v1.4 | August 2021 | BP-2079

BEST PRACTICES

# Citrix Virtual Apps and Desktops on Nutanix

NUTANIX
YOUR ENTERPRISE CLOUD

# Copyright

# Contents

# 1. Executive Summary

This best practice guide discusses the best practices for running Citrix Virtual Apps and Desktops on Nutanix. Nutanix offers a powerful, flexible, and reliable platform for the full spectrum of desktop virtualization requirements, with unrivaled uptime and the freedom to mix and match workloads to fulfill your enterprise needs and the objectives of your operators, whether they're task workers or power users.

Nutanix includes Acropolis distributed storage, which offers a range of advantages for Citrix Virtual Apps and Desktops deployments:

- Optimized data path that easily handles increased read I/O.
- Data avoidance technologies that you can implement on a fit-for-purpose basis.
- A single datastore, which dramatically reduces administrative overhead and update time.
- Storage efficiency techniques, such as deduplication, that can reduce the storage footprint of Virtual Apps and Desktops deployments.
- Nutanix Shadow Clones, which cut network latency and improve user experience.

Delivering applications through Citrix Virtual Apps (previously XenApp) or Virtual Desktops (previously XenDesktop) on Nutanix means you can easily deploy thousands of desktops with an optimal user experience across multiple use cases.

 

# 2. Introduction

## Audience

This hypervisor-agnostic best practice guide is part of the Nutanix Solutions Library. We wrote it for individuals responsible for designing, building, managing, and supporting Citrix Virtual Apps and Desktops on Nutanix infrastructures. Readers should be familiar with Nutanix AOS, Prism, AHV, Files, Citrix Virtual Apps and Desktops, Microsoft Hyper-V, and VMware vSphere.

## Purpose

This document covers the following subject areas:

- Overview of the Nutanix solution.
- Best practices for delivering Citrix Virtual Apps and Desktops on Nutanix.

Unless otherwise stated, the solution described in this document is valid on all supported AOS releases.

Table 1: Document Version History

| Version Number | Published | Notes |
|---|---|---|
| 1.0 | April 2017 | Original publication. |
| 1.1 | March 2018 | Updated platform overview. |
| 1.2 | October 2019 | Added NGT information for Citrix App Layering, MCS, and Provisioning base images. |
| 1.3 | April 2021 | Refreshed and updated document. |
| 1.4 | August 2021 | Added additional Provisioning content and updates. |

 

# 3. Nutanix Enterprise Cloud Overview

Nutanix delivers a web-scale, hyperconverged infrastructure solution purpose-built for virtualization and both containerized and private cloud environments. This solution brings the scale, resilience, and economic benefits of web-scale architecture to the enterprise through the Nutanix enterprise cloud platform, which combines the core HCI product families—Nutanix AOS and Nutanix Prism management—along with other software products that automate, secure, and back up cost-optimized infrastructure.

Available attributes of the Nutanix enterprise cloud OS stack include:

- Optimized for storage and compute resources.
- Machine learning to plan for and adapt to changing conditions automatically.
- Intrinsic security features and functions for data protection and cyberthreat defense.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.
- Simplified one-click upgrades and software life cycle management.
- Native file services for user and application data.
- Native backup and disaster recovery solutions.
- Powerful and feature-rich virtualization.
- Flexible virtual networking for visualization, automation, and security.
- Cloud automation and life cycle management.

Nutanix provides services and can be broken down into three main components: an HCI-based distributed storage fabric, management and operational intelligence from Prism, and AHV virtualization. Nutanix Prism furnishes one-click infrastructure management for virtual environments running on AOS. AOS is hypervisor agnostic, supporting two third-party hypervisors

 

—VMware ESXi and Microsoft Hyper-V—in addition to the native Nutanix hypervisor, AHV.

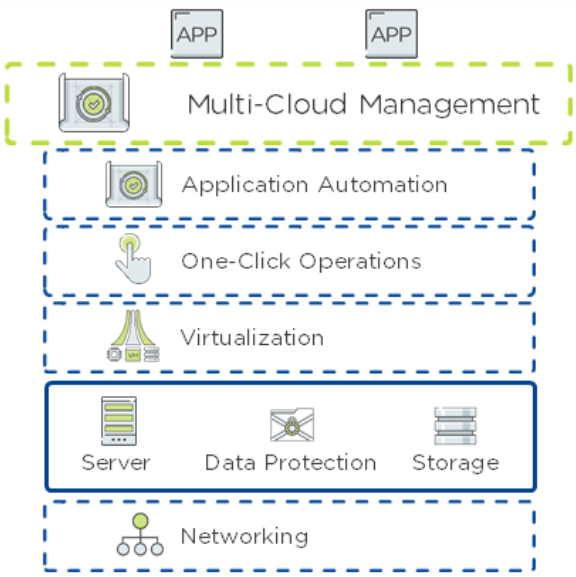


Figure 1: Nutanix Enterprise Cloud OS Stack

## Nutanix HCI Architecture

Nutanix doesn't rely on traditional SAN or network-attached storage (NAS) 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 delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution requires no SAN constructs, such as LUNs, RAID groups, or expensive storage switches. All storage management is VM-centric, and I/O is optimized at the VM virtual disk level. The software solution runs on nodes from a variety of manufacturers that are either entirely solid-state storage with NVMe for optimal performance or a hybrid combination of SSD and HDD storage that provides a combination of performance and additional capacity. The storage fabric automatically tiers data across the cluster to different classes of storage devices using intelligent data placement algorithms. For best

 

performance, algorithms make sure the most frequently used data is available in memory or in flash on the node local to the VM.

To learn more about Nutanix enterprise cloud software, visit the Nutanix Bible and Nutanix.com.

 

# 4. Best Practice Checklist

We can summarize the best practices for deploying Citrix Virtual Apps and Desktops on Nutanix into the following items.

## General

- Perform a current state analysis to identify workloads and sizing for the desktops and applications you plan to virtualize.
- Gather and document functional and technical requirements for the virtual desktop solution.
- Spend time up-front to create a solution that meets both current and future needs.
- Design for end-user experience to deliver consistent performance, reliability, and scale.
- Start with a PoC, then test, optimize, iterate, and scale.

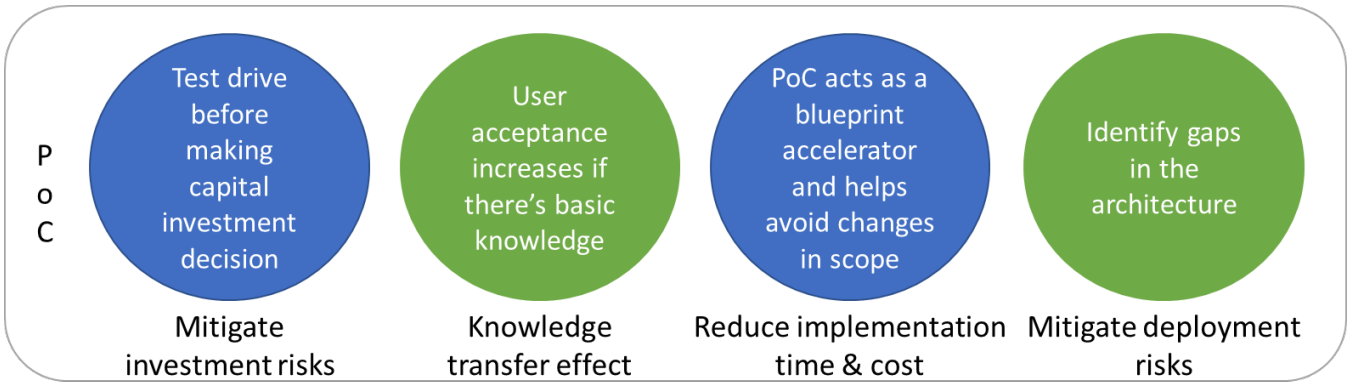


Figure 2: Benefits of Running a PoC

 

## Core Components

### Workloads

- Size workloads appropriately for each particular use case.
- Use a mix of application virtualization and applications installed in gold images, depending on the scenario.
- Disable unnecessary OS services and applications.
- Use optimization tools to further enhance the workloads.
- Configure proper antivirus and security exclusions.
- Implement GPUs when necessary to improve user experience.

### Citrix Virtual Apps and Desktops

- Design for reliability and scale.
- Redirect the home directory or use a profile management tool (Citrix Workspace Environment Management (WEM) or Profile Management, or Microsoft FSLogix, for example) for user profiles and documents.
- Use Citrix High Definition Experience (HDX) Adaptive Transport and apply the relevant HDX policies in Virtual Desktops. Optimize and customize the HDX templates to fit your environment before you apply them.

### Hypervisor

- Follow vendor best practices.
- Keep track of the CPU ready times to ensure that CPU overcommit ratios remain within acceptable thresholds.
- Don’t overcommit RAM.

### Nutanix

- Use a single container and datastore for virtual desktops and Virtual Apps–based VMs.
- Increase Nutanix CVM memory per the sizing tables in this document.

 

- Configure storage containers based on the workload delivery method per the storage best practices in this document.

## Supporting Components

### Active Directory

- Have local global catalogs and DNS (Domain Name System) servers at each site.
- Redirect home directories for users.
- Configure DNS scavenging.
- Configure location-specific organizational units so workloads can use local resources.

### DHCP

- For the Citrix Virtual Apps and Desktops infrastructure scope, reduce default DHCP lease times from eight days to one hour.
- Configure proper DHCP options for Citrix Provisioning.
- If you use AHV, set up and configure Nutanix IP address management (IPAM) for built-in DHCP services.

### File Services

> Note: Nutanix Files supports distributed single namespace to store user and profile data for all Citrix users.

> Note: Nutanix Files provides key analytics data through the File Analytics view and provides access to configure ransomware protection, file-blocking features, and other important features.

- Map home directory redirection to the Nutanix Files namespace.
- Deploy distributed shares (previously called home shares) for user data and profiles.
- Configure connected shares for profile solutions (Citrix User Layer or User Personalization Layer, for example) that have subfolders below the distributed share top-level directory.

 

- Enable SMB transparent failover for containerized profile solutions (Citrix User Layer or User Personalization Layer and Profile Management Containers along with Microsoft FSLogix, for example) for full nondisruptive operations against SMB shares (also known as continuously available shares).
- Deploy Citrix Profile Management user stores inside the users' top-level directory for Director Profile Reset to work properly (\\server\share\%username%.%userdomain%\!CTX_OSNAME!!CTX_PROFILEVER!, for example).

### Virus Scan

- Schedule scans to run outside business hours.
- Stagger system scans in phases.

### Network

- Use and optimize QoS (quality of service) for HDX to prioritize Citrix network traffic.
- Use at least 1 GbE access ports for end-user LAN connectivity.
- Ensure adequate bandwidth for WAN and VPN clients.

### OS and Application Updates

- Apply updates outside business hours to avoid performance impact.
- Stagger updates in phases.
- Use image management technologies such as Citrix App Layering, Machine Creation Services (MCS), or Citrix Provisioning.

 

# 5. Desktops

## Sizing

### Compute

For desktop-based operating systems, Nutanix typically recommends at least 2 vCPU per VM so the system can run multiple threads simultaneously. If you assign a single vCPU for light workloads, the associated desktops are more likely to experience session or application interruptions.

> Tip: Assign 2 vCPU per VDI desktop so the system can run multiple threads simultaneously.

> Note: Moving from 1 vCPU to 2 vCPU reduces desktop density by approximately 20 percent, not by 50 percent.

For server-based operating systems, Nutanix and Citrix Consulting typically recommend 8 vCPU for Microsoft Server 2016 and Microsoft Server 2019.

> Tip: Assign 8 vCPU per 2016 or 2019 VM to optimize CPU-to-user ratios.

Sizing physical CPU cores differs for VDI and RDSH (Remote Desktop Session Host) because of the difference in CPU overcommit ratios and the number of VMs required to host user workloads. The following table provides guidance on the number of virtual apps (RDSH) servers and virtual desktops supported per physical core for light, medium, and heavy workloads, and shows the CPU overcommit ratio of physical CPU to virtual CPU.

