

Citrix XenDesktop on AHV

Nutanix Reference Architecture

Version 3.0 • May 2017 • RA-2035

Copyright

Copyright 2017 Nutanix, Inc.

Nutanix, Inc.
1740 Technology Drive, Suite 150
San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws.

Nutanix is a trademark of Nutanix, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Contents

1. Executive Summary.....	5
2. Introduction.....	6
2.1. Login VSI and EUC Environment Validation.....	6
2.2. Audience.....	6
2.3. Purpose.....	6
3. Nutanix Enterprise Cloud Platform Overview.....	8
3.1. Nutanix Acropolis Overview.....	8
3.2. Distributed Storage Fabric.....	9
3.3. App Mobility Fabric.....	9
3.4. AHV.....	10
3.5. Third-Party Hypervisors.....	10
3.6. Nutanix Acropolis Architecture.....	10
3.7. Nutanix Acropolis File Services.....	12
4. Application Overview.....	14
4.1. What Is Citrix XenDesktop?.....	14
4.2. Deployment Scenario: Automated via the MCS Integration.....	15
4.3. Deployment: Infrastructure Components.....	16
4.4. Citrix XenDesktop the Nutanix Way.....	22
5. Solution Design.....	25
5.1. Desktop Optimizations.....	29
6. XenDesktop and Windows 10 Best Practices.....	31
6.1. Optimizations.....	32
6.2. XenDesktop on AHV.....	34
6.3. Nutanix: Compute and Storage.....	35
6.4. Network.....	36

7. Solution Application.....	38
7.1. Scenario: 600 Desktops.....	38
7.2. Scenario: 900 Desktops.....	38
7.3. Scenario: 2,400 Desktops.....	39
7.4. Scenario: 4,800 Desktops.....	39
7.5. Scenario: 9,600 Desktops.....	40
7.6. Scenario: 19,200 Desktops.....	41
7.7. Scenario: 38,400 Desktops.....	41
8. Validation and Benchmarking.....	42
8.1. Environment Overview.....	42
8.2. Test Environment Configuration.....	42
8.3. XenDesktop Configuration.....	43
8.4. Login VSI Benchmark.....	44
8.5. How to Interpret the Results.....	45
9. Results.....	48
9.1. 600 Knowledge Worker Desktops (Four Nodes).....	48
9.2. 900 Knowledge Worker Desktops (Six Nodes).....	49
9.3. 1,200 Knowledge Worker Desktops (Eight Nodes).....	51
10. Conclusion.....	52
Appendix.....	53
Configuration.....	53
Performance.....	54
About the Author.....	54
About Nutanix.....	55
List of Figures.....	56
List of Tables.....	58

1. Executive Summary

Citrix XenDesktop on Nutanix is a powerful solution that offers unrivaled user experience, simple administration, and web-scale flexibility and economics. In this reference architecture, we make recommendations for the design, optimization, and scaling of XenDesktop deployments on AHV. We used Login VSI and an intelligent scripting framework on Nutanix to simulate real-world workloads in a XenDesktop environment. The Nutanix Enterprise Cloud Platform offers high IOPS and low latency. Our testing demonstrates that host CPU resources drive desktop user density, rather than concerns about I/O or resource bottlenecks—you can easily deploy 150 knowledge worker desktops per node, with four nodes per 2RU appliance. In addition to desktop and application performance reliability, when deploying XenDesktop on Nutanix you get unlimited scalability, data locality, AHV clones, and a single datastore. Nutanix also takes the Citrix commitment to simplicity to another level, with streamlined management, reduced rollout time, and lower operating expenses.

2. Introduction

2.1. Login VSI and EUC Environment Validation

Login VSI is the industry-standard performance validation tool that simplifies load testing, benchmarking, and capacity planning. It emulates real-world conditions using synthetic user logons with predefined profiles to ensure end users get the best possible experience.

Because Login VSI is agentless and has minimal infrastructure requirements, it works with any mainstream broker, Windows operating system, and Microsoft Office version.

2.2. Audience

This reference architecture document is part of the Nutanix Solutions Library and is intended for architecting, designing, managing, and supporting Nutanix infrastructures. Readers should be familiar with Nutanix Acropolis, Prism, AHV, and Citrix XenDesktop. The key items and roles that we address focus on enabling a successful design, implementation, and transition to operation.

2.3. 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.
- Architecting a complete Citrix XenDesktop solution on the Nutanix platform.
- 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	December 2015	Original publication.

Version Number	Published	Notes
2.0	May 2016	Updated platform overview and added a description of the full integration for Citrix MCS on Nutanix AHV.
2.1	November 2016	Addressed performance bottleneck due to a bug in qemu and USB-tablet CPU consumption.
3.0	May 2017	Results validated by LoginVSI.

