v1.0 | August 2022 | NVD-2158

NUTANIX VALIDATED DESIGN

Nutanix Citrix DaaS Validated Design

NUTANIX
YOUR ENTERPRISE CLOUD

Copyright

Contents

1. Executive Summary

Nutanix continues to innovate and engineer solutions that are simple to deploy and operate. To further improve customer experience and add value for customers, Nutanix uses robust validation to simplify the process of architecting and deploying solutions. This document details the design decisions that support the deployment of a scalable, resilient, and secure private cloud solution for Citrix Desktop as a Service (DaaS) using Machine Creation Services (MCS) or Provisioning (PVS) in a single datacenter.

Nutanix can deliver this Nutanix Validated Design (NVD), based on our Citrix DaaS Reference Architecture and Citrix DaaS Best Practices Guides, as a bundled solution for end user computing that includes hardware, software, and services to accelerate and simplify the deployment and implementation process.

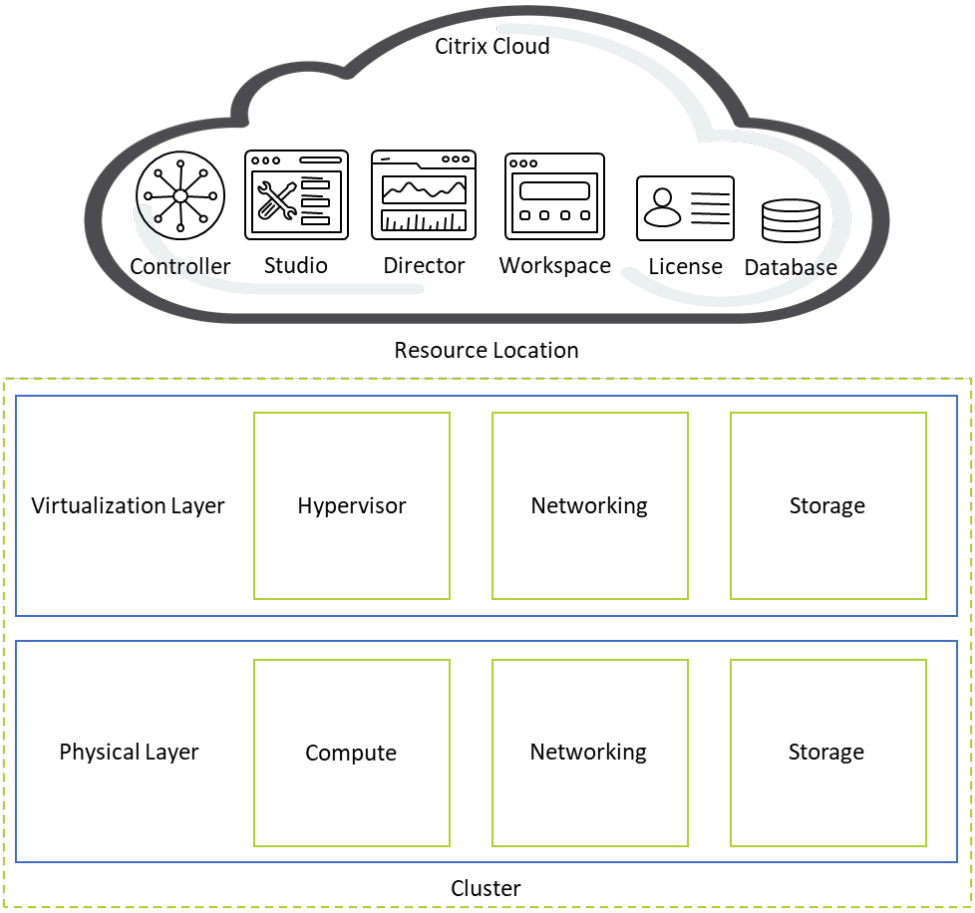

Figure 1: Architectural Layers of the Nutanix Validated Design

This scalable modular design, based on the Nutanix block-and-pod architecture, is well suited for Citrix DaaS use cases of all sizes. Some highlights of the NVD include:

- Solution for Citrix DaaS deployments that integrate multiple products including AOS, AHV, and Files.
- Single-datacenter design built for failure tolerance and 99.999 percent availability.
- Single datacenter with single resource location run under 85 percent capacity to allow for in-datacenter reliability.
- Accelerated customer time-to-value and reduced risk.

- Orderable as a solution with a fully elaborated BOM for hardware, software, and services.
- This NVD focuses on Nutanix NX nodes, but you can use any Nutanix-supported hardware vendor with the same or better hardware specifications.