Table 2: CPU Overcommit Ratios

| Deployment | pCPU | vCPU |
|---|---|---|
| Virtual Apps | 1 | 1–2 |
| Virtual Desktops: Light users | 1 | 11–12 |
| Virtual Desktops: Medium users | 1 | 7–10 |

 

| Deployment | pCPU | vCPU |
|---|---|---|
| Virtual Desktops: Heavy users | 1 | 4–6 |
| Note: Processor architecture and speed have a direct impact on the number of users the system can support. We base these estimates on Intel Broadwell processor architecture. | | |

These scalability estimates take the performance benefits of hyperthreading into account. As noted in the Citrix VDI Handbook and Best Practices, hyperthreading can improve user density per VM (server-based computing (SBC)) or VM density per host (SBC and VDI), and typically provides a performance boost of between 20 and 30 percent. These numbers vary based on the operating system (OS) you deploy, optimizations you configure, and applications you install. In addition, you need to consider other factors, such as CPU clock speed versus more CPU cores and single-threaded versus multithreaded applications.

## Memory

In general, don't overcommit memory for desktop virtualization, as doing so can impact performance and the overall user experience. Nutanix AHV doesn't allow you to overcommit memory. VMware vSphere versions 6 and later disable memory overcommit by default; although you can enable it, doing so requires manual actions.

> Tip: Don't overcommit memory when you virtualize desktops.

VMware vSphere allows you to assign and reserve a certain amount of memory per VM. Performance testing indicates that using such a reservation can reduce the platform's user density up to 18 percent.

> Tip: Don't use memory reservations when you virtualize desktops.

Using GPUs allocates all VM guest memory and automatically enables the option to reserve all guest memory (all locked) on the VM. You can review your memory reservation selections in the VM memory settings on VMware vSphere.

Size memory for Virtual Apps VMs based on the memory requirements for the applications you use on those VMs. As a guideline, we advise assigning at least 3 GB of RAM per vCPU. In a Windows 2016 or Windows 2019 deployment, this

 

guideline results in 8 × 3 GB = 24 GB of RAM per VM. In this case, we advise assigning at least 24 GB of RAM per VM.

Table 3: Memory Assignment

| Deployment | Memory |
|---|---|
| Virtual Apps based on Windows Server | 24 GB and up |
| Virtual Desktops: Light user | 1.5–2 GB |
| Virtual Desktops: Medium user | 2–4 GB |
| Virtual Desktops: Heavy user | 4 GB and up |

## VM Configuration

Use settings from the following table when you configure a base VM.

Table 4: VM Parameters

| | VMware Best Practice | Hyper-V Best Practice | AHV Best Practice |
|---|---|---|---|
| SCSI controller | LSI Logic SAS controller / PVSCSI | SCSI | SCSI |
| Hard disk | Thin provisioning | Default | Default |
| Video card | Automatically detect | Configure according to MS guidelines | N/A |
| Floppy | Remove | Remove | N/A |
| NIC | VMXNET3 | Synthetic NIC | Default |
| BIOS disable ports | Disable LPT and COM ports | Disable LPT and COM ports | N/A |
| Disable HotAdd and HotPlug | Disable HotPlug | Disable HotPlug | N/A |

Tip: Follow best practices for VM configuration.

 

Note: Disk controllers can influence the performance users perceive.

## OS Optimization

Configure your Windows image to the specifications outlined in the Citrix Windows 10 Optimization Guide (Citrix recommends that you use the Citrix Optimizer). Here is a summary of the optimizations:

- Set display to Adjust for best performance.
- Disable unnecessary services and remove unused components.
- Antivirus:
  - Full clones: Run hypervisor-level antivirus scans. If you need to run OS-level antivirus full scan jobs, do so during off hours and in phases.
  - MCS or Provisioning clones: Run hypervisor-level antivirus scan jobs during off hours and in phases. If you need to run OS-level antivirus jobs, do so during off hours and in phases, and run a full scan before sealing your image.
- Updates:
  - Full clones: Update the OS during off hours and in phases.
  - MCS or Provisioning clones: Update the base image during off hours and recompose in phases.

Nutanix recommends that you use the Citrix Optimizer, VMware OS Optimization Tool (TargetOSOptimizer), or the Best Practice Analyzer tools to prepare your gold image.

## Application Delivery

### Application Virtualization Solutions

Application virtualization encapsulates application software from the underlying OS it runs on. This solution doesn't install a fully virtualized application in the traditional sense: at runtime, the application

 

behaves as if it's directly interfacing with the original OS and all the resources that OS manages, but you can isolate or sandbox it to varying degrees.

Available application distribution solutions include:

- Microsoft App-V
- Third-party application virtualization solutions

### Layering Solutions

Layering solutions offer a solution based on separate disks that you can layer over the base image. Layering technology allows you to segment security and isolate user settings, applications, and environment configuration.

Available layering solutions include:

- Citrix App Layering (App Layers or Elastic Layers)
- Microsoft FSLogix (App Masking)
- Microsoft MSIX App Attach
- Third-party layering solutions

> Tip: Test your solutions with real-world applications before you choose between application virtualization or layering. Application layering offers different benefits than application virtualization. You can combine them to achieve the benefits of both, but such combinations increase complexity.

## Profile Management and User Data

To explain the differences between profile management, user environment management, and application virtualization, we first need to define what a user profile does and what profile options are available.

User profiles provide the user with a consistent experience. The following user profile types are available:

- Local
- Roaming
- Mandatory
- Virtual disk–based

 

### Local Profiles

Local profiles are available on a per-computer basis. When a user logs on, the system creates a user profile directory in the default profile directory. For example, for Windows 7 and later, the default directory is C:\Users\%username%. When a user saves a file on their desktop, Windows stores it in C:\Users\%username%\Desktop. This directory includes folders such as Favorites and Documents as well. Windows stores user settings in the registry HKEY_CURRENT_USER on the local machine.

> Note: Local profiles require dedicated desktops to allow user settings and documents to persist after logoff. Additionally, using local profiles without a backup solution for each individual desktop may result in data loss if a user's desktop is corrupted or destroyed.

### Roaming Profiles

The roaming option allows users to roam across different computers and take their settings with them. When a user logs on to a computer, the folder containing the roaming profile downloads to the local computer. Any modifications the user makes are stored locally and synchronized at logoff.

> Tip: Using roaming profiles can result in poor logon and logoff times for users, so you shouldn't use them. Look into using the profile solutions Citrix includes with their solutions or third-party solutions such as hybrid or virtual disk–based solutions.

### Mandatory Profiles

Often used in SBC scenarios, this profile type provides the user with the same group of base settings after each logon. During logon, the system copies the profile to the C:\Users\%username% folder, and the user can make modifications. At logoff, the system deletes the profile and all changes made to it.

### Virtual Disk–Based Profiles

You can also store the user profile in a virtual disk. Each user has their own virtual disk stored in the datacenter close to the virtual desktops and RDSH servers. When the user logs on, the OS learns the location of the user's profile and redirects reads and writes going to the user's profile to that virtual disk. Because the profile doesn't need to be copied at logon, the user experiences

 

a fast logon, as if the profile is already available locally to the virtual desktop. Available virtual disk–based profile solutions are:

- Citrix App Layering (User Layers)
- Citrix Profile Management Profile Containers
- Citrix User Personalization Layers
- Microsoft FSLogix Profile Containers
- Third-party containerized profile solutions

Microsoft acquired FSLogix in 2018 and the solution is included in the following licenses:

- Microsoft 365 E3/E5
- Microsoft 365 A3/A5/Student Use Benefits
- Microsoft 365 F1/F3
- Microsoft 365 Business
- Windows 10 Enterprise E3/E5
- Windows 10 Education A3/A5
- Windows 10 VDA per user
- Remote Desktop Services (RDS) Client Access License (CAL)
- Remote Desktop Services (RDS) Subscriber Access License (SAL)

### Profile Management

Profile management delivers personal settings to the user's (virtual) desktops and applications, often through a profile optimization service that provides these settings in an easy, reliable way to ensure a consistent experience for the user. Although it's a common solution, profile management limits itself to the user profile. To provide the user with drive mappings or printers (for example), you must use a different solution, such as logon scripts or group policy preferences. Available profile management solutions include:

- Microsoft UE-V
- Citrix Profile Management

 

- Third-party profile management solutions

> Note: Nutanix Files is built to support all forms of profile management with key optimizations for faster VDI logon times by efficient metadata caching techniques.

## User Environment Management

User environment management (UEM) allows you to manage the user experience in an enterprise environment, using a traditional desktop infrastructure solution, a virtual desktop infrastructure solution, or even a mobility-based solution.

Today's workplace gives employees one or more devices for accessing IT services, as well as the applications they need for their roles. Employee access to these services and applications operates within the boundaries of a corporate policy to ensure that each individual has sufficient access rights. The IT services that users access include objects such as:

- Network drivers
- Printers
- Applications

Users customize their way of working and make changes to their system within the limits of organizational boundaries. Common customizations could be email signatures, web browser favorites, and shortcuts. The combination of the corporate policy with user preferences is the user persona, which you can manage with a UEM solution.

UEM is a good addition to local profiles and mandatory profiles, creating a form of hybrid profile that blends the speed of local and mandatory profiles with the ability to have user settings roam across devices and operating systems. Available UEM solutions include:

- Citrix Workspace Environment Management
- Microsoft GPO, GP preferences, and user state virtualization
- Third-party UEM solutions

> Tip: Select a UEM solution or a profile management solution or use a combination of the two to manage your profiles and user settings. Managing Microsoft native profiles works as a solution but isn't cost efficient.

 

# 6. Citrix Virtual Apps and Desktops on Nutanix

## Architecture

Nutanix allows organizations to start small and scale from hundreds to thousands of desktops. To enable this kind of growth, you must design a solution with scalability in mind. The following figure presents an example architecture using a pod methodology to design any size of Citrix Virtual Apps and Desktops infrastructure over multiple sites with Nutanix.

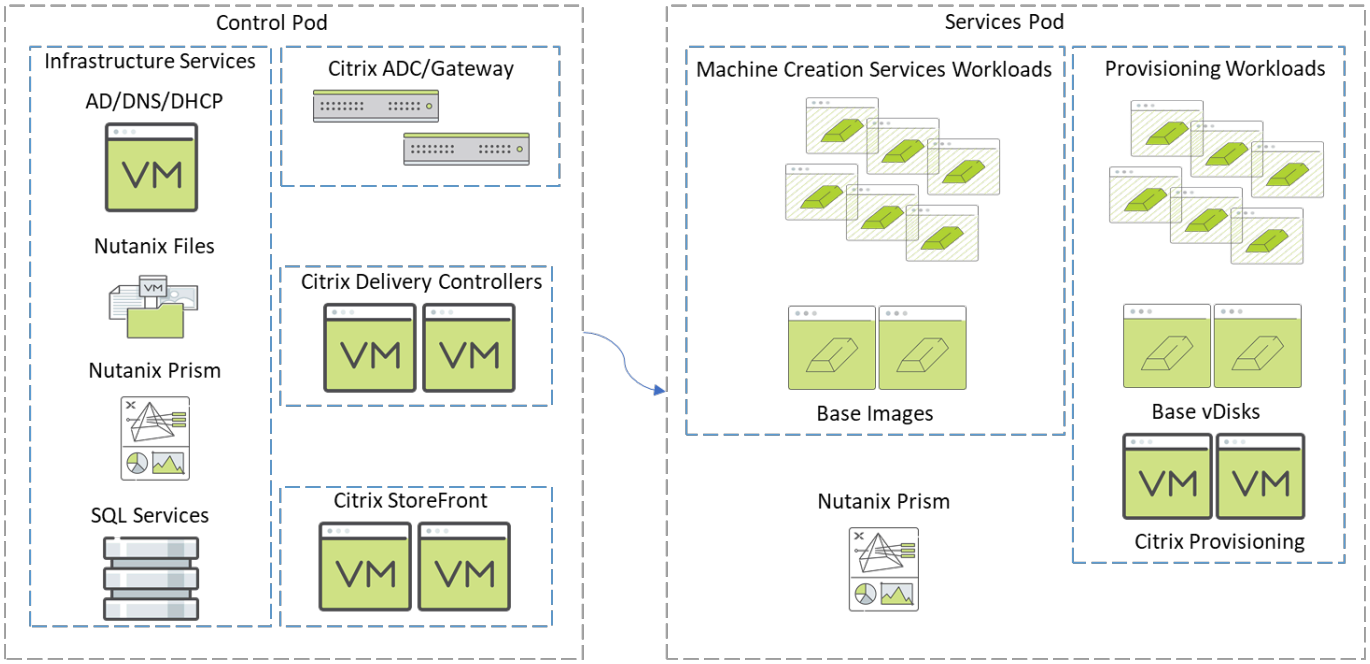


Figure 3: Virtual Apps and Desktops Pod Overview

A control pod contains all the components for the Citrix Virtual Apps and Desktops infrastructure services, and a services pod contains one or more clusters hosting Virtual Apps and Desktops workloads. A single control pod can manage multiple services pods. For more information on the Citrix Virtual Apps

 

and Desktops on Nutanix design, refer to the Citrix Virtual Apps and Desktops reference architectures on the Nutanix Portal.