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

3.7. Nutanix Acropolis File Services

Acropolis File Services (AFS) is a scale-out approach that provides Server Message Block (SMB) file services to Windows clients using SMB 2.1. Acropolis file servers are composed of a set of VMs (called FSVMs) on top of the hypervisor. AFS requires at least three FSVMs running on three nodes to satisfy a quorum for high availability.




Figure 5: Acropolis File Servers Run as VMs to Isolate from DSF

Acropolis File Services (AFS) delivers on-demand performance and automated provisioning to provide highly scalable file management. It reduces the administration and configuration time needed to deploy and maintain your environment, providing a public cloud experience within your private cloud.

The Acropolis DSF and Prism make hosted file services highly resilient and easy to use—for example, you can configure the built-in data efficiency features of DSF, such as compression, deduplication, and EC-X, for each individual file server. You can also use the Prism UI to administer network and resource management, Active Directory, fault tolerance, and share creation all in one place, vastly improving operational efficiency.

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 even 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 (DDC)

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. In some editions, the Citrix license needed to run XenDesktop also includes profile management 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 DDC that automates and orchestrates desktop deployment 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 Virtual Desktop Agent (VDA) enables direct FMA (FlexCast Management Architecture) 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.

4.2. Deployment Scenario: Automated via the MCS Integration

Citrix MCS for AHV provides the simplest way to create a machine catalog. In XenDesktop 7.7 and later, MCS tools and services are baked into a default XenDesktop installation and are ready for immediate use.

When creating a machine catalog with MCS:

- A master image is prepared and created based on a standard VM that contains all the software and customization choices that an administrator wants to make available in users' virtual desktops, as well as the Citrix VDA registered to the Citrix DDC(s).
- The administrator defines the size of the VMs by choosing how many vCPUs and how much vMemory each desktop VM should receive.
- The administrator chooses how many VMs to create and add to the machine catalog and defines a naming convention, such as "Win81-MCS-###" (where ### automatically increments as 001, 002, and so on).
- VMs are all created as "linked clones."

Each VM created has two or three disks assigned to it (see the following figure for reference).




Figure 6: MCS Disk Layout

- **Base VM:** A single base disk residing on a shared NFS datastore that has been mounted on the hosts in the cluster configured with the Studio host connection.
- **Identity disk (or ID disk):** A very small disk (16 MB maximum) that contains identity information; this information provides a unique name for the VM and allows it to join Active Directory (AD).
- **Difference disk:** Also known as the write cache, this disk separates the writes from the master disk, while the system still acts as if the write has been committed to the master disk.

The identity disk and the difference disk together make the VM unique. We offer a more comprehensive discussion of the inner workings of MCS and PVS in a [Nutanix Connect blog entry](#).

4.3. Deployment: Infrastructure Components

Desktop Delivery Controllers

When extending your environment, you can choose between scaling up or out. Scaling up lets you provide services to more users per site, and scaling out means adding more sites. The following constraints may influence which option you choose:

- **Failure domains:** Creating multiple sites in turn creates more failure domains, which are smaller, so outages affect fewer users.

- **Geolocation:** You can split XenDesktop sites into zones. A zone can be a physical location; this capability allows for more flexibility, because you can utilize one XenDesktop site (and thus one configuration) over multiple zones.

Redundancy

Nutanix recommends having at least two DDCs per XenDesktop site for redundancy.

Sizing for Scaling

To ensure that your DDCs aren't a bottleneck in your infrastructure, assign enough resources to account for peak endurance moments, such as boot storms.

Citrix [tested](#) the DDC role's scalability and found that a virtual XenDesktop controller (2 vCPUs, 4 GB RAM) can handle more than 2,500 virtual desktop logons within a four-hour period. For more information, refer to the [Citrix XenDesktop Design Handbook](#).

SQL Database and XenDesktop

SQL Database Redundancy

With the move from IMA (Independent Management Architecture) to FMA, SQL database availability became increasingly important. The DDCs use the centralized databases to store static configurations and more dynamic session-related information. Although current running sessions aren't affected when the database becomes unavailable, it is impossible to start new sessions. For this reason, it is important that the databases are highly available via mirroring, clustering, or AlwaysOn Availability Groups. For more information, please refer to the [Citrix XenDesktop and XenApp best practices document](#).

Backups

Back up the databases at regular intervals to mitigate the impact from disasters and to reduce the size of the SQL transaction log.

Scalability

To be sure that sufficient resources are assigned to the SQL Server environment, Nutanix recommends using the [Citrix XenDesktop Database Sizing document](#).

License Server

Citrix License Server Redundancy

As every Citrix product supports a license server outage of up to 30 days without any decrease in functionality, it's a best practice to deploy a single Citrix license server and, when using a virtualization platform, to rely on the HA functionality for these particular VMs.

Scalability

Both Citrix and Nutanix recommend scalability tests to ensure that you size your license server accurately. Citrix's scalability testing indicates that a single Citrix license server (2 cores, 2 GB RAM) can issue approximately 170 licenses per second, or 306,000 licenses per 30 minutes.

