v1.4 | September 2022 | BP-2060

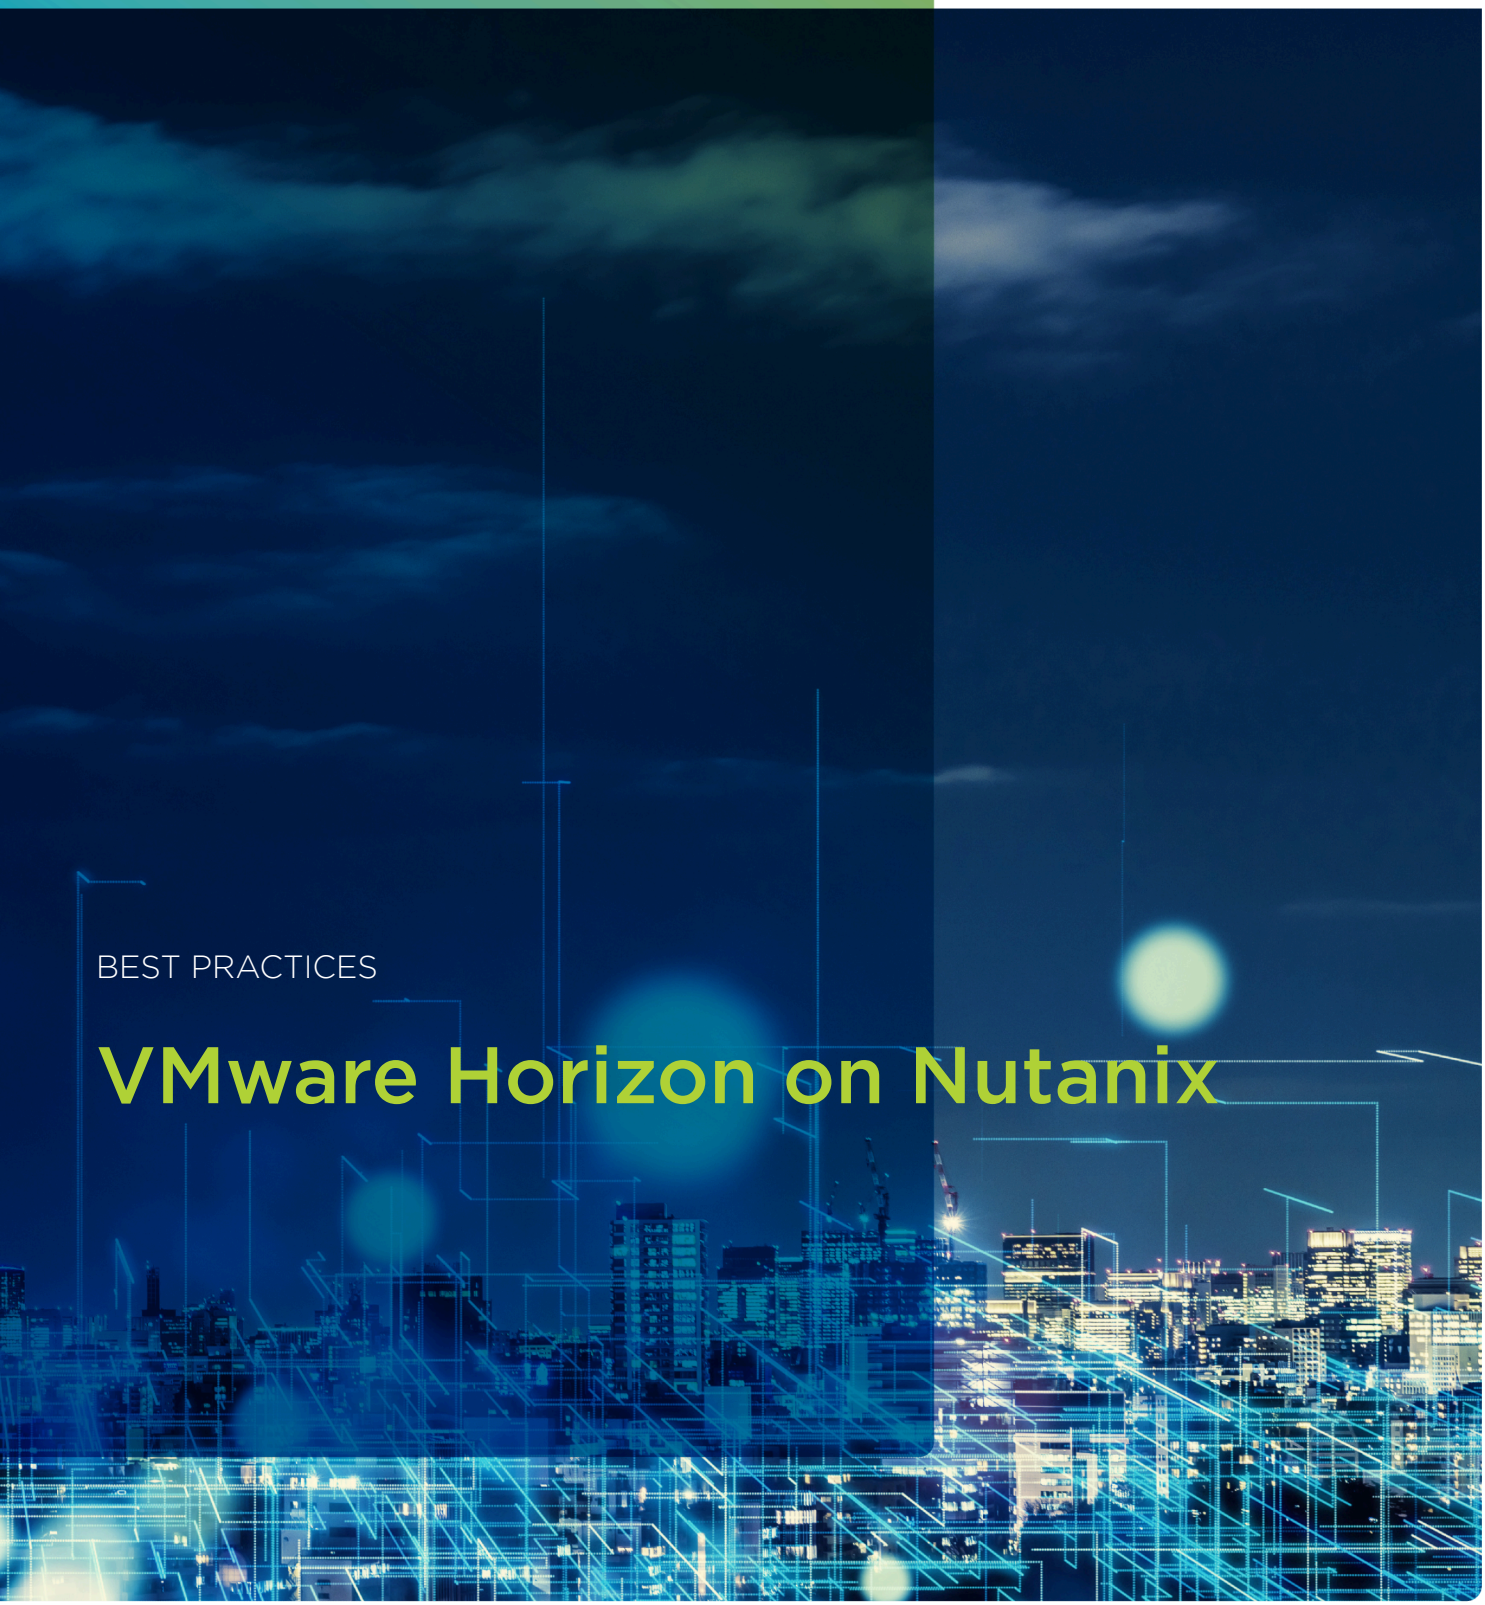

BEST PRACTICES

VMware Horizon on Nutanix

NUTANIX
YOUR ENTERPRISE CLOUD

Copyright

Contents

1. Executive Summary

This guide discusses the best practices for running VMware Horizon 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 fit your enterprise needs and the objectives of your operators, whether they're task workers or power users.

