Citrix XenDesktop on vSphere

Nutanix Reference Architecture

Version 1.2 • November 2017 • RA-2022

Copyright

Contents

1. Executive Summary

This document makes recommendations for the design, optimization, and scaling of Citrix XenDesktop deployments on the Nutanix Enterprise Cloud Platform, including detailed performance and configuration information. We used Login VSI to simulate real-world workloads and a XenDesktop environment using MCS and PVS on Nutanix. The sizing data and recommendations in this document derive from multiple testing iterations and thorough technical validation.

We completed the solution and testing data with Citrix XenDesktop deployed on VMware vSphere. When deploying Citrix XenDesktop on Nutanix, host CPU resources determine desktop user density rather than concerns about I/O or resource bottlenecks for MCS and PVS. The Login VSI test results confirm that Nutanix supports densities of over 120 Office Worker desktops per node and more than 100 Knowledge Worker desktops per node (counting four per 2RU appliance). We determined sizing for the pods after carefully considering performance and accounting for the additional resources needed for N+1 failover capabilities.

2. Introduction

2.1. Audience

This reference architecture is part of the Nutanix Solutions Library and is intended for architecting, designing, managing, and supporting Nutanix infrastructures. Consumers of this document should be familiar with VMware vSphere, Citrix XenDesktop, and the Nutanix Enterprise Cloud Platform.

We addresses key items for each role, focusing on enabling a successful design, implementation, and transition to operation.

2.2. Purpose

This document covers the following subject areas:

- Overview of the Nutanix solution.
- Overview of Citrix XenDesktop and its use cases.
- The benefits of Citrix XenDesktop on Nutanix.
- Design and configuration considerations when architecting a Citrix XenDesktop solution on Nutanix.
- Benchmarking Citrix XenDesktop performance on Nutanix.

Table 1: Document Version History

Version Number	Published	Notes
1.0	June 2015	Original publication.
1.1	July 2016	Updated platform overview.
1.2	November 2017	Updated platform overview

3. Nutanix Enterprise Cloud Platform Overview

3.1. Nutanix Acropolis Overview

Nutanix delivers a hyperconverged infrastructure solution purpose-built for virtualization and cloud environments. This solution brings the performance and economic benefits of web-scale architecture to the enterprise through the Nutanix Enterprise Cloud Platform, which is composed of two product families—Nutanix Acropolis and Nutanix Prism.

Attributes of this solution include:

- Storage and compute resources hyperconverged on x86 servers.
- System intelligence located in software.
- Data, metadata, and operations fully distributed across entire cluster of x86 servers.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.
- Simplified one-click upgrade.
- Native file services for hosting user profiles.
- Native backup and disaster recovery solutions.

Nutanix Acropolis provides data services and can be broken down into three foundational components: the Distributed Storage Fabric (DSF), the App Mobility Fabric (AMF), and AHV. Prism furnishes one-click infrastructure management for virtual environments running on Acropolis. Acropolis is hypervisor agnostic, supporting three third-party hypervisors—ESXi, Hyper-V, and XenServer—in addition to the native Nutanix hypervisor, AHV.

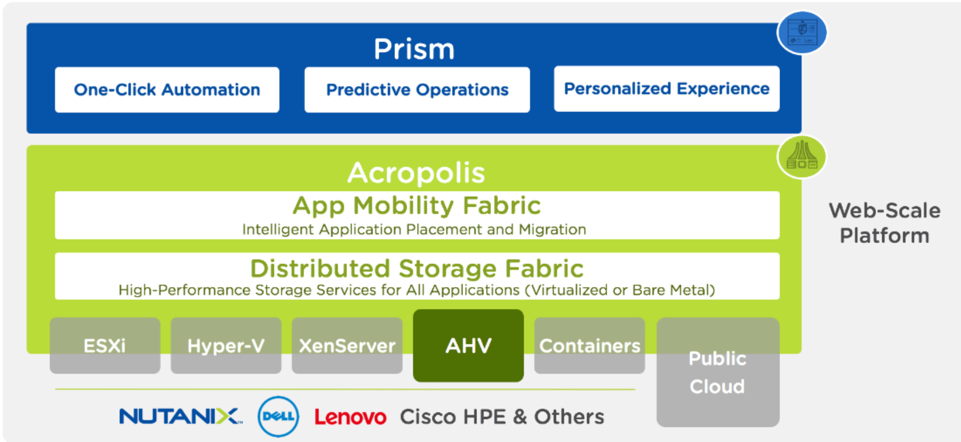

Figure 1: Nutanix Enterprise Cloud Platform

3.2. Distributed Storage Fabric

The Distributed Storage Fabric (DSF) delivers enterprise data storage as an on-demand service by employing a highly distributed software architecture. Nutanix eliminates the need for traditional SAN and NAS solutions while delivering a rich set of VM-centric software-defined services. Specifically, the DSF handles the data path of such features as snapshots, clones, high availability, disaster recovery, deduplication, compression, and erasure coding.

The DSF operates via an interconnected network of Controller VMs (CVMs) that form a Nutanix cluster, and every node in the cluster has access to data from shared SSD, HDD, and cloud resources. The hypervisors and the DSF communicate using the industry-standard NFS, iSCSI, or SMB3 protocols, depending on the hypervisor in use.

3.3. App Mobility Fabric

The Acropolis App Mobility Fabric (AMF) collects powerful technologies that give IT professionals the freedom to choose the best environment for their enterprise applications. The AMF encompasses a broad range of capabilities for allowing applications and data to move freely between runtime environments, including between Nutanix systems supporting different hypervisors, and from Nutanix to public clouds. When VMs can migrate between hypervisors (for example, between VMware ESXi and AHV), administrators can host production and development or test environments concurrently on different hypervisors and shift workloads between them as needed. AMF is implemented via a distributed, scale-out service that runs inside the CVM on every node within a Nutanix cluster.

3.4. AHV

Nutanix ships with AHV, a built-in enterprise-ready hypervisor based on a hardened version of proven open source technology. AHV is managed with the Prism interface, a robust REST API, and an interactive command-line interface called aCLI (Acropolis CLI). These tools combine to eliminate the management complexity typically associated with open source environments and allow out-of-the-box virtualization on Nutanix—all without the licensing fees associated with other hypervisors.

3.5. Third-Party Hypervisors

In addition to AHV, Nutanix Acropolis fully supports Citrix XenServer, Microsoft Hyper-V, and VMware vSphere. These options give administrators the flexibility to choose a hypervisor that aligns with the existing skillset and hypervisor-specific toolset within their organization. Unlike AHV, however, these hypervisors may require additional licensing and, by extension, incur additional costs.

3.6. Nutanix Acropolis Architecture

Acropolis does not rely on traditional SAN or NAS storage or expensive storage network interconnects. It combines highly dense storage and server compute (CPU and RAM) into a single platform building block. Each building block is based on industry-standard Intel processor technology and delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution has no LUNs to manage, no RAID groups to configure, and no complicated storage multipathing to set up. All storage management is VM-centric, and the DSF optimizes I/O at the VM virtual disk level. There is one shared pool of storage composed of either all-flash or a combination of flash-based SSDs for high performance and HDDs for affordable capacity. The file system automatically tiers data across different types of storage devices using intelligent data placement algorithms. These algorithms make sure that the most frequently used data is available in memory or in flash for optimal performance. Organizations can also choose flash-only storage for the fastest possible storage performance. The following figure illustrates the data I/O path for a write in a hybrid model (mix of SSD and HDD disks).

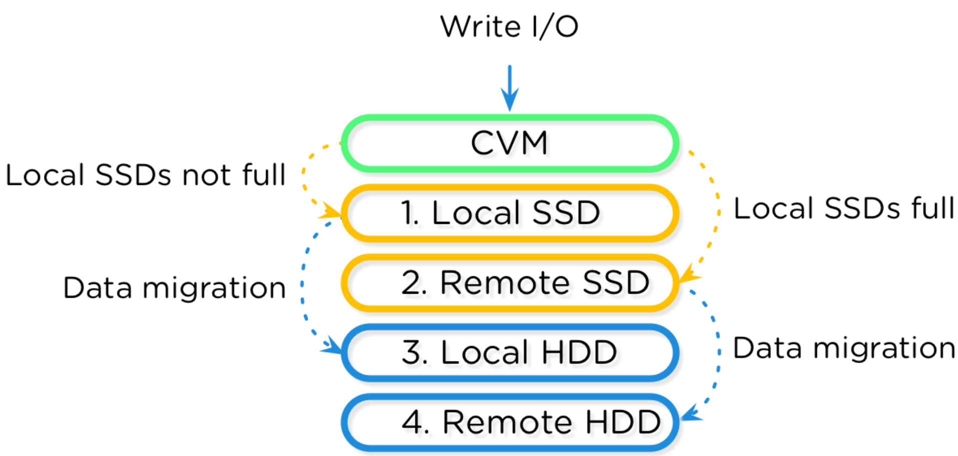

Figure 2: Information Life Cycle Management

The figure below shows an overview of the Nutanix architecture, including the hypervisor of your choice (AHV, ESXi, Hyper-V, or XenServer), user VMs, the Nutanix storage CVM, and its local disk devices. Each CVM connects directly to the local storage controller and its associated disks. Using local storage controllers on each host localizes access to data through the DSF, thereby reducing storage I/O latency. Moreover, having a local storage controller on each node ensures that storage performance as well as storage capacity increase linearly with node addition. The DSF replicates writes synchronously to at least one other Nutanix node in the system, distributing data throughout the cluster for resiliency and availability. Replication factor 2 creates two identical data copies in the cluster, and replication factor 3 creates three identical data copies.

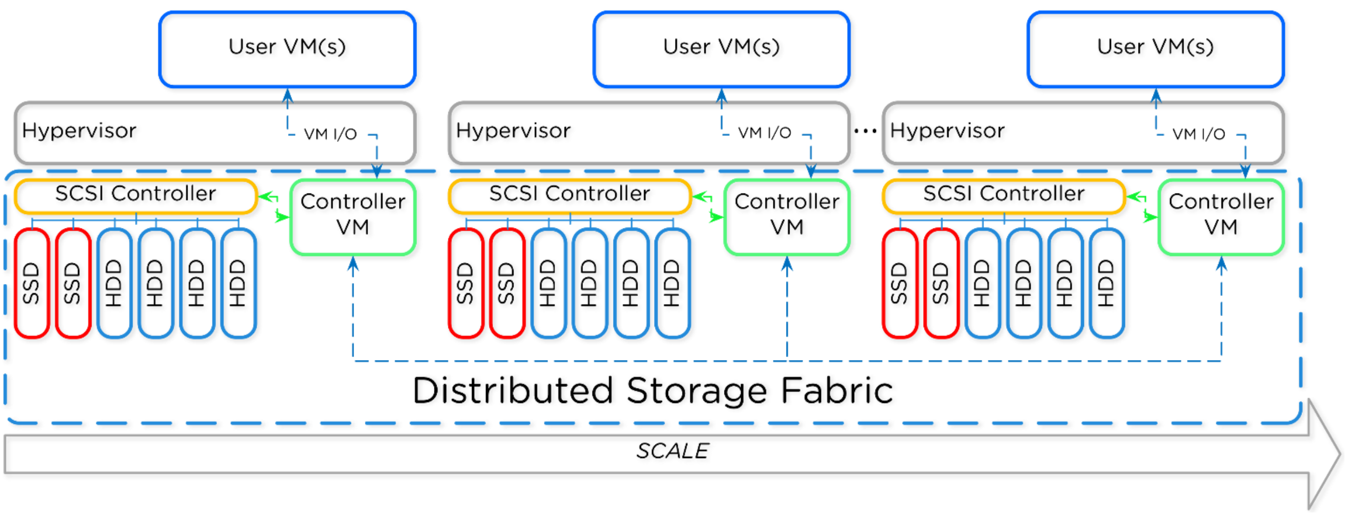

Figure 3: Overview of the Nutanix Architecture

Local storage for each Nutanix node in the architecture appears to the hypervisor as one large pool of shared storage. This allows the DSF to support all key virtualization features. Data localization maintains performance and quality of service (QoS) on each host, minimizing the effect noisy VMs have on their neighbors' performance. This functionality allows for large, mixed-workload clusters that are more efficient and more resilient to failure than traditional architectures with standalone, shared, and dual-controller storage arrays.

When VMs move from one hypervisor to another, such as during live migration or a high availability (HA) event, the now local CVM serves a newly migrated VM's data. While all write I/O occurs locally, when the local CVM reads old data stored on the now remote CVM, the local CVM forwards the I/O request to the remote CVM. The DSF detects that I/O is occurring from a different node and migrates the data to the local node in the background, ensuring that all read I/O is served locally as well.

The next figure shows how data follows the VM as it moves between hypervisor nodes.

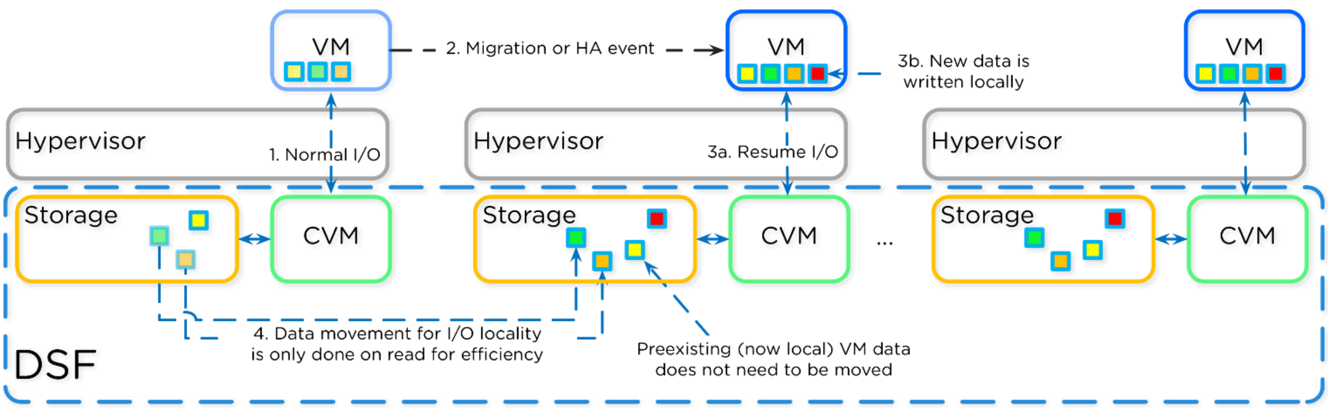

Figure 4: Data Locality and Live Migration

4. Application Overview

4.1. What Is Citrix XenDesktop?

Citrix XenDesktop is a desktop virtualization solution that transforms desktops and applications into a secure, on-demand service available to any user, anywhere, on any device. With XenDesktop, you can deliver individual Windows, web, and SaaS applications, or full virtual desktops, to PCs, Macs, tablets, smartphones, laptops, and thin clients with a high-definition user experience.

Citrix XenDesktop provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure.

The core components of XenDesktop are:

- Desktop Delivery Controller

  Installed on servers in the datacenter, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops. It controls the state of the desktops, starting and stopping them based on demand and administrative configuration. The Citrix license needed to run XenDesktop also includes profile management in some editions to manage user personalization settings in virtualized or physical Windows environments.