Microsoft License Server Redundancy

Nutanix recommends implementing two Remote Desktop Services (RDS) license servers. If the first license server is not available, the second license server can be contacted.

StoreFront

Scaling

Because StoreFront relies on Microsoft IIS (Internet Information Services), scaling out is preferable to scaling up. Together with a load balancing service, scaling out ensures service availability combined with a predictable user experience.

Redundancy

Citrix recommends having at least two StoreFront Servers per site for redundancy. Although you can join them in a server group, this doesn't mean they are utilized equally. You need additional load balancing for an active/active configuration. The Citrix NetScaler appliance is the most common load balancing appliance in Citrix XenApp and XenDesktop environments.

Sizing for Scaling

To ensure that your DDCs aren't a bottleneck in your infrastructure, assign enough resources to account for peak endurance moments, such as boot storms.

User Profile Management

Profile management distributes personal settings to the user's virtual desktops and applications. Often, a profile optimization service that can deliver these settings in an easy, reliable way provides profile management, to ensure a consistent user experience.

Profile management services are limited to the user profile. To provide the user with drive mappings or printers, for example, you should use another solution, such as login scripts or group policy preferences.

Following are some integrated profile management solutions that are available for your XenApp and XenDesktop deployments:

- Microsoft UE-V
- Citrix User Profile Management

- Citrix WEM (User Environment Management)

All of these solutions need a centralized repository for files and folders; Nutanix offers Acropolis File Services (AFS) for this specific use case. See the [Nutanix AFS data sheet](#) and the blog entry [Ten Things You Need to Know about Nutanix Acropolis File Services](#) for more information. Utilizing AFS for user profiles gives you a scalable, highly available file server solution based on the Nutanix Enterprise Cloud Platform's web-scale architecture, ensuring predictable logon times for your end users.

Home Directories

One other common use case for integrated file services are user home directories, typically in a nonpersistent VDI environment or shared hosted desktops (such as with XenApp, which is nonpersistent by nature). All user-created files are stored in either the profile, departmental shared folders, or user home directories. Nutanix AFS offers scalability, predictability, and high availability for your home directories as well.

AOS 4.6 and Later

Nutanix introduced new APIs with AOS 4.6, some of which are directly related to our MCS integration. Among other advantages, these APIs provide Citrix MCS with a way to read and write to the small (16 MB) identity disk associated with every VM.

Provisioning SDK

The Citrix Provisioning SDK is a new addition to the XenDesktop and XenApp platforms for developers and technology partners. It applies the power and flexibility of Citrix MCS to any hypervisor or cloud infrastructure service that you choose.

It does so using the same tried-and-tested technologies that have made MCS successful on Citrix XenServer, Microsoft SCVMM, VMware vSphere, and, more recently, Amazon Web Services (AWS) and Citrix CloudPlatform. Whether your infrastructure resides in your data center or in the cloud, the Citrix Provisioning SDK is your first step toward an integrated provisioning and management solution for desktops and applications.

And integration is only the beginning. Perhaps your infrastructure or cloud service has unique and differentiating features, setting it apart from other offerings in the market. With the SDK, you can expose your chosen functionality and bring it right to the administrator's fingertips in the Citrix Studio management console.

The SDK enables you to create your own provisioning plug-in, which you can add alongside the plug-ins that are installed by default with the XenDesktop and XenApp products. Once your plug-in is installed, the services on the delivery controller discover and load it automatically. It then appears as a new connection type in Citrix Studio, allowing administrators to easily connect, configure, and provision on your chosen infrastructure platform using two key features:

- A set of .NET programming interfaces that are used to call your provisioning plug-in whenever it needs to take action. Your plug-in takes the form of a .NET assembly (DLL) that implements these interfaces. There are several .NET interfaces that a plug-in must implement, but each is designed to be small and easy to understand. Most interfaces are defined with both a synchronous and an asynchronous variant, allowing you to choose the programming pattern that works best. We provide further information about individual interfaces in the Functional Areas section below.
- The Citrix Common Capability Framework, which allows the rest of the product to understand the specific custom features of your infrastructure and how those features should be displayed in Citrix Studio. The framework uses an XML-based high-level description language. Your plug-in produces specifications using this language, allowing Citrix Studio to intelligently adapt its task flows and wizard pages.

Citrix MCS for AHV

The Citrix Provisioning SDK allows us to create the connection between Studio and AHV and gives us access to all the APIs offered through AHV. This integration results in the architecture shown in the following figure.




Figure 7: MCS on AHV: Architecture

By loading the plug-in on an existing XenDesktop 7.7 or later installation, you receive the additional services to both deploy your desktops from the DDC on a Nutanix AHV cluster and have the full native integration that Nutanix supports for SCVMM and vCenter.




