: Cisco UCS service profile configuration

# VMware vSphere Configuration and Tuning

---

In this section, we discuss the ESXi 6.0 cluster configuration, network configuration and ESXi tuning for the hypervisor layer configuration.

## ESXi Cluster and Storage Configuration

A datacenter and a cluster with up to twenty-four hosts were configured with VMware High Availability (HA) clustering and Distributed Resource Scheduling (DRS) features. DRS was set to fully automated mode with power management turned off. The host EVC policy was set to Intel Haswell. The default BIOS for B200-M4 was chosen for all the service profiles. We created three datastores for the ESXi cluster for making the HA cluster datastore heartbeat work correctly. Note that DRS and HA were not active for our single host testing as a cluster size larger than that is required for those two features to be enabled.

Due to the simplicity of both the Pure Storage FlashArray and the Cisco UCS, configuration of VMware ESXi best practices are accordingly simple. ESXi uses its Native Multipathing Plugin architecture to manage I/O multipathing to underlying SAN storage volumes. Pure Storage FlashArray volumes (while not actually an ALUA array—it indeed is active/active) volumes are claimed by default by the Storage Array Type Plugin (SATP) for ALUA devices. Therefore all devices (by default) would inherit the Most Recently Used (MRU) Path Selection Policy (PSP). This would limit I/O to a single path and would be a colossal detriment to performance, as only leveraging a single path/port to the array would remove the active/active nature and performance advantage of the FlashArray.

All the ESXi servers were configured to change the default PSP for Pure devices from MRU to Round Robin (with advanced configuration to alternate paths after every I/O). The following command was run on each ESXi server prior to the presentation of FlashArray devices:

```
esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -V "PURE" -M "FlashArray" -P "VMW_PSP_RR" -O "iops=1"
```

[Figure 10](#) shows a properly configured Pure Storage LUN.




Figure 10: Properly configured Pure Storage Datastore

## ESXi Network Configuration

Two virtual switches each containing two vmnics were used for each host. We went with a standard vSwitch for this design. The redundant NICs were teamed in active/active mode and VLAN configurations were done on the upstream Cisco Nexus 9396 switches. The virtual switch configuration and properties are shown in [Figure 11](#) and [Figure 12](#) on the next page.




Figure 11: ESXi server network configuration on all servers (vSwitch1 for XenDesktop desktops)




Figure 12: ESXi server network configuration on all servers (vSwitch0 for host management)

# Citrix XenDesktop 7.7 Configuration and Tuning

Citrix XenDesktop 7.7 customizations were quite minimal; some of the tuning is highlighted in the section.

## XenDesktop Delivery Controller

Tune maximum simultaneous actions on Delivery Controller

The default concurrent XenDesktop concurrent operations are defined in the Hosting Configuration section of Citrix Studio. These default values are quite conservative and can be increased to higher values. Pure Storage FlashArray can withstand more operations including:

- Max new actions per minute (recommended value  $\geq 50$ )
- Max Simultaneous actions (all types) (recommended value  $\geq 50$ )
- The higher values will drastically cut down the amount of time spent for operations such as creating and updating your machine catalogs.




Figure 13: Advanced Connection settings in Citrix XenDesktop

Some caveats include –

- These settings are global settings and will affect all machine catalogs and delivery groups. Catalogs using other disk arrays will suffer if you set these values higher, so enabling these will have adverse effects.

- vCenter and Delivery Controller configurations, especially number of vCPUs, amount of memory, and the backing storage have implications from these settings. In order to attain the performance levels we have described in this white paper, it is important to note the ESXi configurations listed above and the virtual server configurations used in our testing.

## Citrix XenDesktop Machines Creation Services Desktop Components

VMware vCenter 6.0 was installed on a Windows 2012 R2 VM with 8 vCPU/32GB of memory. For deployments using sixteen or fewer UCS ESXi hosts a 6vCPU/16GB configuration is sufficient. For deployments using 4 UCS hosts or fewer, a vCenter instance with 4 vCPUs and 8GB of memory can be used.

For resiliency and failover, two paired Delivery Controller instances were installed on separate Windows 2012 R2 VMs with 4 vCPU/16 GB of memory for MCS desktop deployment and was used to create the MCS pools used for Login VSI testing. These servers were also used to provide connection brokering and power management for both MCS and PVS-based virtual desktops. Lastly, the Storefront service was also installed on each server.

The MCS-based Machine Catalogs in our test environment all had the following characteristics:

- Windows 7x64 Enterprise Desktop OS
- Power-managed VMs using Citrix Machine Creation Services
- Users connect to a new (random) desktop at each login and the machines were used to deliver Desktops only (not applications). Changes were discarded upon user log out and the delta disk was deleted.
- All desktop OS optimizations were made by the Virtual Delivery Agent upon installation.
- MCS Desktops were configured with 2vCPUs, 2GB of RAM and a 32GB local hard drive. A page file of 1024MB was statically set on the C: drive.
- We did not use a NetScaler Gateway for our testing as all components were within the Pure Storage lab environment.

## Citrix XenDesktop Provisioning Services Desktop Components

For our PVS testing, we utilized the vCenter and Delivery Controller/StoreFront servers described above to provide management for the virtual desktops.

In addition, we built two Provisioning Servers that each featured 4vCPUs and 8GB of memory. The Provisioning Servers each had 3 VMXNET3 NICs, one for management while the other two were teamed as a bonded NIC to enable the streaming service with a robust level of performance. We followed the best practices outlined in the [Best Practices Guide for Provisioning Services and XenApp](#) document and this [CTX document](#) from Citrix to further optimize our Provisioning Servers.

The PVS-based Machine Catalogs in our test environment all had the following characteristics:

- Windows 7x64 Enterprise Desktop OS
- Power-managed VMs using the Delivery Controllers
- Users connect to a new (random) desktop at each login and the machines were used to deliver Desktops only (not applications). Changes were discarded upon user log out.
- Target Devices were created using the XenDesktop Setup wizard in the Provisioning Services Console.
- Target Devices were configured with 2vCPUs, 2GB of RAM and a 32GB local hard drive. A 10GB write-cache was created and designated as the D:\ partition and a page file of 1024MB was set and pointed to D:\. The write-cache was thin-provisioned and attached to the target device as an entirely separate VMDK upon Target Device creation.
- A single Win7x64 vDisk was set in Standard Image mode and was set to 'cached on device hard disk' within the Provisioning Services Console. The vDisk was housed and streamed from each Provisioning Server and new versions were manually replicated between two Provisioning Servers. Automated solutions such as Microsoft's DFS or Robocopy for vDisk replication are relatively easy to script and implement in production implementations.
- All OS optimizations were made by the Virtual Delivery Agent and Provisioning Services Imaging wizard upon installation. The same networking optimizations as were made on our Provisioning Servers were also set on our target devices as described in the PVS best practice documentation in the links above.