> Tip: Design and build your environments in pods to make scaling easy and create smaller failure domains.

## Control Layer

### Delivery Controllers

To extend your Virtual Apps and Desktops environment, you can scale up (provide services to more users per site or zone) or scale out (add more sites or zones). The choice between scaling up and scaling out may be subject to the following constraints:

- Failure domains: Creating multiple sites produces more and smaller failure domains, so outages impact only a subset of the total users.
- Geographical location: You can split Virtual Apps and Desktops sites into zones. A zone can be a physical location, offering you more flexibility, as you can use one Citrix Virtual Apps and Desktops site (and thus one configuration) over multiple zones. The number of Delivery Controllers in a site can affect the performance of some operations, so we recommend limiting the number of zones in a Virtual Apps and Desktops site to no more than 50.

#### Delivery Controller Redundancy

Nutanix recommends having at least two Delivery Controllers per Virtual Apps and Desktops site for redundancy and high availability. The primary zone that contains the SQL Server site database should have at least two Delivery Controllers for redundancy and high availability. Depending on their size, satellite zones may also require two Delivery Controllers for redundancy and high availability. If a customer needs infrastructure redundancy during maintenance or a failure, you should configure at least three Citrix Virtual Apps and Desktops Delivery Controllers per site in the primary zone. The number of Delivery Controllers for the satellite zone depends on size and requirements.

 

#### Delivery Controller Scalability

To ensure that the Delivery Controllers aren't the bottleneck in your infrastructure, assign sufficient resources to each VM. Scale these resources for peak endurance moments like periods of boot (or logon) storms. Citrix performed internal testing on the scalability of the Delivery Controller role and report in their Citrix VDI Handbook and Best Practices document that a Virtual Apps and Desktops controller (4 vCPU and 4 GB of RAM) can support more than 5,000 virtual desktops. The Local Host Cache requires additional RAM in an environment with many logons occurring during an extended outage. Consider increasing the total amount of RAM capacity to support the Local Host Cache during an extended outage.

### StoreFront

#### StoreFront Redundancy

Citrix recommends that you have at least two StoreFront servers per site for redundancy. Keep in mind that joining them in a server group doesn't necessarily mean that the system uses them equally. You need additional load balancing for an active-active configuration. The Citrix ADC appliance is the most common load balancing appliance in Citrix Virtual Apps and Desktops environments.

#### StoreFront Scalability

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

### Provisioning

#### Provisioning Redundancy

Nutanix recommends that you have at least two Provisioning servers per Provisioning site in a farm for redundancy and high availability. Plan for high availability and redundancy in the design so that a single Provisioning server failure doesn't reduce the amount of target devices supported in a site. The number of Provisioning servers for each farm site depends on

 

size and requirements. Configure the Provisioning boot file with multiple Provisioning servers in a site for high availability (the Provisioning boot file can have up to four Provisioning servers). Each Provisioning server should have itself listed highest in the boot file order for the local boot file configuration. Enable Provisioning load balancing on the vDisk for load distribution across Provisioning servers in the site.

### Provisioning Scalability

To ensure that the Provisioning servers aren't the bottleneck in your infrastructure, assign sufficient resources to each VM. Scale these resources for peak endurance moments like periods of boot (or logon) storms. Citrix recommends 4 vCPU per Provisioning server in small environments (up to approximately 500 target devices) and 8 vCPU in large environments. Because the Provisioning Streaming Service is configured for 20 sequential network ports and 8 threads per port by default, a Provisioning server can support 160 concurrent target devices by default. If you need more than 160 streams, Provisioning continuously switches streaming between different target devices.

To support more than 160 concurrent target devices, you can adjust the number of ports and threads per port for each Provisioning server in the Provisioning console. A Provisioning server has the best performance when the threads per port don't exceed the number of cores or vCPU assigned to the Provisioning Server. Higher CPU utilization occurs when a Provisioning server doesn't have sufficient cores or vCPU, which causes increased read latency in target devices waiting for requests to be processed.

A Provisioning server's Windows OS partially caches vDisks in memory (system cache), which reduces the number of reads required from the storage. Allocate memory to the Provisioning servers to maximize the benefit of caching vDisks in memory. Citrix recommends setting total Provisioning server memory to 2 GB + (number of vDisks × 2 GB).

 

# SQL Database

## SQL Database Redundancy

With the transition from the Independent Management Architecture (IMA) to the FlexCast Management Architecture (FMA), SQL database availability has become increasingly important. Delivery Controllers use centralized databases to store both static configuration and more dynamic session information. Currently running sessions aren't affected until a user either logs off or disconnects when the database becomes unavailable. The Local Host Cache (LHC) feature allows connection-brokering operations in the Citrix Virtual Apps and Desktops site to continue for server-hosted applications and desktops and static (assigned) desktops when the database becomes unavailable. To access pooled desktops when the database becomes unavailable, you must run PowerShell commands to reuse machines without shutdown during an outage. When the database becomes unavailable, site management is unavailable. Therefore, it's very important that the databases be highly available through mirroring, clustering, or Always On availability groups.

Because Citrix Provisioning uses a database to store all Provisioning farm configuration details and other information, SQL database availability has become increasingly important. Citrix Provisioning has an offline database support feature that allows Provisioning for target devices to remain operational when a database outage occurs. When the Provisioning database is offline, management functions and the management console become unavailable. The Provisioning Stream process uses a database snapshot that is created and initialized at Provisioning server startup and continually updates the snapshot. The offline database snapshot is stored in Provisioning server memory.

Offline database support isn't enabled by default; if you want to use it, enable the feature after you set up and configure the Provisioning farm. To enable offline database support, you must restart the Stream services on Provisioning servers in the farm. For more information, consult the Citrix Provisioning Managing for highly available implementations documentation.

 

> Tip: Ensure that your database is highly available through Always On failover cluster instances or Always On Availability Groups (including Basic Availability Groups). Follow the Microsoft SQL Server best practices for Nutanix to achieve optimal performance.

### SQL Database Backups

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

> Tip: Back up your Citrix databases at regular intervals to minimize the impact of a disaster on your environment.

### SQL Database Scalability

To ensure that you've assigned sufficient resources to the Microsoft SQL Server environment, Nutanix recommends that you use the Database Sizing Guidance for XenApp and XenDesktop versions 7.6 through Current Release document.

> Tip: Size your SQL Server based on the database sizing document Citrix provides.

For specific details regarding SQL implementations on Nutanix, refer to the Nutanix best practices guide for Microsoft SQL.

## License Server

### Citrix License Server Redundancy

Because 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. When you use a virtualization platform, rely on the HA functionality for this specific VM.

> Tip: Deploy a single Citrix License Server and use HA when you use a hypervisor.

### License Server Scalability

Nutanix recommends that you conduct scalability tests to size your license server correctly. Citrix scalability testing indicates that a single Citrix License Server (with two cores and 2 GB of RAM) can issue approximately 170 licenses per second, or 306,000 licenses every 30 minutes.

 

## Microsoft License Server Redundancy

Nutanix recommends that you implement at least two Remote Desktop Services (RDS) license servers. When the first license server isn't available, the system can contact the second license server. The Group Policy Object (GPO) setting is located on the following path:

```
Computer Configuration\Policies\Administrative Templates\Windows Components
\Remote Desktop Services\Remote Desktop Session Host\Licensing
```

Tip: Deploy two Microsoft RDS license servers for redundancy.

 

# 7. Citrix Delivery Options

## Citrix Provisioning Solution

After they install Citrix Provisioning, an administrator or consultant prepares a base target device for imaging. This process installs all required software on that device (for example, MS Office, PDF readers and writers, and the Citrix Provisioning target device tools) along with all the optimizations you need for this particular image.

The administrator then creates a vDisk image from the base target device and saves it on the Citrix Provisioning server or storage device—a file share or storage system the Citrix Provisioning server can communicate with using iSCSI, SAN, NAS, or CIFS. Once the vDisk is available from the network, the image can stream from that location to a target device VM, allowing the target device VM to boot directly across the network. The Citrix Provisioning server streams the vDisk content to the target device on demand, and the target device acts as if it's running from a local drive.

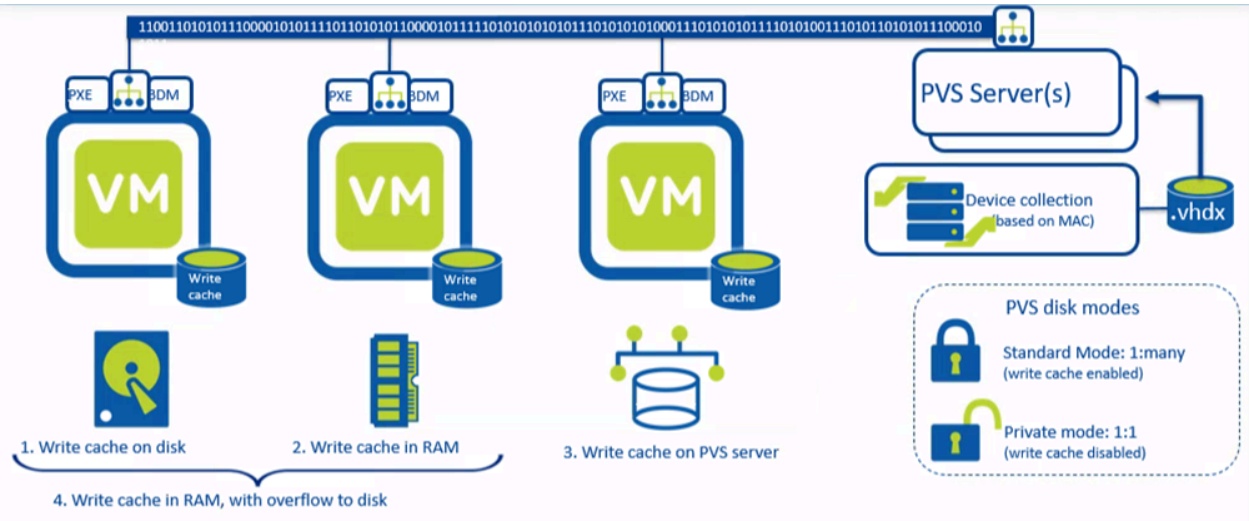


Figure 4: Provisioning vDisk Architecture

 

The administrator can assign a vDisk to either a single target device in Private Image mode, which allows them to make changes to the image, or to multiple target devices in Standard Image mode, which refreshes the default image when one of the target devices restarts.

## Provisioning Target Device Write Cache

While Citrix Provisioning is great for image management, there are certain items in Windows configurations that benefit from persistency, such as event logs, antivirus definitions, App-V cache, and other logs. The write cache is the Citrix Provisioning feature that allows you to save differential writes for persistent items in the Windows configuration. When data is written to the image with a configured write cache, it goes to the write cache file rather than to the base image itself. The write cache has the following options:

- Cache in device RAM.
- Cache on device RAM with overflow on hard disk.
- Cache on device hard disk.
- Cache on server.

When the target device boots to a vDisk in standard mode, the software on the virtual disk streams to the target device as needed. The target device write cache information is checked to determine the location of the write cache. When the target device restarts, it deletes the write cache, so the write cache is clean and doesn't contain anything from previous sessions.