Nutanix includes the Nutanix storage method, which offers a range of advantages for VMware Horizon deployments:

- An 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 VMware Horizon deployments.
- Nutanix Shadow Clones, which cut network latency and improve user experience.

Delivering applications using VMware Horizon on Nutanix means you can easily deploy thousands of desktops with an optimal user experience across multiple use cases.

2. Introduction

Audience

This best practice guide is part of the Nutanix Solutions Library. We wrote it for IT architects and administrators as a technical introduction to Nutanix and VMware Horizon.

Purpose

This document covers the following subject areas:

- Overview of the Nutanix solution.
- Best practices for delivering VMware Horizon on Nutanix.

Document Version History

Version Number	Published	Notes
1.0	April 2019	Original publication.
1.1	August 2019	Added VCAI information.
1.2	September 2020	Updated with VMware Horizon 2006 information.
1.3	September 2021	Refreshed content.
1.4	September 2022	Updated Desktops section.

3. Best Practice Checklist

We can categorize the best practices for deploying VMware Horizon on Nutanix into the following areas.

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 upfront to design and build a solution that meets both current and future needs.
- Design for a user experience that delivers consistent performance, reliability, and scale.
- Start with a proof of concept (PoC), then test, optimize, iterate, and scale.

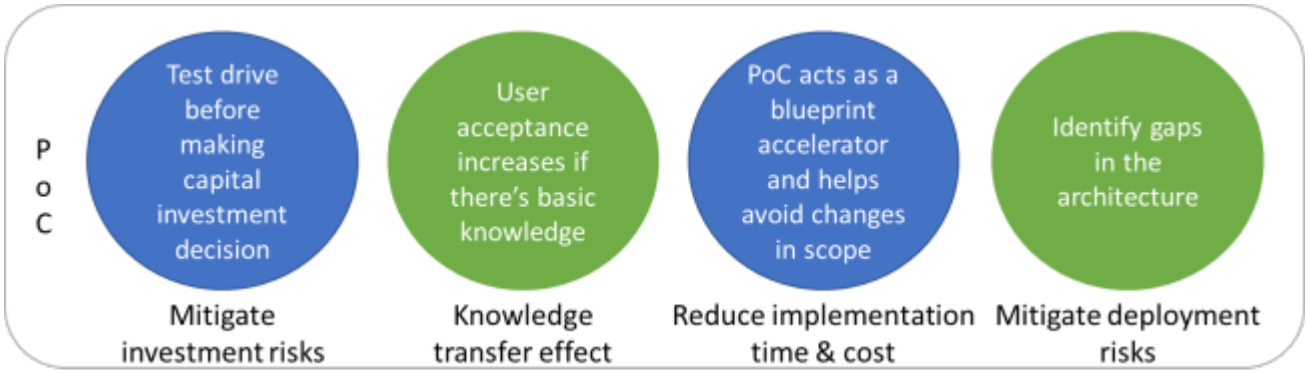

Figure 1: Benefits of Running a PoC

Core Components

Desktops:

- Size desktops appropriately for each use case.
- Use a mix of virtualized applications and applications installed in gold images, depending on the scenario.
- Disable unnecessary OS services and applications.

VMware Horizon:

- Design for reliability and scale.
- Redirect home directory or use a profile management tool (VMware Dynamic Environment Manager (DEM), for example) for user profiles and documents.
- Use the Blast protocol and apply the relevant Blast policies through VMware DEM or Group Policies. Optimize and customize the Blast settings 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 Remote Desktop Session Host (RDSH) based VMs.
- Use a separate container for App Volumes AppStacks.
- Increase Nutanix Controller Virtual Machine (CVM) memory per the sizing tables in the online documentation.

Supporting Components

Active Directory:

- Have local GCs (global catalogs) and DNS (Domain Name System) servers at each site.

- Redirect home directories for users.
- Configure DNS scavenging.

DHCP:

- For the virtual desktop infrastructure (VDI) scope, reduce default DHCP lease times from eight days to one hour.

Files:

- Use Nutanix Files for user data and user profile settings.
- Map home directory redirection to the Nutanix Files namespace.

Virus scan:

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

Network:

- Use and optimize network quality of service (QoS) for Blast traffic.
- Use at least 1 GbE access ports for 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.

4. Desktops

Sizing

Compute

For desktop-based operating systems, Nutanix typically recommends two or more 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.

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

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

For server-based operating systems, Nutanix typically recommends 8 vCPU for Microsoft Server 2019 and Microsoft Server 2022.

> Note: Assign 8 vCPU per 2019 or 2022 VM to optimize CPU-to-user ratios.

Sizing physical CPU cores differs for VDI and RDSH 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 desktops supported per physical core for light, medium, and heavy workloads.

Table: CPU Overcommit Ratios

Deployment	pCPU	vCPU
Horizon RDSH	1	1–2
Horizon desktop: Light users	1	11–12
Horizon desktop: Medium users	1	7–10
Horizon desktop: Heavy users	1	4–6

Note: Processor architecture and speed have a direct impact on the number of users you can support.

These scalability estimates take the performance benefits of hyperthreading into account. Hyperthreading can improve user density per VM (RDSH) or VM density per host (RDSH and VDI) by 20 to 30 percent.

Memory

In general, don't overcommit memory for desktop virtualization, as doing so can impact performance and thus the overall user experience.

Note: Don't overcommit memory when you virtualize desktops.

Sizing memory for RDSH-based VMs depends on the memory requirements for the applications used on those VMs. As a guideline, we advise assigning at least 4 GB of RAM per vCPU. In a Windows 2019 or 2022 deployment, this guideline results in 8 x 4 GB = 32 GB of RAM per VM. In this case, we advise assigning 32 GB of RAM or more per VM (depending on the applications that are used).

Table: Memory Assignment

Deployment	vRAM
RDSH based on Windows 2019 and 2022	32–42 GB
Windows 10 VDI: Light user	3–4 GB
Windows 10 VDI: Medium user	4–6 GB
Windows 10 VDI: Heavy user	8 GB and up

VM Configuration

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

Table: VM Parameters

	VMware Best Practice
SCSI controller	LSI Logic SAS controller / PVSCSI
Hard disk	Thin provisioning
Video card	Automatically detect

	VMware Best Practice
Floppy	Remove
NIC	VMXNET3
BIOS disable ports	Disable LPT and COM ports
Disable HotAdd and HotPlug	Disable HotPlug

Note: Follow best practices for VM configuration.