Figure 8: MCS on AHV: Site Setup




Figure 9: MCS on AHV: Site Setup—Network




Figure 10: MCS on AHV: Hosting Connection

Power Management for XenDesktop 7.6

Nutanix offers a XenDesktop power management plug-in that allows you to automatically control desktop power state from within XenDesktop tools. The plug-in is available for XenDesktop 7.6 at http://download.nutanix.com/utills/NutanixPlugin-PM-XD76_1.0.msi.

4.4. Citrix XenDesktop the Nutanix Way

The Nutanix modular web-scale architecture lets you start small and then expand to meet increasing demand—a node, a block, or multiple blocks at a time—with no impact on performance. This design removes the hurdle of a large initial infrastructure purchase, generating faster time-to-value for your XenDesktop implementation.

The figure below presents a typical XenDesktop or XenApp deployment on Nutanix.




Figure 11: XenDesktop on Nutanix Conceptual Architecture

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: With the Nutanix solution you can start small and scale up. A single Nutanix block provides dozens of TB of storage and over 600 virtual desktops in a compact footprint. Given the modularity of the solution, you can granularly scale by node, by block, or with multiple blocks, giving you the ability to accurately match supply with demand and minimize the upfront CapEx.

- High performance: By using system memory caching for read I/O and flash storage for read and write I/O, you can deliver high performance throughput in a compact form factor.
- 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 and need enterprise-grade data management features, including backup and disaster recovery.
- Data efficiency: The DSF offers both compression and deduplication to help reduce the storage footprint. The compression functionality is truly VM-centric. Unlike traditional solutions that perform compression mainly at the LUN level, the Nutanix solution provides all of these capabilities at the VM and file level, greatly increasing efficiency and simplicity. These capabilities ensure the highest possible compression and decompression performance even below the block level. By allowing for both inline and post-process compression capabilities, the Nutanix solution breaks 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 the bonjour mechanism to automatically detect new nodes in the cluster. As a result, you can spend more time enhancing your environment rather than maintaining it.
- High-density architecture: Nutanix uses an advanced server architecture that, when using the NX-3000 series as an example, can house eight Intel CPUs (up to 144 cores) and up to 4 TB of memory in a single 2RU appliance. Coupled with data archiving and compression, Nutanix can reduce desktop hardware footprint 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 cloud. Get the efficiency and savings you require with a converged cloud on a truly unified architecture.

5. Solution Design

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

Table 2: Solution Design Decisions

Item	Detail	Rationale
General		
Software Version	Citrix XenDesktop 7.12	
Software Version	Acropolis Operating System (AOS) 5.0.2	
Minimum Size	6 Nutanix Acropolis hosts running AHV (NX-3000: 2 blocks)	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 capacity demands Scale in n x node increments
Nutanix AHV		
Cluster Size	As many as 12–32 hosts (starting with a minimum of 3 hosts)	Isolated fault domains best practice
Datastore(s)	1x Nutanix DSF datastore per pod (XenDesktop Server VMs, VM clones, and so on)	Nutanix handles I/O distribution and localization n-Controller model
Infrastructure Services	Small deployments: Shared cluster Large deployments: Dedicated cluster	Dedicated infrastructure cluster for larger deployments (best practice)

Item	Detail	Rationale
Nutanix Enterprise Cloud Platform		
Cluster Size	As many as 32 nodes	Isolated fault domains best practice
Storage Pool(s)	1x storage pool (SSD, SATA SSD, SATA HDD)	Standard practice Intelligent tiering handles data locality
Container(s)	1x container for VMs 1x container for data (not used here)	Standard practice
Features / Enhancements	Increase CVM memory to 32 GB Turn on deduplication and compression for persistent desktops. Turn on compression only for nonpersistent desktops	Deduplication needs 32 GB of RAM to be enabled
Citrix XenDesktop		
XenDesktop Controllers	Minimum: 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 (including NetScaler VPX)	Ensures availability of controllers Balances load between controllers and pods
Desktop Power Management	Nutanix plug-in for controlling desktop power	Allows clean reboot and power down
Citrix StoreFront		
StoreFront Servers	Minimum: 2 (n+1)	HA for StoreFront servers
Load Balancing	Citrix NetScaler (including NetScaler VPX)	Ensures availability of StoreFront servers Balances load between StoreFront servers

Item	Detail	Rationale
		Citrix NetScaler including NetScaler VPX (if used)
NetScaler Servers	Minimum: 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

The next table shows a high-level summary of the pod design for Citrix XenDesktop on Nutanix.

Table 3: XenDesktop on AHV Highlights

Item	Quantity
Control Pod	
# of XenDesktop Controller(s)	2
# of XenDesktop StoreFront Server(s)	2
Services Pod	
# of AHV Hosts	Up to 32
# of Nutanix Cluster(s)	1
# of Datastore(s)	1




Figure 12: XenDesktop Pod Overview

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

The following table contains examples of typical scenarios for desktop deployment and use (drawn from Login VSI definitions).

Table 4: Desktop Scenarios: Typical 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 login 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 table below proposes initial sizing recommendations for a workload for these typical scenarios.