This validated design is just one example of a supported Citrix DaaS configuration. There are many ways to design and build a Citrix DaaS solution on Nutanix, and you can deviate from this specific configuration while still following Nutanix best practices.

As a member of the Citrix Ready program, Nutanix has completed a rigorous verification process established by Citrix to ensure compatibility with its products and services. By successfully completing this series of tests, Nutanix has proven it can run Citrix DaaS seamlessly in a hybrid cloud environment.

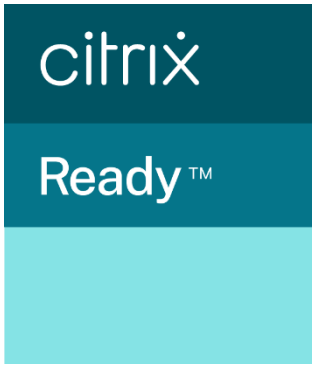

Figure 2: Citrix Ready Verified

You can have this validated solution up and running in weeks with minimal burden on your internal teams, allowing you to realize the full value of your infrastructure quickly. After you place your order, Nutanix takes care of the rest.

Audience

This guide is part of the Nutanix Solutions Library, intended for architects and engineers responsible for scoping, designing, installing, and testing server virtualization solutions. Readers of this document should already be familiar with the Nutanix Hybrid Cloud Reference Architecture.

Purpose

This document describes the components, integration, and configuration for the NVD packaged hybrid cloud solution and covers the following topics:

- Core Nutanix infrastructure and related technology.
- Citrix DaaS on the Nutanix platform for multi-session or single-session nonpersistent workloads using MCS or PVS.
- Test plan.
- Bill of materials.

Table 1: Document Version History

Version Number	Published	Notes
1.0	August 2022	Original publication.

2. Core Infrastructure Design

The following tables provide core infrastructure design requirements, assumptions, risks, and constraints.

Table 2: Core Infrastructure Design Requirements

Component	Description
Citrix Cloud Connectors	Deploy Citrix Cloud Connectors on the dedicated management cluster.
Citrix Cloud Connectors	Deploy Citrix Cloud Connectors in a highly available configuration.
Microsoft SQL Services	Deploy Microsoft SQL Services for Citrix Provisioning on the dedicated management cluster.
Microsoft SQL Services	Deploy Microsoft SQL Services for Citrix Provisioning in a highly available configuration.
Microsoft Licensing	Deploy Microsoft RDS Licensing for multi-session workloads on the dedicated management cluster.
Citrix Licensing	Deploy Citrix Licensing for Provisioning on the dedicated management cluster.
Nutanix Files	Deploy Nutanix Files on the workload cluster.
Nutanix Files	Deploy Nutanix Files in a highly availability configuration.
Citrix Provisioning	Deploy Citrix Provisioning on the workload cluster.
Citrix Provisioning	Deploy Citrix Provisioning in a highly available configuration on the workload cluster.
Virtual Machines	Support at least three virtual machine (VM) sizes: small, medium, and large.

Component	Description
Virtual Machines	Support Windows Server 2019 and Windows 10 as VM operating systems.
Virtual Machines	For Windows Server 2019, limit virtual CPU overcommitment to 8:1, or 8 vCPU per physical CPU core. For Windows 10, limit it to 2:1.
Monitoring	Enable platform fault monitoring and use email to send alerts.
Monitoring	Monitor performance metrics and store historical data for the past 12 months.
Monitoring	Keep resource usage under 85 percent; usage over 85 percent generates an email alert.
Monitoring	Monitor resources critical to Nutanix AOS operations (for example, CPU, memory, storage, and network resources); resource usage that exceeds configured limits generates an alert.
Monitoring	For resources that have high availability reservations, measure the resource utilization threshold against the usable capacity after subtracting the capacity reserved for high availability.
Monitoring	Monitor all network links (including host-switch and switch-switch) for bandwidth utilization and store historical data for the past 12 months.
Monitoring	Use email as the primary channel for event monitoring alerts.

Table 3: Core Infrastructure Design Assumptions

Component	Description
Clusters	The maximum number of Windows 10 VMs per workload cluster is 2,475 (165 VMs per usable node). The maximum number of Windows Server 2019 VMs per workload cluster is 120 (8 VMs per usable node).