Note: Disk controllers can affect performance for users.

OS Optimization

Configure your Windows image to the specifications outlined in the Windows 10 Optimization for Horizon document. 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.
  - Linked 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.
  - Linked clones: Update the base image during off hours and recompose in phases.

Nutanix recommends that you use a script or tool to optimize the Windows OS for a virtualized environment. You can also use the VMware OS Optimization Tool or the Best Practice Analyzer tools to prepare your gold image.

Note: Use the optimization guides or tools to optimize 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:

- VMware ThinApp
- Microsoft App-V

Layering Solutions

Layering solutions are 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:

- VMware App Volumes
- Liquidware FlexApp
- Citrix App Layering

> Note: 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 are a way to 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 also includes folders such as Favorites and Documents. 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. Using either Active Directory User Object configuration or Group Policies, you can alter the roaming user profile path to copy changes in the profile to the network at logoff. 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.

> Note: Using roaming profiles can result in poor logon and logoff times for users.

Mandatory Profiles

Often used in RDSH 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 profile and all changes made to it are deleted.

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 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:

- Microsoft FSLogix Profile Container
- Liquidware ProfileDisk
- VMware App Volumes (writeable volume)

Microsoft acquired FSLogix in 2018, and the following licensing includes the solution:

- Microsoft 365 E/A3+
- Windows 10 E/A3+
- Win10 VDA
- RDS CAL

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:

- VMware Dynamic Environment Manager
- Ivanti Workspace Control
- Ivanti User Workspace Manager
- Liquidware ProfileUnity Pro
- Citrix WEM
- Scense Workspace Management
- Microsoft GPO, GP preferences, and user state virtualization

Note: Select a UEM solution or a profile management solution to manage your profiles and user settings. Managing Microsoft native profiles works as a solution but isn't cost efficient.

5. VMware Horizon 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 for designing either a small or massive VMware Horizon infrastructure over multiple sites with Nutanix.

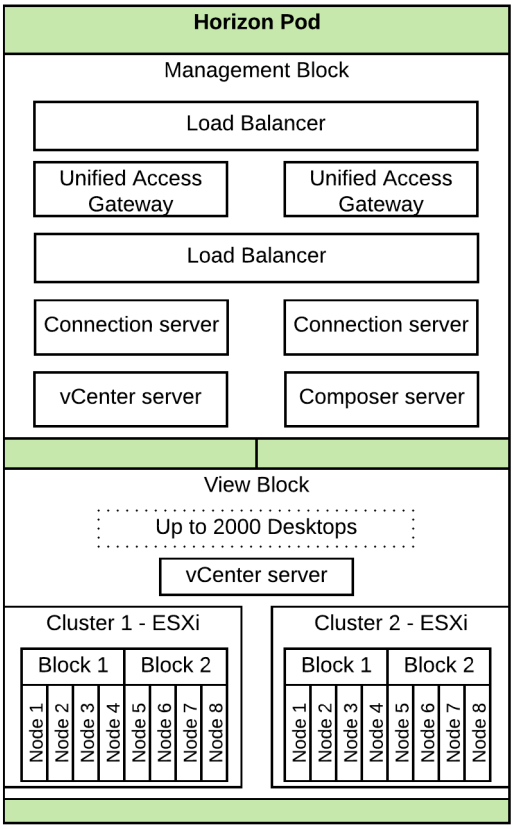

Figure 2: VMware Horizon Pod Overview

A management block contains all the components for the VMware Horizon services, and a view block contains one or more clusters hosting the virtual desktops. A single management block can manage multiple view blocks. Although a single vCenter Server can manage up to 8,000 instant-clone desktops, we advise you to limit the maximum number of desktops per vCenter to 2,000. For more information on the VMware Horizon on Nutanix design, refer to the VMware Horizon on Nutanix reference architecture.

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

Management Layer

Horizon Connection Servers

Scale Up or Scale Out

To extend your VDI or application virtualization environment, you could choose to scale up (provide services to more users per pod) or scale out (add more pods). The choice between scaling up and scaling out may be subject to the following constraints:

- Failure domains: Creating multiple pods creates more and smaller failure domains, so outages impact only a subset of the total users.
- Geographical location: In a Cloud Pod Architecture (CPA) environment, a site is a collection of well-connected pods in the same physical location, typically in a single datacenter. A pod can't span multiple physical locations. With multiple datacenters you need at least one pod per datacenter and preferably one site per datacenter. Depending on the Horizon version you're using, you can create up to 15 sites and 50 pods in a CPA, for a maximum of 250,000 concurrent sessions.

Redundancy

Nutanix recommends that you have at least two VMware Horizon Connection Servers per Horizon CPA for redundancy and high availability. If you need your infrastructure to be redundant during maintenance or a failure, you should configure at least three VMware Horizon Connection Servers per CPA.

Scalability

To ensure that the VMware Horizon Connection Servers aren't the bottleneck in your infrastructure, assign sufficient resources to each Connection Server. Scale these resources for peak endurance moments like periods of boot (or logon) storms.

VMware's minimum requirement for a Horizon Connection Server is 4 vCPU and 10 GB of RAM. Although a Connection Server can handle 4,000 connections, we recommend a maximum of 2,000 connections per Connection Server.

For further information, refer to the VMware Horizon limits and recommendations KB article.

SQL Database

SQL Database Redundancy

The requirement for a SQL database in a VMware Horizon infrastructure depends on the components used. You can use a SQL database for:

Event logging

VMware Horizon continues to work if the SQL database is unavailable for event logging. Alternatively, you can log events to a Syslog server.

Horizon Composer

The Horizon Composer service stores the configuration of the linked-clone desktops in a SQL database. If the SQL database is unavailable, the Composer service fails. Users can continue to work and log on to available desktops but can't create new desktops or refresh existing desktops after they log off. Therefore, we recommend that you use a highly available SQL database if you use linked-clone desktops.

VMware App Volumes

App Volumes stores its configuration in a SQL database. If the SQL database is unavailable, existing AppStack usage isn't interrupted, but new user logons can't use AppStack. We recommend that you use a highly available SQL database for VMware App Volumes.

> Note: Ensure that your database is highly available via mirroring, clustering, or Always On. 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.

> Note: Back up your VMware Horizon databases at regular intervals to minimize the impact of a disaster.

SQL Database Scalability

To ensure that you've assigned sufficient resources to the Microsoft SQL Server environment, Nutanix recommends that you use the VMware App Volumes database best practices document.

> Note: Size your SQL Server based on the database sizing document VMware provides.

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

License Server

VMware Horizon Licensing

VMware Horizon licenses are applied on the Connection server and the App Volumes server. The licenses are stored in the configuration of these services. No separate license server is required.

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 here:

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

> Note: Deploy two Microsoft RDS license servers for redundancy.

6. Nutanix Storage Optimization Technologies

Capacity Optimization

Nutanix incorporates a wide range of storage optimization technologies that work in concert to efficiently use the capacity available for any workload. These technologies are intelligent and adaptive to workload characteristics, eliminating the need for manual configuration and fine-tuning.

Nutanix offers the following optimizations, which we discuss in the following sections:

- Compression
- Deduplication
- Erasure coding

Compression

The Nutanix Capacity Optimization Engine (COE) performs data transformations to increase data efficiency on disk. Currently, compression is one of the features the COE uses to optimize data. Nutanix storage provides both inline and post-process compression to suit the customer's specific needs and types of data.