Note: These are recommendations 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	3 GB	35 GB (OS)
Power Users	2	4 GB	35 GB+ (OS)

5.1. Desktop Optimizations

We generated our design with the following high-level desktop guidelines in mind:

- 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.

For more detail on desktop optimizations, refer to the [Citrix XenDesktop Windows 10 Optimization Guide](#).

6. XenDesktop and Windows 10 Best Practices

Windows 10 includes the concepts of versions and branches, which are essential to your Windows 10 implementation. Picking the right version and branch determines your update and validation strategy, which can directly affect your user experience.

There are multiple versions of Windows 10. For this document we address versions applicable to your XenDesktop implementation: Professional, Education, and Enterprise. We do not address Windows 10 Core IoT (specific for Internet of Things features running on a Raspberry PI or Arduino) or Home (no domain join). Microsoft has published a [guide](#) to help you choose the right version for your environment.

With the release of Windows 10, Microsoft also introduced the notion of branch releases, which enables Windows as a service. This model focuses on adding new capabilities and updates by releasing software in a more predefined and predictable cadence. Microsoft plans for this delivery model to replace the traditional Windows upgrade projects, as Windows 10 would update to the latest version within the branch without any additional administrative overhead.

There are four types of branches:

- Windows Insider
- Current Branch (CB)
- Current Branch for Business (CBB)
- Long Term Servicing Branch (LTSB)

The CB pushes feature updates when Microsoft releases them. With Windows 10 version 1607, [which Microsoft recommends](#), you can use the servicing tools to delay CB feature updates up to 180 days before implementing them. The CB is a good fit for pilot deployments and early adopters that need the latest features directly.

While the CB suits early adopters by providing a test cycle before the major branch roll out, the CBB is the branch typically deployed for the actual major roll out. The CBB gets the same version numbers as the CB, but at a later point in time, to make sure that organizations are ready for the wider distribution of the release.

Systems that require a longer-term servicing branch due to their specialized use cases benefit from the LTSB. The LTSB targets systems with devices that have limited changes and need to be stable and always available, such as medical equipment or ATMs.

6.1. Optimizations

LoginVSI has a detailed document on how to optimize Windows 10 for XenDesktop, which they have consolidated into a template for the [VMware OS Optimization Tool](#). Applying these best practices significantly increases user and desktop per host ratio density. LoginVSI has also shown these optimizations improve [VSIMax 44 percent over a vanilla Windows 10 installation](#). For more detailed information about the optimizations, see [this recorded session](#).

Most of these best practice elements are handled by the LoginVSI template for the VMware OS optimization Tool, which provides the best density while maintaining the optimal user experience. For this reference architecture, we followed the [Citrix Windows 10 Optimization Guide](#), BPAnalyzer, and the LoginVSI Windows 10 template. The main items of all three sources are:

- Features
- Scheduled tasks
- Services
- Modern apps
- Microsoft OneDrive
- Hardware acceleration

Features

The new features in Windows 10 that have garnered the most discussion are those that relate to privacy and telemetry. When Citrix XenDesktop delivers Windows 10, these features negatively affect performance, while providing only minimal gains for user experience. In light of this fact, disabling these new features enhances security and privacy, while also improving user density.



Tip: Disabling new features such as boot logging and call home, as well as features related to telemetry, improves security and user density.

Scheduled Tasks

Any default Windows client operating system schedules tasks by default. While these scheduled tasks are useful and have little effect on a physical deployment, they can have a significant effect on your XenDesktop image. The newly introduced Indexer Automatic Maintenance task, for example, scans your Windows operating system anytime the machine is idle to check for errors, defragment the hard drive, and run optimization services. Running this task on hundreds or thousands of desktops in a centralized environment negatively affects your system's active users. Other scheduled tasks magnify this effect, which is why we recommend disabling scheduled tasks whenever possible.



Tip: When implementing XenDesktop, disable scheduled tasks whenever possible, including defragmentation, automatic maintenance, and Bluetooth scanning.

Services

XenDesktop lets you consolidate your desktop centralization into one or more datacenters, which means that some of the default services that come with Windows 10 add little benefit. Other default services add limited value when desktops are virtualized (rather than deployed on PCs or laptops). Services like BranchCache, Internet Connection Sharing, and Geolocation don't add value when XenDesktop delivers the desktop services from a central location. Accordingly, disabling these services has the advantage of hardening your operating system and improving user experience by lowering the operating system overhead.



Tip: Disabling unused services hardens your operating system and lowers the operating system overhead, which improves user experience.

Consumer-Focused Apps

The Windows 10 default deployment includes a list of preinstalled applications, many of which are consumer focused, such as Twitter, Asphalt 8, Airborn, and FarmVille 2. Removing those consumer-market applications frees up resources and helps limit distractions for end users.



Tip: Removing consumer-focused apps frees up machine resources and limits end-user distraction.