### Target Device Write Cache Best Practices

AOS distributed storage simplifies write cache placement: There's no local disk management and the solution isn't typically constrained by IOPS. As a result, you can greatly simplify your architecture by directing the base-image VM's write cache to the Nutanix datastore. Place the write cache on the target device hard disk to obtain the maximum benefit using one of the following options:

- Cache on device RAM with overflow on hard disk: Set RAM to zero so the write cache goes directly to the target device's hard disk.

 

- Cache on device hard disk: Citrix plans to deprecate this option and remove it from the vDisk properties screen. However, if you need to, you can configure it using the Provisioning API.

### Best Practices for Provisioning Target Device Deployment

When you deploy provisioned target devices on AHV, consider the following:

- AHV only supports the Citrix Virtual Apps and Desktops Setup Wizard, not the Streamed VM Wizard.
- AHV uses snapshots for Provisioning Target Device templates. Create a template VM with the settings you want, then take a snapshot to use with the Citrix Virtual Apps and Desktops Setup Wizard.
  - Before you create the snapshot, use the aCLI to set the PXE boot settings in the template VM.
  - Before you create the snapshot, mount the Boot Device Manager (BDM) ISO in the template VM. You can't select the BDM option in the setup wizard on AHV. If you select PXE in the setup wizard, the provisioned target devices honor the template VM snapshot settings. If the template VM snapshot contains the BDM ISO and has the default boot order, you don't need to do anything else.
- AHV snapshots with disks attached aren't removed when target devices are provisioned using the setup wizard.
- You can only add provisioned target devices to a machine catalog using the Citrix Virtual Apps and Desktops Setup Wizard. As of version 2.7.0.0 of the Nutanix AHV Plugin for Citrix, you can add existing PVS VMs to Citrix Studio using the Citrix Provisioning option.
- Don't remove and re-add the NIC of a provisioned VM.
- AHV only supports Windows target devices at this time.
- AHV doesn't support BDM partitions at this time.
- Set UEFI on the VM before you install Windows to create a provisioned vDisk image. Set UEFI in the aCLI after you create the VM and before you install the guest OS.

 

- AHV doesn't support the Export Devices Wizard at this time.

## Citrix Machine Creation Services Solution

Citrix Virtual Apps and Desktops offers a fully integrated desktop virtualization suite for both single-session and multisession desktops. The Virtual Desktops broker console, Citrix Studio, lets users deploy all types of desktop and application workloads, whether persistent or nonpersistent, and the built-in Citrix Machine Creation Services (MCS) can derive each of these workloads from base images and clone them on the spot.

When nonpersistent environments use MCS, the broker copies the base image to each configured datastore that the Studio host connection specifies. This configured datastore can either be local on each host or shared on a SAN or NAS. The administrator then selects the available datastores, which VMware vCenter, Microsoft SCVMM, Citrix XenCenter, or the Nutanix Prism interface reads from the hypervisor cluster. After this copy finishes (which can take some time, depending on the number of datastores configured), the broker points all the VMs in the catalog to these local copies.

MCS works as shown in the following figure. Each supported hypervisor has its own specific MCS disk management implementation, but the net effect is the same.

 

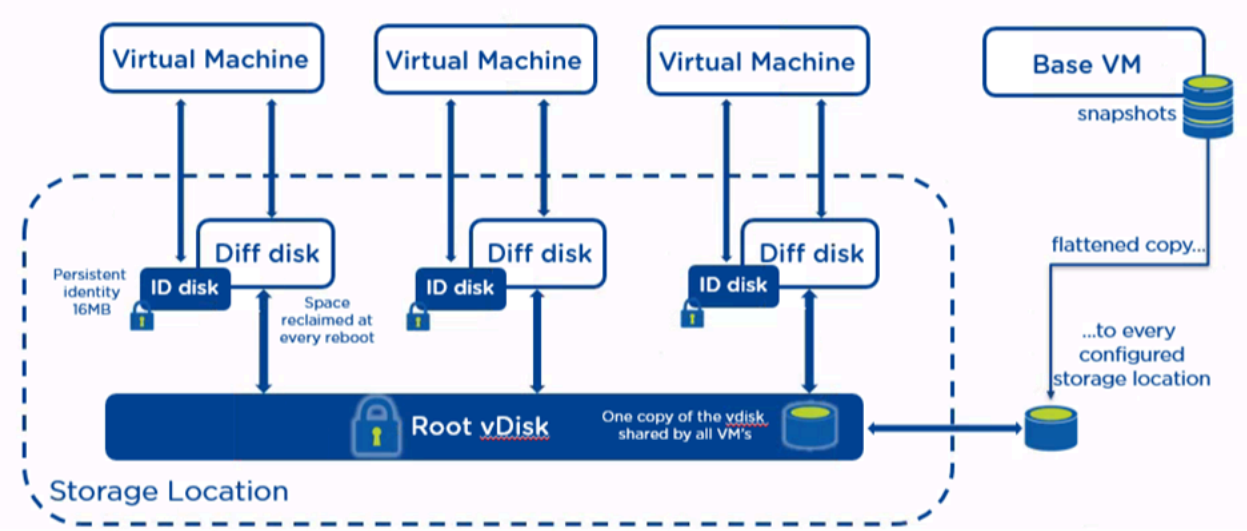


Figure 5: Citrix MCS Architecture

To make each VM unique and able to write data, MCS uses two disks in addition to the primary disk.

The identity disk, or ID disk, is 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. The broker fully manages this process; the administrator only needs to provide Active Directory accounts the VMs can use. The broker then creates a unique ID disk for every VM.

The difference disk, also known as the write cache, separates the writes from the primary disk, while the system still functions as if the write has been committed to the primary disk. VMware environments don't write changes to the difference disk file; instead, MCS on VMware uses a VMDK disk chain with multiple child disks. On Hyper-V and Citrix Hypervisor, MCS uses VHD chaining, an approach similar to VMware's, though slightly different in implementation and disk naming. Nutanix AHV uses copy-on-write, which avoids disk chaining and the potential disk corruption and performance issues associated with it.

Tip: When you use MCS with Virtual Desktops 7.9 or later on Nutanix, don't enable MCS Storage Optimization (MCS I/O).

Tip: Use vStorage API for Array Integration (VAAI) on vSphere or Offloaded Data Transfer (ODX) on Hyper-V when you use Citrix MCS. These selections speed up machine catalog deployment.

 

## Citrix Virtual Apps and Desktops and Nutanix Guest Tools on AHV

Nutanix Guest Tools (NGT) is a software bundle you can install in a guest VM to enable advanced VM functionality. The NGT installer contains the following components:

- Nutanix Guest Agent service.
- Self-Service Restore (SSR), also known as File-Level Restore (FLR), CLI.
- VM Mobility Drivers for cross-hypervisor disaster recovery between AHV and ESXi.
- VSS requestor and hardware provider for Windows VMs.
- Application-consistent snapshot for Linux VMs.

The NGT guest agent service uses SSL to communicate with the guest tools services through the Nutanix cluster IP address. Each VM running NGT needs a unique certificate pair to communicate with guest tools services.

> Tip: When you use Citrix App Layering, MCS, or PVS on Nutanix, don't install and enable NGT on base images. Nutanix recommends that you only install VirtIO drivers in base images with Citrix App Layering, MCS, or PVS nonpersistent images. If you use Citrix MCS to deploy a full clone image, you can install NGT after the MCS full clone deployment finishes.

## Citrix Virtual Apps and Desktops CPU and Core Assignment on AHV

When you create VMs for Citrix Virtual Apps and Desktops on AHV you can select vCPU and cores per CPU. This Nutanix best practice guide shows that there is no performance impact when you configure VMs with multiple CPUs or cores. Using the VDI and RDSH benchmarking tool LoginVSI, we show that on AHV, from a performance perspective, it doesn't matter if you add more vCPU or more CPU cores to the VMs; the user density and user experience is the same. We also compared the impact of using 4 vCPU with Windows 10 versus using 2 vCPU VMs. The results show that user density decreases but user experience improves.

 

## Assign CPUs and Cores on AHV

When you create virtual desktops or Remote Session Hosts with Citrix Virtual Apps and Desktops, you can assign the number of CPU sockets and the number of cores per CPU socket.

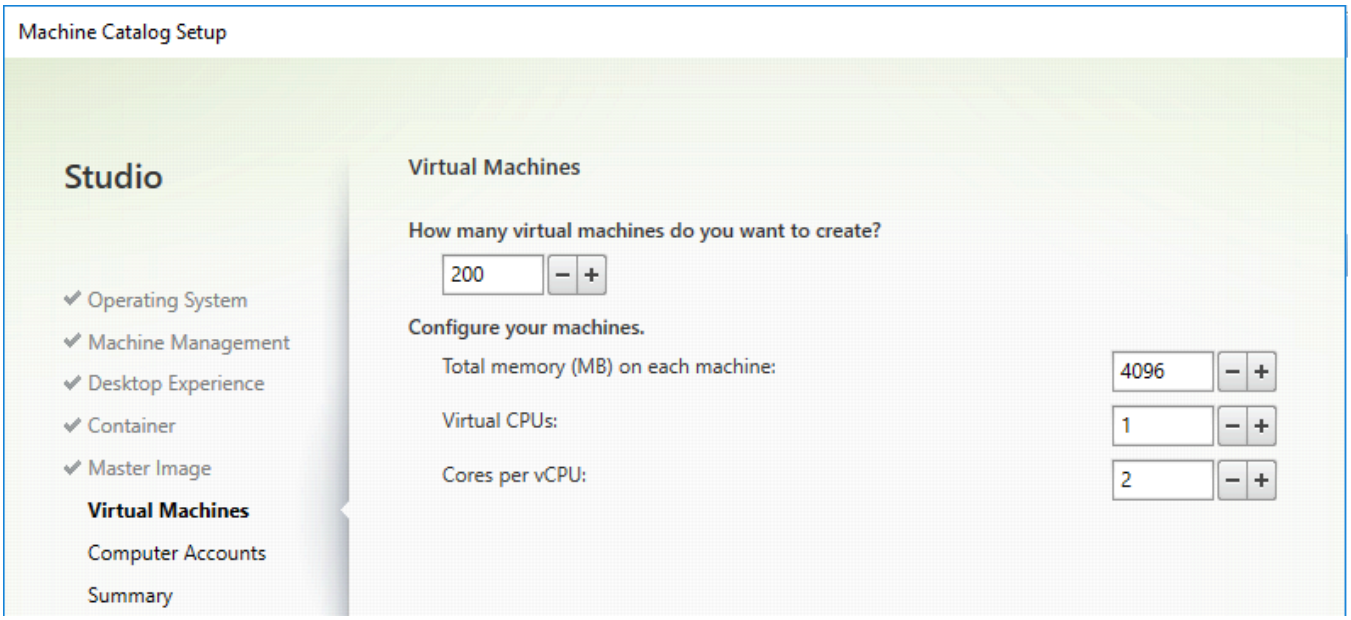


Figure 6: Machine Catalog Setup Virtual Machines

In a Virtual Apps or Desktops environment, you often use more than one CPU per VM. If the VMs need 2 vCPU, you can assign 1 vCPU with two cores or you can assign 2 vCPU with a single core each. If you run these VMs on AHV, there should be no differences in performance or density between the two options. We used LoginVSI to test performance and density using VMs with the following configurations.

Table 5: LoginVSI VM Configurations

| Parameter | Virtual Desktops | Virtual Apps |
|---|---|---|
| Operating system | Windows 10<br>Build 17134.254 | Windows Server 2016<br>Build 14393.2941 |

 

| Parameter | Virtual Desktops | Virtual Apps |
|---|---|---|
| Assigned memory per VM | 2.5 GB | 42 GB |
| CPU configurations | 2 CPU:<br>— 1 CPU, 2 cores<br>— 2 CPUs, 1 core<br>4 CPU:<br>— 1 CPU, 4 cores<br>— 4 CPU, 1 core<br>— 2 CPU, 2 cores | 8 CPU:<br>— 1 CPU, 8 cores<br>— 8 CPU, 1 core<br>— 2 CPU, 4 cores |
| Number of VMs | 200 | 10 |
| Launched sessions | 200 | 300 |
| LoginVSI workload | Knowledge Worker | |
| Display protocol | Citrix HDX | |
| AOS | 5.11 | |
| AHV | 20170830.301 | |

## Test Results

We used LoginVSI to simulate the user workload so we can compare the results of the metrics VSImax and VSIbaseline. The VSImax indicates user capacity, where a higher VSImax is better. The VSIbaseline represents the user experience by calculating an average response time for a predefined set of actions (lower VSIbaseline is better).