- Studio

  Citrix Studio is the management console that allows you to configure and manage your Citrix XenDesktop environment. It provides different wizard-based deployment or configuration scenarios to publish resources using desktops or applications.

- Virtual Desktop Provisioning powered by Citrix Machine Creation Services

  Machine Creation Services (MCS) is the building mechanism of the Citrix Desktop Delivery Controller that automates and orchestrates the deployment of desktops using a single image. MCS communicates with the orchestration layer of your hypervisor, providing a robust and flexible method of image management.

- Virtual Desktop Provisioning powered by Citrix Provisioning Services

  Provisioning Services (PVS) creates and provisions virtual desktops from a single desktop image on demand, optimizing storage utilization and providing a pristine virtual desktop to each user every time they log on. Desktop provisioning also simplifies desktop images, provides optimal flexibility, and offers fewer points of desktop management for both applications and desktops.

- Virtual Desktop Agent

  Installed on virtual desktops, the agent enables direct FlexCast Management Architecture (FMA) connections between the virtual desktop and user devices.

- Citrix Receiver

  Installed on user devices, the Citrix Desktop Receiver enables direct ICA connections from user devices to virtual desktops.

- Citrix FlexCast

  Citrix XenDesktop with FlexCast delivery technology lets you deliver virtual desktops and applications tailored to meet the diverse performance, security, and flexibility requirements of every worker in your organization through a single solution. Centralized, single-instance management helps you deploy, manage, and secure user desktops more easily and efficiently.

Deployment Scenario: Machine Creation Services

Machine Creation Services (MCS) provides images only to desktops virtualized on a hypervisor. The images are contained within the hypervisor pool and then thin-provisioned as needed. The thin-provisioned virtual desktops use identity management functionality to overcome the new security identity (SID) requirements typical of cloning. XenDesktop Controllers integrate with and manage MCS, which uses the underlying hypervisor's capabilities.

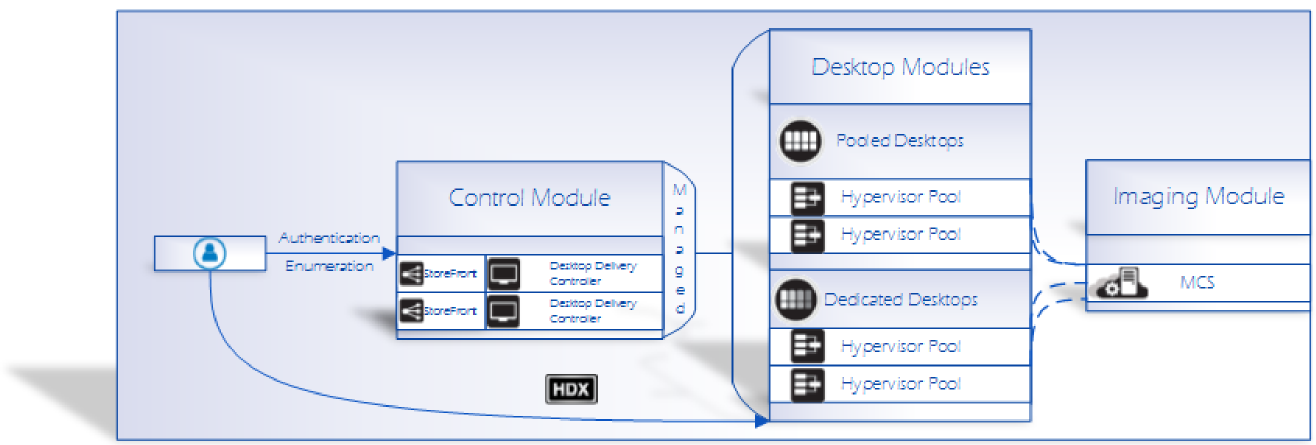

Figure 5: Machine Creation Services

MCS does not require additional servers; it uses integrated functionality built into Citrix XenServer, Microsoft Hyper-V, and VMware vSphere. As MCS uses hypervisor functionality, it is only a viable option for desktops virtualized on a hypervisor. A master desktop image is created and maintained within the hypervisor pool. The XenDesktop Controller instructs the hypervisor to create a snapshot of the base image and thin-provision new VMs through the built-in hypervisor functions. However, thin provisioning images often results in cloning issues, as each provisioned desktop has the same identity as the master. MCS uses special functionality within the XenDesktop Controller and XenDesktop Virtual Desktop Agent (installed within the virtual desktop image) to build unique identities for each VM; these identities are stored within the virtual desktop's identity disk. This functionality allows each virtual desktop to be unique even though it uses the same base image.

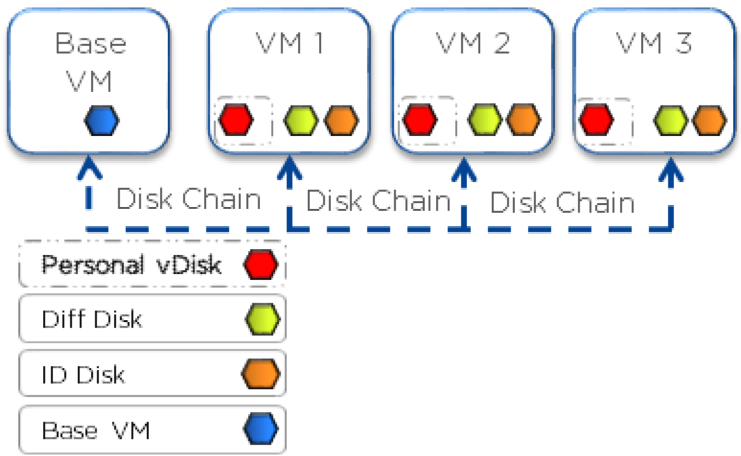

Figure 6: Machine Creation Services: vDisks

Deployment Scenario: Provisioning Services

Provisioning Services (PVS) streaming technology allows computers to be provisioned in real-time from a single shared-disk image. Administrators manage all images on the master image instead of managing and patching individual systems. The local hard-disk drive of each system may be used for runtime data caching or, in some scenarios, removed from the system entirely, which reduces power usage, system failure rates, and security risks. PVS can stream these images to both virtual and physical devices.

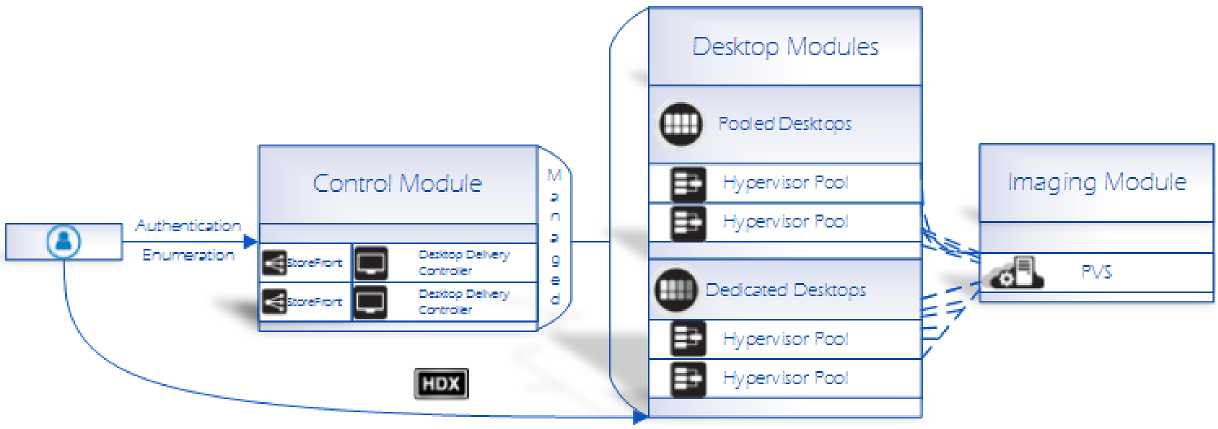

Figure 7: Provisioning Services

The PVS solution’s infrastructure is based on software-streaming technology. After installing and configuring PVS components, a vdisk is created from a device’s hard drive by taking a snapshot of the OS and application image and then storing that image as a vdisk file on the network. A device that is used during this process is called a master target device. The devices that use those vdisks are called target devices.

vDisks can exist on a provisioning server, file share, or (in larger deployments) on a storage system with which the provisioning server can communicate (iSCSI, SAN, NAS, and CIFS). vDisks can be assigned to a single target device (Private Image Mode) or to multiple target devices (Standard Image Mode).

When a target device is turned on, it starts up from the network and communicates with a provisioning server. Unlike thin-client technology, processing takes place on the target device (see the following figure).

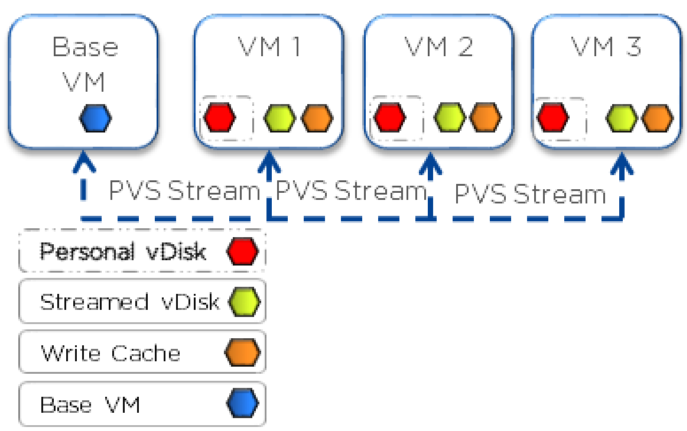

Figure 8: Provisioning Services: vDisks

4.2. Citrix XenDesktop the Nutanix Way

The Nutanix platform operates and scales Citrix XenDesktop MCS and PVS. The figure below illustrates the XenDesktop on Nutanix solution.

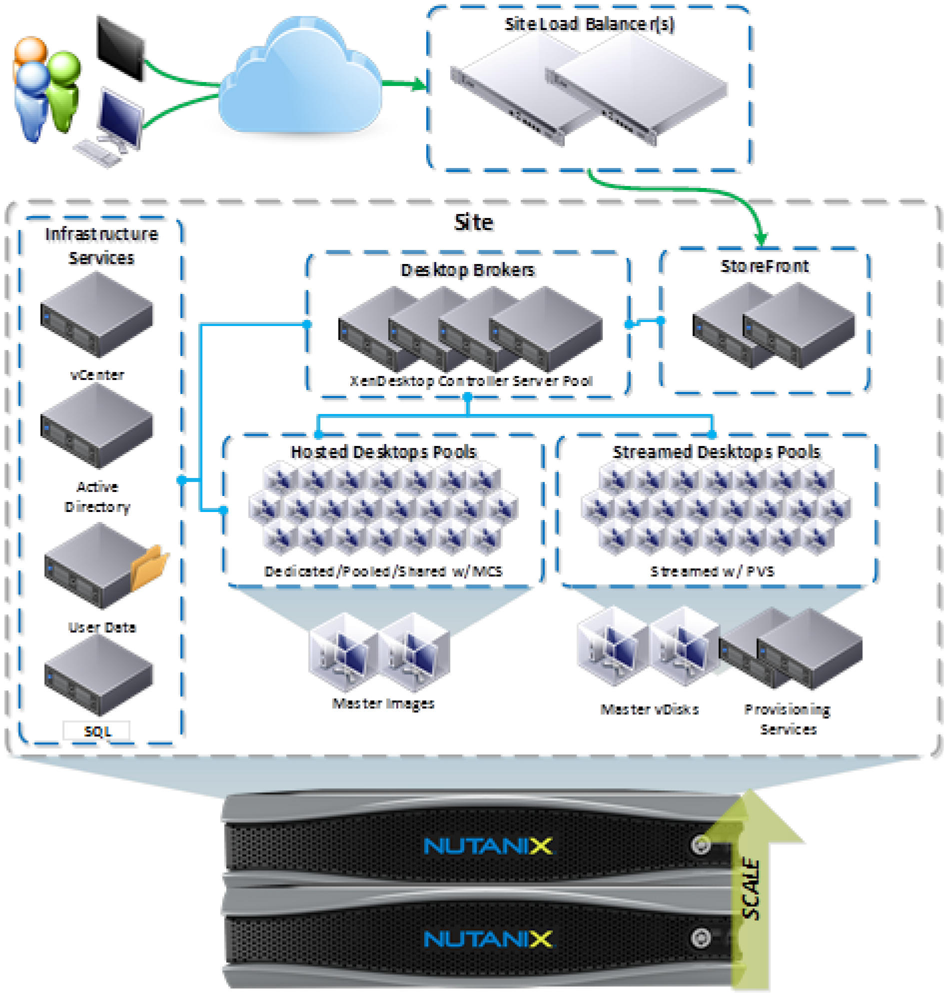

Figure 9: XenDesktop on Nutanix Conceptual Arch

Nutanix modular scale-out design enables customers to select any initial deployment size and then grow in granular data and desktop increments. Customers can realize a faster time-to-value for their XenDesktop implementation because it removes the hurdle of a large initial infrastructure purchase.

The Nutanix solution is fully integrated with the VMware APIs for Array Integration (VAAI) and provides high-performance SSD flash for optimal end user experience along with the flexibility of a single modular platform.

Running Citrix XenDesktop on Nutanix enables you to run multiple workloads all on the same, scalable converged infrastructure while achieving these benefits:

- Modular incremental scale: A single Nutanix block provides up to 20 TB storage and 400 desktops in a compact 2RU footprint. Nutanix modular design lets you scale by node (up to approximately 5 TB per 100 desktops); by block (up to approximately 20 TB per 400 desktops); or with multiple blocks, matching supply with demand and minimizing the upfront CapEx.
- Integrated: The Nutanix platform provides full support for VAAI, allowing you to take advantage of the latest advancements from VMware and to optimize your VDI solution.
- High performance: By using memory caching for read I/O and flash storage for write I/O, you can deliver high performance throughput in a compact 2RU four-node cluster.
- Change management: Maintain environmental control and separation between development, test, staging, and production environments. Snapshots and fast clones can help in sharing production data with nonproduction jobs, without requiring full copies and unnecessary data duplication.
- Business continuity and data protection: User data and desktops are mission critical; Nutanix protects this vital information with enterprise-grade data management features, such as backup and DR.
- Data efficiency: The Nutanix solution offers truly VM-centric compression policies. Unlike traditional solutions that perform compression mainly at the LUN level, Nutanix provides these capabilities at the VM and file level, greatly increasing efficiency and simplicity and ensuring the highest possible compression and decompression performance on a subblock level. Nutanix allows both inline and post-process compression, breaking the bounds set by traditional compression solutions.
- Enterprise-grade cluster management: A simplified and intuitive approach to managing large clusters, including a converged GUI that serves as a central point for servers and storage, alert notifications, and a bonjour mechanism that automatically detects new nodes in the cluster. These features let you spend more time enhancing your environment rather than maintaining it.
- High-density architecture: Nutanix uses an advanced server architecture in which eight Intel CPUs (up to 96 cores) and up to 3 TB of memory are integrated into a single 2RU appliance. Coupled with data archiving and compression, Nutanix can reduce desktop hardware footprints by up to 5x.
- Time-sliced clusters: Like public cloud EC2 environments, Nutanix can provide a truly converged cloud infrastructure, allowing you to run your server and desktop virtualization on a

single converged cloud. Get the efficiency and savings you require with a converged cloud on a truly converged architecture.