Microsoft OneDrive

Microsoft OneDrive is a file synchronization solution that allows you to connect to your Office 365 account (for example) and store your files both locally and in the cloud. Microsoft's OneDrive, however, is not built for virtualized desktops, as it syncs the complete OneDrive content locally. If your XenDesktop deployment contains nonpersistent desktops, OneDrive syncs the complete folder per user after every new logon to the system. Accordingly, we recommend removing Microsoft OneDrive to avoid wasting disk capacity and to prevent network congestion.



Tip: Uninstalling Microsoft OneDrive eliminates network congestion due to user folder synchronization actions with Office 365.

Hardware Acceleration

Since the release of Microsoft Office 2013, Office has used a hardware-accelerated method for presenting Office programs such as Word and Excel. The Microsoft operating system offers this feature, but the operating system itself relies on the graphics processing unit (GPU) and the display drivers. If there is no GPU available, the CPU performs the rendering, which uses additional system resources.



Tip: If you have no GPU available, make sure to enable software rendering for Internet Explorer and to disable hardware acceleration for Office 2010, Office 2013, and Office 2016 to improve user experience.

6.2. XenDesktop on AHV

Using a base or “Master VM,” the script framework creates full clone VMs that duplicate the base image, which provides both data locality and persistent desktops.

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




Figure 13: Communication Path

AHV Pod Design

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

Table 6: AHV Pod Detail

Item	Quantity
Control Pod	
# of XenDesktop Controller(s)	2
# of XenDesktop StoreFront Server(s)	2
Services Pod	
# of AHV Hosts	Up to 32
# of Nutanix Cluster(s)	1
# of Datastore(s)	1
# of Desktops	Up to 4,800

6.3. Nutanix: Compute and Storage

The Nutanix Enterprise Cloud Platform provides an ideal combination of both high-performance compute with localized storage to meet any demand. True to this capability, this reference architecture contains no 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 the Nutanix storage pool and containers.




Figure 14: Nutanix Logical Storage Configuration

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

Table 7: Nutanix Storage Configuration

Name	Role	Details
SP01	Main storage pool for all data	SSD and HDD
CTR VSI	Container for all VMs	AHV Datastore
CTR-RF2-DATA-01	Container for all data (not used here)	AHV Datastore

6.4. Network

Designed for true linear scaling, Nutanix uses 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 due 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 scale from one Nutanix block to thousands without any impact to available bandwidth.




Figure 15: Leaf-Spine Network Architecture

7. Solution Application

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



Note: We provide detailed hardware configuration and product models in the appendix.

7.1. Scenario: 600 Desktops

Table 8: Detailed Component Breakdown: 600 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix desktop pods	1 (partial)	# of AHV hosts	4
# of Nutanix chassis (blocks)	1	# of datastore(s)	1
# of RU (Nutanix)	2		
# of 10 GbE ports	8		
# of 100/1000 ports (IPMI)	4		
# of L2 leaf switches	2		
# of L3 spine switches	1		

7.2. Scenario: 900 Desktops

Table 9: Detailed Component Breakdown: 900 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix desktop pods	1 (partial)	# of AHV hosts	6

Item	Value	Item	Value
# of Nutanix blocks	2	# of datastore(s)	1
# of RU (Nutanix)	4		
# of 10 GbE ports	16		
# of 100/1000 ports (IPMI)	8		
# of L2 leaf switches	2		
# of L3 spine switches	1		

7.3. Scenario: 2,400 Desktops

Table 10: Detailed Component Breakdown: 2,400 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix desktop pods	1	# of AHV hosts	16
# of Nutanix blocks	4	# of datastore(s)	1
# of RU (Nutanix)	8		
# of 10 GbE ports	32		
# of 100/1000 ports (IPMI)	16		
# of L2 leaf switches	2		
# of L3 spine switches	2		

7.4. Scenario: 4,800 Desktops

Table 11: Detailed Component Breakdown: 4,800 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	

Item	Value	Item	Value
# of Nutanix desktop pods	1	# of AHV hosts	32
# of Nutanix blocks	8	# of datastore(s)	1
# of RU (Nutanix)	16		
# of 10 GbE ports	64		
# of 100/1000 ports (IPMI)	32		
# of L2 leaf switches	2		
# of L3 spine switches	2		

7.5. Scenario: 9,600 Desktops

Table 12: Detailed Component Breakdown: 9,600 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix desktop pods	2	# of AHV hosts	64
# of Nutanix blocks	16	# of datastore(s)	2
# of RU (Nutanix)	32		
# of 10 GbE ports	128		
# of 100/1000 ports (IPMI)	64		
# of L2 leaf switches	4		
# of L3 spine switches	2		

7.6. Scenario: 19,200 Desktops

Table 13: Detailed Component Breakdown: 19,200 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix desktop pods	4	# of AHV hosts	128
# of Nutanix blocks	32	# of datastore(s)	4
# of RU (Nutanix)	64		
# of 10 GbE ports	256		
# of 100/1000 ports (IPMI)	128		
# of L2 leaf switches	8		
# of L3 spine switches	2		