Component	Description
Clusters	When using FSLogix with Nutanix Files on the workload cluster or on the infrastructure cluster, the maximum number for Windows 10 VMs per workload cluster is 2,250 (150 VMs per usable node). The maximum number of Windows Server 2019 VMs per workload cluster is not impacted.
Monitoring	IT operations teams can continuously staff the mailbox that receives monitoring alerts to address critical issues in a timely manner.
Monitoring	IT operations teams can provide email infrastructure with sufficient resilience to send, receive, and access emails even during critical outages.
Infrastructure	IT operations teams can deploy Active Directory and DNS in a highly available configuration on the dedicated management cluster.
Infrastructure	IT operations teams can deploy Nutanix Files in a highly available configuration on the workload cluster for profile and user data.

Table 4: Core Infrastructure Design Risks

Component	Description
Infrastructure	IT operations teams can deploy Nutanix Files in a highly available configuration on the dedicated management cluster, workload cluster, or on a dedicated Nutanix Files cluster depending on size, scalability, and feature requirements. In this NVD, Nutanix Files is hosted on the workload cluster for FSLogix Profile Containers using VHD Locations.
Infrastructure	IT operations teams can deploy Citrix Provisioning in a highly available configuration on the dedicated management cluster or on the workload clusters. In this NVD, Citrix Provisioning is hosted on the workload cluster.

Table 5: Core Infrastructure Design Constraints

Component	Description
Clusters	The number of VMs per workload cluster doesn't exceed 2,475 (165 Windows 10 desktops is the max number tested per node in our reference architecture with a steady-state CPU usage maximum of 85 percent). The number of workload VMs per pod doesn't exceed 10,000 (Citrix DaaS maximum single-session VDAs per resource location).
Monitoring	SMTP is an available channel in the environment that can receive event monitoring alerts. Syslog captures logs but does not generate alerts on events.
Monitoring	SNMP polls monitor performance metrics using an existing Prometheus deployment.

Core Infrastructure Conceptual Design

The conceptual pod design has the following features:

- A single datacenter.
- A small management cluster that hosts services such as Active Directory, DNS, SQL, Cloud Connectors, and Licensing.
- An instance of Nutanix Files on the workload cluster for FSLogix Profile Containers.
- A highly available instance of Citrix Provisioning on the workload cluster.
- A workload cluster that hosts the Citrix DaaS workloads.

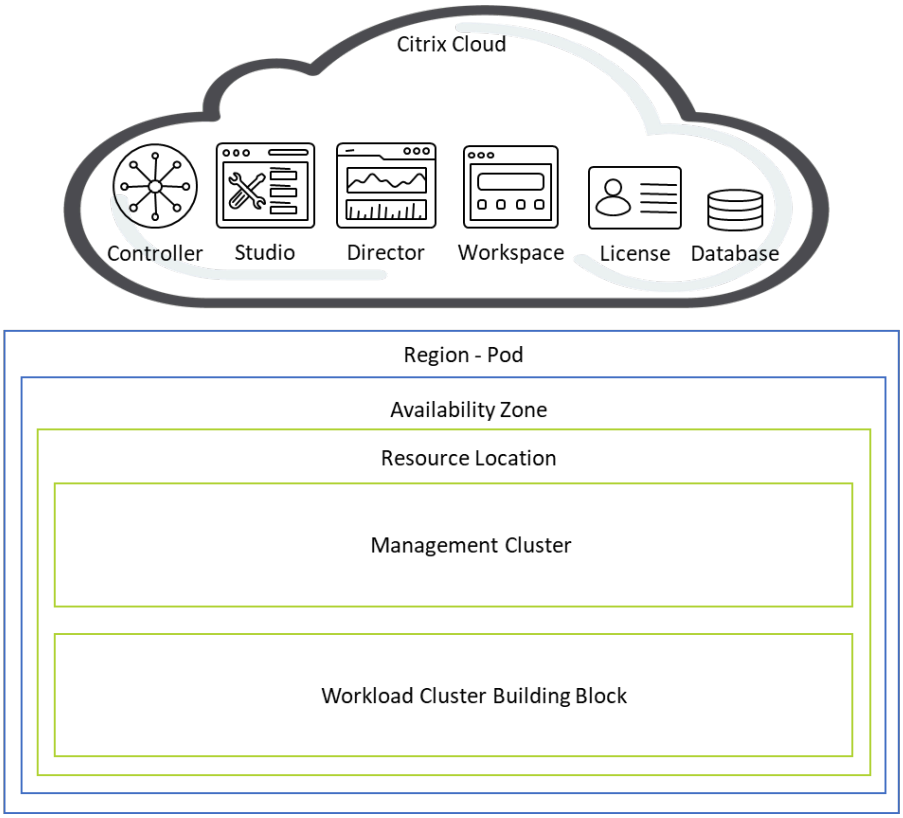

Figure 3: Conceptual Pod Design

Scalability

Scalability is one of the core concepts of the Nutanix platform and refers to the ability to increase storage and compute capacity to meet both current and future workload demands. A well-designed cluster meets current requirements while providing a path to support future growth.