> Note: You can't use the VSImax results in this document for sizing purposes.

We performed at least five tests with each configuration. The following table shows the results of the tests performed with different Virtual Desktops configurations.

 

Table 6: LoginVSI VM Configurations and Results for Virtual Desktops

| Windows 10 – 1803 | VSImax | % Difference | VSIbaseline | % Difference |
|---|---|---|---|---|
| | | 2 vCPU VMs | | |
| 1 CPU, 2 cores | 187 | 0.0% | 880 | 0.0% |
| 2 CPU, 1 core | 186 | –0.5% | 884 | 0.5% |
| | | 4 vCPU VMs | | |
| 1 CPU, 4 cores | 144 | –23.0% | 812 | –7.7% |
| 4 CPU, 1 core | 146 | –21.9% | 809 | –8.1% |
| 2 CPU, 2 cores | 147 | –21.4% | 812 | –7.7% |

If you configure your Windows 10 VMs with 2 vCPU, it doesn't matter if you select 1 CPU with two cores or 2 CPU with one core each. The density and user performance are the same.

If you decide to configure the VMs for 4 vCPU, the density is around 22 percent lower compared to 2 vCPU (using the same workload). The user experience improves, decreasing the response times by around 8 percent. If you use 4 CPU, it doesn't matter if you configured the VMs with 1 CPU and four cores, 4 CPU with one core, or 2 CPU with two cores each; the density and performance are the same.

The following table shows the results of the tests performed with Virtual Apps workloads.

Table 7: LoginVSI VM Configurations and Results for Virtual Apps

| Windows Server 2016 | VSImax | % Difference | VSIbaseline | % Difference |
|---|---|---|---|---|
| 1 CPU, 8 cores | 244 | –1.2% | 731 | 0.6% |
| 8 CPU, 1 core | 247 | 0.0% | 727 | 0.0% |
| 2 CPU, 4 cores | 247 | 0.0% | 725 | –0.3% |

 

Again, using different configurations for Virtual Apps VMs doesn't impact performance. The differences are minimal and within the margins commonly seen using LoginVSI.

## Nutanix AHV Plugin for Citrix

The Nutanix AHV Plugin for Citrix Virtual Apps and Desktops allows Delivery Controllers to manage workloads running on Nutanix AHV. Install the Nutanix AHV plugin for Citrix Virtual Apps and Desktops on all Delivery Controllers in the site for single-zone sites or all Delivery Controllers in the same zone for multizone sites. For multizone sites, you must install the Nutanix AHV plugin for Citrix Virtual Apps and Desktops on the primary-zone Delivery Controllers and the satellite-zone Delivery Controllers where you plan to deploy Nutanix.

Nutanix frequently updates the AHV Plugin for Citrix Virtual Apps and Desktops with fixes and new functionalities and to accommodate changes between Citrix Virtual Apps and Desktops versions. Nutanix recommends that customers running Citrix Virtual Apps and Desktops on Nutanix AHV update their plugin installs whenever Nutanix updates the plugin to ensure continued functionality.

You can use the AHV Plugin for Citrix Virtual Apps and Desktops with Citrix Virtual Apps and Desktops 7.9 and later. The AHV Plugin for Citrix Virtual Apps and Desktops requires administrator privileges for the AHV cluster. Nutanix recommends that customers use a local cluster account for the AHV Plugin for Citrix Virtual Apps and Desktops, so create a local cluster account on the Nutanix AHV cluster with cluster administrator privileges.

> Note: Nutanix recommends using local authentication, but you can use Active Directory authentication. With Active Directory authentication, Nutanix recommends having a user or a user in a group with Cluster Admin permissions and the Authentication Configuration search type set to Non Recursive (default).

 

# 8. Nutanix Storage Configuration

Table 8: Storage Best Practices for Citrix Virtual Apps and Desktops

| Delivery Method | Compression | Elastic Deduplication Engine | Erasure Coding |
|---|---|---|---|
| Full clones | X | X | |
| Citrix MCS | X | | |
| Citrix Provisioning | X | | |

Tip: Enabling compression for Virtual Apps or Virtual Desktops is a general best practice; only enable the Elastic Deduplication Engine for full clones. Erasure coding isn't a suitable data reduction technology for desktop virtualization.

## Capacity Optimization

Nutanix hybrid multicloud software offers capacity optimization features that improve storage utilization and performance. The three key features are:

1. Compression
2. Deduplication
3. Erasure coding (EC-X)

For more information on Nutanix capacity optimization techniques, read the Data Efficiency tech note.

Tip: As of AOS 4.5, you can enable both deduplication and compression on the same container. However, unless the data can be deduplicated (conditions explained later), only use compression.

### Compression

The Nutanix Capacity Optimization Engine (COE) performs data transformations to increase data efficiency on disk. Compression is one of the key features of the COE. AOS storage provides inline, oplog, and offline (post-process) compression to best suit the customer's specific needs and types of data. As of AOS 5.1, offline compression is enabled by default.

 

### Inline Compression

The system compresses data synchronously as it's written to optimize capacity and maintain high performance for sequential I/O operations. Inline compression only compresses sequential I/O to avoid degrading performance for random write I/O. The following figure shows an example of how inline compression interacts with the AOS distributed storage write I/O path.

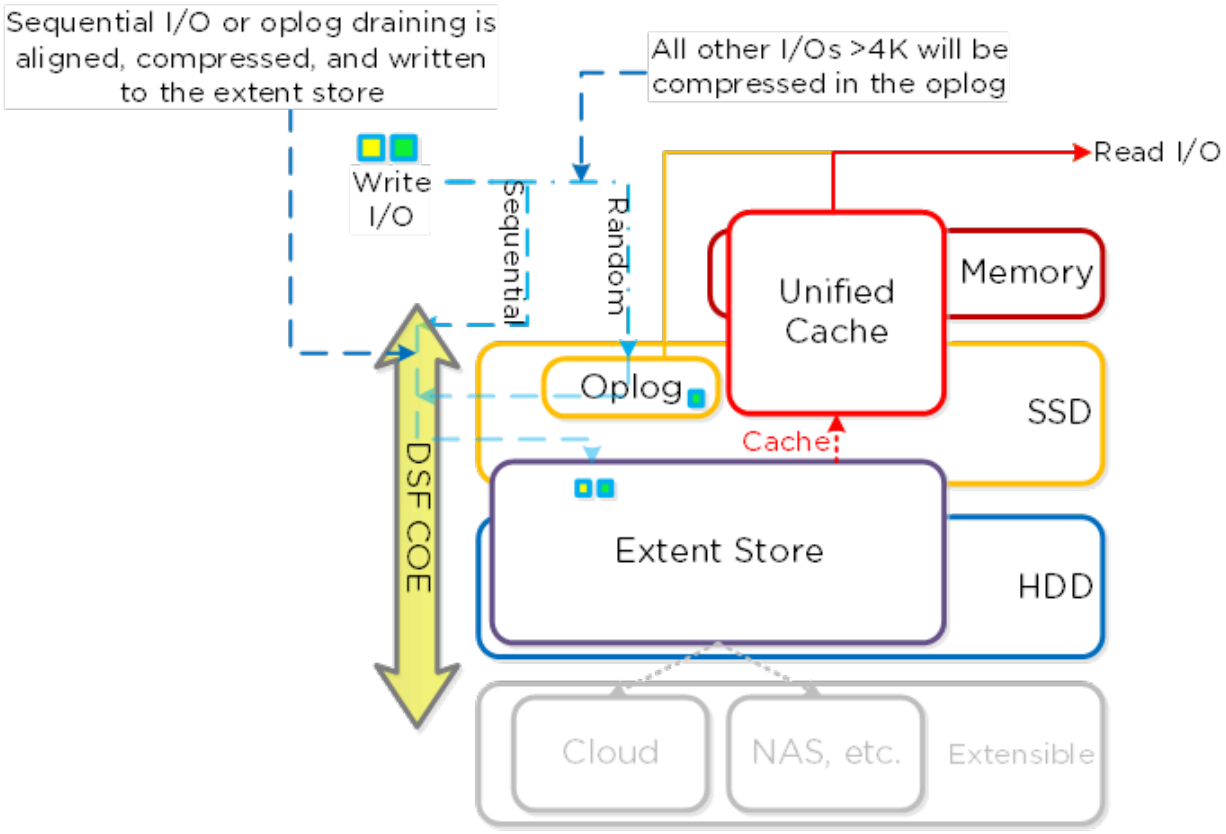


Figure 7: Inline Compression I/O Path

Tip: Almost always use inline compression (compression delay = 0), as it only compresses larger or sequential writes and doesn't impact random write performance.

Inline compression also increases the usable size of the SSD tier, which increases effective performance and allows more data to sit in the SSD tier. Also, for larger or sequential data written and compressed inline, the replication to maintain the replication factor ships the compressed data, further increasing performance by sending less data across the wire.

 

Inline compressions pairs perfectly with erasure coding.

When you enable inline compression but the I/O operations are random, the data is written uncompressed in the oplog, coalesced, then compressed in memory before being written to the extent store.

### Oplog Compression

As of AOS 5.0, the oplog compresses all incoming writes greater than 4 KB that show good compression to enable more efficient usage of the oplog capacity and help sustained performance. When the data drains from the oplog to the extent store, it's decompressed, aligned, then recompressed at a 32 KB aligned unit size (as of AOS 5.1). This feature is on by default and no user configuration is necessary.

### Post-Process Compression

Post-process compression writes all new write I/O in an uncompressed state and sends it on the normal AOS storage I/O path. After the data meets the configurable compression delay and becomes cold, the data is eligible for compression, which can occur through Information Life Cycle Management (ILM). Post-process compression uses the Curator MapReduce framework and all nodes can perform compression tasks. Chronos throttles compression tasks.

The following figure shows an example of how post-process compression interacts with the AOS storage write I/O path.