## Microsoft Windows 7 Desktop Configuration

---

Login VSI provides guidelines for the suggested hardware configuration of the base Windows 7 image for each level of workload that is to be exercised. Since we wanted to mirror production customer environments as closely as possible in our testing, we elected to use Windows 7 Enterprise 64-bit, Office 2013 and the Knowledge Worker workload as we find that environment is most common in our present customer base.

To optimize the desktop OS we first installed VMware Tools with a typical installation. As we were using VMXNET3 network adapters, we also applied Microsoft Hotfix [2550978](#). At this point we cloned the desktop VM so that one could be used as the template device for MCS and the other could be used for the PVS vDisk. The vast majority of remaining OS-level Windows7 optimizations were accomplished via the Virtual Delivery Agent installer for MCS and a combination of the Virtual Delivery Agent and Provisioning Services Imaging Wizard for PVS.

The Win7 64-bit OS configuration for MCS and PVS can be seen in more detail in [Table 3](#) and [Table 4](#) below, respectively.

Component	Description
Desktop	Windows 7 Enterprise (64-bit)
Hardware Version	11
vCPU	2 (1 Socket, 2 Cores)
Memory	2 GB
vNIC Adapter	VMXNET 3
SCSI Controller	Paravirtual
Virtual Disk	32 GB Thin Provision
XenDesktop Virtual Delivery Agent (VDA)	Citrix VDA 7.7
Installed Applications	Microsoft Office 2013, Adobe Reader 11, Flash Player 11 Active X, Doro 1.82, Internet Explorer, Archive-7Zip, Windows Media Player

Table 3: Windows 7 64-bit MCS Desktop Configuration

Component	Description
Desktop	Windows 7 Enterprise (64-bit)
Hardware Version	11
vCPU	2 (1 Socket, 2 Cores)
Memory	2 GB
vNIC Adapter	VMXNET 3
SCSI Controller	Paravirtual
Virtual Disk	32 GB Thin Provision
XenDesktop Virtual Delivery Agent (VDA)	Citrix VDA 7.7 Provisioning Services Imaging Wizard 7.7
Installed Applications	Microsoft Office 2013, Adobe Reader 11, Flash Player 11 Active X, Doro 1.82, Internet Explorer, Archive-7Zip, Windows Media Player

Table 4: Windows 7 64-bit PVS Desktop Configuration

## Desktop Testing Tool – Login VSI 4.1.4

Login VSI is a 3<sup>rd</sup> party tool and the industry-standard designed to simulate real-world deployments of virtualized desktop systems and study its effects on an entire virtualized infrastructure. The tool is scalable from a few virtual machines running on one VMware vSphere host up to tens or even hundreds of thousands of virtual machines distributed across clusters of vSphere hosts.

Login VSI runs a set of operations selected to be representative of real-world user applications, and reports data on the latencies of those operations. In our tests, we used this tool to simulate a real world scenario, and then accepted the resultant application latency as a metric to measure end user experience.

Login VSI provides performance insights for virtualized desktop and server environments. Enterprise IT departments use Login VSI products in all phases of their virtual desktop deployment—from planning to deployment to change management—for more predictable performance, higher availability and a more consistent end user experience. The world's leading virtualization vendors use the flagship product, Login VSI, to benchmark performance. With minimal configuration, Login VSI products works in VMware Horizon View, Citrix XenDesktop and XenApp, Microsoft Remote Desktop Services (Terminal Services) and any other Windows-based virtual desktop solution.

For more information, download a trial at [www.loginvsi.com](http://www.loginvsi.com).

For this validation, we chose to use the Knowledge Worker workload in order to closely emulate a real-world VDI deployment. Those workloads simulate the following applications with varying degrees of intensity found in almost every company:

- Microsoft Word
- Microsoft Excel
- Microsoft PowerPoint
- Microsoft Outlook
- Microsoft Internet Explorer
- Document browse
- Picture album browse
- Adobe Reader
- Archiving software
- Video playback software
- Doro PDF Writer
- Photo Viewer/Edit

The simulated desktop workload is scripted in a 48 minute loop when a simulated Login VSI user is logged on, performing typical user activities. After the loop is finished it will restart automatically. Within each loop the response times of five specific operations are measured in a regular interval: twelve times in within each loop. The response times of these five operations are used to determine VSI<sub>max</sub>. VSI<sub>max</sub> is the “Virtual Session Index” that provides a useful aggregate single score that can be used to easily compare results when environment changes are being considered or different cluster configurations are tested. When the VSI<sub>max</sub> threshold is reached, that value shows the maximum number of concurrent VDI

sessions that can be successfully run on a given infrastructure before some resource (e.g. RAM) becomes saturated and the end-user experience becomes degraded.

The five operations from which the response times are measured are:

### Notepad File Open (NFO)

Loading and initiating VSINotepad.exe and opening the open file dialog. This operation is handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

### Notepad Start Load (NSLD)

Loading and initiating VSINotepad.exe and opening a file. This operation is also handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

### Zip High Compression (ZHC)

This action copy's a random file and compresses it (with 7zip) with high compression enabled. The compression will very briefly spike CPU and disk IO.

### Zip Low Compression (ZLC)

This action copy's a random file and compresses it (with 7zip) with low compression enabled. The compression will very briefly disk IO and creates some load on the CPU as well.

### CPU

Calculates a large array of random data and spikes the CPU for a short period of time.

The specific workload virtual machine configuration, applications open in parallel and IOPs generated per VM are shown below in [Table 5](#). As we will be comparing XenDesktop MCS and PVS, we elected to only use the Knowledge Workload in order to limit variables within our simulations.

Workload Name	Guest OS	Apps Open	IOPS (per VM)	Memory	vCPU
Office	Windows Enterprise 7 32-bit	5-8	8.1	1.5GB	1 vCPU
Knowledge	Windows 7 Enterprise 64-bit	5-9	8.5	2GB	2 vCPU
Power	Windows Enterprise 7 64-bit	8-12	10.8	2GB	2 vCPU

Table 5: Login VSI Workload characteristics

## Pure Storage FlashArray Configuration

---

The FlashArray//m20 contains no special configurations or value changes from the default values. The FlashArray contains twenty drive bays fully populated with two 256 GB SSDs each with two NVRAM devices for 10TB of raw space in total.

The UCS hosts are redundantly connected to the controllers with two FC connections to each controller from two HBAs on each host over the Fibre Channel protocol for a total of eight logical paths.

A cluster group was configured with all the ESXi hosts and a private volume was created for boot from SAN for each host. One 30TB and two 50 TB LUNs were shared across the entire host group for hosting the desktops; one for XenDesktop infrastructure, one for MCS and one for PVS, respectively.