5. Solution Design

With the Citrix XenDesktop on Nutanix solution, you gain the flexibility to start small with a single block and scale up incrementally a node, a block, or multiple blocks at a time, without any impact to performance.

In the following section we cover the design decisions and rationale for XenDesktop deployments on the Nutanix Enterprise Cloud Platform.

Table 2: Solution Design Decisions

Item	Detail	Rationale
	General	
Minimum Size	1x Nutanix block (4 ESXi hosts)	Minimum size requirement
Scale Approach	Incremental modular scale	Allow for growth from PoC (hundreds of desktops) to massive scale (thousands of desktops)
Scale Unit	Node(s), block(s), or pod(s)	Granular scale to precisely meet the capacity demands Scale in nx node increments
	VMware vSphere	
Cluster Size	Up to 12–32 vSphere hosts (minimum of 3 hosts	Isolated fault domains VMware best practice
Clusters per vCenter	Up to 2x24 or 4x12 host clusters	Task parallelization
Datastore(s)	1x Nutanix DSF datastore per pod (XenDesktop Server VMs, Provisioning Services Store, VM clones, VAAI clones, and so on) (Max: 2,000 machines per container)	Nutanix handles I/O distribution/ localization n-Controller model

Item	Detail	Rationale
Infrastructure Services	Small deployments: shared cluster Large deployments: dedicated cluster	Dedicated infrastructure cluster for larger deployments (best practice)
	Nutanix	
Cluster Size	Up to 16 nodes	Isolated fault domains
Storage Pool(s)	1x storage pool (SSD, SATA SSD, SATA HDD)	Standard practice ILM handles tiering
Container(s)	1x container for VMs 1x container for data (not used here)	Standard practice
Features and Enhancements	Increase CVM Memory to 32 GB Turn on MapReduce Dedupe	MapReduce dedupe needs 32 GB of RAM to be enabled
	Citrix XenDesktop	
XenDesktop Controllers	Min: 2 (n+1) Scale: 1 per additional pod	HA for XenDesktop Controllers
Users per Controller	Up to 5,000 users	XenDesktop best practice
Load Balancing	Citrix NetScaler	Ensures availability of controllers Balances load between controllers and pods
	Citrix Provisioning Services	
PVS Servers	Min: 2 (n+1) Scale: 1 per additional pod	HA for PVS servers
Users per PVS Server	Up to 1,500 streams	PVS best practice

Item	Detail	Rationale
Load Balancing	Provisioning services farm	Ensures availability of PVS servers Balances load between PVS servers and pods
vDisk Store	Dedicated disk on Nutanix	Standard practice
Write Cache	On local hard drive	Best practice if the storage can provide enough I/O
	Citrix StoreFront	
StoreFront Servers	Min: 2 (n+1)	HA for StoreFront servers
Load Balancing	Citrix NetScaler	Ensures availability of StoreFront servers Balances load between StoreFront servers
	Citrix NetScaler (If Used)	
NetScaler Servers	Min: 2	HA for NetScaler (active-passive)
Users per NetScaler Server	See product data sheet	Varies per model
Load Balancing	NetScaler HA	Ensures availability of NetScaler servers Balances load between NetScaler servers and pods

Highlights from a high-level snapshot of the Citrix XenDesktop on Nutanix Pod are shown in the table below.

Table 3: Pod Highlights

Item	Quantity
Control Pod	
# of vCenter Server(s)	1
# of XenDesktop Controller(s)	2

Item	Quantity
# of XenDesktop StoreFront Server(s)	2
Services Pod	
# of Nutanix Blocks	Up to 4
# of ESXi Hosts	Up to 32
# of Nutanix Cluster(s)	1
# of Datastore(s)	1

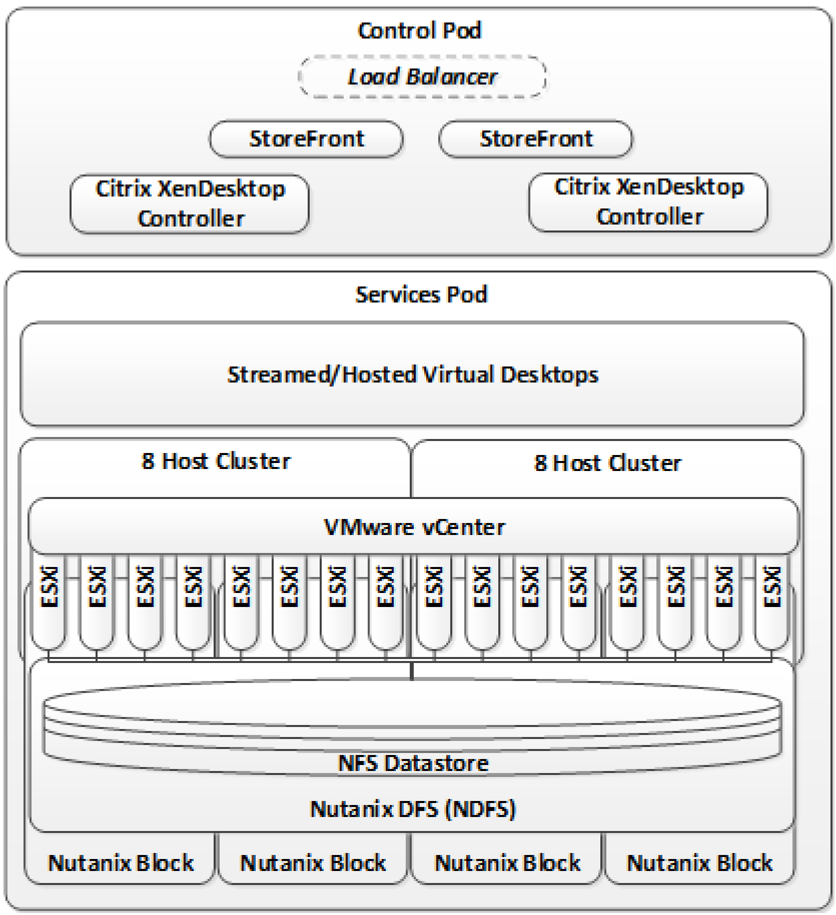

Figure 10: XenDesktop Pod Overview

The section below describes the desktop sizing and considerations for hosted virtual and streamed desktops.

The following are examples of typical scenarios for desktop deployment and use based on the Login VSI definition.

Table 4: Desktop Scenario Definition

Scenario	Definition
Task Workers	Task workers and administrative workers perform repetitive tasks within a small set of applications, usually at a stationary computer. The applications are usually not as CPU and memory intensive as the applications used by knowledge workers. Task workers who work specific shifts might all log in to their virtual desktops at the same time. Task workers include call center analysts, retail employees, and warehouse workers.
Knowledge Workers	Knowledge workers' daily tasks include accessing the Internet, using email, and creating complex documents, presentations, and spreadsheets. Knowledge workers include accountants, sales managers, and marketing research analysts.
Power Users	Power users include application developers and people who use graphics-intensive applications.

The following table proposes some initial recommendations for desktop sizing for a Windows 7 desktop.

> **Note:** These are recommendations for sizing and should be modified after a current state analysis.

Table 5: Desktop Scenario Sizing

Scenario	vCPU	Memory	Disks
Task Workers	1	1.5 GB	35 GB (OS)
Knowledge Workers	2	2 GB	35 GB (OS)
Power Users	2	4 GB	35 GB+ (OS)

5.1. Desktop Optimizations

We used the following high-level desktop optimizations for this design:

- Size desktops appropriately for each particular use case.
- Use a mix of applications installed in gold images and application virtualization, depending on the scenario.
- Disable unnecessary OS services and applications.
- Redirect home directories or use a profile management tool for user profiles and documents.

For more detail on desktop optimizations refer to the Citrix XenDesktop Windows 7 Optimization Guide.

5.2. XenDesktop Machine Creation Services

Citrix MCS uses a standardized model for hosted virtual desktop creation. Using a base or master VM, MCS creates clone VMs that consist of a delta and identity disk and that link back to the base VM's disks.

The figure below shows the main architectural components of an MCS deployment on Nutanix and the communication path between services.

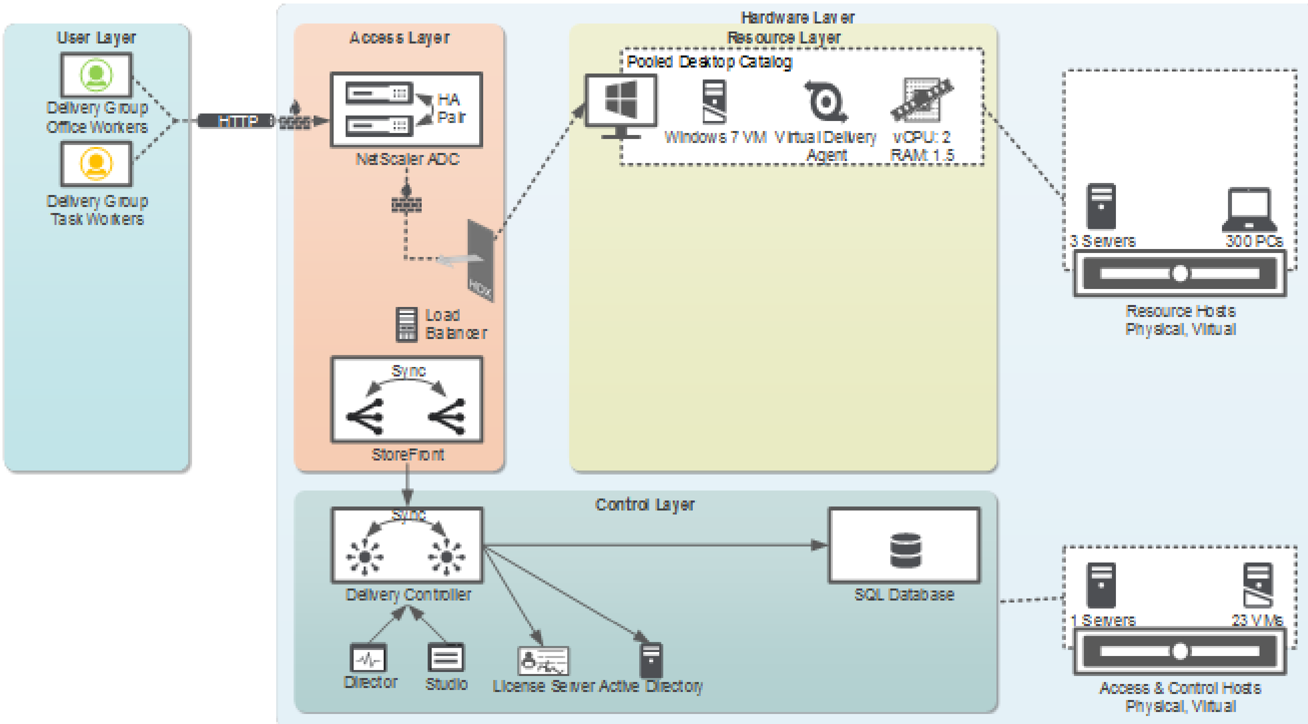

Figure 11: MCS Communication

MCS Pod Design

The following table shows highlights from a high-level snapshot of the Citrix XenDesktop on Nutanix Hosted Virtual Desktop Pod.

Table 6: MCS Pod Detail

Item	Quantity
Control Pod	
# of vCenter Server(s)	1
# of XenDesktop Controller(s)	2
# of XenDesktop StoreFront Server(s)	2
Services Pod	
# of Nutanix Blocks	Up to 4

Item	Quantity
# of ESXi Hosts	Up to 16
# of Nutanix Cluster(s)	1
# of Datastore(s)	1
# of Desktops	Up to 1,200

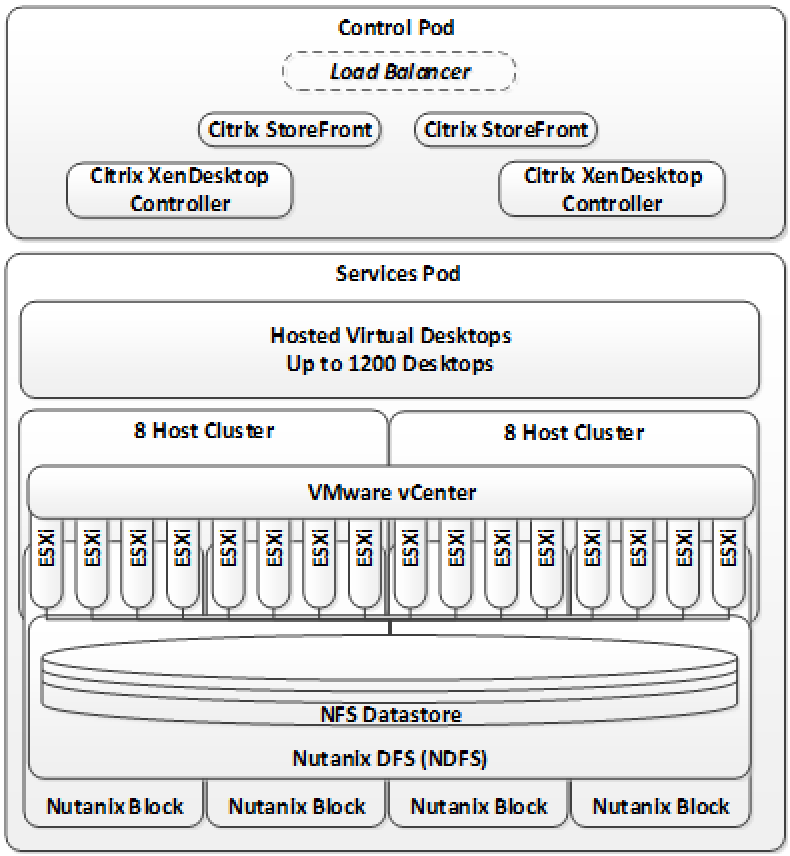

Figure 12: MCS Pod Detail

Hosted Virtual Desktop I/O Path with MCS

The figure below describes the high-level I/O path for an MCS-based desktop on Nutanix. As shown, all I/O operations are handled by the DSF and occur on the local node to provide the highest possible I/O performance. Read requests for the master VM occur locally for desktops hosted on the same ESXi node and over 10 GbE for desktops hosted on another node.