 

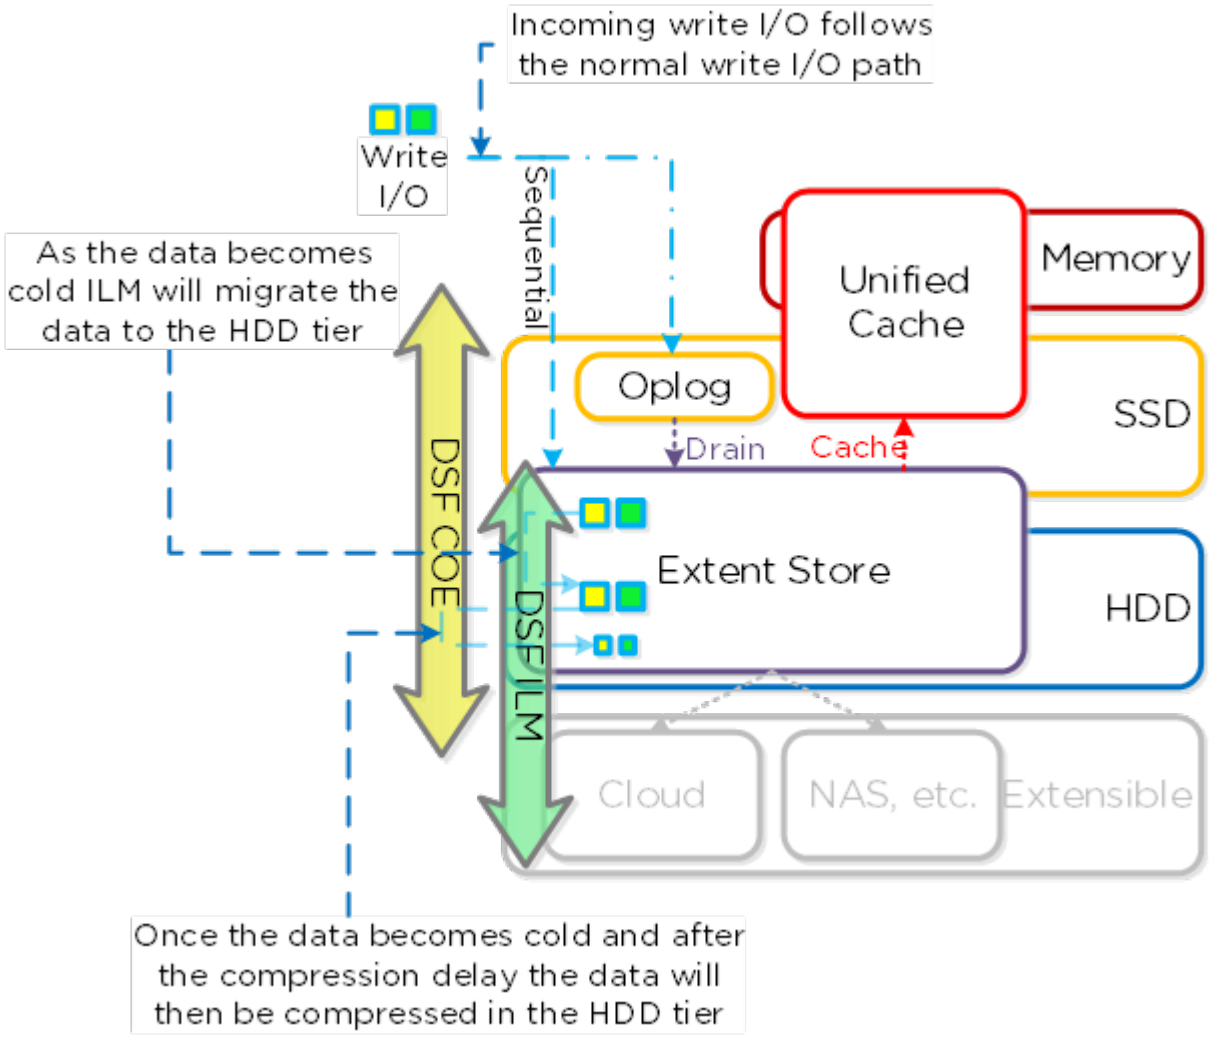


Figure 8: Post-Process Compression I/O Path

For read I/O, the system first decompresses the data in memory, then serves the I/O. Heavily accessed data is decompressed in the HDD their and then uses ILM to move up to the SSD tier and the cache.

The following figure shoes an example of how decompression interacts with the distributed storage I/O path during reads.

 

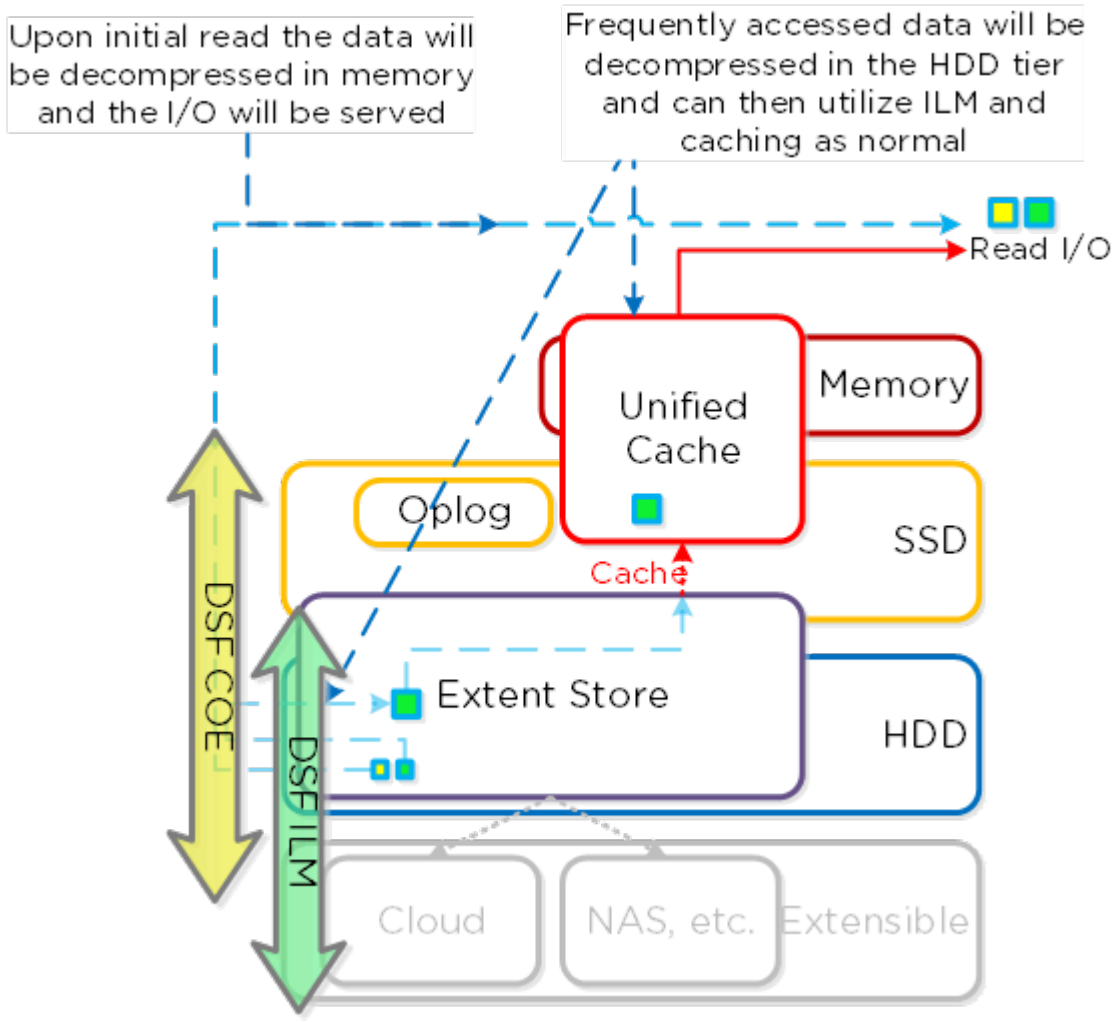


Figure 9: Decompression I/O Path

## Elastic Deduplication Engine

Note: Only enable the Elastic Deduplication Engine for full clones.

The Elastic Deduplication Engine is a software-based feature that deduplicates data in the capacity (extent store) and performance (unified cache) tiers. The system fingerprints streams of data during ingest using a SHA-1 hash at an 8 KB granularity. This fingerprint only occurs on data ingest and is then stored persistently as part of the written block's metadata. The stored fingerprints

 

allow the Elastic Deduplication Engine to detect and remove duplicate copies easily, without scanning or reading the data again.

To make metadata overhead more efficient, Nutanix monitors fingerprint refcounts to track dedupability. The system discards fingerprints with low refcounts to minimize the metadata overhead. Capacity-tier deduplication prefers full extents to minimize fragmentation.

> Tip: Use performance-tier deduplication on your base images (you can manually fingerprint them using vdisk_manipulator) to take advantage of the unified cache. Use capacity-tier deduplication for P2V and V2V, when you use Hyper-V (because ODX does a full data copy), or when you do cross-container clones (not usually recommended). In most other cases compression yields the highest capacity savings and should be used instead.

The following figure shows an example of how the Elastic Deduplication Engine scales and handles local VM I/O requests.

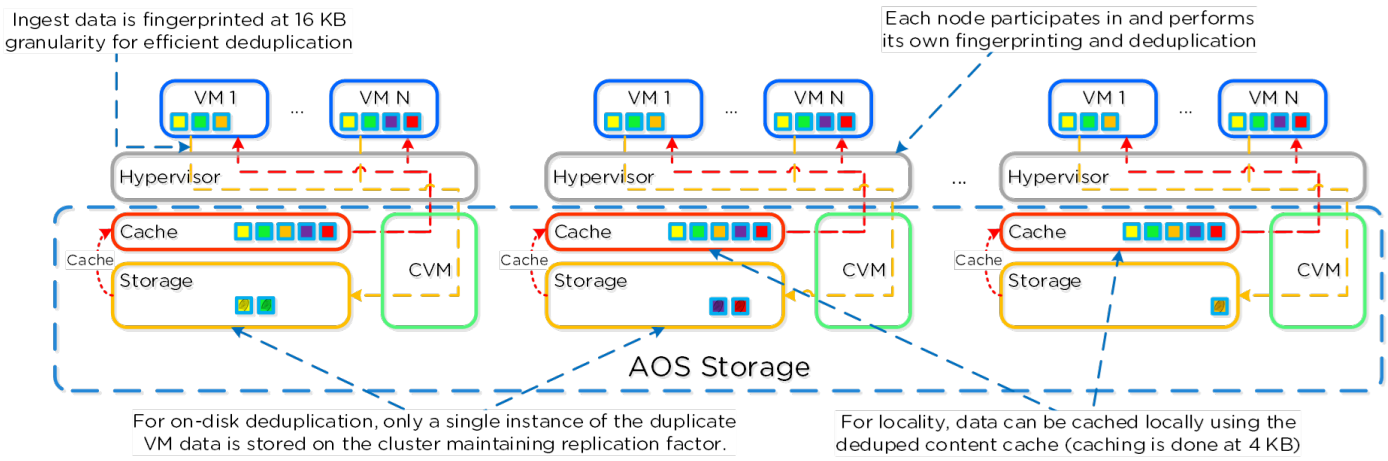


Figure 10: Elastic Deduplication Engine: Scale

Fingerprinting occurs during data ingest with an I/O size of 64 KB or greater (either initial I/O or when draining from the oplog). The engine uses Intel acceleration for the SHA-1 computation, which creates minimal CPU overhead. In cases where fingerprinting doesn't occur during ingest (for example, with smaller I/O sizes), it can run as a background process.

As the engine identifies duplicate data (multiple copies of the same fingerprints), a background process removes the duplicate data using the Curator MapReduce framework. Data being read is pulled into the unified

 

cache, which is a multitier or pool cache. Any subsequent requests for data with the same fingerprint are pulled directly from the cache.

> Tip: AOS versions 4.6.1 and later have no limit to fingerprinted vDisk offsets, so you can fingerprint or dedupe the full vDisk.

The following figure shows how the Elastic Deduplication Engine interacts with the AOS I/O path.

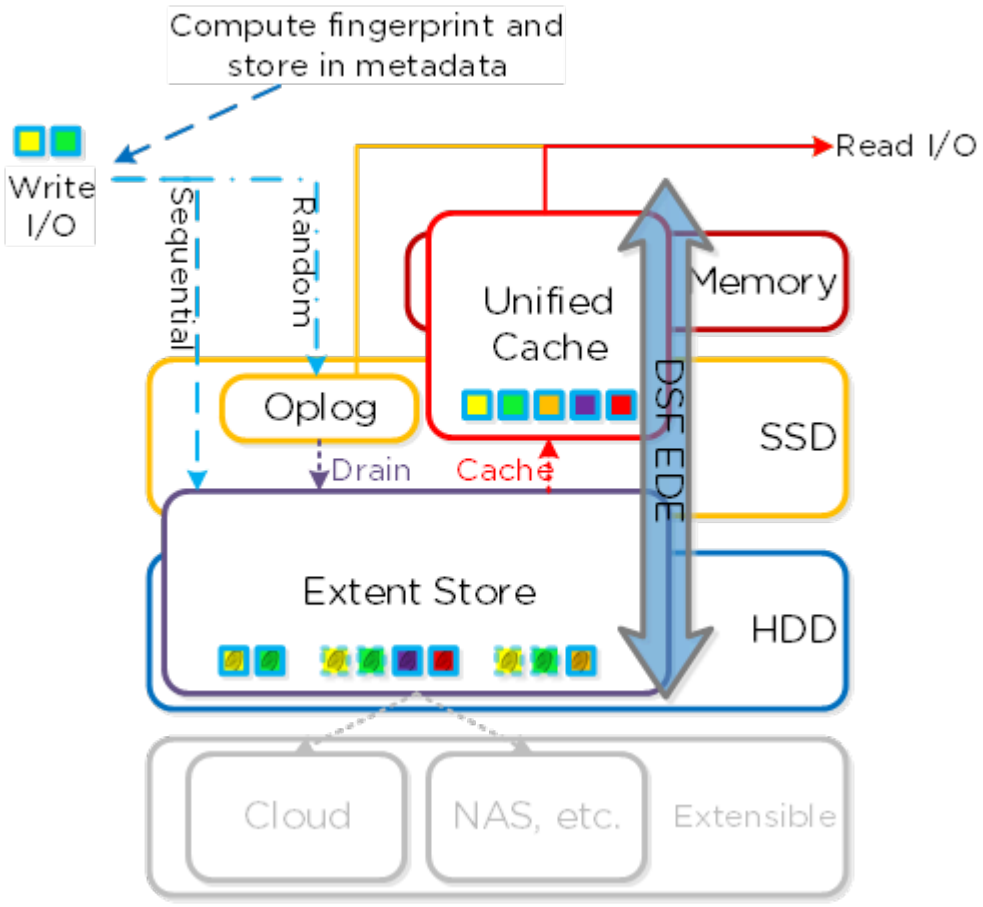


Figure 11: Elastic Deduplication Engine I/O Path

## Erasure Coding

> Note: Erasure coding isn't a suitable data reduction technology for desktop virtualization.