Scalability Conceptual Design

This NVD allows horizontal and vertical scaling within the boundaries set by running workloads in a single rack per datacenter in a single resource location. If the workload grows, you can add nodes and storage capacity to the cluster. This design has a maximum of 16 nodes per cluster; if you need to scale beyond that number, you can create additional Nutanix clusters.

Note: If workloads change or grow, you must account for that in the workload cluster and in any additional workload clusters in the resource locations.

Because this NVD supports three general VM sizes, each node's memory is fully populated to accommodate the resulting mixed memory requirements. This approach also provides maximum memory performance, even if you don't need it. If memory pressure increases, add more nodes. The design uses all-flash disks to accommodate peak workload demands.

We selected a Nutanix model that puts four workload nodes in two rack units to maximize workload density. The design uses a single rack in the datacenter, with redundant top-of-rack network switches. This approach reduces operational complexity but constrains the number of nodes in a rack, as only a certain number of network ports are available. Datacenter power and cooling limitations may introduce further constraints; refer to the Datacenter Infrastructure section for more information.

When you scale VM workloads, cluster design is the biggest constraint.

Table 6: Scalability Design Decisions

Decision Name	Decision
Node memory population	Fully populate node memory
Node drive type	Use all-flash drives
Node drive population	Do not fully populate nodes with disk drives
Single rack	Use one rack per cluster
Establish scalability boundaries	Workload load per node is established with reference architecture testing (See the Appendix for links to Nutanix's scalability results and testing methodologies)
Rack availability use	Don't use rack availability

Nutanix recommends balanced memory configurations:

- Use identical CPUs.
- Keep the memory configuration identical across CPUs.

- Keep the memory channel configuration identical.
- Use identical DIMMs.

> Note: Improper memory configurations will impact system performance.

For supported memory configurations in Nutanix G8 platforms, see Supported Memory Configurations in the NX Series Hardware Administration Guide and Nutanix Physical Memory Configuration on the Nutanix Portal.

Configuration maximums also constrain solution scalability. For the latest limits, refer to the configuration maximums or the maximum system values on the Nutanix Support Portal (portal account required). Note that you can reach a constraint before you reach a configuration maximum. For example, a workload node that contains only Windows Server 2019 multisession VMs could theoretically not hold more than eight VMs, assuming we stay within 80 percent of the host CPU usage.

Table 7: Configuration Maximums or Maximum System Values

Entity	Decision
VMs or volume groups	Async DR: 200 VMs or volume groups for each protection domain or consistency group.
vDisks (including snapshots)	600,000

Resilience

Nutanix provides many resilience features, including storage replication, snapshots, block awareness, degraded node detection, and self-healing. These capabilities increase the resilience of all workloads, even if the application itself has limited resilience options. Nutanix layers these software features on hardware designed with resilience in mind (for example, with redundant physical components and power supplies, many of which are hot-swappable or otherwise easily serviceable). Running workloads in a virtualized environment adds another kind of resilience, as you can perform many maintenance operations without application downtime. A resilient network fabric that

can sustain individual link, node, or block failures without significant impact completes the architecture.

Resilience Conceptual Design

All components are physically redundant. The physical components include the top-of-rack switches, the nodes and their internal parts, and the datacenter itself in case of a disaster.

To protect workloads to meet or exceed SLAs, this NVD separates the workload clusters from the management clusters. The workload cluster sizing allows for n + 1 failure redundancy. Monitoring and alerting ensure that any issues result in an alert; consistently monitoring workload growth ensures that sufficient headroom is available at any time.

There is no ideal cluster size for a generic workload. This NVD uses 16-node building blocks to take advantage of block awareness, a key platform resilience feature.