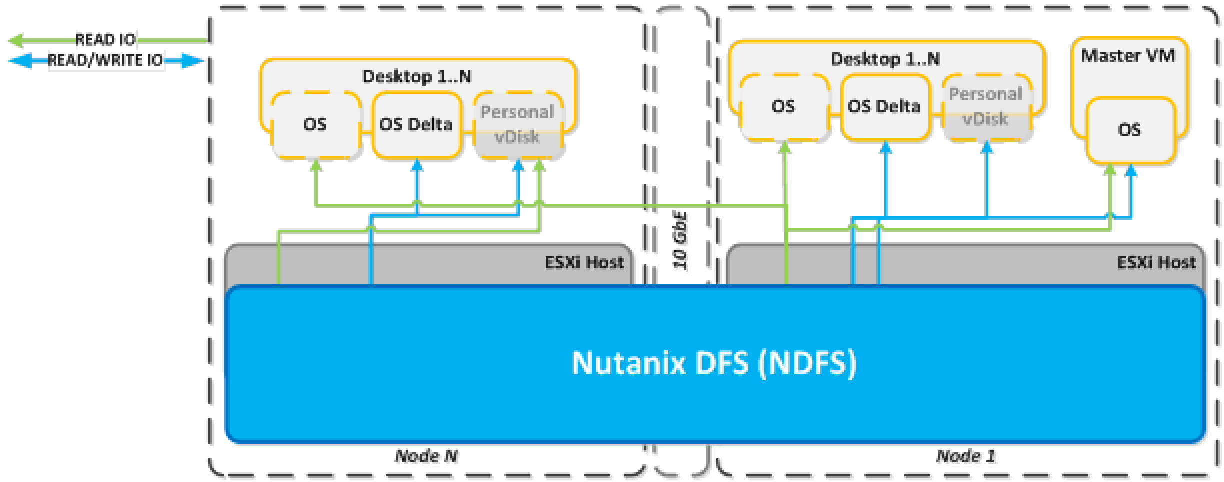

Figure 13: MCS I/O Overview

The following figure describes the detailed I/O path for an MCS-based desktop on Nutanix. All write I/O occurs locally on the local node's SSD tier to provide the highest possible performance. Read requests for the master VM occur locally for desktops hosted on the same ESXi node and over 10 GbE for desktops hosted on another node. These reads are served from the high-performance read cache (if cached) or the SSD tier. Each node also saves frequently accessed data in the read cache for any local data (delta disks and, if used, personal vdisks). Nutanix ILM continually monitors data and the I/O patterns to choose the appropriate tier placement.

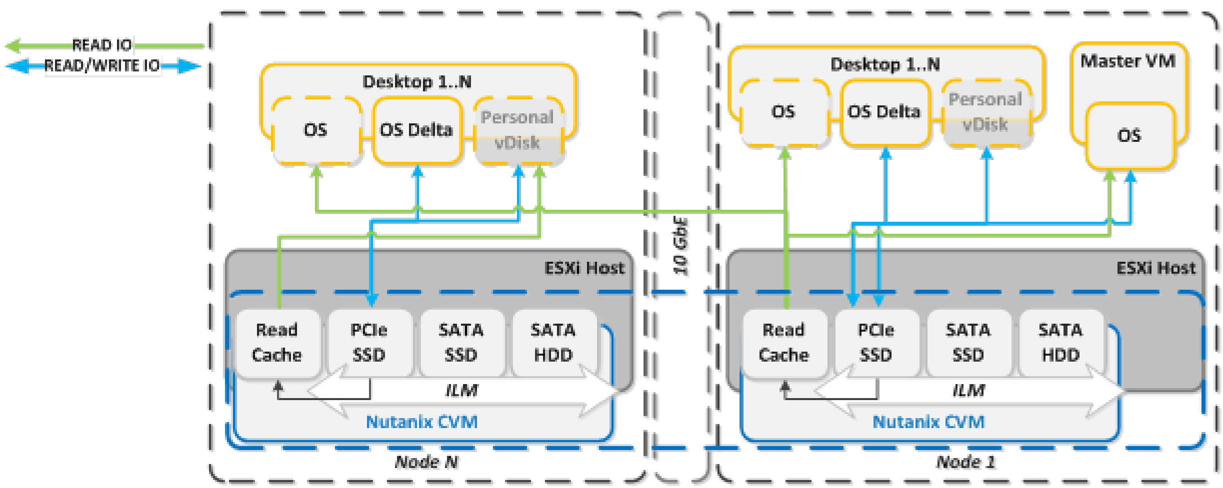

Figure 14: MCS I/O Detail

5.3. XenDesktop Provisioning Services

Citrix PVS streams desktops over the network from a centralized store of master vdisks (OS images). The PVS server stores the vdisks and the Citrix Stream service delivers them. During startup, the streamed desktop pulls the configuration using PXE/TFTP and then initiates communication with the PVS server to continue starting the vdisk.

The next figure shows the main architectural components of a PVS deployment on Nutanix and the communication path between services.

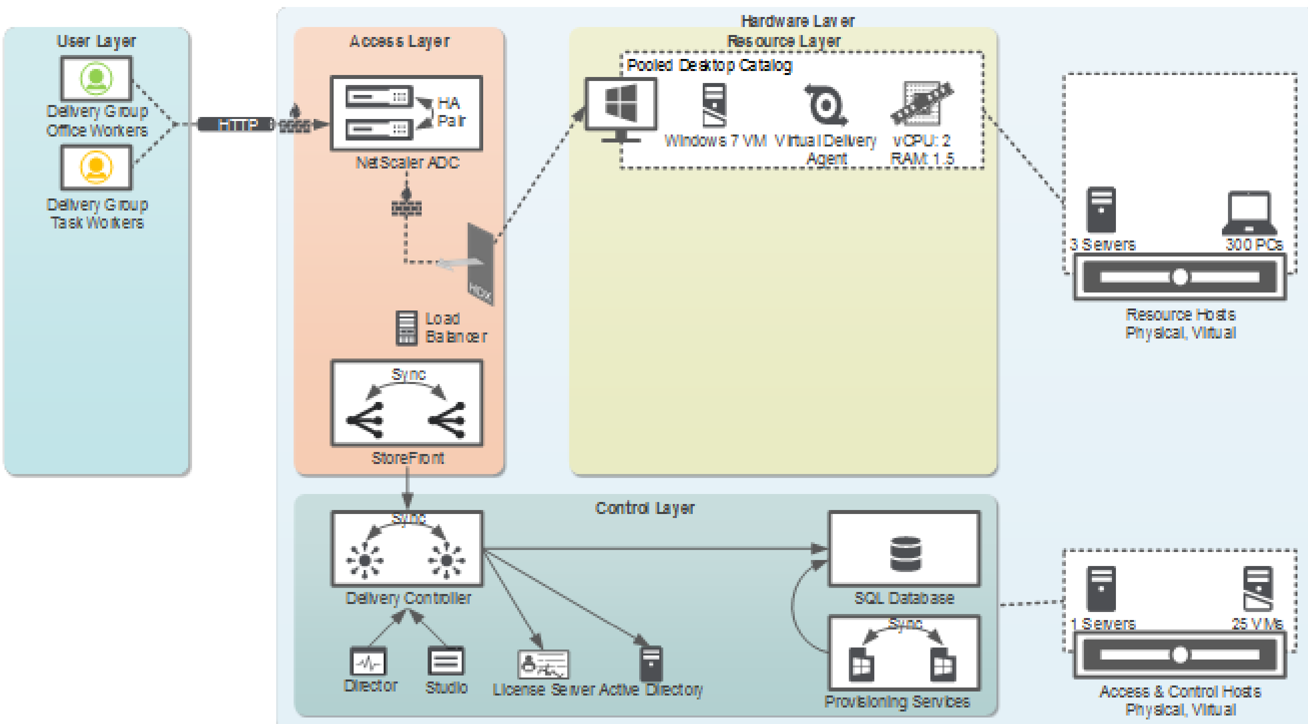

Figure 15: PVS Communication

PVS Pod Design

The table below highlights a high-level snapshot of the Citrix XenDesktop on Nutanix Streamed Desktop Pod.

Table 7: PVS Pod Detail

Item	Quantity
Control Pod	
# of vCenter Server(s)	1
# of XenDesktop Controller(s)	2
# of XenDesktop StoreFront Server(s)	2
# of PVS Server(s)	2
Services Pod	
# of Nutanix Blocks	Up to 4
# of ESXi Hosts	Up to 16
# of Nutanix Cluster(s)	1
# of Datastore(s)	1
# of Desktops	Up to 1,200

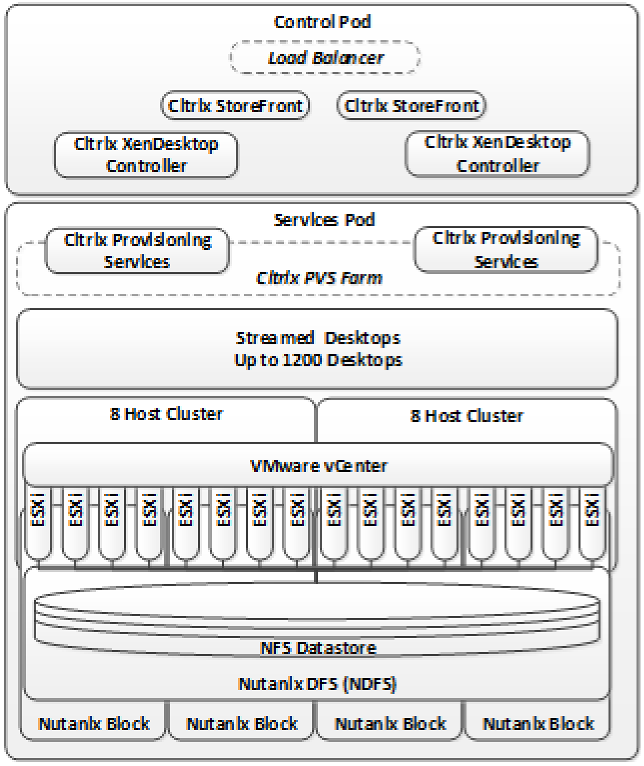

Figure 16: PVS Pod Detail

PVS Store and Network Mapping

The following figure shows the mapping for the PVS server's storage and network. In this case we used dedicated interfaces for both PVS server management and stream services.

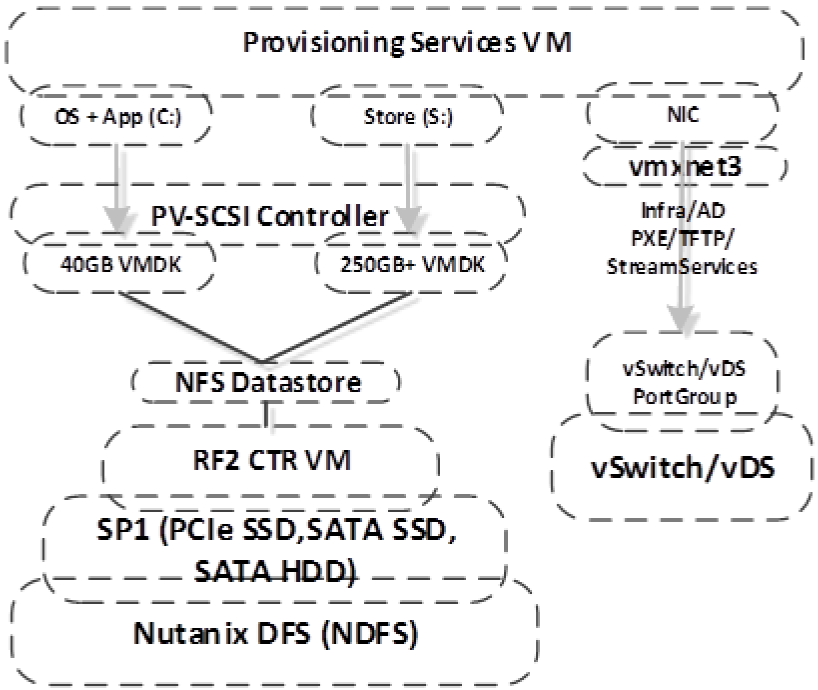

Figure 17: PVS Component Mapping

Streamed Desktop I/O Path with PVS

The next figure describes the high-level I/O path for a streamed desktop on Nutanix. All write I/O operations are handled by the DSF and occur on the local node to provide the highest possible I/O performance. Streamed desktops hosted on the same server as the PVS host are handled by the host's local vSwitch and do not use the external network.

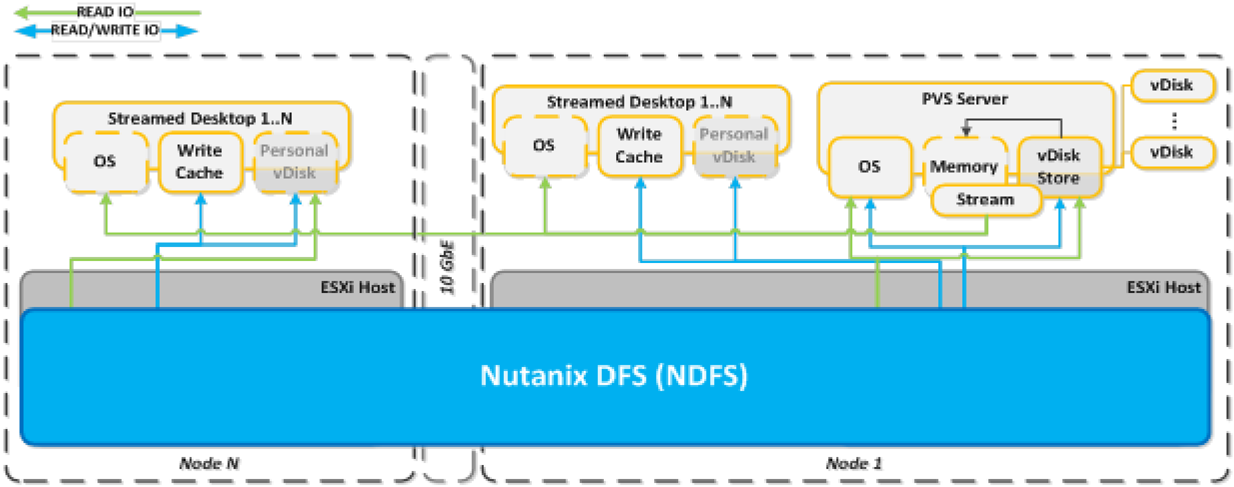

Figure 18: PVS I/O Overview

The figure below describes the detailed I/O path for a streamed desktop on Nutanix. All write I/O (write cache or, if used, personal vdisks) occurs locally on the local node's SSD tier to provide the highest possible performance. The PVS server's vdisk store is hosted on the local node's SSD tier and is also cached in memory. All read requests from the streamed desktop are then streamed either from the PVS server's memory or its vdisk store, which is hosted on the DSF. Each node saves frequently accessed data in the read cache for any local data (write cache, personal vdisks). Nutanix ILM continually monitor data and I/O patterns to choose the appropriate tier placement.