The Nutanix platform relies on a replication factor for data protection and availability. This method provides the highest degree of availability because

 

it doesn't require reading from more than one storage location or data recomputation on failure. However, because this feature requires full copies, it uses additional storage resources. The distributed storage minimizes the required storage from this feature by encoding data using erasure codes (EC-X).

Similar to the concept of RAID (levels 4, 5, 6, and so on), EC-X encodes a strip of data blocks on different nodes and calculates parity. In the event of a host or disk failure, the system can use the parity to decode any missing data blocks. In AOS distributed storage, the data block is an extent group.

## Networking, I/O, and Data Locality

The Nutanix platform relies on a standard 10 GbE or higher network (rather than a backplane) for communication between nodes. The hypervisor handles all storage I/O for VMs running on a Nutanix node on a dedicated private network. The hypervisor forwards I/O requests to the private IP on the local CVM. The CVM then uses its external IP over the public network to perform remote replication with other Nutanix nodes. Because the system can serve nearly all read requests locally, the public 10 GbE network is generally reserved for remote replication traffic and VM network I/O. Exceptions can occur when a CVM goes down and its requests are forwarded to other CVMs in the cluster, or when the system must pull remote data. Cluster-wide tasks, such as disk balancing, can also temporarily generate I/O on the 10 GbE network.

The following figure shows how the VM's I/O path interacts with the private and public 10 GbE network.

 

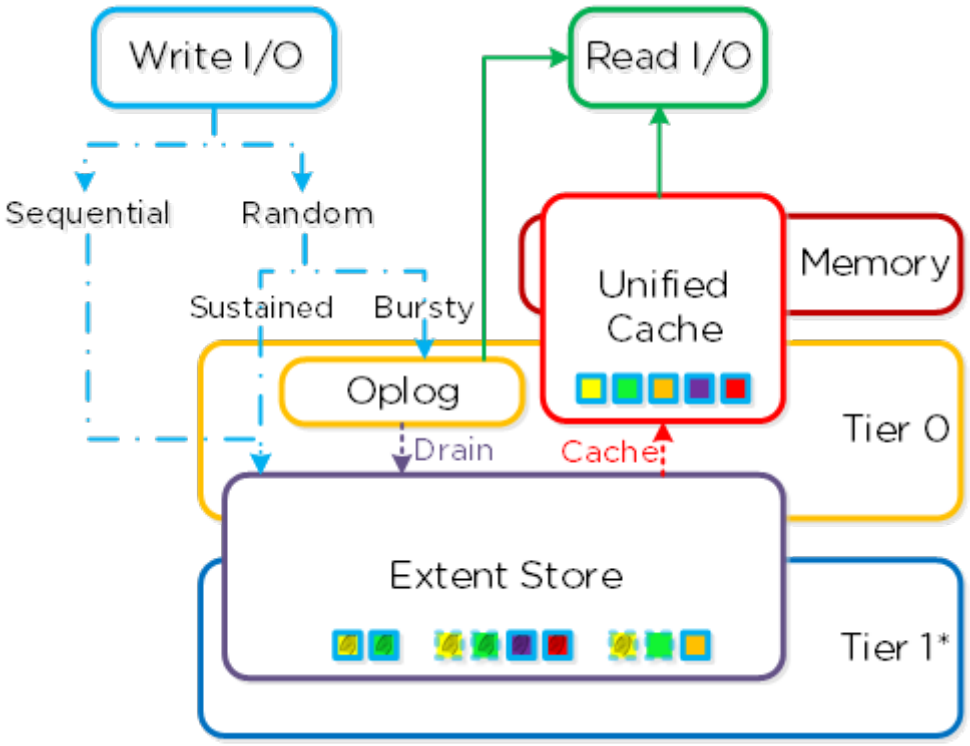


Figure 12: VM I/O Path

As a converged (compute and storage) platform, Nutanix relies heavily on I/O and data locality for cluster and VM performance. As described above, the local CVM controls its local disks and serves all read/write I/O operations. When a VM moves from one hypervisor node to another (or during a HA event), the CVM that's now local to the newly migrated VM takes over serving its data. When reading old data (stored on the node that is now remote), the now-local CVM forwards I/O to the remote CVM. All write I/O occurs locally right away. The system detects that I/O is now occurring on a different node and migrates the data in the background, which allows all read I/O operations to be served locally once again. To keep migration tasks from flooding the network, the data only migrates when it's read.

Data locality takes two main forms:

1. Cache locality: vDisk data is stored locally in the unified cache. vDisk extents may be remote to the node.
2. Extent locality: vDisk extents are local to the same node that hosts the VM.

 

Cache locality occurs in real time and is determined by vDisk ownership. When a vDisk or VM moves from one node to another, ownership of those vDisks or VMs transfers to the new local CVM. Once ownership transfers, the data can be cached locally in the unified cache. In the interim, the cache is located in the now-remote host, which holds ownership. The instance of Stargate that previously hosted the vDisk token relinquishes the vDisk token when it sees remote I/O operations for 300+ seconds, at which point the new local Stargate instances takes the token. Cache coherence is enforced because ownership is required to cache the vDisk data.

Extent group locality is a sampled operation. Extent groups are migrated when the following occurs: 3 touches for random I/O or 10 touches for sequential I/O in a 10-minute window (multiple reads in a 10-second sampling count as a single touch).

The following figure shows how data follows a VM as it moves between nodes.

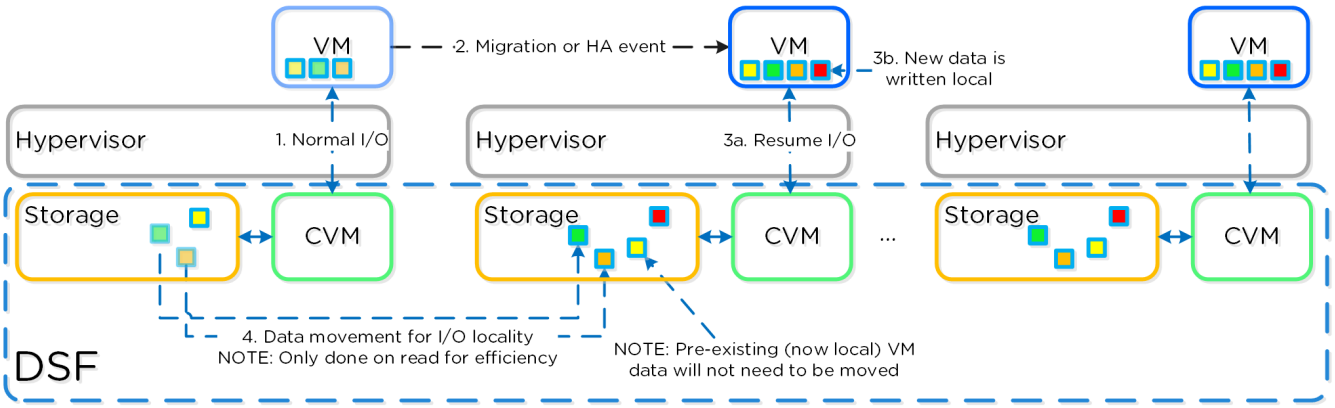


Figure 13: Data Locality

For more details on networking, data locality, and I/O with AHV, read the AHV best practice guide or the AHV Networking best practice guide .

## I/O Path Details

In the CVM, the Stargate process handles all I/O coming from user VMs and persistence (replication factor). When a write request comes to Stargate, a write characterizer determines whether the write is persisted to the oplog, extent store, or autonomous extent store. For reads, the read characterizer handles reads and manages caching and readahead.

 

### OpLog

The oplog is a persistent write buffer similar to a filesystem journal. It's built as a staging area to handle bursts of random writes, coalesce them, then sequentially drain the data to the extent store. On a write, the oplog is synchronously replicated to another number of CVMs before the write is acknowledged for data availability purposes. All CVM oplogs partake in replication and are dynamically chosen based on load.

The oplog is stored on the SSD tier of the CVM to provide extremely fast write I/O performance, especially for random I/O workloads. All SSD devices participate and handle a portion of oplog storage. For sequential workloads, the oplog is bypassed and the writes go directly to the extent store. If data is sitting in the oplog and hasn't been drained, all read requests are directly fulfilled from the oplog until the data has been drained, at which point read requests are served by the extent store or unified cache.

### Extent Store

The extent store provides the bulk of persistent data storage. It spans all device tiers (PCIe SSD, SATA SSD, HDD) and can extend to facilitate additional devices and tiers. Data entering the extent store is either draining from the oplog or sequential or sustained in nature and bypassed the oplog. Nutanix ILM determines tier placement dynamically and moves data between tiers based on I/O patterns.

> Note: Write I/O is sequential when there is more than 1.5 MB of outstanding write I/O going to a vDisk. Sequential write I/O bypasses the oplog and goes directly to the extent store because it's already a large chunk of data and doesn't benefit from coalescing. All other I/O types, including large I/O (greater than 64 KB), are still handled by the oplog.

In all-flash configurations, the extent store consists of SSD devices, and no tier ILM occurs because there's only a single flash tier. In hybrid all-flash configurations (for example, NVMe or Intel Optane with a SATA SSD), the highest performance media is Tier 0 and the lower performance media is Tier 1. For hybrid configurations that aren't all flash, the SSD is Tier 0 and the HDD is Tier 1.

> Note: As of AOS 5.10, the autonomous extent store (AES) can handle sustained random workloads when requisite conditions are met.

 

### Unified Cache

The unified cache acts as the dynamic read cache. It's used for data, metadata, and deduplication, and is stored in the CVM's memory. On a read request for data not in the cache (or based on a particular fingerprint), the system places data into the single-touch pool of the unified cache. The system assigns the data an LRU (least recently used) counter, which it uses to sequentially evict data from the unified cache. Any subsequent read request for data in the single-touch pool moves the data to the multitouch pool (no actual data is moved, just cache metadata). Any read request for data in the multitouch pool moves the data to the peak of the multitouch pool, where it receives a new LRU counter.

You can calculate cache size as ((CVM memory – 12 GB) × 0.45). For example, a 32 GB CVM has the following cache size: ((32 – 12) × 0.45) = 9 GB.

The following figure shows a high-level overview of the unified cache data path.