Table 8: Resilience Design Decisions

Decision Name	Decision
Full redundancy of all components	Ensure the full redundancy of all components in the datacenter
Established resilience boundaries	Test infrastructure and workload clusters to find resilience constraints

Nutanix Files Design

Note: There are multiple ways to manage the user profiles in a nonpersistent desktop virtualization environment. Whichever one you choose, you need to store user profile data and user data on a file share. This NVD uses Nutanix Files for storing user profile data and user data. FSLogix Profile Container and Office Container store the user profiles on a Nutanix Files share. Depending on the size and contents of the FSLogix containers, you might encounter lower storage performance when using FSLogix Cloud Cache. If you use Cloud Cache for more than 100 sessions per node, we recommend using either four SSD disks per hybrid node or all-flash nodes with a minimum of four disks per node for the workload cluster. A high sign-in rate over a short period can also decrease storage performance. While your mileage may vary, Cloud Cache can greatly decrease storage performance.

You can deploy Nutanix Files on a dedicated Nutanix cluster, on the management clusters, or on the workload cluster. For this NVD, Nutanix Files is deployed on the workload cluster.

Nutanix Files Configuration

Size and deploy Nutanix Files file server VMs (FSVMs) by the number of connections and resources required. Each connection to a file share or to a different top-level directory (TLD) within a file share is a connection. If a user connects to two file shares, each with a single TLD, that counts as two connections for that user. If that same user accesses another TLD within one of those previous shares, that is an additional connection, giving the user three connections total to Nutanix Files.

> Note: Have users access a single share with a single TLD for user profiles unless different settings are required for the user profile data. For example, if FSLogix Profile Containers are replicated at the share level, but the FSLogix Office Containers are not, two separate shares are required.

Nutanix Files has built-in high availability and resilience to recover from a range of service disruptions. For containerized solutions such as FSLogix Containers and other similar solutions, Nutanix Files also supports SMB 3.0 Transparent Failover, also known as continuously available file shares, for nondisruptive operations. You can enable continuous availability on a per share basis, which needs to be enabled after the share is created. In this NVD, continuous availability is enabled on the file share for FSLogix Profile Containers and Office Containers.

Table 9: Nutanix Files Configuration

Item	Detail
Version	3.8.1.3
Cluster size	3 VMs
vCPUs per VM	6
Memory per VM	16 GB

Item	Detail
Share settings	Distributed share Enable file system compression Enable access-based enumeration Enable continuous availability (enabled after share deployment using AFS CLI commands)

Note: Although this NVD uses Nutanix Files 3.8.1.3, you can use newer versions of Nutanix Files. Scalability and sizing can change with newer versions of Nutanix Files.

This NVD discusses an initial six-node workload cluster deployment to support up to 1,500 sessions for Windows Server 2019 or 1,200 sessions for Windows 10 using FSLogix Profile Containers on Nutanix Files. If you expand the workload cluster past six nodes, then you will need to adjust your Nutanix Files resources to account for the additional usage capacity increase. For example, if you deploy all 16 nodes for the maximum of 2,250 Windows 10 users when using FSLogix with Nutanix Files, then Nutanix Files will need to be expanded to five FSVMs.

Note: To expand Nutanix Files, you can scale up the existing FSVMs with more compute resources or scale out by adding more FSVMs to the deployment. Scaling up compute resources and scaling out FSVMs is determined by the number of additional connections to Nutanix Files. Refer to Nutanix Files Sizing Guide for more information.

Citrix DaaS Design

This NVD's desktop virtualization solution uses Citrix DaaS to deploy nonpersistent workloads on the workload cluster. Citrix DaaS infrastructure components are hosted and maintained by Citrix, with the Cloud Connectors hosted on the management cluster to connect with Citrix DaaS.

Citrix DaaS Decisions

This NVD supports the Citrix DaaS platform configuration detailed in the following table.

Table 10: Citrix DaaS Platform Configuration

Item	Detail	Rationale
General		
Software versions	Citrix DaaS	-
	Citrix Virtual Delivery Agent 1912 CU5	You can use any version of Citrix VDA 7.15 or higher with Citrix DaaS.
	Citrix Provisioning 1912 CU5	You can use any version of Citrix Provisioning 7.18 or higher with Citrix DaaS.
Citrix DaaS		
Cloud Connectors	Minimum: 3 (n + 1) Scale: 3 per additional resource location	High availability for Cloud Connectors
Users per resource location	Up to 10,000 single-session VDAs Up to 1,000 multisession VDAs	Citrix DaaS resource location limits
Active Directory domains per resource location	1 Active Directory domain	Citrix DaaS resource location limit
Clusters per resource location	Up to 20 host connections	Citrix DaaS resource location limits
Load balancing	Built into Cloud Connectors	Ensures availability of Cloud Connectors Balances load between Cloud Connectors