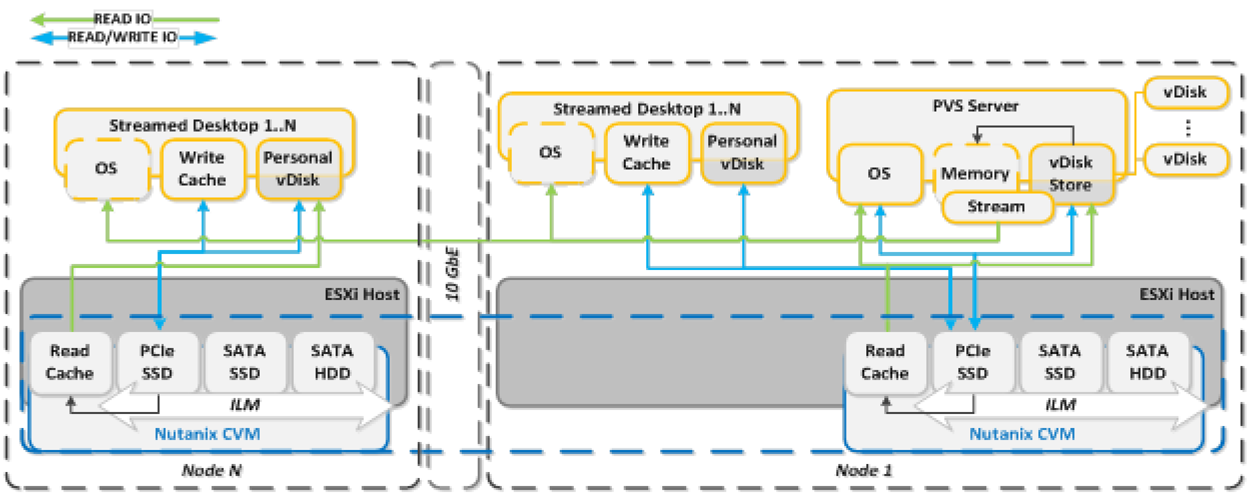

Figure 19: PVS I/O Detail

5.4. Nutanix: Compute and Storage

The Nutanix Enterprise Cloud Platform provides an ideal combination of high-performance compute with localized storage to meet any demand. True to this capability, this reference architecture contains zero reconfiguration of or customization to the Nutanix product to optimize for this use case.

The following figure shows a high-level example of the relationship between a Nutanix block, node, storage pool, and container.

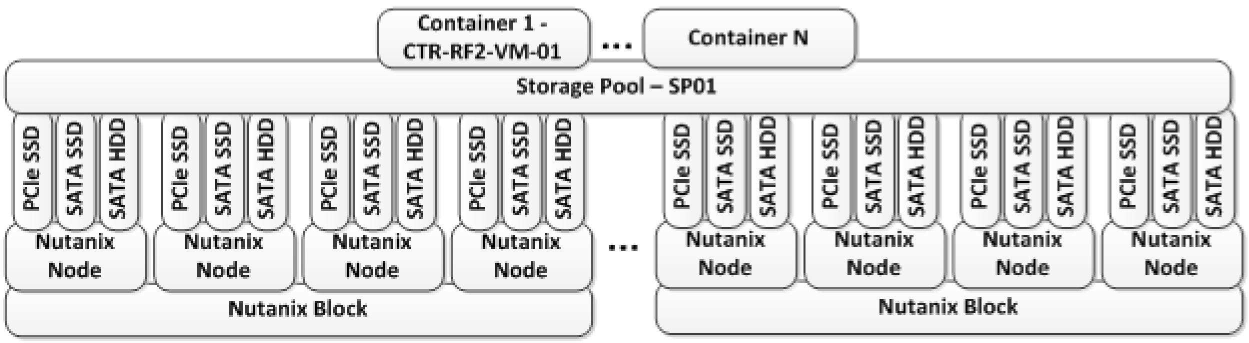

Figure 20: Nutanix Component Architecture

The table below shows the Nutanix storage pool and container configuration.

Table 8: Nutanix Storage Configuration

Name	Role	Details
SP01	Main storage pool for all data	PCI-e SSD, SATA, SSD, SATA-HDD
CTR-RF2-VM-01	Container for all VMs	ESXi: Datastore
CTR-RF2-DATA-01	Container for all data (not used here)	ESXi: Datastore

5.5. Network

Designed for true linear scaling, we use a leaf-spine network architecture. A leaf-spine architecture consists of two network tiers: an L2 leaf and an L3 spine based on 40 GbE and nonblocking switches. This architecture maintains consistent performance without any throughput reduction thanks to a static maximum of three hops from any node in the network.

The figure below shows a design of a scaled-out leaf-spine network architecture that provides 20 Gb active throughput from each node to its L2 leaf and scalable 80 Gb active throughput from each leaf-to-spine switch, providing the ability to scale from one Nutanix block to thousands without any impact to available bandwidth.

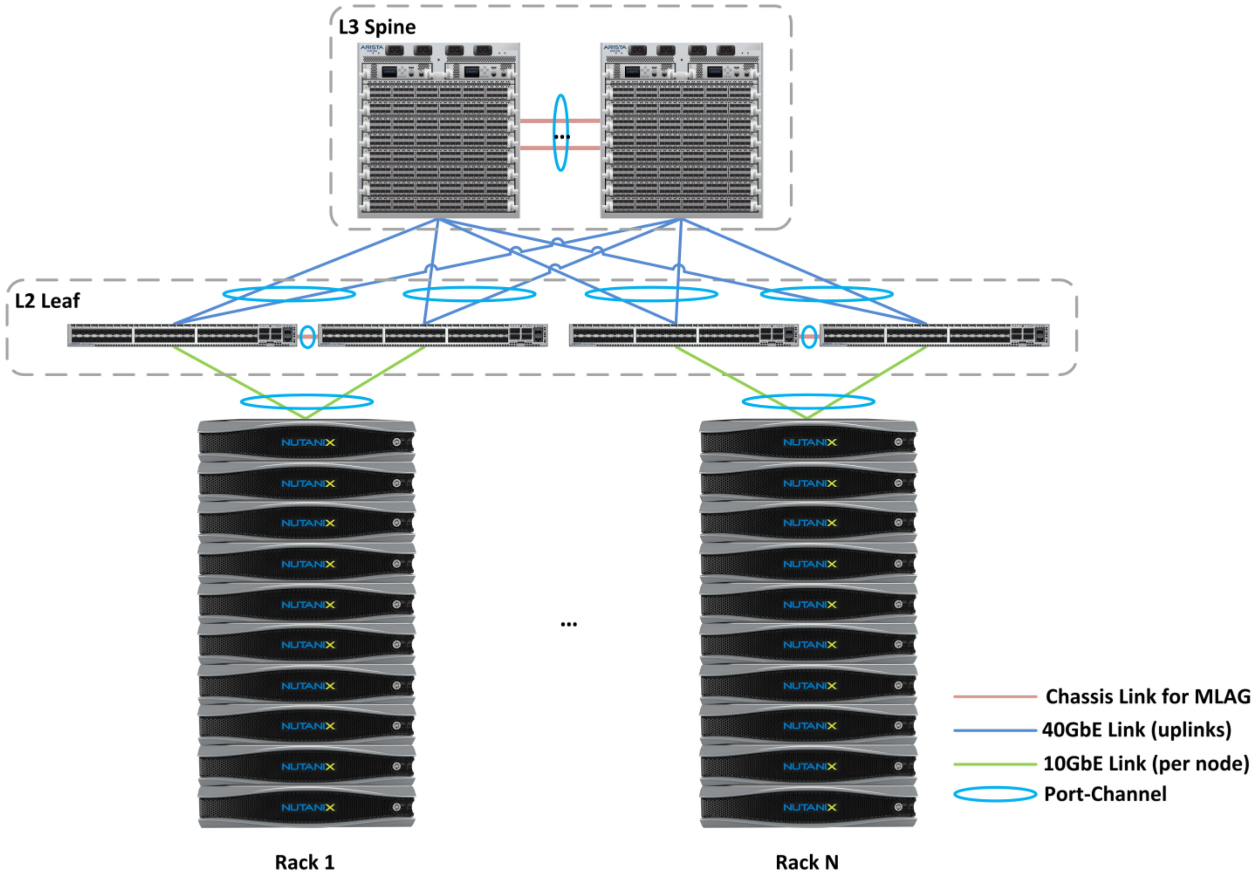

Figure 21: Leaf-Spine Network Architecture

6. Solution Application

This section applies the recommended pod-based reference architecture to real-world scenarios and outlines the sizing metrics and components.

Note: Detailed hardware configuration and product models can be found in the appendix.

6.1. Scenario: 400 Desktops

Table 9: Detailed Component Breakdown: 400 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	1 (partial)	# of vCenter Servers	1
# of Nutanix Blocks	1	# of ESXi Hosts	4
# of RU (Nutanix)	2	# of vSphere Clusters	1
# of 10 GbE Ports	8	# of Datastore(s)	1
# of 100/1000 Ports (IPMI)	4		
# of L2 Leaf Switches	2		
# of L3 Spine Switches	1		

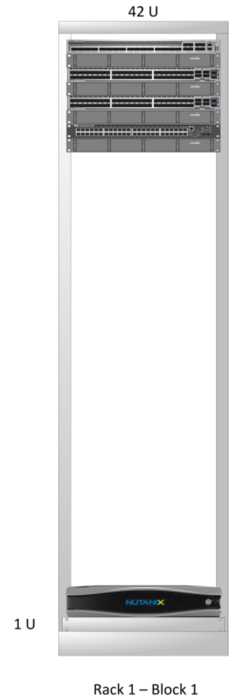

Figure 22: Rack Layout: 400 Desktops

6.2. Scenario: 800 Desktops

Table 10: Detailed Component Breakdown: 800 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	1 (partial)	# of vCenter Servers	1
# of Nutanix Blocks	2	# of ESXi Hosts	8
# of RU (Nutanix)	4	# of vSphere Clusters	1
# of 10 GbE Ports	16	# of Datastore(s)	1
# of 100/1000 Ports (IPMI)	8		
# of L2 Leaf Switches	2		
# of L3 Spine Switches	1		

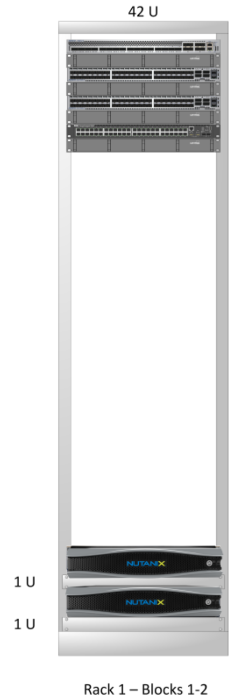

Figure 23: Rack Layout: 800 Desktops

6.3. Scenario: 1,600 Desktops

Table 11: Detailed Component Breakdown: 1,600 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	1	# of vCenter Servers	1
# of Nutanix Blocks	4	# of ESXi Hosts	16
# of RU (Nutanix)	8	# of vSphere Clusters	2
# of 10 GbE Ports	32	# of Datastore(s)	1
# of 100/1000 Ports (IPMI)	16		
# of L2 Leaf Switches	2		
# of L3 Spine Switches	2		

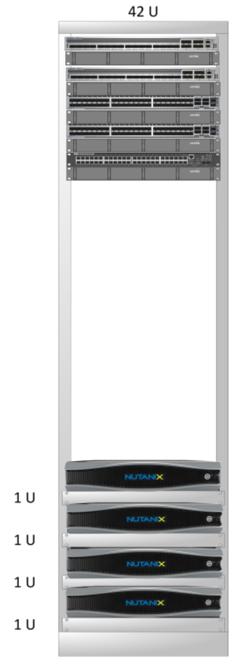

Figure 24: Rack Layout: 1,600 Desktops

6.4. Scenario: 3,200 Desktops

Table 12: Detailed Component Breakdown: 3,200 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	2	# of vCenter Servers	1
# of Nutanix Blocks	8	# of ESXi Hosts	32
# of RU (Nutanix)	16	# of vSphere Clusters	1
# of 10 GbE Ports	64	# of Datastore(s)	2
# of 100/1000 Ports (IPMI)	32		
# of L2 Leaf Switches	2		
# of L3 Spine Switches	2		

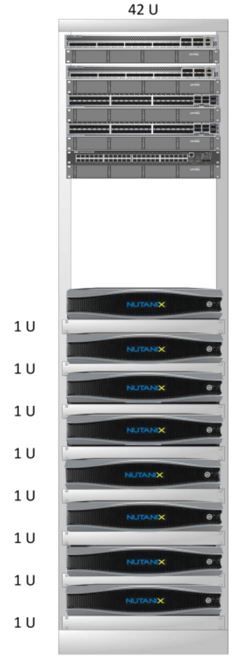

Figure 25: Rack Layout: 3,200 Desktops

6.5. Scenario: 6,400 Desktops

Table 13: Detailed Component Breakdown: 6,400 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	4	# of vCenter Servers	1
# of Nutanix Blocks	16	# of ESXi Hosts	64
# of RU (Nutanix)	32	# of vSphere Clusters	2
# of 10 GbE Ports	128	# of Datastore(s)	4
# of 100/1000 Ports (IPMI)	64		
# of L2 Leaf Switches	4		
# of L3 Spine Switches	2		

Figure 26: Rack Layout: 6,400 Desktops

6.6. Scenario: 12,800 Desktops

Table 14: Detailed Component Breakdown: 12,800 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	8	# of vCenter Servers	2
# of Nutanix Blocks	32	# of ESXi Hosts	128
# of RU (Nutanix)	64	# of vSphere Clusters	8
# of 10 GbE Ports	256	# of Datastore(s)	8
# of 100/1000 Ports (IPMI)	128		
# of L2 Leaf Switches	8		
# of L3 Spine Switches	2		

Figure 27: Rack Layout: 12,800 Desktops

6.7. Scenario: 25,600 Desktops

Table 15: Detailed Component Breakdown: 25,600 Desktops

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix Desktop Pods	16	# of vCenter Servers	2
# of Nutanix Blocks	64	# of ESXi Hosts	256
# of RU (Nutanix)	128	# of vSphere Clusters	16
# of 10 GbE Ports	512	# of Datastore(s)	16
# of 100/1000 Ports (IPMI)	256		
# of L2 Leaf Switches	14		
# of L3 Spine Switches	2		

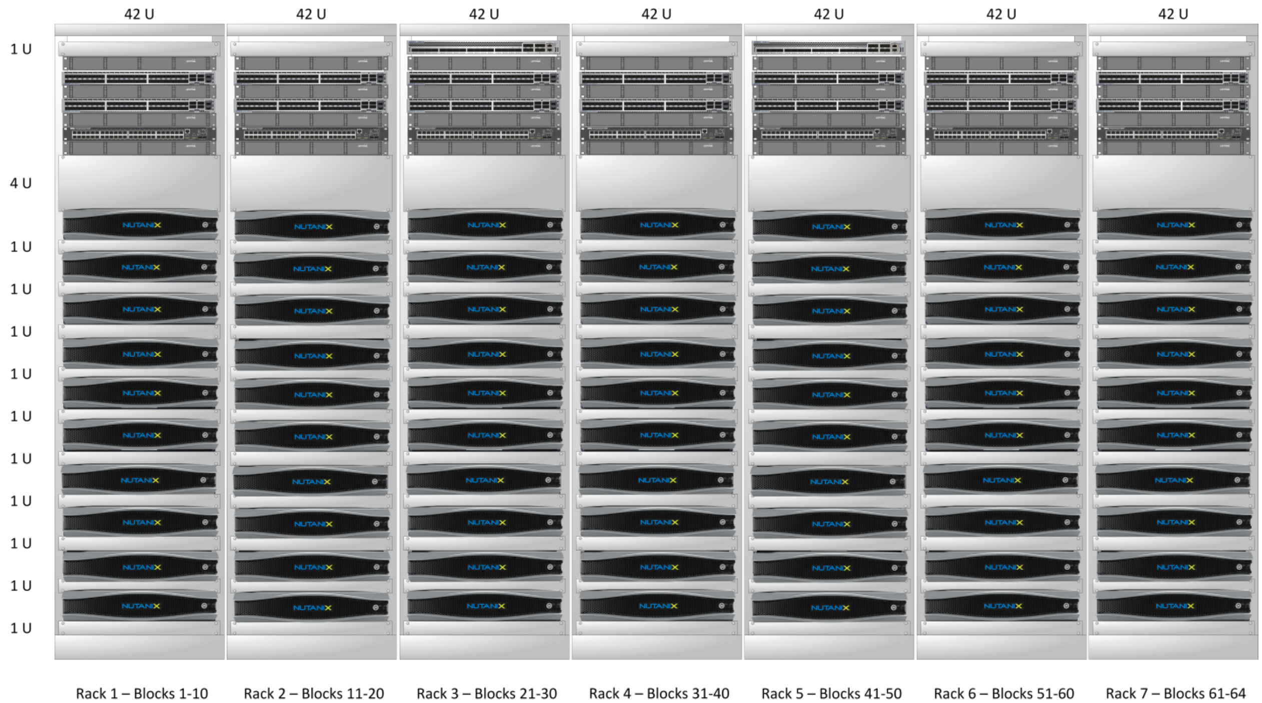

Figure 28: Rack Layout: 25,600 Desktops

7. Validation and Benchmarking

We conducted the solution and testing in this document with Citrix XenDesktop 7.6 deployed on VMware vSphere 5.5, on the Nutanix Enterprise Cloud Platform. Login VSI Office Worker and Knowledge Worker benchmarks supplied the model for the Knowledge Worker on the Nutanix appliance.