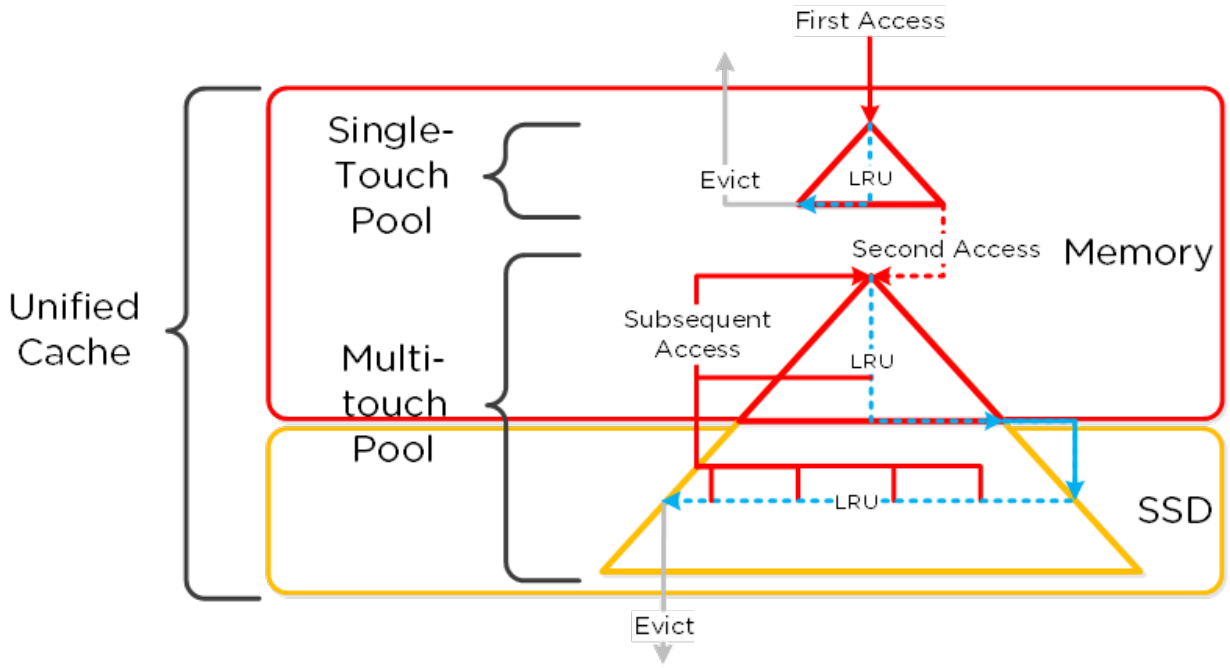


Figure 14: Unified Cache Data Path

Each CVM has its own local cache that it manages for the vDisks it hosts (for example, VMs running on the same node). When you clone a vDisk (new clones,

 

snapshots, and so on), each new vDisk has its own block map and the original vDisk is marked as immutable. This method ensures that each CVM can have its own cached copy of the base vDisk with cache coherency. In the event of an overwrite, that copy is redirected to a new extent in the VM's own block map to protect the cache from corruption.

## Shadow Clones

Shadow Clones are a unique feature of AOS distributed storage that enables distributed caching of vDisks across Nutanix. Shadow Clones provide effective caching optimization in distributed multireader scenarios, including large VDI and cloud deployments, where VMs on multiple nodes (VM hosts) in a Nutanix cluster read from the same set of base data.

Continuing with the VDI example, when you deploy desktops as linked clones, the system forwards some of the read requests to a central leader or base VM for the clone pool. In VMware Horizon, the base VM is called the replica disk and all linked or instant clones can read it. In Citrix Virtual Apps and Desktops deployments that use Machine Creation Services (MCS), the base VM is called the MCS Base VM. Shadow Clone optimization works in either of these cases. Shadow Clones help decrease read latency in any scenario with distributed multireader access, not just in VMware and Citrix VDI environments.

Systems that run on Nutanix AOS use data and I/O locality to obtain the highest possible VM performance. With Shadow Clones, distributed storage monitors vDisk access patterns to determine whether VMs are frequently reading the same data set from multiple nodes in the cluster. If it detects this situation, distributed storage marks the vDisk as immutable. The remote CVM then creates a local cached copy, which serves subsequent read requests for that node. The remote read-only cached copy of the vDisk is a Nutanix Shadow Clone. The process of determining the need for and creating a Shadow Clone is transparent to the hypervisor and VMs and is entirely handled by AOS distributed storage. When the base VM or vDisk is modified, distributed storage removes the remote Shadow Clone. This procedure ensures that the VMs receive current data. Depending on future read patterns, the process starts over again.

 

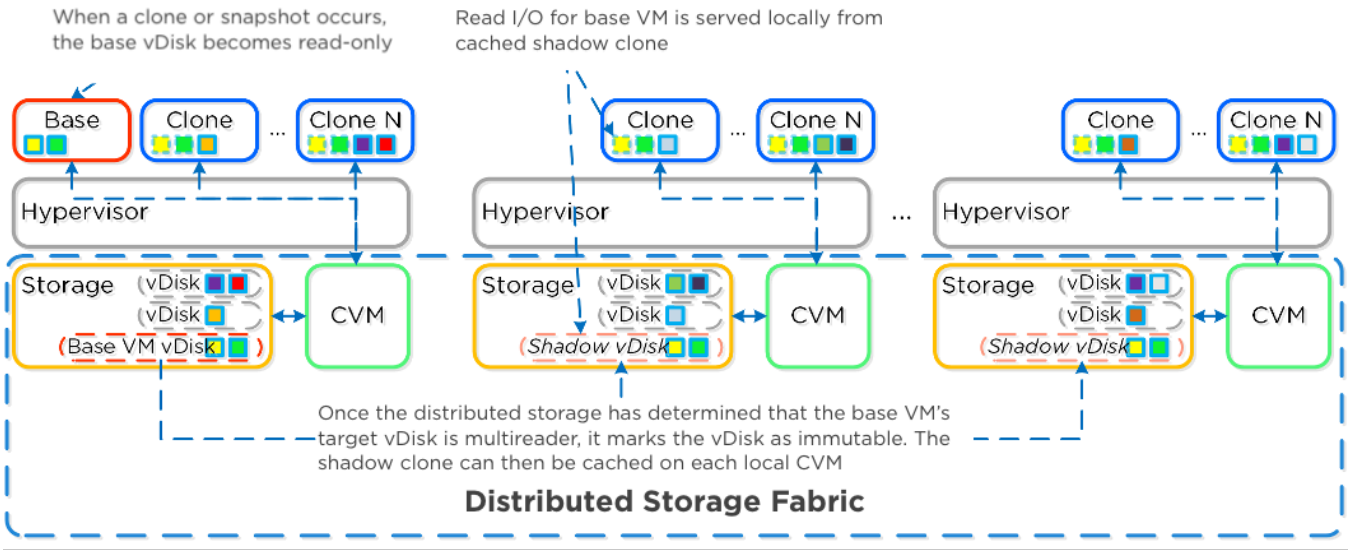


Figure 15: Nutanix Shadow Clones

For additional information on Nutanix Shadow Clones, see the Performance Analysis of Nutanix Shadow Clones tech note.

## Nutanix Controller VM

Each node runs an industry-standard hypervisor (ESXi, AHV, or Hyper-V) and the Nutanix CVM. The Nutanix CVM runs the Nutanix software and serves all I/O for the chosen hypervisor and the VMs running on that host. For Nutanix units that run VMware vSphere, the SCSI controller—which manages the SSD and HDD devices—directly passes to the CVM through VMDirectPath (Intel VT-d). In Hyper-V, storage devices pass through to the CVM.

The following figure provides an example of what a typical node looks like logically.

 

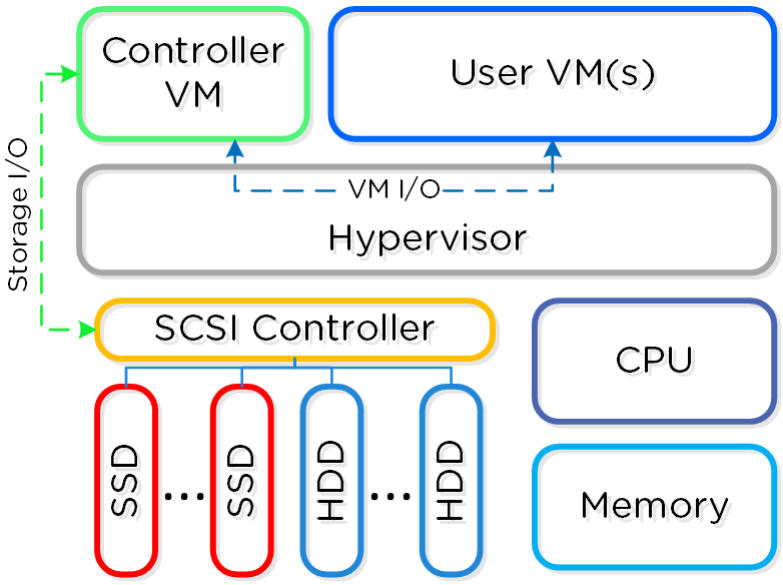


Figure 16: Hyperconverged Platform

The CVM runs on the selected hypervisor and has assigned vCPU and RAM resources; the ideal configuration depends on the platform and enabled features. The following table lists the most popular Nutanix platforms for desktop virtualization.

Table 9: Nutanix Models

| Hardware Platform | Default Memory | Default vCPU |
|---|---|---|
| NX-3060-G7 | 32 GB | 12 |
| NX-3155-G7 | 32 GB | 12 |
| NX-3170-G7 | 32 GB | 12 |

The CVM uses assigned CPU cycles to handle storage I/O. Because desktop virtualization I/O requirements don't push AOS storage to its limits, the CVM typically has unused resources that it actively pushes back to user VMs.

Note: Don't change the CPU sizing of your CVMs unless Nutanix Support tells you to do so.

The following table lists the minimum amount of memory required to add each of the given features. These memory requirements are in addition to the default

 

memory available. Add the memory required for each desired feature to the amount for your model type (specified in the previous table) to determine the total memory the CVM needs.

Note: For erasure coding and compression, you don't need to add memory to the base CVM configuration. The additional memory can't exceed 16 GB for the features in the following table.

Table 10: Memory Assignment per Feature

| Features | Memory |
|---|---|
| Capacity tier deduplication (includes performance tier deduplication) | 12 GB |
| Replication factor 3 | 8 GB |
| Performance tier deduplication | 8 GB |

Tip: Assign memory according to the required features. Always discuss memory allocation with your Nutanix representative.

 

# 9. Conclusion

Our extensive testing of Citrix MCS and Provisioning deployments on Nutanix demonstrates that desktop user density is based primarily on the available host CPU resources, not on any I/O or resource constraints. The Nutanix data reduction technologies maximize available capacity and improve performance. These features have minimal impact on the total available resources and thus minimal impact on user VMs and user performance.

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

 

# Appendix

## Best Practices Checklist

- Perform a current state analysis to identify workloads and sizing for the desktops and applications you plan to virtualize.
- Gather and document functional and technical requirements for the virtual desktop solution.
- Spend time upfront to architect a solution that meets both current and future needs.
- Design for end-user experience to deliver consistent performance, reliability, and scale.
- Start with a PoC, then test, optimize, iterate, and scale.
- Size workloads appropriately for each particular use case.
- Use a mix of application virtualization and applications installed in gold images, depending on the scenario.
- Optimize images using the tools listed in this document.
- Design for reliability and scale.
- Don't overcommit RAM.
- Use a single container and datastore for virtual desktops and Virtual Apps–based VMs.
- Configure storage containers based on the workload delivery method per the storage best practices in this document.
- Configure proper infrastructure components per the supporting components best practices in this document.
- Deploy and configure Nutanix AHV Plugins for Citrix per the Citrix delivery options best practices in this document.

 

## References


1. Citrix VDI Handbook and Best Practices document
2. Citrix Tech Zone
3. VMware OS Optimization Tool
4. Best Practice Analyzer
5. Citrix Optimizer
6. Microsoft SQL Server best practices guide
7. Database Sizing Guidance for XenApp and XenDesktop versions 7.6 through Current Release document
8. Citrix Provisioning Managing for highly available implementations
9. Windows 10 VDI Performance Impact Analysis
10. Nutanix Bible


## About the Author

Jarian Gibson is a staff solutions architect on the End-User Computing Engineering team at Nutanix. Follow Jarian on Twitter @JarianGibson.

## About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix enterprise cloud software 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 us on Twitter @nutanix.

 

# List of Figures

# List of Tables