The inline method compresses sequential streams of data or large I/O sizes in memory before writing them to disk, while post-process compression initially writes the data as usual (in an uncompressed state), then uses the Curator MapReduce framework to compress the data cluster-wide. When you've enabled inline compression but the I/O is random, the system writes data in the oplog uncompressed, coalesces it, then compresses it in memory before writing it to the extent store. The COE uses LZ4 and LZ4HC for data compression, providing good compression ratios with minimal computational overhead and extremely fast compression and decompression rates.

The following figure shows an example of how inline compression interacts with the distributed storage write I/O path.

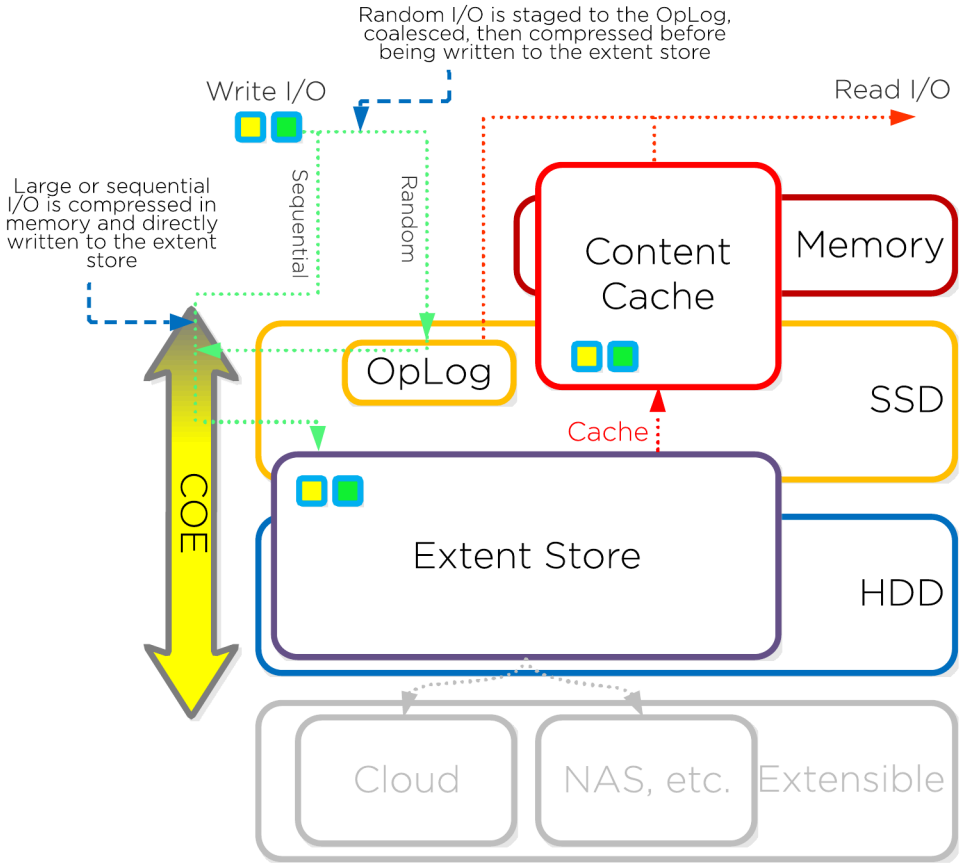

Figure 3: Inline Compression I/O Path

Post-process compression writes all new write I/O in an uncompressed state and follows the normal distributed storage I/O path. After the data has met the configurable compression delay and become cold, it migrates down to the HDD tier through Information Life Cycle Management (ILM) and becomes eligible for compression. All nodes perform compression tasks using the Curator MapReduce framework.

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

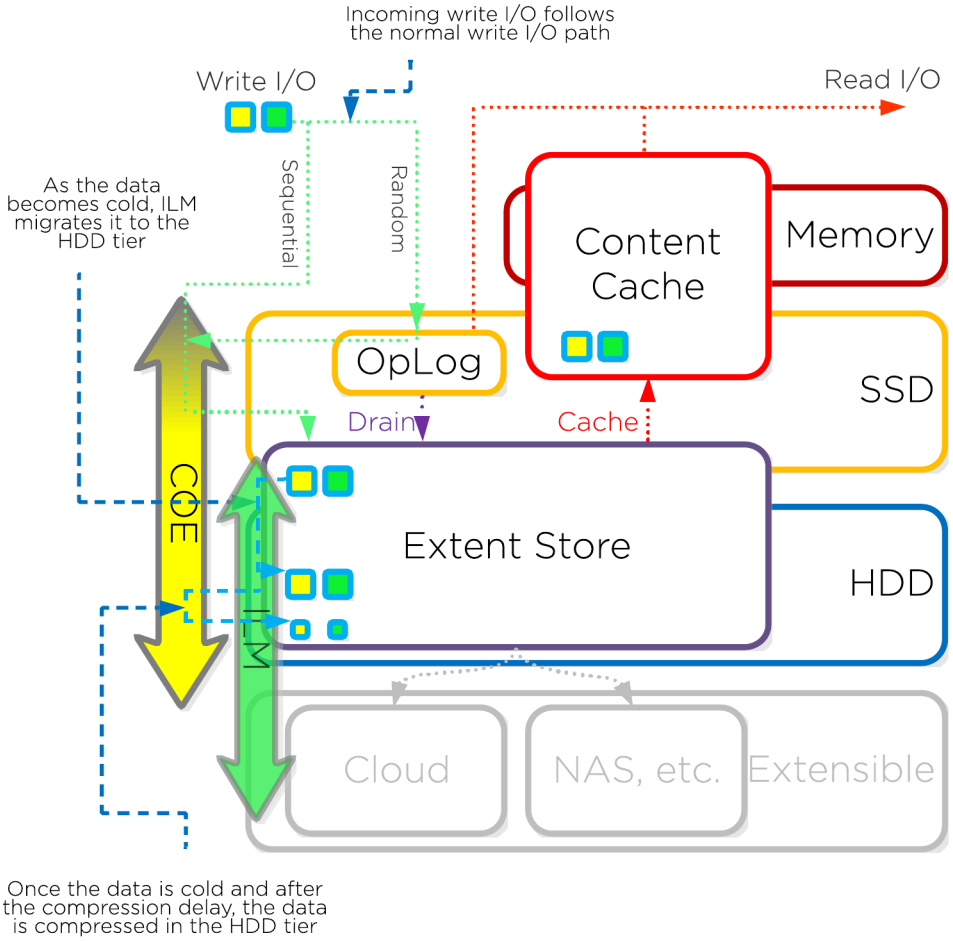

Figure 4: Post-Process Compression I/O Path

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 tier and then uses ILM to move up to the SSD tier and the cache.