## Solution Validation

---

In order to deploy and scale from hundreds to thousands of desktops, proper hardware and software design, a good test plan and success criteria are all required. This section talks about the test infrastructure, hardware configuration and software configuration we had in place for this reference architecture.

### Test Setup

The core infrastructure components including our Login VSI environment, Active Directory, DNS and DHCP servers were placed on a dedicated infrastructure cluster that was completely separated from the FlashStack other than running under the same vCenter instance and on the same network so that we could focus specifically on VDI performance on the FlashStack environment.

The infrastructure cluster included:

- Eight dedicated infrastructure servers (B200-M3 UCS servers) in an HA and DRS-enabled cluster were used to host the all of the infrastructure virtual machines:
  - 2 paired Active Directory, DNS and DHCP Windows 2012R2 Servers
  - 1 VMware vSphere Virtual Center Windows 2012R2 Server
  - 1 Microsoft SQL Windows 2012R2 Server for vCenter
  - 1 Login VSI Management Console on Windows Server 2012R2
  - 120 Login VSI Launcher Windows 7x64 Desktop VMs that could each support up to 25 concurrent VDI sessions
- One 5.5 TB (raw) FlashArray FA-405
  - 8 x 50GB ESXi boot volume for the 8 infrastructure Cisco UCS blade servers
  - 1 x 20 TB volume for the virtual server core infrastructure components listed above
  - 1 x 20 TB volume for the 120 Login VSI Launcher VMs

The FlashStack VDI test cluster included:

- One Pure Storage FlashArray//m20 in HA configuration, including two controllers and two 5TB data packs for 10TB raw:
  - 1 x 20TB volume was provisioned for the XenDesktop infrastructure components, including:
    - 2 XenDesktop Delivery Controller/StoreFront Servers
    - 2 XenDesktop Provisioning Servers
    - 1 XenDesktop Licensing Server
    - 1 XenDesktop Database Server
  - 1 x 40 TB volume was provisioned for the XenDesktop MCS desktops
  - 1 x 40 TB volume was provisioned for the XenDesktop PVS desktops
  - 24 x 50 GB ESXi boot volume for the twenty four Cisco UCS Blade servers
- We started with one, then incrementally scaled up to 24 Cisco UCS B-series blade server based on dual socket Intel E5-2670v3 @ 2.3GHz processor with 392 GB of memory running ESXi 6.0 as the hypervisor used to host the MCS and then later the PVS virtual desktops.

## Test Plan and Success Criteria

The test procedure was broken up into the following segments:

1. A single ESXi host comprised of 100% Knowledge Worker desktop testing using Machine Creation Services (MCS). This was performed in order to provide scaling guidelines for that VDI technology on the FlashStack Converged Infrastructure.
2. A single ESXi host comprised of 100% Knowledge Worker desktop testing using Provisioning Services (PVS). This was performed in order to provide scaling guidelines for that VDI technology on the FlashStack Converged Infrastructure.
3. MCS and PVS relative performance on a single host is exhibited in order to provide a comparison point between the two technologies to provide insight on how each functions and behaves on FlashStack.
4. Cluster sizes of 4, 8, 16 and 20 hosts were tested and performance was characterized with the Knowledge workload using XenDesktop MCS. We elected to show MCS at this scale as it is generally used in smaller VDI projects due to simplicity of setup and management.
5. Cluster sizes of 20 and 24 hosts were tested and performance was characterized with the Knowledge workload using XenDesktop PVS. We chose to show PVS for the bigger cluster sizes as most large Citrix deployments in our customer base use this technology over MCS due to easier manageability at scale.

6. MCS and PVS relative performance on a 20 host cluster is shown in order to provide a comparison point between the two technologies on FlashStack at a larger scale for customers looking for guidance or curious about how each of those two technologies best fits their environment.

The success criteria is as follows:

1. The Average Index score must show steady-state simulation performance below the VSI<sub>max</sub> threshold amount during the entire simulation. Additionally, ESXi host memory consumption must show no memory ballooning or swapping for our recommended number of virtual machines per host. Since the Login VSI tool uses a relatively small set of applications to exercise the environment, we purposefully targeted concurrent VDI sessions that showed conservative CPU and memory utilization during our testing to allow for confidence in using additional applications and more stressful workloads used in production environments.
2. The test must be repeated at least three times with minimal variability in order to show consistent performance for the specified VDI deployment size.
3. Login VSI provides a Baseline Performance Score as listed below in addition to the VSI<sub>max</sub> value. Successful tests VSI Baseline and Average scores must fall into the 'Good' category in order to confirm both good end-user experience and adequate environment performance throughout the entire simulation.

Baseline from	Baseline to	Name
0	1500	Good
1501	3000	Average
3001	4500	Below Average
4501	9999	Poor

Table 6: Login VSI Baseline Scoring Values

4. The backend Pure Storage FlashArray is able to service the I/O needs and is getting below 1 millisecond latency for over 99% of the simulation run.

Once we determine the single server scalability numbers we will then scale the number of servers from 4 to 8 servers (1 full chassis) to 16 servers (two full chassis) to 20 servers and finally to 24 servers (three full chassis) and linearly scale the concurrent number of VDI sessions in parallel. At each point we captured the ESXi server CPU/Memory utilization along with the Pure Storage FlashArray utilization data. Login VSI provided a detailed report on each of the test runs which would clearly show if our environment was properly configured and if we were seeing linear scaling of VMs per host as expected. The Login VSI report is additionally a measurement of VDI responsiveness and overall end-user experience throughout the simulation as shown by the VSI Index Average metric.

## Scalability Results

---

This section highlights the test results for the scalability testing starting with single server scale testing which would be used in the subsequent four, eight, sixteen, twenty and twenty-four server testing.