7.1. Environment Overview

A single Nutanix NX-3400 node hosted all infrastructure and XenDesktop services, as well as the Login VSI test harness. The three remaining nodes in the Nutanix NX-3400 served as the target environment and provided all desktop hosting. We connected the Nutanix block to an Arista 7050S top-of-rack switch using 10 GbE.

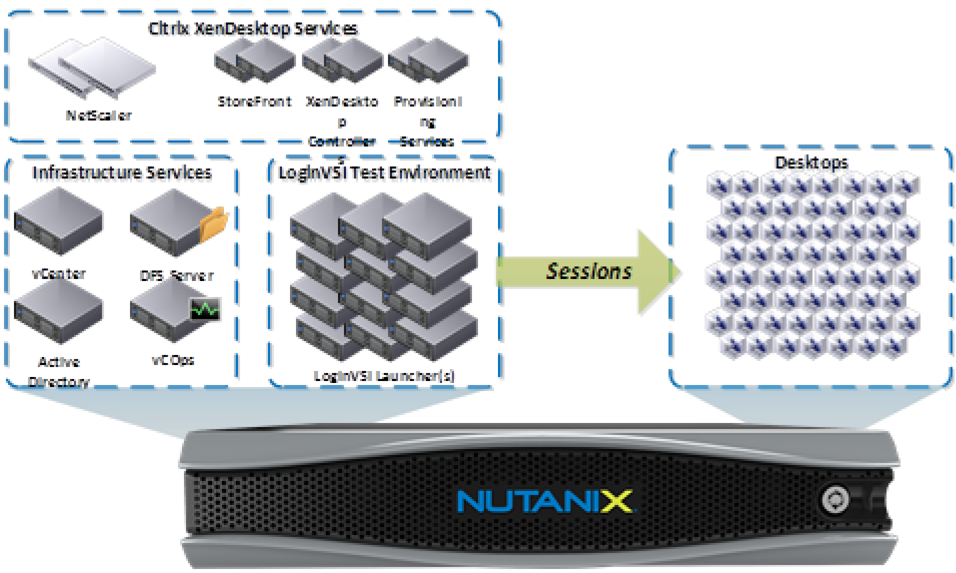

Figure 29: Test Environment Overview

7.2. Test Environment Configuration

Assumptions:

- Knowledge Worker use case

- Per-desktop IOPS (Office Worker): 5 sustained, 70 peak (startup)
- Per-desktop IOPS (Knowledge Worker): 10 sustained, 70 peak (startup)
- Using both MCS and PVS

Hardware:

- Storage and compute: 1 Nutanix NX-3400
- Network: Arista 7050Q (L3 spine) and 7050S (L2 leaf) series switches

Desktop configuration:

- OS: Windows 7 SP1 x86
- 2 vCPU and 2 GB memory
- 1x 35 GB OS disk
- Applications:
  - Microsoft Office 2013
  - Adobe Acrobat Reader XI
  - Internet Explorer
  - Flash Video

Login VSI:

- Login VSI 4.1 Professional

7.3. XenDesktop Configuration:

The table below shows the XenDesktop configuration used in the test environment.

Table 16: XenDesktop Configuration

VM	Quantity	vCPU	Memory	Disks
XenDesktop Controller(s)	2	4	8	1x 40 GB (OS)
PVS Server(s)	2	4	16	1x 40 GB (OS) 1x 250 GB (Store)
StoreFront Server(s)	2	4	4	1x 40 GB (OS)

Test image preparation: MCS

1. Create base VM.
2. Install Windows 7.
3. Install standard software.
4. Optimize Windows 7.
5. Add machine to domain.
6. Install Citrix VDA.
7. Install Login VSI components.
8. Create snapshot.
9. Create clones using XenDesktop Create Machine Catalog Wizard.

Test image preparation: PVS

1. Create base VM.
2. Install Windows 7.
3. Install standard software.
4. Optimize Windows 7.
5. Install PVS target device.
6. Create vDisk.
7. Set bios to start up from PXE.
8. Remove VMDK.
9. Boot VM from vDisk (private mode).
10. Add machine to domain.
11. Install Citrix VDA.
12. Install Login VSI components.
13. Create disk for write cache.
14. Convert to template.
15. Convert vDisk (standard mode).
16. Set cache to local disk.
17. Create clones using XenDesktop setup wizard.

Test execution:

1. Restart or turn on desktops.
2. Restart or start Login VSI launcher(s).
3. Log on to VSI Management Console.
4. Set test parameters and number of sessions.
5. Start test.
6. Wait for test execution to finish.

7. Analyze results (Login VSI).

7.4. Login VSI Benchmark

Login Virtual Session Indexer (Login VSI) is the de-facto industry standard benchmarking tool for testing the performance and scalability of centralized Windows desktop environments like server-based computing (SBC) and virtual desktop infrastructures (VDI).

Login VSI is 100 percent vendor independent and tests virtual desktop environments like Citrix XenDesktop and XenApp, Microsoft VDI, and Remote Desktop Services, VMware View or any other Windows-based SBC or VDI solution. It works with standardized user workloads and thus conclusions based on Login VSI test data are objective, verifiable, and replicable.

For more information about Login VSI visit http://www.loginvsi.com/.

The following table includes all four workloads available on Login VSI 4.1.

Table 17: Login VSI 4.1 Workloads

Task Worker	Office Worker	Knowledge Worker	Power User
Light	Medium	Medium	Heavy
1 vCPU	1 vCPU	2 vCPUs	2–4 vCPUs
2–3 apps	4–6 apps	4–7 apps	5–9 apps
No video	240p video	360p video	720p video

Login VSI Workflows

The Login VSI Workflow base layout is captured in the Login VSI 4.1 Workloads document, which also documents the changes from previous versions of Login VSI to version 4.1 in great detail.

Table 18: Login VSI Workload Definitions

Workload Name	Light	Medium	Heavy	Task Worker	Office Worker	Knowledge Worker	Power User
VSI Version	4	4	4	4.1	4.1	4.1	4.1
Apps Open	2	5–7	8–10	2–7	5–8	5–9	8–12

Workload Name	Light	Medium	Heavy	Task Worker	Office Worker	Knowledge Worker	Power User
CPU Usage	66%	99%	124%	70%	82%	100%	119%
Disk Reads	52%	93%	89%	79%	90%	100%	133%
Disk Writes	65%	97%	94%	77%	101%	100%	123%
IOPS	5.2	7.4	7	6	8.1	8.5	10.8
Memory	1 GB	1 GB	1 GB	1 GB	1.5 GB	1.5 GB	2 GB
vCPU	1 vCPU	2 vCPU	2 vCPU	1 vCPU	1 vCPU	2 vCPU	2 vCPU+

7.5. How to Interpret the Results

Login VSI

Login VSI is a test benchmark used to simulate real-world user workload on a desktop. These values represent the time it takes for an application or task to complete (launching Outlook, for example) and is not in addition to traditional desktop response times. These figures do not refer to the round-trip time (RTT) for network I/O; rather, they refer to the total time to perform an action on the desktop.

During the test, all VMs are turned on and the workload is started on a new desktop every 30 seconds, until all sessions and workloads are active.

We quantified the evaluation using the following metrics:

- Minimum Response: The minimum application response time.
- Average Response: The average application response time.
- Maximum Response: The maximum application response time.
- VSI Baseline: Average application response time of the first 15 sessions.
- VSI Index Average: The VSI index average is the average response time, dropping the highest and lowest two percent.
- VSImax: If reached, the maximum value of sessions launched before the VSI index average gets above the VSI baseline x 125 percent + 3,000 ms.

Based on user experience and industry standards, we recommend keeping ranges for these values below those stated in the table below.

Table 19: Login VSI Metric Values

Metric	Value (ms)	Rationale
Minimum Response	<1,000	Acceptable ideal response time
Average Response	<2,000	Acceptable average response time
Maximum Response	<3,000	Acceptable peak response time
VSI Baseline	<1,000	Acceptable ideal response time
VSI Index Average	<2,000	Acceptable average response time

Login VSI Graphs

The Login VSI graphs show the values defined in the table above during the launching of each desktop session. The following figure shows an example graph of the test data. The y-axis is the response time in milliseconds and the x-axis is the number of active sessions.

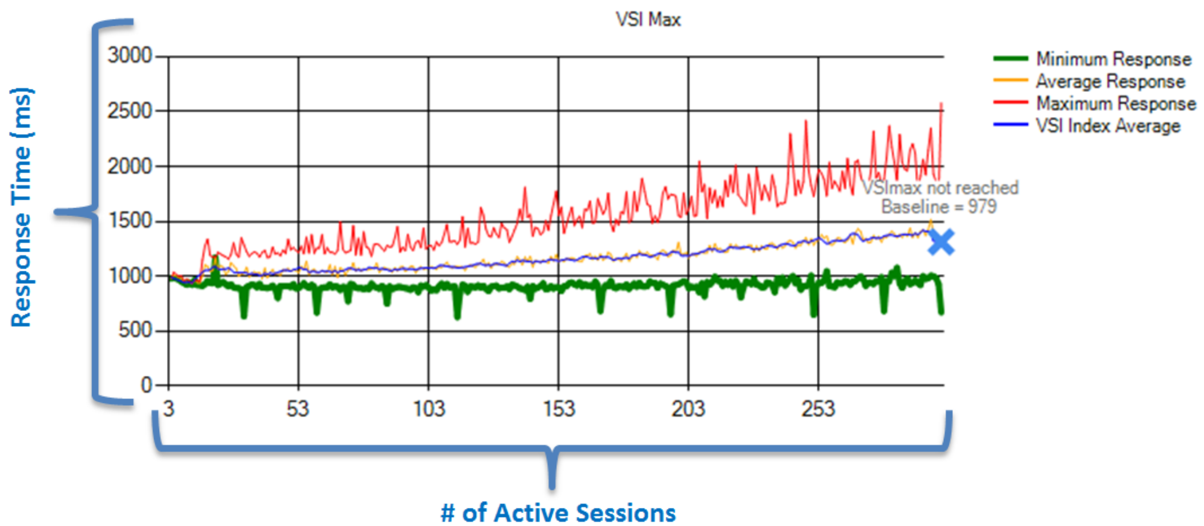

Figure 30: Example Graph of Login VSI Test

8. Validation Results

8.1. MCS: 360 Office Worker Desktops

During the testing with 360 Office Worker desktops, VSImax was not reached, with a baseline of 1,096 ms and average VSI index of 2,006 ms.

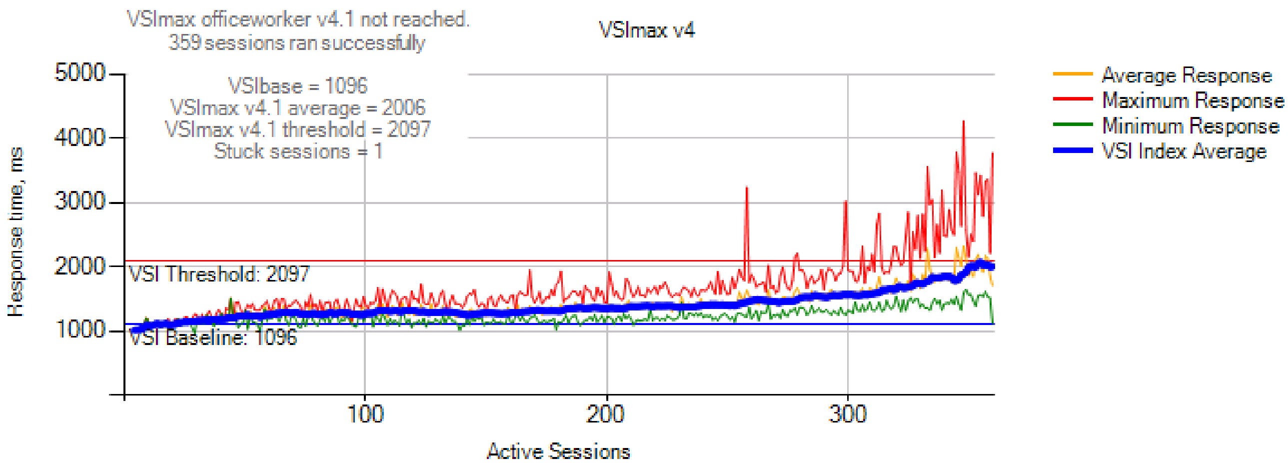

Figure 31: Login VSI 360 Office Worker Desktops with MCS

8.1. Cluster Metrics

The figure below shows user sessions over time, measured with Splunk and UberAgent.

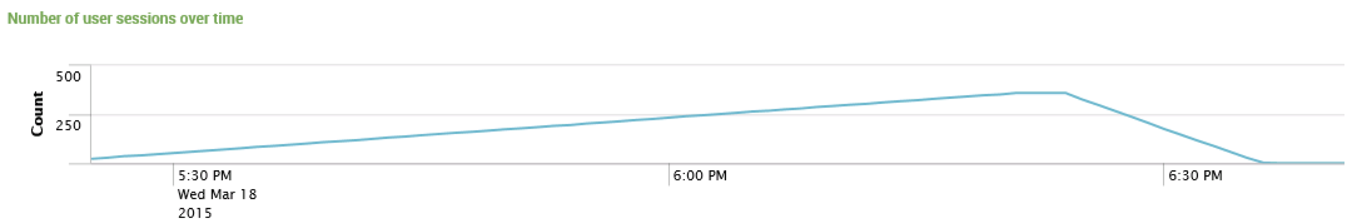

Figure 32: 360 Office Worker User Sessions over Time with MCS

Average logon duration over time during the test was measured with Splunk and UberAgent. The scale is from 0 to 20 seconds.