Item	Detail	Rationale
Virtual hardware specs	vCPU: 4 Memory: 8 GB+ (Local Host Cache) Disk: 60 GB vDisk	Standard sizing practice
VDA management	Nutanix AHV plug-in VDA management	Allows deployment and management
Citrix Provisioning		
Provisioning servers	Minimum: 2 (n + 1) Scale: 2 per additional resource location	High availability for Provisioning servers
Load balancing	Built into Provisioning servers	Balances load between Provisioning servers
Virtual hardware specs	vCPU: 4 Memory: 12 GB+ (number of vDisks) Disk: 60 GB vDisk	Standard sizing practice
VDA provisioning	Nutanix AHV plug-in VDA provisioning	Allows deployment

Note: Ensure that the same version of the Nutanix AHV plug-in for Citrix is installed on all Citrix Cloud Connectors in the DaaS instance in resource locations where AHV is deployed.

This NVD supports the Citrix DaaS infrastructure configuration detailed in the following table.

Table 11: Citrix DaaS Infrastructure Configuration

Item	Detail	Rationale
Active Directory		

Item	Detail	Rationale
Global catalog and DNS	Minimum: 2 (n + 1)	High availability for global catalog and DNS Microsoft best practice
DHCP		
DHCP servers	Nutanix IPAM	High availability for Nutanix IPAM is built in
Load balancing	Built in	Ensures availability of DHCP
SQL Server (Citrix Provisioning only)		
SQL Servers	Minimum: 2 (n + 1) per resource location Scale: 2 per additional resource location	High availability for SQL Servers
SQL Server database versions	Microsoft SQL Server 2008 SP3 and later	Citrix Provisioning database requirements
Data protection	SQL Server Clustering, Mirroring, or Always On availability groups (including Basic Always On availability groups)	Ensures availability of SQL Servers
Virtual hardware specs	vCPU: 4 Memory: 8 GB+ Disk: 3 × 60 GB vDisk	Standard sizing practice Separate disks for OS, data, logs

Virtual Machine Design

As the overall objective is to provide a hybrid cloud environment for desktop virtualization workloads, this NVD establishes three standard VM sizes to facilitate consistent deployment, automation, sizing, and capacity planning for the environment. The Cluster Design section specifies the maximums for each

VM size to help with capacity planning, but you can combine any number of VMs of any size up to the maximums Nutanix designed this architecture to support.

Virtual Machine Names

Nutanix recommends that you keep the VM name and the guest OS host name the same. This approach streamlines operational and support requirements and minimizes confusion when you identify systems in the environment.

Virtual Machine Guest Clustering

You can use VM guest clustering to form failover clusters using shared disk devices with both Windows and Linux guest operating systems. Nutanix AHV allows you to use a shared volume group between multiple VMs as part of a failover cluster—just connect the shared volume group to the VMs and install the necessary guest software. Nutanix natively integrates SCSI-based fencing using persistent reservations and does not require any complex configuration.

Virtual Machine Standard Deployment Sizes

This NVD supports the VM configurations detailed in the following table.

Table 12: Supported VM Configurations

VM Size	Small	Medium	Large
Virtual CPU	2	3	8
Virtual memory	4 GB	6 GB	42 GB
Virtual storage	60 GB	60 GB	80 GB
Virtual NIC	1	1	1
Virtual CD-ROM	1	1	1
Volume groups	No	No	No
Maximum VM instances per node	165	150	8

Note: This design targets an oversubscription ratio of eight or fewer virtual CPUs per physical CPU.

Windows Virtual Machines

All Windows VMs in this NVD are based on Windows 10 or Windows Server 2019. Windows VMs use the standard templates detailed in the following table when provisioned with MCS or PVS.

Table 13: Standard Templates for Windows VMs

Template	Windows 10	Windows 10 Power User	Windows Server 2019
Base template size	Small	Medium	Large
Virtual CPU	2 per VM	3 per VM	8 per VM
Virtual memory	4 GB per VM	6 GB per VM	42 GB per VM
Virtual storage	60 GB per VM (VirtIO-SCSI)	60 GB per VM (VirtIO-SCSI)	80 GB per VM (VirtIO-SCSI)
Virtual NIC	1 (VirtIO-Net: kNormal)	1 (VirtIO-Net: kNormal)	1 (VirtIO-Net: kNormal)
Virtual CD-ROM	1	1	1

The Windows templates have all necessary application components preinstalled and ready to deploy on demand as a single VM through MCS and PVS.