For each UCS cluster size, the results will be broken up as follows:

1. Recommended number of concurrently active Machine Creation Services (MCS) VDI sessions per UCS Cluster Sizes of one, four, sixteen, and twenty hosts
  - I. The Login VSI VSI<sub>maxv4</sub> value will not be reached in this simulation, but the VSI Baseline Performance Score and VSI<sub>maxv4</sub> chart showing the environment's ability to handle the recommended number of concurrent VDI desktops will be used to confirm our recommended sizing guidelines as well as provide independent, 3<sup>rd</sup> party proof of an excellent end-user experience.
  - II. Charts from the ESXi cluster showing: CPU utilization, overall memory utilization (including any minimal ballooning) to confirm ample headroom to further scale and perform common administrative tasks without user disruption.
  - III. Pure Storage Dashboard showing the array performance during the entire simulation run including latency, IOPs, bandwidth, storage space used and data reduction. Specific utilization metrics were captured when the last Login VSI sessions launched during the simulation run to show peak values during the run.
2. Recommended number of concurrently active Provisioning Services (PVS) VDI sessions per UCS Cluster Sizes of one, twenty and twenty-four hosts
  - I. The Login VSI VSI<sub>maxv4</sub> value will not be reached in this simulation, but the VSI Baseline Performance Score and VSI<sub>maxv4</sub> chart showing the environment's ability to handle the recommended number of concurrent VDI desktops will be used to confirm our recommended sizing guidelines as well as provide independent, 3<sup>rd</sup> party proof of an excellent end-user experience.
  - II. Charts from the ESXi cluster showing: CPU utilization, overall memory utilization (including any minimal ballooning) to confirm ample headroom to further scale and perform common administrative tasks without user disruption.
  - III. Pure Storage Dashboard showing the array performance during the entire simulation run including latency, IOPs, bandwidth, storage space used and data reduction. Specific utilization metrics were captured when the last Login VSI sessions launched during the simulation run to show peak values during the run.
3. Relative performance of PVS and MCS for both single server testing as well as at a large scale of 20 servers.
  - I. Compare and contrast Login VSI performance metrics for each scenario, including overall score as well as important metrics such as average logon time.

- II. Compare and contrast charts from the ESXi cluster showing: CPU utilization, overall memory utilization (including any ballooning or swapping).
- III. Compare and contrast utilization details from the Pure Storage Dashboard throughout both simulation runs.

## XenDesktop MCS and PVS Single Server Testing

Figure 14 below describes the test setup for a single UCS server in the FlashStack topology. A single Cisco UCS B200-M4 blade server is deployed and concurrent VDI sessions are executed on the single host until a recommended number of VMs per single host is determined using our standard VDI desktop configuration described in the Windows7 Desktop OS section. Storage performance is also closely monitored and captured. Using Login VSI as the workload generator we will show outstanding end-user experience while still leaving server host resource headroom for additional operations to take place and for additional applications to be installed.




Figure 14: Single server scale test configuration

## Test Results: 100 MCS Knowledge Workers on a Single Server

Running the Login VSI Knowledge Workload Benchmark on a single server cluster (note that vSphere HA and DRS were not enabled since a single host does not meet minimum requirements for those features) we were able to run 100 desktops while achieving a Login VSI baseline performance score of 897 and an average score of 1174. The resource utilization of the ESXi server and storage array is shown below as well.

The below chart in [Figure 15](#) shows a VSI Index Average well beneath the threshold level and we can also see that the FlashArray//m20 was not taxed in any regard during this simulation.






Figure 15: VSImax chart showing performance characteristics of 100 MCS Knowledge Workers



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	10.21.32.31	Usage	Average	%	0.36	84.52	0.36	35.927
■	10.21.32.31	Usage in MHz	Average	MHz	205	46551	205	19790.714

Figure 16: Single ESXi host CPU utilization during 100 MCS Knowledge Worker simulation



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	10.21.32.31	Balloon	Average	KB	0	0	0	0
■	10.21.32.31	Consumed	Average	KB	5290724	2234508...	5290724	1650078...
■	10.21.32.31	Usage	Average	%	1.31	55.52	1.31	40.998

Figure 17: Single ESXi host memory utilization during 100 MCS Knowledge Worker simulation showing no ballooning or swapping and a maximum of 55.5% memory used during the simulation




Figure 18: Pure Storage Dashboard during 100 MCS Knowledge Worker testing



Figure 19: Pure Storage Array values during peak simulation workload

## Test Results: 100 PVS Knowledge Workers on a Single Server

Running the Login VSI Knowledge Workload simulation on a single server cluster we were able to run 100 PVS-based desktops on a single host while achieving a Login VSI baseline performance score of 822 and an average score of 1153, maintaining a good end-user experience throughout. The resource utilization of the ESXi server and storage array is also shown below. The below chart in [Figure 20](#) shows a VSI Index Average well beneath the threshold level and we can also see that the FlashArray//m20 was not taxed in any regard during this simulation. Further comparison of a single host for MCS and PVS can be found in the results discussion section.




Figure 20: VSImax chart showing performance characteristics of 100 Knowledge Workers





Figure 21: Single ESXi host CPU utilization during 100 Knowledge Worker simulation.



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	10.21.32.32	Balloon	Average	KB	0	0	0	0
■	10.21.32.32	Consumed	Average	KB	60019284	2178035...	60019284	1913608...
■	10.21.32.32	Usage	Average	%	14.91	54.12	14.91	47.548

Figure 22: Single ESXi host memory utilization during 100 Knowledge Worker simulation showing no ballooning or swapping




Figure 23: Pure Storage Dashboard during 100 Knowledge Worker simulation



Figure 24: Pure Storage Array values during peak simulation workload

## Comparing PVS and MCS on a Single Server

At this small scale we found that there was practically no difference between the MCS and PVS technologies on any component of the FlashStack. ESXi CPU and memory utilization was close to identical with sufficient headroom.

Figure 25 shows an overlay of the two single host Login VSI charts and we can see that while the PVS-based test (purple line) shows a slight performance advantage over MCS (blue line), it would not be discernible to the end-user and they are in effect basically equivalent.




Figure 25: Login VSI performance chart overlay of 100 PVS and MCS simulations showing very similar performance

The most visible metric of performance for any VDI project is the amount of time it takes from when the user clicks 'connect' from within Citrix Receiver to when they have a usable desktop in front of them. For our comparison between MCS and PVS we again can see very similar performance with less than a second separating the two technologies (note that the y-axis is measured in milliseconds).




Figure 26: Login VSI Logon Timer for 100 PVS Desktops




Figure 27: Login VSI Logon Timer for 100 MCS Desktops

The only item of notice on the FlashArray//m was that MCS requires more read-IOPS relative to PVS, which is not a concern for an AFA as there are hundreds of thousands of read IOPs available. Below in [Table 7](#) we captured peak storage metrics as the 100<sup>th</sup> desktop was launched during the simulation run.

# of Hosts	MCS or PVS	# of VMs	Read Latency (ms)	Write Latency (ms)	Read IOPS (K)	Write IOPS (K)	Read Bandwidth (MB/s)	Write Bandwidth (MB/s)	Total Bandwidth (MB/s)
1	MCS	100	0.33	0.28	4.89	2.54	108.84	61.04	169.88
1	PVS	100	0.34	0.29	0.54	3.48	18.47	73.09	91.57

Table 7: Single host MCS and PVS observed storage metrics

## XenDesktop MCS Four Server Testing

[Figure 28](#) below describes the setup and connectivity for our four UCS server FlashStack. Four Cisco UCS B200-M4 blade servers were deployed in a single cluster with HA and DRS enabled. Specific desktop configurations and workload profiles were discussed at length in the Windows 7 Desktop Configuration and Test Plan sections. Login VSI, ESXi host and Pure Storage performance is closely monitored and captured. Keeping with linear scalability, we will show outstanding end-user experience while leaving ample server, network and storage resource headroom for additional operations to take place with 400 concurrent VDI users.