7.7. Scenario: 38,400 Desktops

Table 14: Detailed Component Breakdown: 38,400 Desktops Using NX-3060-G5

Item	Value	Item	Value
Components		Infrastructure	
# of Nutanix desktop pods	8	# of AHV hosts	256
# of Nutanix blocks	64	# of datastore(s)	8
# of RU (Nutanix)	128		
# of 10 GbE ports	512		
# of 100/1000 ports (IPMI)	256		
# of L2 leaf switches	14		
# of L3 spine switches	2		

8. Validation and Benchmarking

We conducted the solution and testing in this document with Citrix XenDesktop 7.12 deployed on AHV and the Nutanix Enterprise Cloud Platform. Login VSI Office Worker and Knowledge Worker benchmarks supplied the model for the Knowledge Worker user on the Nutanix appliance.

8.1. Environment Overview

We used two nodes of an existing Nutanix NX-8050 to host all infrastructure and XenDesktop services and a Nutanix NX-3460 to host the Login VSI test harness. The six nodes in the eight node Nutanix NX-3060-G5 functioned as the target environment and provided all desktop hosting. The Nutanix block was connected to an Arista 7050S top-of-rack switch using 10 GbE.




Figure 16: Test Environment Overview

8.2. Test Environment Configuration

Assumptions:

- Knowledge Worker use case
- Per-desktop IOPS (Knowledge Worker): 10 sustained and 70 peak (startup)
- Using MCS persistent Machine Catalogs

Hardware:

- Storage and compute: 2x Nutanix NX-3460-G5
- Network: Arista 7050Q (L3 spine) and 7050S (L2 leaf) series switches

Desktop configuration:

- OS: Windows 10 x64
- 2 vCPU and 3 GB memory
- 1x 32 GB OS Disk

Applications:

- Microsoft Office 2013
- Adobe Acrobat Reader XI
- Internet Explorer
- Flash video

Login VSI:

- Login VSI 4.1 Professional

8.3. XenDesktop Configuration

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

Table 15: XenDesktop Configuration

VM	Quantity	vCPU	Memory	Disks
XenDesktop Controller(s)	1	4	8	1 x 40 GB (OS)
StoreFront Server(s)	1	4	4	1 x 40 GB (OS)

Test Image Preparation

1. Create base VM.
2. Install Windows 10.
3. Install standard software.
4. Optimize Windows 10.
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.
10. Create clones using XenDesktop Setup Wizard.

Test Execution




Figure 17: Test Execution Flowchart

8.4. 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 such as 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, so 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 16: 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 4.1 Workloads](#) document captures the Login VSI Workflow base layout and documents the changes from previous versions of Login VSI to version 4.1 in great detail.

Table 17: 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
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+

8.5. 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 18: 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 launch for 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 18: Example Graph of Login VSI Test

9. Results

9.1. 600 Knowledge Worker Desktops (Four Nodes)

During the testing with 600 desktops, VSImax was not reached, with a threshold of 1,705 ms, VSI average of 1,251 ms, and a VSIBase of 705 ms.




Figure 19: 600 Knowledge Worker Desktops

9.2. 900 Knowledge Worker Desktops (Six Nodes)

During the testing with 900 desktops, VSImax was not reached, with a threshold of 1,718 ms, VSI average of 1,230 ms, and a VSIBase of 718 ms.




Figure 20: 900 Knowledge Worker Desktops

CPU utilization for the AHV hosts during the tests peaked at 88.73 percent.




Figure 21: Peak CPU Utilization

Nutanix Metrics

IOPS peaked at approximately 17,927 during the high-volume startup period to refresh the desktops; the peak IOPS value during the test was 6,510.




Figure 22: Peak IOPS




Figure 23: Command Latency Peak

9.3. 1,200 Knowledge Worker Desktops (Eight Nodes)

During the testing with 1,200 desktops, VSI_{max} was not reached, with a threshold of 1,720 ms, VSI average of 1,334 ms, and a VSIBase of 720 ms.




Figure 24: 1,200 Knowledge Worker Desktops

10. Conclusion

Our extensive testing of Citrix XenDesktop on Nutanix revealed a light I/O footprint on the Nutanix Enterprise Cloud Platform. Aggregate IOPS peaked at approximately 18,000 during the high-volume startup periods, and sustained IOPS ranged from 5,000–6,000. I/O latencies averaged less than 2 ms for read and less than 5 ms for write during peak load. Nutanix offers the performance and low latency necessary for uncompromised end-user experience, even during boot storms, patch and update operations, and application and OS upgrades.