Table 14: Standard Templates for Windows VM Applications

Parameter	Setting
Operating system	Windows 10 21H2 (x64) and Windows Server 2019
Windows Updates	Windows 10 1/11/22 and Windows Server 2019 3/8/22

Parameter	Setting
Applications	Adobe Acrobat DC Adobe Flash Player 11 Doro PDF 1.82 FreeMind Internet Explorer 11 Microsoft Edge Browser Microsoft Office 2019
Citrix Virtual Delivery Agent	7.1912 CU5
Citrix Provisioning Target Device Agent (PVS workloads only)	7.1912 CU5
Microsoft FSLogix Apps	2.9.7979.62170
Optimizations	Citrix Optimizer

Note: This NVD uses the applications in the standard templates for Windows VMs in scale testing. Your applications might vary. A single image was used for each provisioning and OS type (MCS and PVS for Windows 10 and Windows Server 2019) in this NVD.

For more detail on desktop optimizations, refer to the Appendix for Windows VM performance tuning recommendations.

Citrix Machine Creation Services Base Image Distribution

With MCS base image distribution across multiple Nutanix clusters (in the same resource location or remote resource locations), Nutanix has built-in data protection. Native data protection uses protection domains to replicate MCS base images across clusters. A protection domain is a defined group of entities (VMs and volume groups) used to take hardware-level snaphosts of a cluster, which you can replicate to one or more remote clusters.

Citrix Provisioning vDisk Image Distribution

With PVS vDisk image distribution across multiple Citrix Provisioning deployments and Nutanix clusters, the vDisk image is replicated at the

application level between PVS vDisk stores. vDisk stores can either be local storage on each PVS server or a shared location, such as Nutanix Files, in each resource location. To replicate PVS vDisk images within or between resource locations for local or shared PVS vDisk stores, consider the following options:

- Manually copy vDisks.
- Use the Citrix vDisk replicator tool .
- Write a script to copy vDisks.

Cluster Design

This design incorporates two distinct cluster types:

1. Management clusters: critical infrastructure and environment management workloads.
2. Workload clusters: the building block for all desktop virtualization workloads.

This section defines the overall high-level cluster design, platform selection, capacity management, scaling, and resilience. This design follows the block-and-pod architecture defined in the Nutanix Hybrid Cloud Reference Architecture.