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

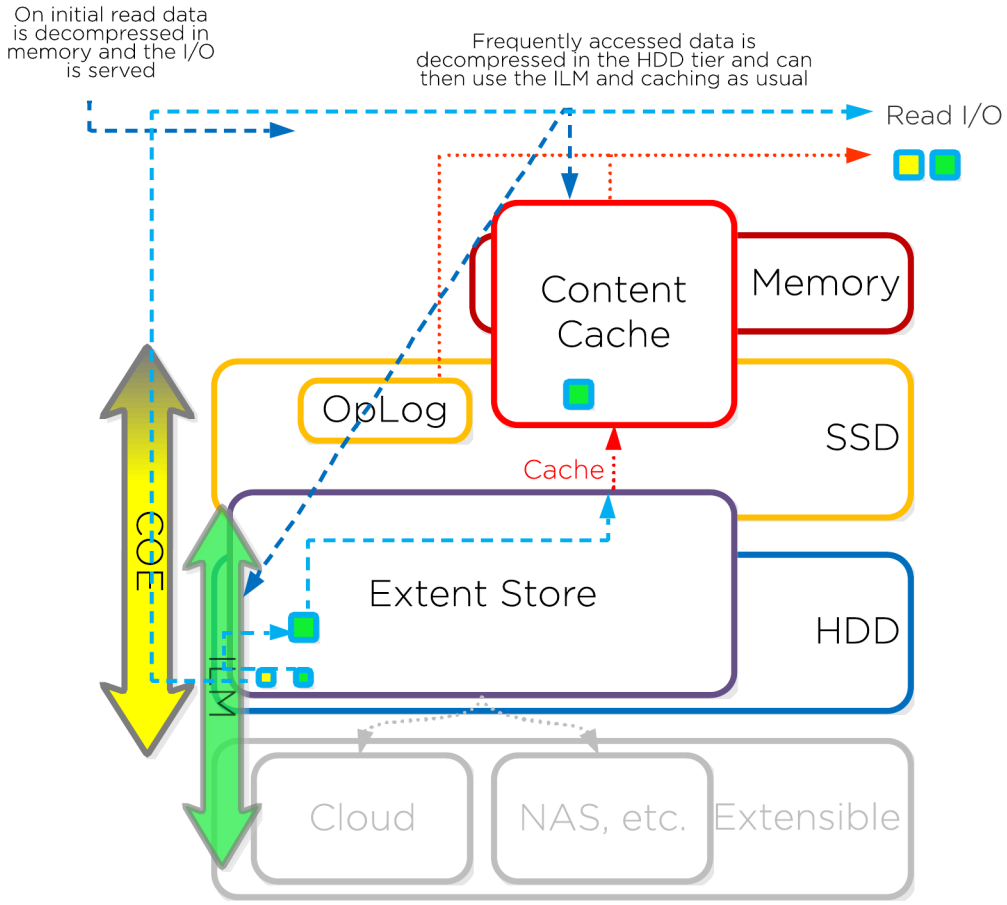

Figure 5: Decompression I/O Path

Elastic Deduplication Engine

The Elastic Deduplication Engine is a software-based feature of the Nutanix storage method that deduplicates data in the capacity (HDD) and performance (SSD and memory) tiers. The system fingerprints streams of data during ingest using a SHA-1 hash at an 8 KB granularity. Unlike traditional approaches that use background scans requiring data to be read again, this fingerprint occurs only 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, minimizing metadata overhead. Capacity tier deduplication prefers full extents to minimize fragmentation.

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

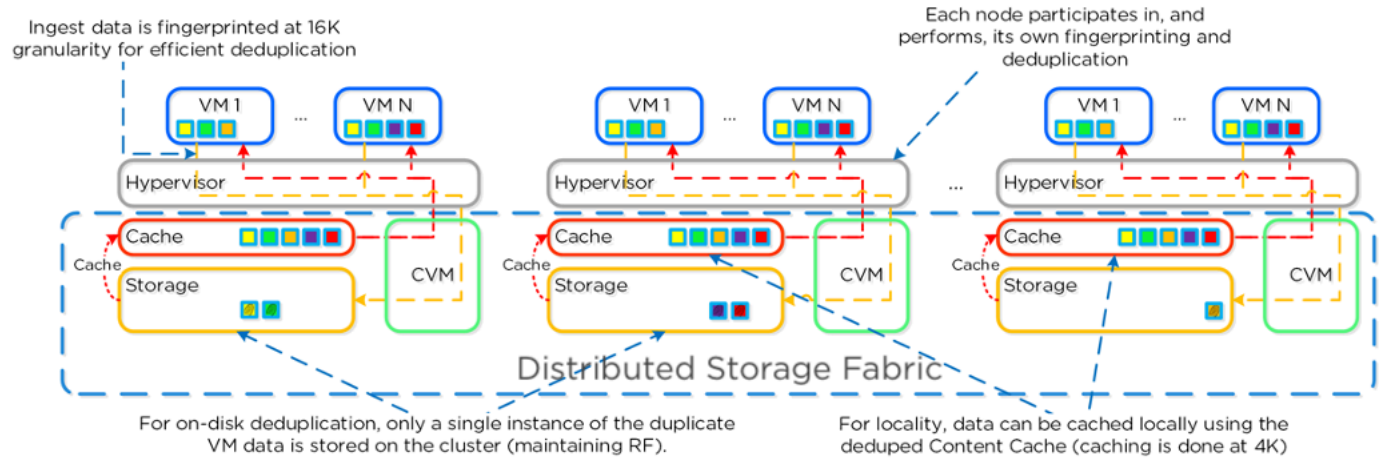

Figure 6: Elastic Deduplication Engine: Scale

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 take place as a background process.

The Elastic Deduplication Engine spans both the capacity disk tier (HDD) and the performance tier (SSD and memory). As the engine identifies duplicate data (based on finding multiple copies of the same fingerprints), a background process removes the duplicate data using the Curator MapReduce framework. Read data is pulled into the unified cache, which is a multitier and pool cache, at a 4 KB granularity. Any subsequent requests for data with the same fingerprint pull directly from the cache.

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

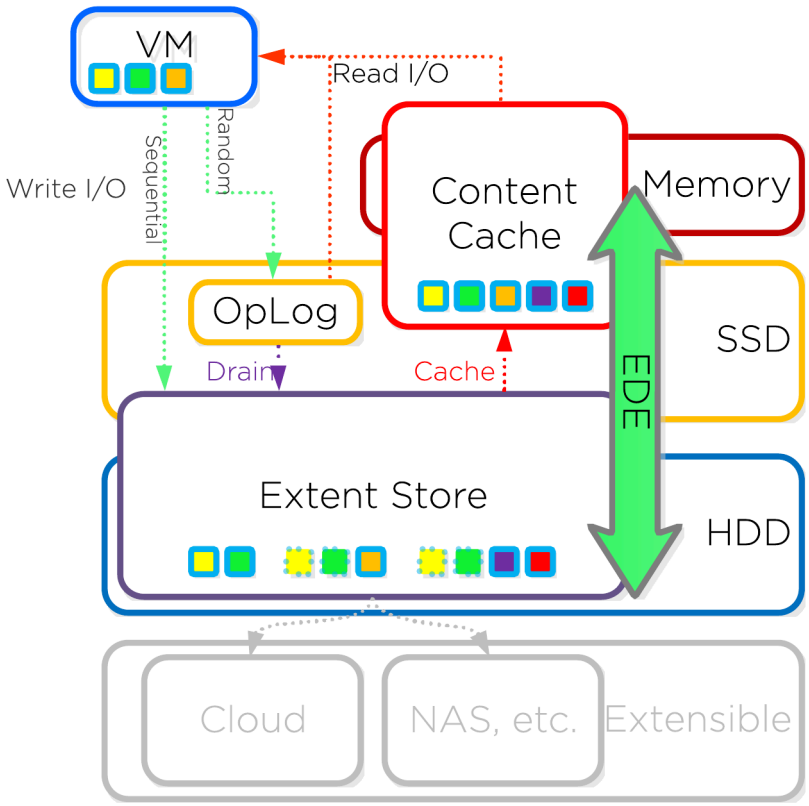

Figure 7: Elastic Deduplication Engine I/O Path

Erasure Coding

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 does occupy additional storage resources. Replication factor 2 stores one additional copy of data; replication factor 3 stores two additional copies. 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 the Nutanix storage method, the data block is an extent group, and each data block must be on a different node and belong to a different vDisk.