Figure 28: Four UCS server scale test configuration

## Test Results: 400 XenDesktop MCS Knowledge Workload on 4 Host Cluster

We continued to observe linear scaling and exceptional performance on a four server cluster via being able to run 400 Knowledge Worker desktops in parallel while achieving a Login VSI baseline performance score of 852 and an average score of 1165. VSImax v4 was not reached during this simulation and the VSI Index Average (blue line) remained well below the VSImax threshold limitation and fell into the 'Good' category throughout. The resource utilization of the ESXi server cluster and storage array is shown below. We can also see that the FlashArray//m20 was not taxed in any regard.






Figure 29: VSImax performance chart for a 400 MCS-based Knowledge Worker simulation



K...	Object	Measurement	Rollup	Units	Latest	Maxim...	Minim...	Average
■	CTX-MCS2	Usage in MHz	Average	MHz	6278	2017...	6278	1030...
■	CTX-MCS2	Usage	Average	%	2.85	91.62	2.85	46.78

Figure 30: 4 ESXi cluster CPU utilization during 400 Knowledge Workload Login VSI simulation



K...	Object	Measurement	Rollup	Units	Latest	Maxim...	Minim...	Average
■	CTX-MCS2	Balloon	Average	KB	0	0	0	0
■	CTX-MCS2	Consumed	Average	KB	8598...	8616...	8597...	8603...
■	CTX-MCS2	Total	Average	MB	1518...	1518...	1518...	1518...

Figure 31: 4 ESXi cluster memory utilization during 400 Knowledge worker Login VSI simulation with no swapping or ballooning noted




Figure 32: Pure Storage dashboard during 400 Knowledge Worker simulation on 4 ESXi hosts




Figure 33: Pure Storage metrics recorded during steady-state of 400 desktop MCS Login VSI simulation

## XenDesktop MCS Eight Server Testing

Figure 34 below describes the setup and connectivity for our eight UCS server FlashStack test. An entire 5108 UCS chassis was populated with Cisco UCS B200-M4 blade servers. Those hosts were deployed in a single cluster with HA and DRS enabled. Login VSI, ESXi host and storage performance is closely monitored and captured and the experiment was repeated 3 separate times in order to confirm no variability and consistent execution. Using Login VSI as the workload generator we will confirm superior end-user operations while leaving headroom for additional operations to take place for 800 concurrent MCS-based sessions.




Figure 34: Eight UCS Server FlashStack Configuration

## Test Results: 800 XenDesktop MCS Knowledge Workload on 8 Host Cluster

We again confirmed linear scaling and very good performance on an eight host cluster via being able to run 800 Knowledge Worker desktops in parallel while achieving a Login VSI baseline performance score of 822 and an average score of 1124. The VSI Index Average (blue line) remained well below the VSImax threshold limitation throughout the run and fell into the 'Good' category throughout. The resource utilization of the ESXi server cluster and storage array is shown below. We can also see that the FlashArray//m continued to provide sub-millisecond latency despite driving a large amount of I/O.




Figure 35: VSImax v4 performance chart during 800 desktop simulation showing excellent performance throughout





Figure 36: Eight host ESXi cluster CPU utilization showing maximum of 64% utilization during 800 desktop simulation



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	CTX-MCS2	Balloon	Average	KB	0	0	0	0
■	CTX-MCS2	Consumed	Average	KB	1720709...	1721156...	1716896...	1718667...
■	CTX-MCS2	Total	Average	MB	3278654	3278805	3278567	3278668...

Figure 37: Eight host ESXi cluster memory utilization showing steady-state usage and no ballooning or swapping during 800 desktop simulation




Figure 38: Pure Storage Dashboard displaying sub-ms latency despite consistent high read and write IO



Figure 39: Pure Storage metrics recorded during end of 800 desktop MCS Login VSI simulation

## XenDesktop MCS Sixteen Server Testing

Figure 40 below graphically shows the setup and connectivity for our sixteen UCS server FlashStack test. Two entire 5108 UCS chassis were populated with Cisco UCS B200-M4 blade servers. Those hosts were deployed in a single cluster with HA and DRS enabled. DRS was set to the most aggressive setting in order to load-balance the servers as much as possible. Login VSI, ESXi host and storage performance is closely monitored and captured and the experiment was repeated 3 separate times in order to confirm little to no variability and consistent execution. In addition to the Login VSI simulation, we also will show performance and array characteristics when we updated this 1600 Machine Catalog with monthly patches from Microsoft.




Figure 40: Sixteen Server FlashStack Configuration

## Test Results: 1600 XenDesktop MCS Knowledge Workload on 16 Host Cluster

We again found linear scaling and very good performance on a sixteen server cluster. This time being able to run 1600 Knowledge Worker desktops in parallel while achieving a Login VSI baseline performance score of 750 and an average score of 918; in fact showing that VDI performance was improving with additional host resources relative to our earlier Login VSI scores. The VSI Index Average (blue line) remained well below the VSImax threshold limitation throughout the run and fell into the 'Good' category throughout. The resource utilization of the ESXi server cluster and storage array is shown below. We can also see that the FlashArray//m continued to provide sub-millisecond latency despite driving a large amount of I/O.






Figure 41: 1600 VSImaxv4 simulation performance chart indicating good user experience throughout



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum
■	CTX-MCS	Usage in MHz	Average	MHz	117760	545852	113379
■	CTX-MCS	Usage	Average	%	13.36	61.96	12.87

Figure 42: 16 host ESXi cluster CPU utilization during 1600 Knowledge worker Login VSI simulation



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	A..
■	CTX-MCS	Balloon	Average	KB	0	0	0	0
■	CTX-MCS	Consumed	Average	KB	3469309696	3472387072	3421845504	3

Figure 43: 16 host ESXi cluster consistent memory utilization during 1600 Knowledge worker Login VSI simulation




Figure 44: Pure Storage Dashboard throughout the 1600 MCS Login VSI simulation



Figure 45: Pure Storage metrics recorded during end of 1600 desktop MCS Login VSI simulation

## Test Result: 1600 XenDesktop MCS Machine Catalog Update on a 16 Host Cluster

Probably the most disruptive common VDI administrative tasks is updating template VDI images with things like monthly Microsoft patches, new versions of applications and entirely new applications and then rolling it out to existing production Machine Catalogs. This task is particularly disruptive for MCS-based desktops as it requires the user logging out so that the desktop can be updated with the required changes.