The Citrix XenDesktop on Nutanix solution provides a single, high-density platform for desktop and application delivery. Running XenDesktop on Nutanix also means that you can:

- Get started quickly, without onerous upfront expenditures, and then scale as you grow.
- Scale out with linear performance.
- Deliver excellent user density, supporting 150 knowledge workers per AHV host.
- Simplify management and administration by eliminating LUNs.
- Significantly lower both acquisition and ongoing costs by running VDI alongside other workloads.




Figure 25: Linear Scale-Out Performance

Appendix

Configuration

Hardware

- Storage and compute
 - Nutanix NX-3460-G5
 - Nutanix NX-3460-G5 (only 2 nodes were used)
- Per-node specs (four nodes per 2RU block):
 - CPU: 2x Intel Xeon E5-2680v4 @ 2.4 Ghz
 - Memory: 512 GB memory
- Network
 - Arista 7050Q L3 spine
 - Arista 7050S L2 leaf

Software

- Nutanix Acropolis
 - AOS 5.0.2
- XenDesktop
 - 7.12
- Desktop
 - Windows 10 x64
 - Office 2013
- Infrastructure
 - AHV

VM

- Desktop
- CPU: 2 vCPU
- Memory: 3 GB

- Storage: 1x 32 GB OS Disk on CTRVSI DSF datastore

Performance

Nutanix has identified a potential performance bottleneck due to a bug in qemu and USB-tablet CPU consumption. The USB host adapter that polls the USB devices may use excessive CPU resources when the host adapter [is emulated](#). While already solved for Linux-based operating systems, this problem still affects Windows, which impacts your XenDesktop implementation.

Nutanix plans to create a fix for the newer releases of our Acropolis Operating System (AOS). For VMs created on AOS 4.7.3 or earlier, an in-guest fix can bring down CPU utilization and increase performance.

Once you confirm that USB type 1 devices are in use, delete these two registry subtrees:

```
HLM\SYSTEM\CurrentControlSet\Control\usbflags
```

```
HLM\SYSTEM\CurrentControlSet\Enum\USB
```

After rebooting the machine, Windows checks the OS descriptors again; at this point the CPU overhead fix is in place.

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, 2014, 2015, and 2016. Citrix also named him a Citrix Technology Professional in 2015 and 2016.

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](https://twitter.com/nutanix).

List of Figures

Figure 1: Nutanix Enterprise Cloud Platform.....	9
Figure 2: Information Life Cycle Management.....	11
Figure 3: Overview of the Nutanix Architecture.....	11
Figure 4: Data Locality and Live Migration.....	12
Figure 5: Acropolis File Servers Run as VMs to Isolate from DSF.....	13
Figure 6: MCS Disk Layout.....	16
Figure 7: MCS on AHV: Architecture.....	20
Figure 8: MCS on AHV: Site Setup.....	21
Figure 9: MCS on AHV: Site Setup—Network.....	21
Figure 10: MCS on AHV: Hosting Connection.....	22
Figure 11: XenDesktop on Nutanix Conceptual Architecture.....	23
Figure 12: XenDesktop Pod Overview.....	28
Figure 13: Communication Path.....	34
Figure 14: Nutanix Logical Storage Configuration.....	35
Figure 15: Leaf-Spine Network Architecture.....	37
Figure 16: Test Environment Overview.....	42
Figure 17: Test Execution Flowchart.....	44
Figure 18: Example Graph of Login VSI Test.....	47
Figure 19: 600 Knowledge Worker Desktops.....	48
Figure 20: 900 Knowledge Worker Desktops.....	49
Figure 21: Peak CPU Utilization.....	50

Figure 22: Peak IOPS.....	50
Figure 23: Command Latency Peak.....	51
Figure 24: 1,200 Knowledge Worker Desktops.....	51
Figure 25: Linear Scale-Out Performance.....	52

List of Tables

Table 1: Document Version History.....	6
Table 2: Solution Design Decisions.....	25
Table 3: XenDesktop on AHV Highlights.....	27
Table 4: Desktop Scenarios: Typical Definition.....	29
Table 5: Desktop Scenario Sizing.....	29
Table 6: AHV Pod Detail.....	35
Table 7: Nutanix Storage Configuration.....	36
Table 8: Detailed Component Breakdown: 600 Desktops Using NX-3060-G5.....	38
Table 9: Detailed Component Breakdown: 900 Desktops Using NX-3060-G5.....	38
Table 10: Detailed Component Breakdown: 2,400 Desktops Using NX-3060-G5.....	39
Table 11: Detailed Component Breakdown: 4,800 Desktops Using NX-3060-G5.....	39
Table 12: Detailed Component Breakdown: 9,600 Desktops Using NX-3060-G5.....	40
Table 13: Detailed Component Breakdown: 19,200 Desktops Using NX-3060-G5.....	41
Table 14: Detailed Component Breakdown: 38,400 Desktops Using NX-3060-G5.....	41
Table 15: XenDesktop Configuration.....	43
Table 16: Login VSI 4.1 Workloads.....	45
Table 17: Login VSI Workload Definitions.....	45
Table 18: Login VSI Metric Values.....	46