You can configure the number of data and parity blocks in a strip based on how many failures you need to tolerate. In most cases, we can think of the configuration as the <number of data blocks> / <number of parity blocks>.

EC-X is a post-process framework that doesn't affect the traditional write I/O path. The encoding uses the Curator MapReduce framework for task distribution.

The following figure depicts a normal environment using replication factors.

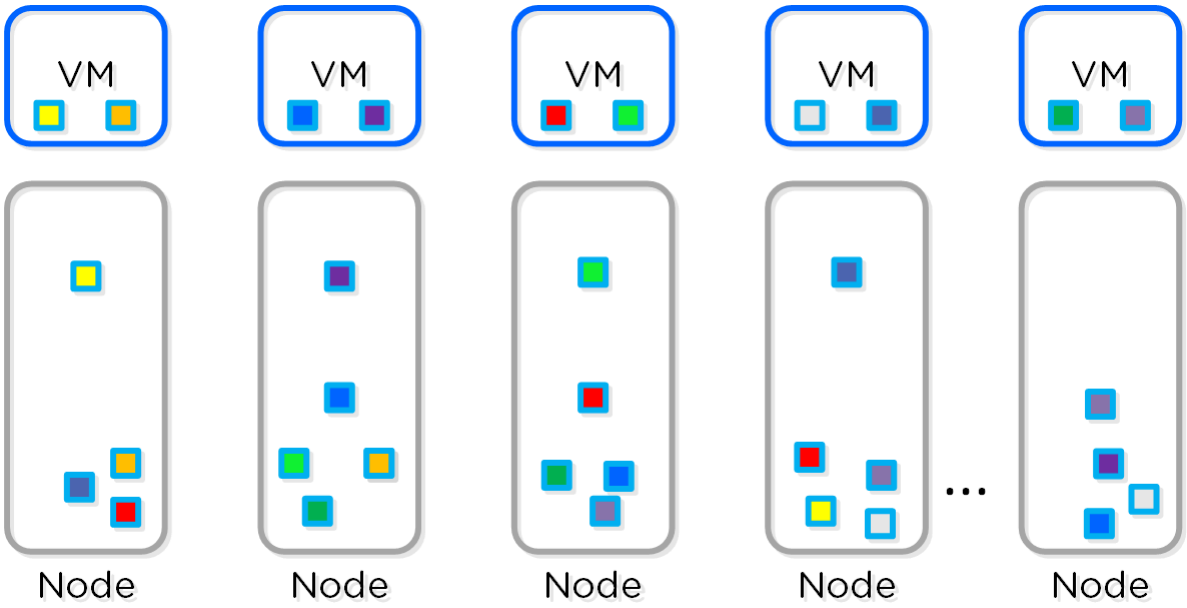

Figure 8: Typical Replication Factor Data Layout

In this scenario, we have a mix of both replication factor 2 and replication factor 3 data with primary copies stored locally and replicas distributed to other nodes throughout the cluster. When a Curator full scan runs, it finds eligible extent groups available for EC-X. Eligible extent groups must be write-cold, meaning nothing has written to them for at least one hour.

The following figure shows an example 4:1 and 3:2 strip.

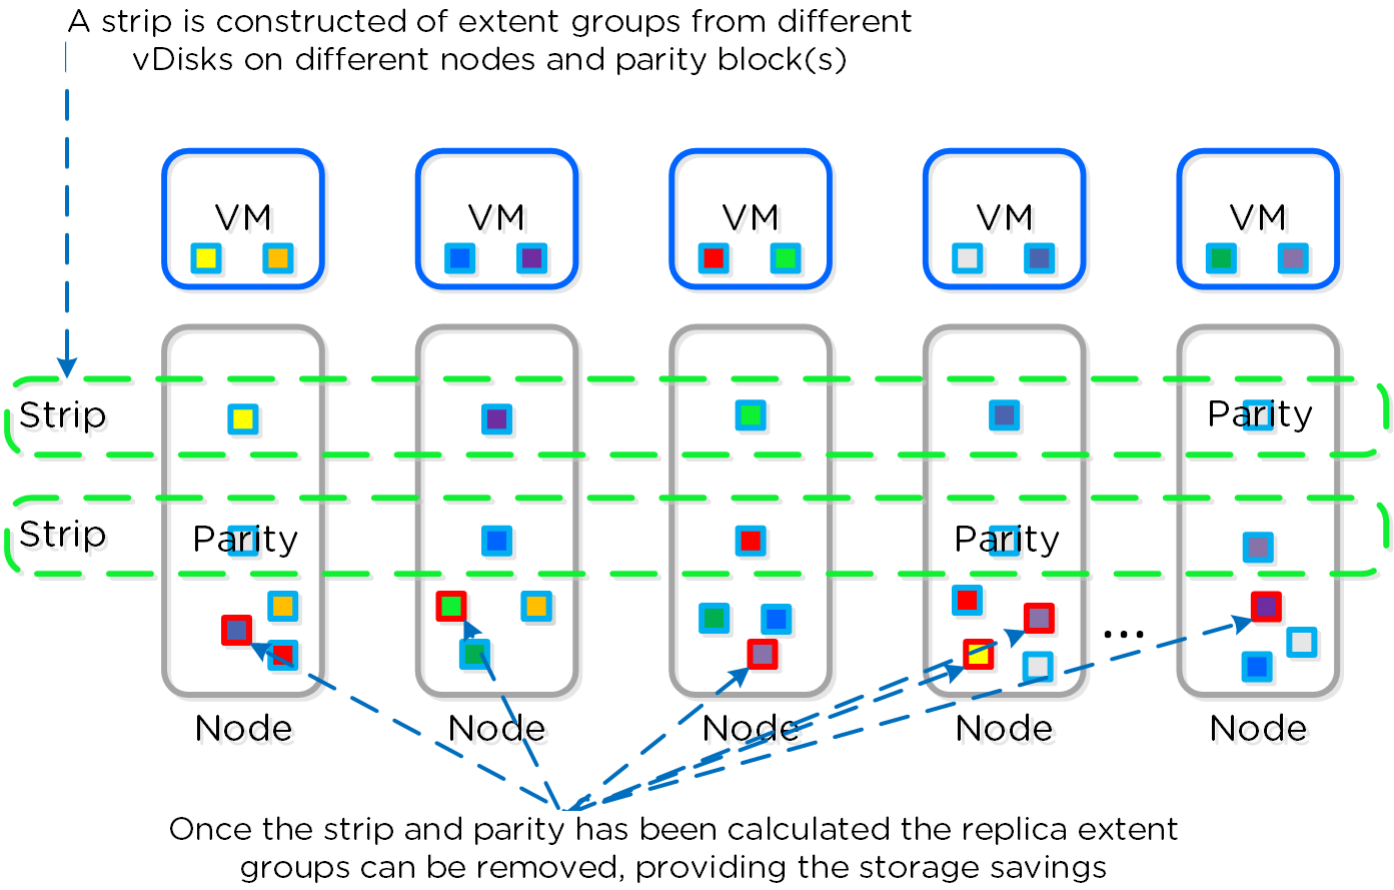

Figure 9: Encoded Strip Before Savings

Once the strips and parity have been successfully calculated, the system removes the replica extent groups. The following figure shows the storage savings in the environment after the system runs EC-X.