To simulate this type of operation we patched our template image with monthly patches from Microsoft and updated all 1600 desktops at once to simulate an overnight update which is commonly employed to minimize user interruption. We were able to perform a Machine Catalog update of 1600 MCS-based desktops in approximately 35 minutes. This includes shutting down all 1600 MCS desktops at the previous image revision, cloning and provisioning the updated image and booting and registering all desktops against the Delivery Controller. We observed peak latency of 0.7ms, IOPS of approximately 100K and bandwidth of 1.5GB/s during this operation which can be seen below in [Figure 46](#).




Figure 46: Pure Storage Dashboard during a 1600 MCS Machine Catalog upgrade which completed in approximately 35 minutes.

## XenDesktop MCS Twenty Server Testing

Figure 47 below visually shows the setup and connectivity for our twenty UCS server FlashStack test. Two entire 5108 UCS chassis and a half of another were populated with Cisco UCS B200-M4 blade servers. Those hosts were deployed in a single cluster with HA and DRS enabled. DRS was set to the most aggressive setting in order to load-balance the virtual desktops across the servers as much as possible. Login VSI, ESXi host and storage performance is closely monitored and captured and the experiment was repeated three separate times in order to confirm little to no variability and consistent execution.




Figure 47: Twenty UCS Server FlashStack Configuration

## Test Results: 2000 XenDesktop MCS Knowledge Workload on 20 Host Cluster

We again found linear scaling and very good performance on a twenty server cluster via being able to run 2000 Knowledge Worker desktops in parallel while achieving a Login VSI baseline performance score of 750 and an average score of 906. The VSI Index Average (blue line) was well below the VSI<sub>max</sub> threshold throughout the run and continually was in the 'Good' category. The resource utilization of the ESXi server cluster and storage array is shown below. We can also see that the FlashArray//m continued to provide sub-millisecond latency despite driving a large amount of I/O.






Figure 48: 2000 MCS VSI<sub>max</sub>v4 simulation performance chart indicating good user experience throughout



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	CTX-MCS	Usage in MHz	Average	MHz	175490	737991	75080	364471...
■	CTX-MCS	Usage	Average	%	15.93	67.02	6.81	33.095

Figure 49: 20 host ESXi cluster CPU utilization during 2000 MCS Knowledge worker Login VSI simulation



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	CTX-MCS	Balloon	Average	KB	0	0	0	0
■	CTX-MCS	Consumed	Average	KB	4441951232	4441951232	1862631424	4062415138.133
■	CTX-MCS	Total	Average	MB	8351373	8351373	8350730	8350952.556

Figure 50: 20 host ESXi cluster memory utilization during 2000 MCS Knowledge worker Login VSI simulation with no swapping or ballooning




Figure 51: Pure Storage Dashboard through the duration of the 2000 XenDesktop MCS Login VSI simulation




Figure 52: Pure Storage IO measurements taken during 2000 desktop MCS simulation

## Test Results: 2000 XenDesktop PVS Knowledge Workload on 20 Host Cluster

At this point we changed from using Machine Creation Services to Provisioning Services but kept the same number of desktops and the exact same host cluster configuration in order to compare the two technologies with minimal differences. We found comparable linear scaling and the same robust performance on a twenty server cluster with 2000 Knowledge Worker desktops achieving a Login VSI baseline performance score of 804 and an average score of 965. The VSI Index Average (blue line) remained well below the VSImax threshold limitation throughout the run. The resource utilization of the ESXi server cluster and storage array is shown below. We can also see that the FlashArray//m continued to provide sub-millisecond latency despite driving a large amount of IO with less read bandwidth for the PVS simulation. Further discussion on the comparison of these two technologies will be in the next section.




Figure 53: 2000 PVS VSImaxv4 simulation performance chart indicating good user experience throughout





Figure 54: 20 host ESXi cluster CPU utilization during 2000 PVS Knowledge worker Login VSI simulation



Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	CTX-PVS	Balloon	Average	KB	0	0	0	0
■	CTX-PVS	Consumed	Average	KB	1486348...	4297960...	1486348...	3950760...

Figure 55: 20 host ESXi cluster memory utilization during 2000 PVS Knowledge worker Login VSI simulation




Figure 56: Pure Storage Dashboard through the duration of the 2000 PVS Login VSI simulation



Figure 57: Pure Storage IO measurements taken during 2000 PVS Knowledge worker simulation

## Comparing 2000 PVS and MCS Desktops on a Twenty Host Cluster

Similar to our single host comparison, we again found that there was no major differences between the MCS and PVS technologies on any component of the FlashStack - even when running 2000 concurrent desktops. ESXi CPU and memory utilization was close to identical with sufficient headroom for additional operations to take place in both cases.

The below figure shows an overlay of the two twenty host Login VSI VSImaxv4 charts and we can see that in this case MCS (purple) provided a slight performance advantage over PVS (blue) but it is not major enough to warrant making a design decision on this factor alone.




Figure 58: VSImax v4 Overlay chart showing relative performance of 2000 Knowledge Worker Login VSI simulations for PVS and MCS

Looking again at the LogonTimes for each virtualization technology we again see near identical performance, though with PVS we did encounter a bit more variability than MCS as we can see in the below two charts for each.




Figure 59: 2000 MCS Desktop Login VSI LogonTimer chart




Figure 60: 2000 PVS Desktop Login VSI LogonTimer chart

Finally, when we compared array performance we do see a substantial amount of read-bandwidth needed for the MCS solution relative to PVS though latency remained at the sub millisecond level throughout the simulation. On the flip side, data reduction rates for MCS fared better than PVS because there are more reducible bits for the separate desktop VMs created by MCS instead of the single streaming vDisk used in our PVS testing.

# of Hosts	MCS or PVS	# of VMs	Read Latency (ms)	Write Latency (ms)	Read IOPS (K)	Write IOPS (K)	Read Bandwidth (MB/s)	Write Bandwidth (MB/s)	Total Bandwidth (MB/s)
20	MCS	2000	0.33	0.31	26.51	17.81	538.89	497.79	1036.68
20	PVS	2000	0.2	0.24	1.4	27.63	24.05	469.13	493.18

Table 8: Twenty host cluster MCS and PVS observed storage metrics during steady-state

From these results, we have found that Pure Storage will work exceedingly well for either PVS or MCS technologies, but customers who are more focused on VM density per array will likely see a slightly better result using PVS over MCS. Meanwhile, those looking to maximize their data reduction would be better served with an MCS-based implementation. As always, we encourage our customers to test with their own unique environments to best understand the implications of each solution to arrive at the optimal solution for their particular project.