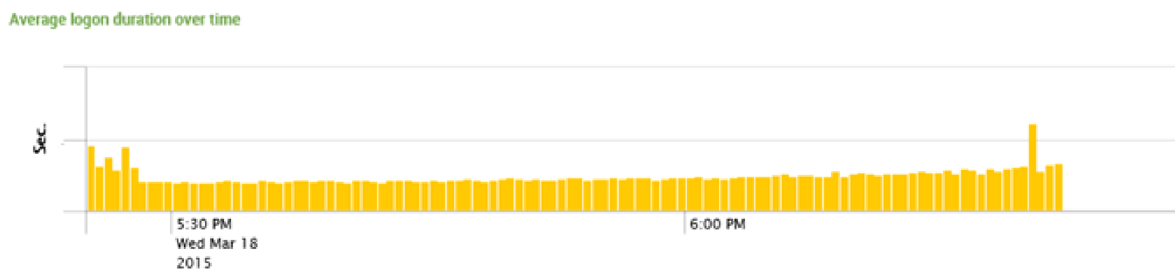

Figure 33: Average Logon Duration for 360 Office Worker Users with MCS

CPU utilization for the ESXi hosts peaked at 99.96 percent and memory utilization peaked at approximately 71.28 percent:

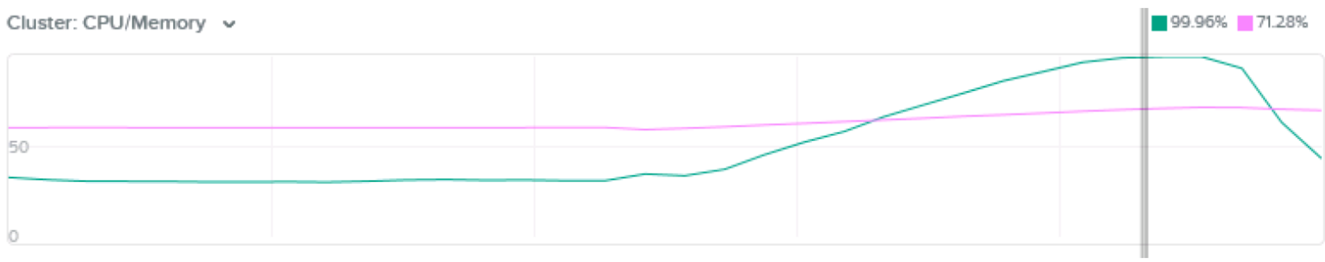

Figure 34: Peak CPU and Memory Utilization for 360 Office Workers with MCS

8.1. Nutanix Datastore Metrics

IOPS peaked at approximately 3,458 during the high-volume startup period to refresh the desktops. The peak IOPS value during the test was a little above 5,059.

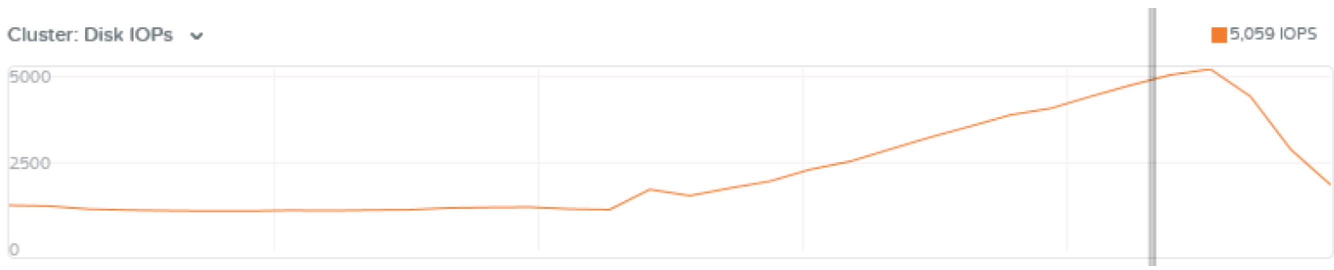

Figure 35: Peak Cluster IOPS for 360 Office Workers MCS

Command latency peaked at approximately 9.91 ms:

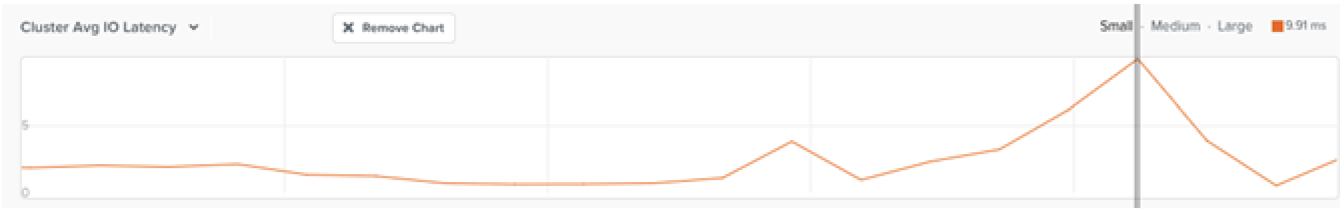

Figure 36: Peak Command Latency for 360 Office Workers with MCS

8.2. MCS: 300 Knowledge Worker Desktops

During the testing with 300 Knowledge Worker desktops, VSImax was not reached, with a baseline of 979 ms and an average VSI index of 1,169 ms.

Figure 37: Login VSI 300 Knowledge Worker Results for MCS

8.2. Cluster Metrics

User sessions over time measured with Splunk and UberAgent:

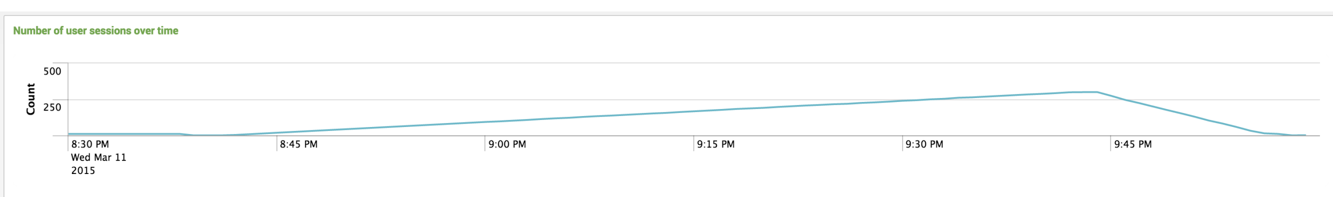

Figure 38: User Sessions over Time for 300 Knowledge Workers with MCS

Average logon duration over time during the test, measured with Splunk and UberAgent. The scale is from 0 to 20 seconds:

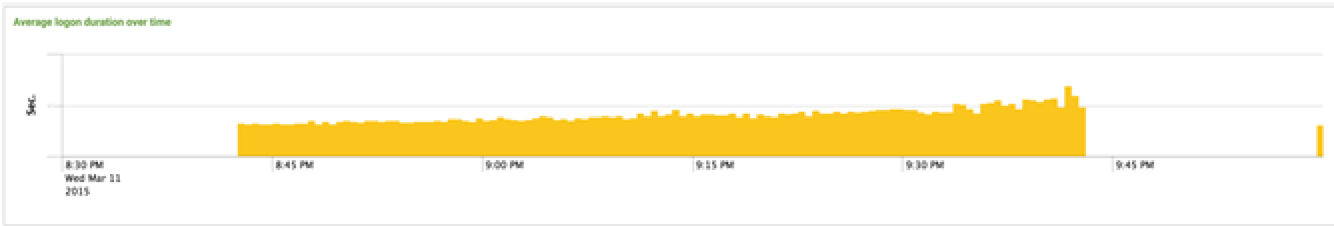

Figure 39: Average Logon Duration over Time for 300 Knowledge Workers with MCS

CPU utilization for the ESXi hosts peaked at 98.2 percent and memory utilization peaked at approximately 82.53 percent.

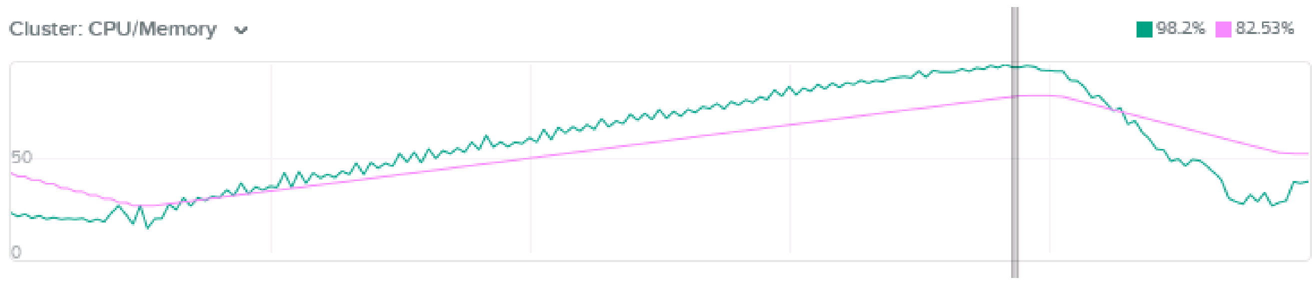

Figure 40: Peak CPU and Memory Utilization for 300 Knowledge Workers with MCS

8.2. Nutanix Datastore Metrics

IOPS peaked at approximately 4,788 during the high-volume startup period to refresh the desktops. The peak IOPS value during the test was 4,837.

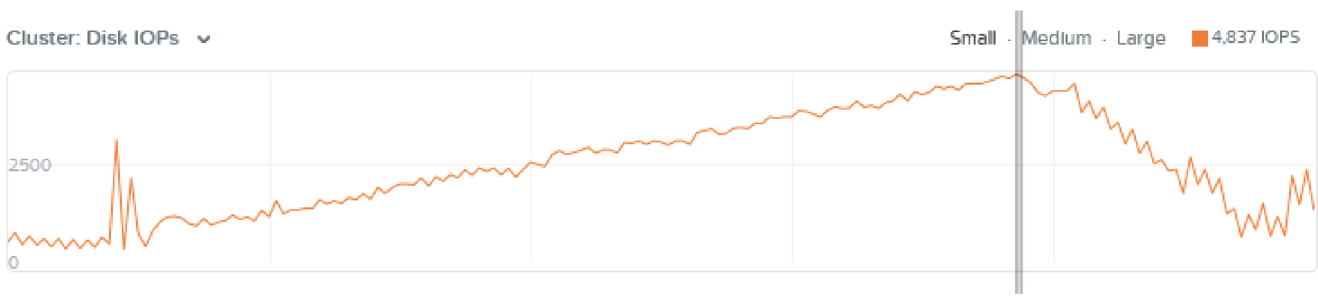

Figure 41: Peak Cluster IOPS for 300 Knowledge Workers with MCS

Command latency peaked at approximately 4.32 ms:

Figure 42: Peak Command Latency for 300 Knowledge Workers with MCS

8.3. PVS: 360 Office Worker Desktops

During the testing with 360 Office Worker desktops, VSImax was not reached, with a baseline of 1,127 ms and average VSI index of 1,791 ms:

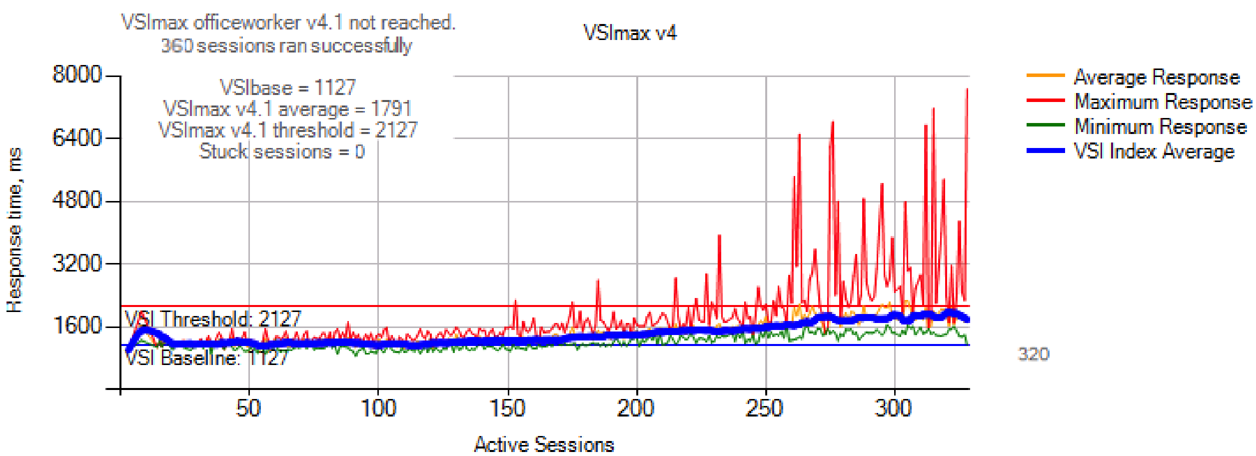

Figure 43: Login VSI 360 Office Worker Results with PVS

8.3. Cluster Metrics

User sessions over time measured with Splunk and UberAgent:

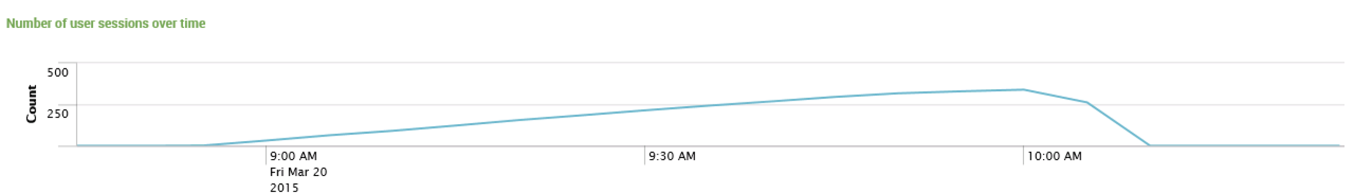

Figure 44: 360 User Sessions over Time with PVS

Average logon duration over time during the test, measured with Splunk and UberAgent. The scale is from 0 to 20 seconds:

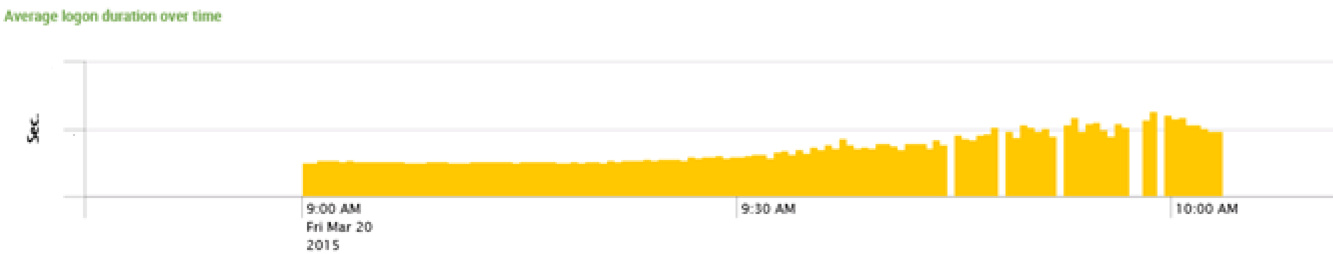

Figure 45: Average Logon Duration for 360 Office Workers with PVS

CPU utilization for the ESXi hosts peaked at 99.36 percent and memory utilization peaked at approximately 71.67 percent:

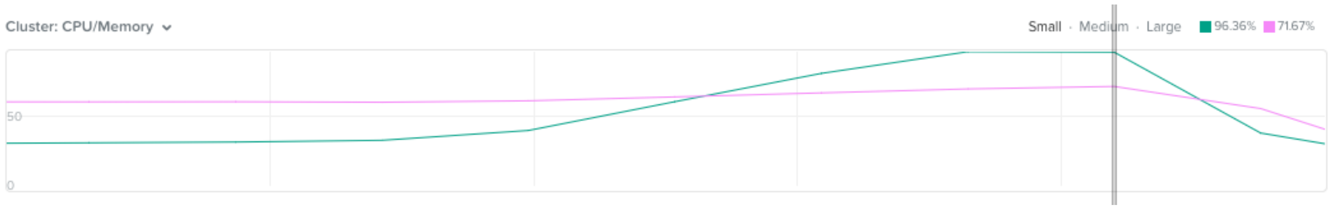

Figure 46: Peak CPU and Memory Utilization for 360 Office Workers with PVS

8.3. Nutanix Datastore Metrics

IOPS peaked at 3,743 during the high-volume startup period to refresh the desktops. The peak IOPS value during the test was a little above 4,807:

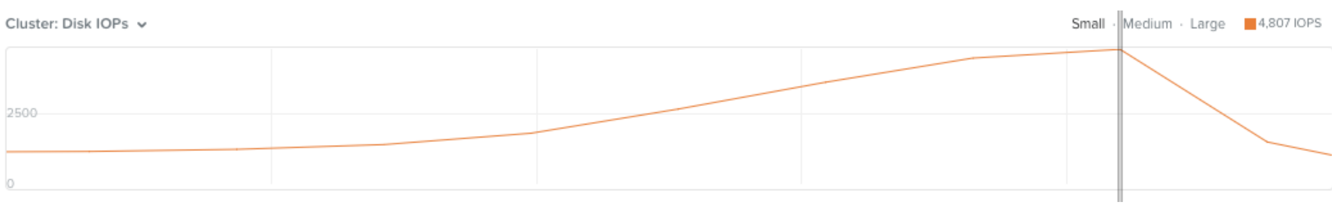

Figure 47: Peak IOPS Volume for 360 Office Works with PVS

Command latency peaked at approximately 10.31 ms:

Figure 48: Command Latency Peak for 360 Office Workers with PVS

8.4. PVS: 300 Knowledge Worker Desktops

During the testing with 300 Knowledge Worker desktops, VSImax was not reached, with a baseline of 721 ms and average VSI index of 1,479 ms:

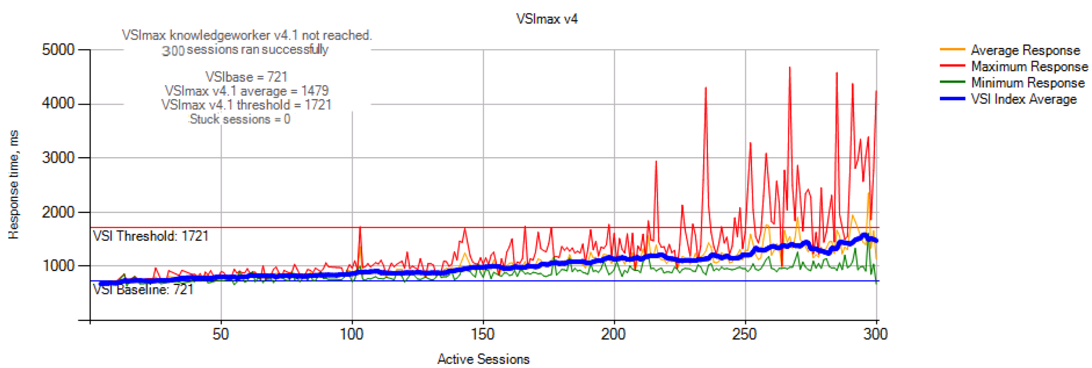

Figure 49: Login VSI 300 Knowledge Worker Results with PVS

8.4. Cluster Metrics

User sessions over time measured with Splunk and UberAgent:

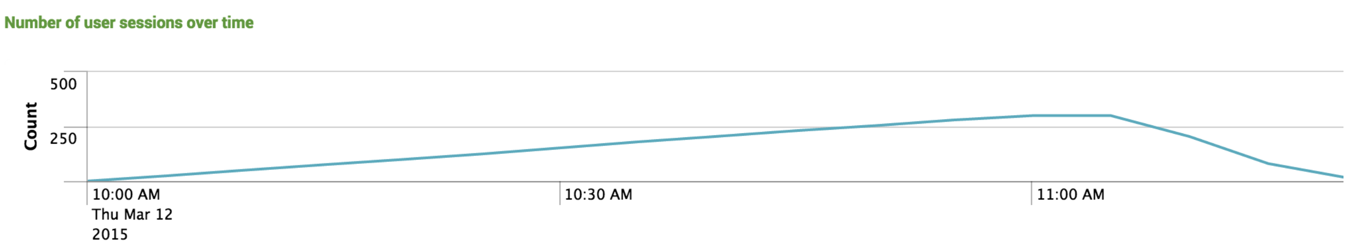

Figure 50: 300 Knowledge Worker User Sessions over Time with PVS

Average logon duration over time during the test, measured with Splunk and UberAgent. The scale is from 0 to 20 seconds:

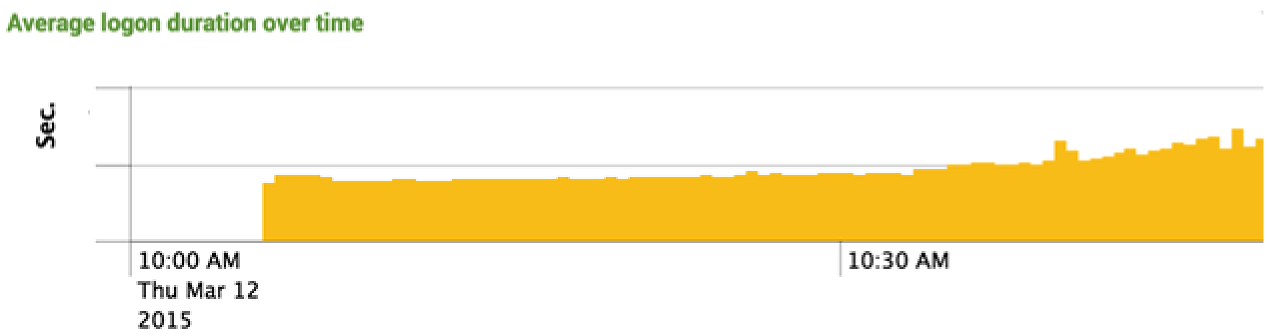

Figure 51: Average Logon Duration for 300 Knowledge Workers with PVS

CPU utilization for the ESXi hosts peaked at 99.42 percent and memory utilization peaked at approximately 83.93 percent:

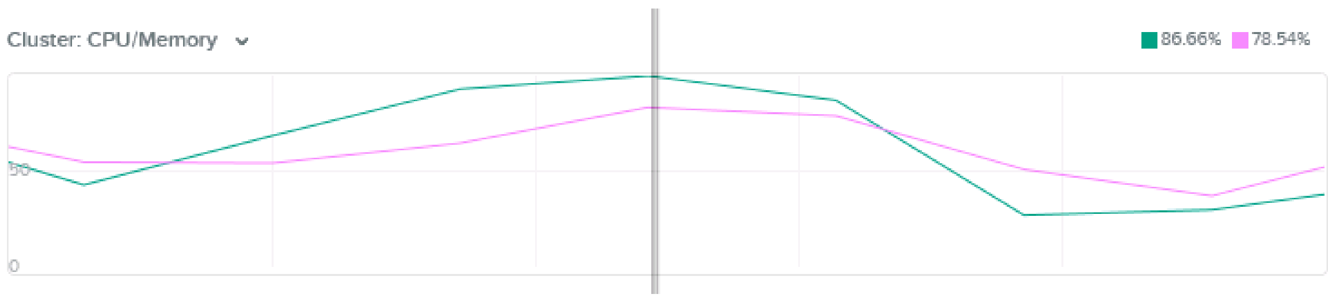

Figure 52: Peak CPU and Memory Utilization for 300 Knowledge Workers with PVS

8.4. Nutanix Datastore Metrics

IOPS peaked at approximately 5,975 during the high-volume startup period to refresh the desktops. The peak IOPS value during the test was 4,542:

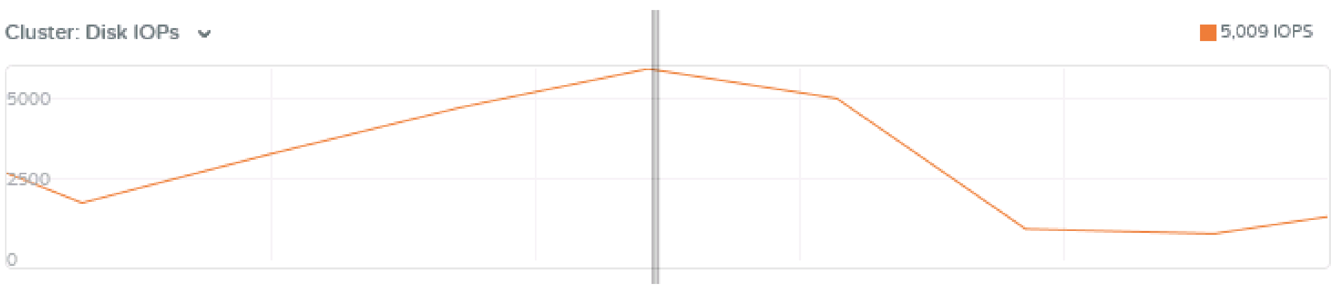

Figure 53: IOPS Volume for 300 Knowledge Workers with PVS

Command latency peaked at approximately 4.18 ms:

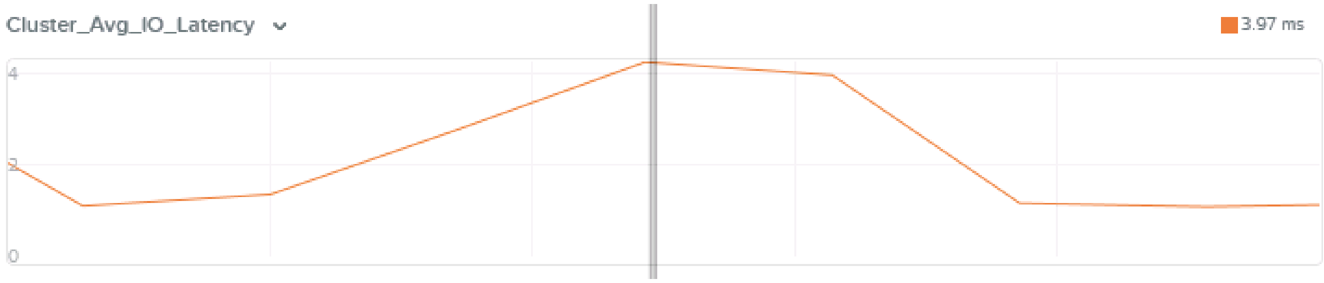

Figure 54: Peak Command Latency for 300 Knowledge Workers with PVS

9. Conclusion

Our extensive testing of MCS and PVS deployments on Nutanix demonstrates that desktop user density is based primarily on the available host CPU resources, not by any I/O or resource constraints. Login VSI Office Worker test results showed densities of more than 480 Office Worker desktops for each 2RU Nutanix appliance. However, most VDI deployments fit within the Knowledge Worker category, which was validated at over 400 desktops for each 2RU appliance. When determining the sizing for the pods, we considered both performance and additional resources for N+1 failover capabilities.

The MCS tests showed light I/O footprints on the Nutanix platform, with a peak of approximately 5,000 aggregate IOPS (during the high-volume startup periods). Sustained IOPS were light, ranging from 4,000 to 5,000. I/O latencies averaged less than 2 ms for read and less than 5 ms for write during peak load.

The PVS tests showed light I/O footprints on the Nutanix platform as well, with a peak of approximately 2,600 aggregate IOPS during the high-volume startup periods. Sustained IOPS were light, ranging from 500 to 2,600. I/O latencies averaged less than 1 ms for read and less than 8 ms for write during peak load. PVS server CPU utilization peaked at approximately 40 percent during the high-volume startup period, with an average steady state at approximately 10 percent.

The Citrix XenDesktop-on-Nutanix solution provides a single, high-density platform for desktop and application delivery. This modular, pod-based approach also enables deployments to scale simply and efficiently with zero downtime.

Appendix

Configuration

Hardware:

- Storage and compute
  - Nutanix NX-3400
- Per-node specs (4 nodes per 2RU block):
  - CPU: 2x Intel Xeon E5-2680
  - Memory: 256 GB
- Network
  - Arista 7050Q: L3 spine
  - Arista 7050S: L2 leaf

Software:

- Nutanix
  - AOS 4.1.1.3
- XenDesktop
  - 7.6
- Provisioning Services
  - 7.6
- Desktop
  - Windows 7 SP1 x86
- Infrastructure
  - ESXi 5.5.2
  - vCenter 5.5.2

VM:

- Desktop
  - CPU: 2 vCPU

    - Memory: 1.5 GB
- Storage
    - 1x 35 GB OS Disk on CTR-RF2-VM-01 DSF-backed NFS datastore

About the Author

Kees Baggerman is a staff solutions architect for End User Computing at Nutanix, Inc. In his role, Kees develops methods for successfully implementing applications on the Nutanix platform. In addition, he delivers customer projects, including defining architectural, business, and technical requirements, creating designs, and implementing the Nutanix solution.

Before working with Nutanix, Kees' main areas of work were migrations and implementations of Microsoft and Citrix infrastructures, writing functional and technical designs for Microsoft infrastructures, Microsoft Terminal Server, or Citrix (Presentation Server, XenApp, XenDesktop, and NetScaler) in combination with RES Workspace Manager and RES Automation Manager.

Kees is a Citrix Certified Integration Architect, Microsoft Certified IT Professional, RES Certified Professional, and RES Certified Trainer. RES Software also named him RES RSVP for seven consecutive years, and Kees was honored as the RES Software Most Valuable Professional of 2011. As a demonstration of his passion for virtualization technology, Kees earned the title of VMware vExpert in 2013–2017. Citrix also named him a Citrix Technology Professional in 2015, 2016 and 2017.

Follow Kees on Twitter at @kbaggerman.

About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix Enterprise Cloud Platform leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization, and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow up on Twitter @nutanix.

List of Figures

List of Tables