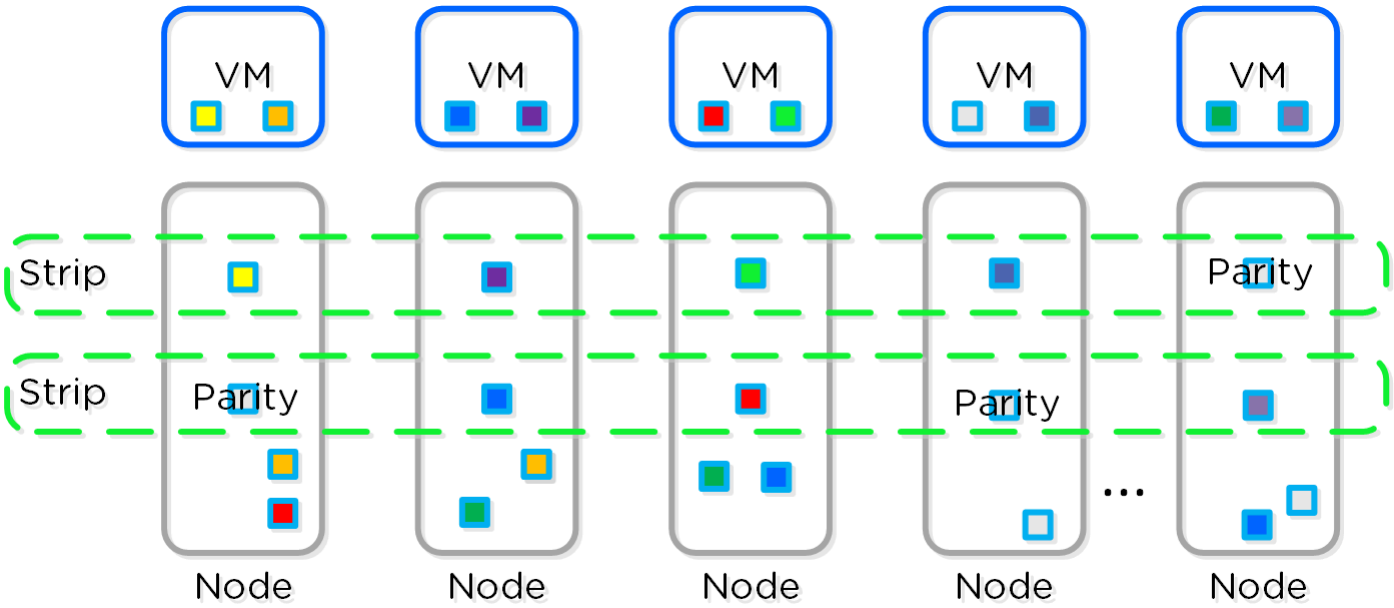

Figure 10: Encoded Strip After Savings

7. VMware Horizon Delivery Options

VMware Horizon has three methods to create virtual desktops:

1. Instant clones
2. Linked clones
3. Full clones

Starting with VMware Horizon version 8 (2006), the linked clones feature is deprecated. The instant clones feature is the preferred method for creating virtual desktops.

Horizon View Instant Clones

Instant-clone provisioning uses built-in technology from VMware vSphere to create virtual desktops on demand. VMware Horizon prepares one replica disk per datastore and one parent VM per host founded on a base VM or gold image. Once you finish the preparation, desktop creation only takes a couple seconds per desktop. You don't need a separate service or database (like Horizon Composer) to use Horizon View instant clones.

VMware Horizon 2006 introduces a new feature called Smart Provisioning. With this feature, Horizon automatically chooses the type of desktop to create, based on the density of VMs per host in the selected cluster. Low-density types are created without parent VMs; high-density types are created with parent VMs.

Horizon View Linked Clones

The Horizon View linked clones feature uses a standardized model for hosted virtual desktop creation. You need a Windows service called Horizon Composer to create linked clones. Starting from a base VM or gold image, Horizon Composer creates a replica that is a locked version of the gold image, so you can make updates to the gold image without affecting production VMs. From there, Horizon Composer creates VM clones that consist of a delta disk and an

identity disk, both linking back to the base VM's disks. These VM clones are also called linked clones.

View Composer Array Integration (VCAI)

When you create a linked-clone desktop pool with Horizon Composer on Nutanix, you can enable the View Composer Array Integration (VCAI) feature. This feature, which uses our native NFS snapshot technology to clone VMs, is the preferred cloning option on Nutanix.

> Note: When you create linked clones on Nutanix with View Composer, enable VCAI for fast provisioning of virtual desktops.

Migrating from Non-VCAI to VCAI

When you need to migrate an existing VMware Horizon environment from a non-Nutanix platform to a Nutanix platform, there are multiple deployment options available. One option is to create new desktop pools, so you can use the best settings based on your use case and the storage capabilities of the Nutanix platform. If you don't want to create new desktop pools, another option is to update the existing desktop pools and select the new Nutanix infrastructure as the target for the virtual desktops. If the existing pool is a View Composer desktop pool with the Reclaim Disk Space setting enabled, you may be unable to select VCAI on it. Use the following procedure to enable VCAI on those desktop pools during migration.

Procedure

Clone the parent VM for each pool from the old platform to Nutanix. During the cloning process, ensure that disk format is set to thin provision.

Create a new snapshot on the new parent VM.

Update the desktop pool in the Horizon administrator console using the following order of operations:

- Select the pool, then Status, then Disable provisioning, and select OK.
- Select the pool, then Edit, then General, and make sure that Select separate datastores for replica and OS disks is deselected.

- In the same window, select vCenter Settings, then select the new parent VM and the new snapshot as well as the Nutanix cluster and datastore.
- Navigate to Advanced Storage Settings, deselect Use View Storage Accelerator, and click OK to save.
- Use the ADSI Edit utility to disable the creation of space-efficient disks in View LDAP:
  - Start the ADSI Edit utility on your View Connection Server system.
  - In the console tree, select Connect to.
  - In the Select or type a Distinguished Name or Naming Context text box, type the distinguished name DC=vdi, DC=vmware, DC=int.
  - In the Select or type a domain or server text box, select or type localhost:389 or the fully qualified domain name (FQDN) of the View Connection Server computer followed by port 389.
  - Expand the ADSI Edit tree, expand DC=vdi, DC=vmware, DC=int, and select OU=Server Groups.
  - In the right pane, select the CN of the desktop pool.
  - In the CN=<pool name> dialog, set the pae-UseSeSparseFormat attribute to 0 and click OK.
  - Restart the VMware Horizon View Connection Server service.
- Return to the Horizon administrator console and select the pool, then Edit, then Advanced Storage Settings, then Other Options. Select Use native NFS snapshots (VAAI) and click OK to save.
- Select the pool, then Status, and select Enable provisioning.
- Select the pool, then View Composer, and select Rebalance.