## XenDesktop PVS Twenty-Four Server Testing

Figure 61 below shows the setup and connectivity for our twenty-four UCS server FlashStack test. Three entire 5108 UCS chassis were fully populated with Cisco UCS B200-M4 blade servers. Those hosts were deployed in a single cluster with HA and DRS enabled. DRS was set to the most aggressive setting in order to load-balance the virtual desktops as much as possible across the hosts. Login VSI, ESXi host and storage performance is closely monitored and captured and the experiment was repeated 3 separate times in order to confirm little to no variability and consistently good performance.

For our last test we elected to simulate a large scale bootstorm in which all PVS desktops were powered off and put into maintenance mode initially. We next changed the PVS Delivery Group power settings to be 100% powered on, turned off maintenance mode and then started a timer to see how long bootup and registering all 2,500 desktops would take.




Figure 61: Twenty-Four UCS Server FlashStack Configuration

## Test Results: 2500 XenDesktop PVS Knowledge Workload on 24 Host Cluster

We again found linear scaling and rock-solid performance on a twenty-four server cluster via being able to run 2500 Knowledge Worker desktops in parallel while achieving a Login VSI baseline performance score of 806 and an average score of 997. The VSI Index Average (blue line) remained well below the VSImax threshold limitation in the run and fell into the 'Good' category throughout. The resource utilization of the ESXi server cluster and storage array is shown below. The FlashArray//m20 continued to prove its performance and resiliency even while running a very large set of PVS-based virtual desktops in parallel with consistent sub-millisecond latency.




Figure 62: 2500 PVS-based VSImaxv4 simulation performance chart indicating excellent user experience throughout




Figure 63: 24 host ESXi cluster CPU utilization during 2500 PVS Knowledge worker Login VSI simulation showing 66% total usage at peak




Figure 64: 24 host ESXi cluster consistent memory usage during 2500 PVS Knowledge worker Login VSI simulation




Figure 65: Pure Storage Dashboard during of the 2500 PVS Knowledge Worker Login VSI simulation




Figure 66: Metrics recorded from Pure Storage GUI while 2500 concurrent PVS sessions were simulated

## Test Result: 2500 XenDesktop PVS Bootstorm on a 24 Host Cluster

Bootstorms are a common occurrence in almost all VDI strategies and are particularly common at the start of the workday when employees begin coming online to start the work day. As such, we wanted to simulate such an event at a large scale to see how long it would take for an environment to recover and pay attention to any performance issues during the operation. We were able to boot 2500 PVS-based desktops in approximately 32 minutes with all VMs showing up in the Delivery Controller as ‘registered’ and available for use within that timeframe. That equates to approximately 83 PVS desktops booting and registering against the Delivery Controller per minute and further illustrates the game-changing agility this Converged Infrastructure provides to VDI administrators.




Figure 67: Pure Storage Dashboard during a 2500 PVS desktop bootstorm which completed in approximately 32 minutes

## Test Summary

Based on the performance data from one, four, eight, sixteen, twenty and twenty-four servers testing, we conclude the following:

1. As we scaled up the number of concurrent desktops in our simulations we expected and observed a corresponding large increase in the amount of I/O on the array. Despite this, we maintained sub-millisecond latency on the array throughout every simulation run all the way up to 2,500 VDI desktops at the same time. This conclusively proves that the FlashStack Converged Infrastructure can support a VDI project from proof-of-concept to pilot to a large scale production environment seamlessly without any disruption or downtime. In [Table 9](#) below we recorded storage metrics of interest within the last 5 minutes of each Login VSI simulation highlighted in the previous results section for each Login VSI simulation.

# of Hosts	MCS or PVS	# of VMs	Read Latency (ms)	Write Latency (ms)	Read IOPS (K)	Write IOPS (K)	Read Bandwidth (MB/s)	Write Bandwidth (MB/s)	Total Bandwidth (MB/s)
1	MCS	100	0.33	0.28	4.89	2.54	108.84	61.04	169.88
1	PVS	100	0.34	0.29	0.54	3.48	18.47	73.09	91.57
4	MCS	400	0.34	0.26	15.02	7.53	377.34	184.58	561.92
8	MCS	800	0.39	0.26	19.26	10.95	431.37	276.03	707.4
16	MCS	1600	0.36	0.28	25.58	14.54	554.35	384.95	939.3
20	MCS	2000	0.33	0.31	26.51	17.81	538.89	497.79	1036.68
20	PVS	2000	0.2	0.24	1.4	27.63	24.05	469.13	493.18
24	PVS	2500	0.21	0.25	1.79	33.42	28.81	575.17	603.98

Table 9: Summary table of Pure Storage metrics during each simulation run discussed in results section

2. We were able to perform a Machine Catalog update of 1600 MCS-based desktops in approximately 35 minutes. This includes shutting down the 1600 desktops at the previous revision, cloning the updated image and booting and registering all desktops against the Delivery Controller. We observed peak latency of 0.7ms, IOPS of approximately 100K and bandwidth of 1.5GB/s during this operation.
3. We were able to start 2500 PVS desktop at once with all of them being registered against the Delivery Controller in approximately 32 minutes. That equates to approximately 78 PVS target devices streaming a vDisk from a Provisioning Server, booting and registering against our Delivery Controllers *per minute*.
4. The Pure Storage FlashArray consistently showed a 0.5 ms latency or less throughout the Login VSI tests for both reads and writes despite a heavy I/O workload being run against it.
5. VM per host density can be further improved by purchasing servers with faster CPUs and/or more memory.
6. The data reduction values ranging from 5:1 to 7.7:1 during these tests was due to our tests using both MCS and PVS-based desktops to showcase the XenDesktop suite of VDI solutions. While we absolutely support both, we generally see a better data reduction ratio with MCS-based clones as those desktops have more Windows bits that we are better able to deduplicate than PVS. For ease of management and upgradability, PVS uses only a single streaming vDisk for the OS-layer and the attached page file/write-cache data is not as reducible. That being said, from a performance perspective we did observe that PVS has less of an I/O requirement than MCS and the array at most was 28% full during our tests, leaving a ton of space available.

## Design Considerations and Sizing Guidelines

The space consumption and the IOPS we saw in the 2500 Knowledge workload desktop deployment could easily have been sustained in the smallest FlashArray configuration. As the deployment grows, it is easy to expand capacity by adding denser SSD drives or more shelves to the array without downtime.

A pilot can be implemented on a two-controller HA system and a single 5TB data pack that populates half of the drive bays on the front of the //m20 array. As the deployment passes out of the pilot phase, you

can upgrade with an additional 5TB data pack and support 2,000 desktops or more depending on the IOPS requirements for your users. As your user data grows, additional shelves can be added. More shelves and/or higher performance controllers can be added without any downtime in the middle of the workday completely transparently while your VDI users are actively using your environment.