When this process is complete, the system deletes unused desktops from the pool and recreates them on the Nutanix cluster you selected. Desktops (whether connected or disconnected) aren't redeployed until after the user logs off from the desktop.

View Storage Accelerator

With VMware Horizon you can configure a feature called View Storage Accelerator (VSA). This feature uses the Content-Based Read Cache (CBRC) feature in ESXi hosts to cache VM disk data. VSA improves the performance during I/O storms, which can take place when many VMs start up or run antivirus scans at once. However, as the Nutanix storage method automatically caches this data, you don't need to enable VSA. For this reason, we advise you to disable VSA. A warning message appears in the management interface when you create an instant-clone desktop pool while VSA is disabled. You can ignore this warning and continue to create the instant-clone desktop pool.

Note: When you use VMware Horizon on Nutanix, disable VSA, as it doesn't provide a performance benefit on Nutanix.

8. Nutanix Storage Configuration

Nutanix Storage Best Practices

Table: Storage Best Practices for VMware Horizon

Delivery Method	Compression	Elastic Deduplication Engine	Erasure Coding
Full clones	Yes	Yes	
Linked clones	Yes		
Instant clones	Yes		
Application AppStacks	Yes		
Writeable AppStacks	Yes	Yes	

Note: Enabling compression for VMware Horizon is a general best practice; only enable the Elastic Deduplication Engine for full clones and writeable AppStacks. Erasure coding isn't a suitable data-reduction technology for desktop virtualization.

Instant-clone desktops and linked-clone desktops are both based on linked-clone technology. The setup process starts by creating a gold image and a snapshot. The user then creates a desktop pool and selects a snapshot. VMware Horizon creates a full copy of the snapshot (replica) as well as a differencing disk and identity disk for every VM.

On traditional systems, this structure produces I/O inefficiencies, leading to higher network latency and decreased user experience. With linked-clone technology, all VM reads go back to the full copy of the snapshot. Thus, in an environment with eight hosts running 100 VMs each, there are 800 VMs returning to a single VMDK for reads. All writes stay local, as they are reverted to the differencing disk.

The Nutanix storage method resolves these inefficiencies with shadow clones. Shadow clones enable distributed caching of vDisks or VM data in a multireader scenario, which is what occurs with a gold image—all reads come from the gold image and all writes go to the differencing disk.

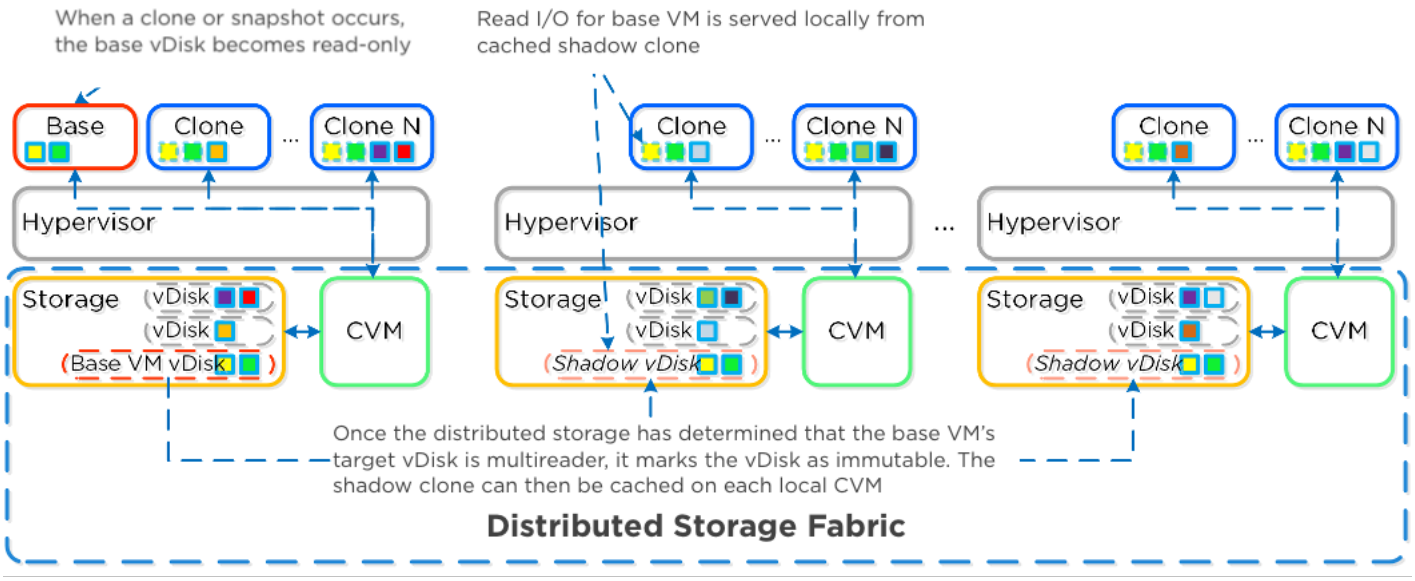

Figure 11: Nutanix Shadow Clones

With shadow clones, the distributed storage monitors vDisk access trends much like it monitors data locality. When requests occur from more than two remote CVMs (as well as the local CVM) and all the requests are read I/O, the distributed storage marks the vDisk as immutable. Then, each CVM can cache the vDisk locally, creating shadow clones of the base vDisk. These shadow clones allow VMs throughout the cluster to read the base VM's vDisk locally. With VDI, each node caches the replica disk and serves all base read requests locally. Data is migrated only on a read to avoid flooding the network and to allow for efficient cache utilization. When anything modifies the base VM, the shadow clones are dropped and the process starts over.

Note: Using shadow clones with VMware Horizon speeds up logons, improves application performance, and prevents boot storms.

Note: In general, use one container per use case. However, if you're provisioning both full clones and linked clones, Nutanix recommends that you use two containers.

9. Conclusion

Our extensive testing of VMware Horizon instant clones and linked clones on Nutanix demonstrates that desktop user density is based primarily on the available host CPU resources, not on I/O or resource constraints. Using our CVM, our 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 VMware Horizon 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.

10. Appendix

References

1. VMware Horizon on Nutanix reference architecture
2. Windows 10 Optimization for Horizon document
3. VMware Horizon limits and recommendations KB article
4. VMware App Volumes database best practices document
5. Microsoft SQL Server on Nutanix best practices guide

About the Authors

Kees Baggerman is a technical director for Solutions Engineering at Nutanix. Follow Kees on Twitter at @kbaggerman.

Sven Huisman is a staff solutions architect for End-User Computing at Nutanix. Follow Sven on Twitter @SvenH.

About Nutanix

Nutanix is a global leader in cloud software and a pioneer in hyperconverged infrastructure solutions, making clouds invisible and freeing customers to focus on their business outcomes. Organizations around the world use Nutanix software to leverage a single platform to manage any app at any location for their hybrid multicloud environments. Learn more at www.nutanix.com or follow us on Twitter @nutanix.

List of Figures