Based on the documented results in this design guide we believe that the //m20 can accommodate up to 3,000 PVS-based and approximately 2500 MCS-based Knowledge Worker Windows 7 desktops from an IOPS perspective. For a 3,000+ desktop deployment, we recommend a fully configured //m50 with 40TB raw space. These sizing guidelines are approximations based upon our internal testing and workload characterization; your actual desktop density may vary depending on how the desktops are configured, whether or not user data is stored in the desktops or the array, and a variety of other factors. Pure Storage recommends a pilot deployment using production applications and desktop images in order to fully-understand space and performance requirements.

Pure Storage highly recommends performing resiliency testing including Machine Catalog updates, storage controller reboots (to simulate a firmware or controller upgrade), power failures, pulling one or two SSD drives and pulling Fibre or iSCSI cables within your own proof-of-concept project as these kinds of interruptions are a reality of the workplace. They can also take place during regular office hours while users are actively logged into your VDI deployment, thus proving environment resiliency should be of high importance. A prime example is when Microsoft releases a critical patch and your desktop pools need to be updated with the update as quickly as possible in order to avoid the possibility of your network being compromised by an attacker. Additionally, IT staff can avoid working off-hours and over the weekends - when they can perform non-disruptive upgrades with no impact to production users.

Adding a new shelf to increase capacity is very straightforward and involves simply connecting SAS cables from the controller to the new shelf that can be done while the array is online. The shelf addition can be done non-disruptively and without any impact to the desktop users (no down time nor any performance impact at all) even when they are actively using their virtual desktops. The Pure Storage FlashArray features stateless controllers, which means all the configuration information is stored on the storage shelves instead of within the controllers themselves. In the event of a controller failure, one can easily swap out a failed controller with a new controller without reconfiguring SAN zoning, which again can be done non-disruptively.

## Summary of Findings

---

- We deployed a highly resilient catalog of 2000 MCS-based Knowledge Workers and later, 2500 PVS-based Knowledge Workers on FlashStack Converged Infrastructure. Using these pools we ran a realistic load generator with Login VSI that simulated Knowledge Workers performing common computing tasks, resulting in **achievement of the highest category of Login VSI Baseline Score in every single test that was performed. VSImax v4 was not reached in any simulation.** This score means that all of the applications used in the Knowledge Workload performed with outstanding performance and response times.
- Throughout the testing the FlashArray delivered up to 100,000 32K IOPS and maintained latency under 1 ms, demonstrating the FlashArray's consistent latency and ability to deliver the best all-flash VDI end-user experience at all times. The FlashArray delivers a better desktop experience for end-users than dedicated laptops with SSDs, and doesn't risk the end-user experience by relying on caching as hybrid flash/disk arrays do.
- In total throughout the testing we deployed more than 2500 desktops (each of 32 GB disk size), together only consuming about 1TB of physical storage on the FlashArray. This massive data reduction is the result of the high-performance inline data reduction (deduplication and compression) delivered by the FlashArray, which enables using any combination of non-persistent or persistent desktops for MCS or PVS – both of which reduce to about the same amount of space on the array.
- As tested, the 10 TB FlashArray//m20 delivered best-in-class VDI performance for 2500 desktops with headroom to scale further with additional UCS chassis and/or servers with more memory and faster CPUs.
- Running Unmap weekly to reclaim dead space is a critical operation for any VDI implementation on Pure Storage and is especially important for non-persistent linked-clones. For more information please see the following primer on Cody Hosterman's [blog](#).
- Running sDelete on the PVS Write-Cache is another important operation to minimize in-guest dead-space for your PVS-based desktops. More info on this can be found on Kyle's [blog](#).
- Throughout the testing we performed common VDI administrator operations and found a drastic reduction in time for cloning persistent desktops, (re)booting desktops, and other day-to-day virtual desktop operations. Taken together these **operational savings deliver substantial efficiency gains for VDI administrators** throughout the VDI day.
- The power footprint for the tested FlashArray//m20 was 8 Amps (110V) which is a fraction of any mechanical disk storage array available in the marketplace. This configuration consumed three rack units (3 RU) in data center space.
- This reference architecture has shown FlashStack CI to be able to provide multiple XenDesktop building blocks so as to better provide sizing and scaling recommendations. Customers can add more server and infrastructure components to scale the architecture out to ten of 1000's of desktops. Based on the results, we believe a single //m20 can support from 2500-3000 MCS or PVS desktops with any mix of non-persistent and/or persistent desktops, depending on the IOPs

requirements per desktop. For a 3000+ desktop deployment, we recommend a fully configured //m50 with 40TB raw space.

## Conclusions

---

We set out to prove that FlashStack Converged Infrastructure with the Pure Storage FlashArray is the ultimate platform for virtual desktop deployment and we achieved consistently impressive results from one server to twenty-four servers while running an industry standard desktop workload generator. The Login VSI scores reflect the latency observations on the Pure Storage dashboard that showed minimal latency throughout the testing and even during boot storm, login storm and test ramp-up scenarios. Factors like rapid deployment, ease of storage management, lower storage cost, lower power, rack space savings, and lower cooling requirements make the TCO for large scale deployments even more attractive.

## About the Author

---



Kyle Grossmiller is a VDI Solutions Architect at Pure where he focuses on helping customers bring their VDI projects to the next level of success using Pure's All-Flash Arrays. He provides technical expertise to help solve pressing client challenges and produces technical collateral that contains insights and guidance on how using Pure Storage delivers the best possible results for VDI.

Prior to joining Pure, Kyle was at Lockheed Martin Space Systems and Enterprise Business Services for over 12 years where he worked in dual IT roles. In this capacity, he supported their engineering user base throughout multiple hardware platform lifecycles and major CAD software upgrades as well as serving as the technical lead for an internal private-cloud VDI project from planning to POC to production. From these experiences he has a unique and deep perspective on the ever-changing nature of IT.

Kyle resides in San Francisco, CA and holds a Bachelor of Science degree in Electrical Engineering from Lafayette College in Easton, PA.

Blog: <http://blog.purestorage.com/author/kyle/>

Twitter: [@kylegrossmiller](https://twitter.com/kylegrossmiller)



Pure Storage, Inc.  
Twitter: [@purestorage](#)  
[www.purestorage.com](#)

650 Castro Street, Suite #260  
Mountain View, CA 94041

T: 650-290-6088  
F: 650-625-9667

Sales: [sales@purestorage.com](mailto:sales@purestorage.com)  
Support: [support@purestorage.com](mailto:support@purestorage.com)  
Media: [pr@purestorage.com](mailto:pr@purestorage.com)  
General: [info@purestorage.com](mailto:info@purestorage.com)