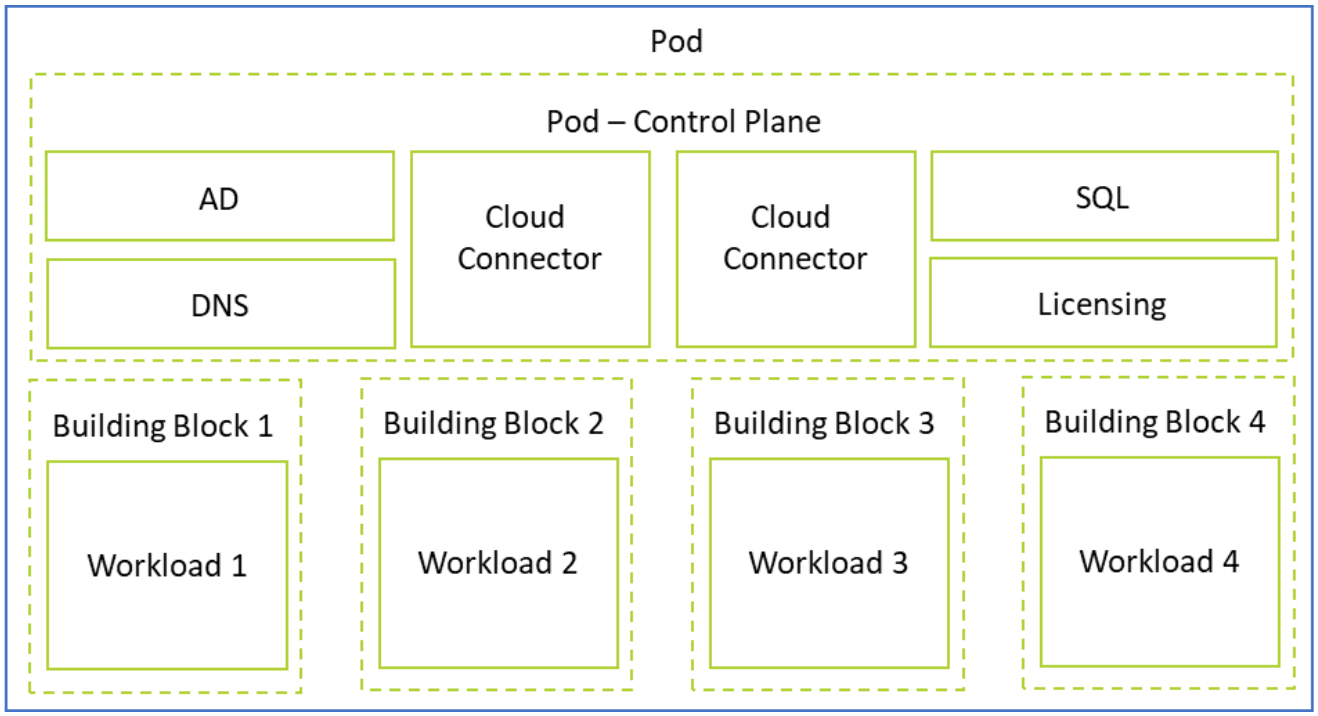

Figure 4: Pod and Building Block Architecture

Cluster Conceptual Design

This NVD solution uses one region with a single availability zone (AZ). The AZ has active workloads provisioned by MCS or PVS using Citrix DaaS with infrastructure for Active Directory, DNS, SQL, Nutanix Files, Licensing, and Provisioning resources.

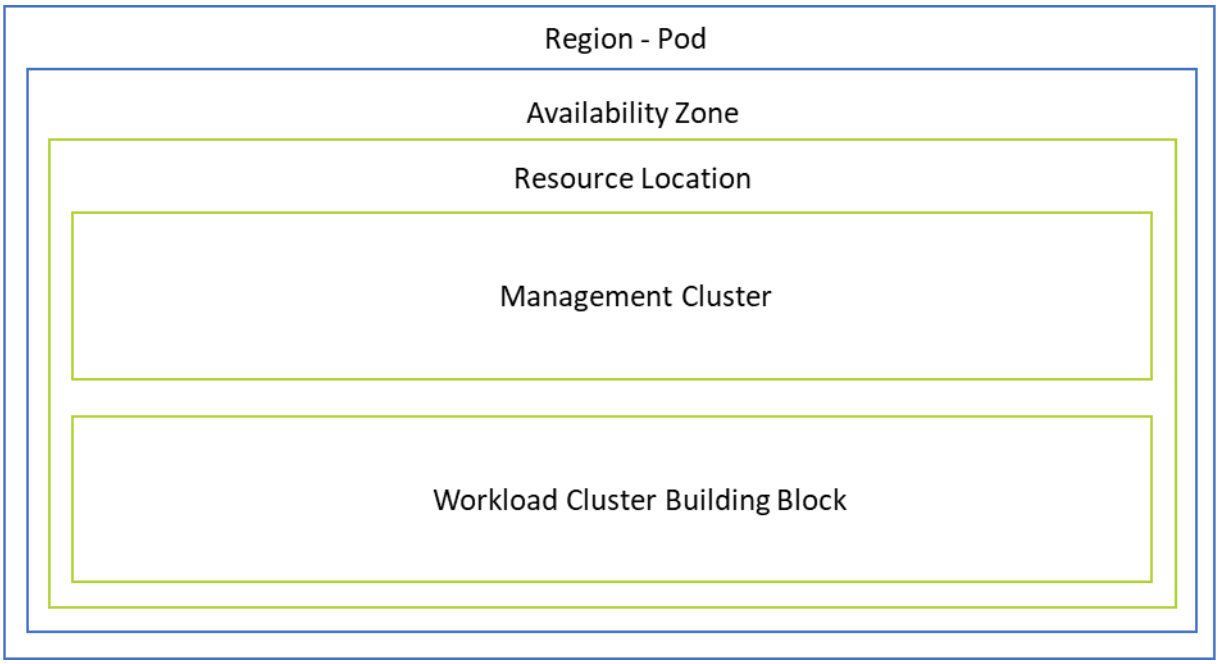

Figure 5: Conceptual Design

Table 15: Cluster Design Decisions

Decision Name	Decision
Number of regions	Use 1 region
Number of AZs	Use 1 AZ
Number of datacenters	Use 1 datacenter: 1 per AZ
Mixed workloads or dedicated workload per cluster	Dedicated workloads per cluster, as this design is for desktop virtualization
Minimum workload cluster building block size	Use at least 4 nodes
Workload cluster building block expansion increments	Use 1 node
Maximum workload cluster building block size for this design	Use at most 16 nodes

Decision Name	Decision
Maximum workload cluster building blocks per pod for this design	Use at most 4 workload cluster building blocks (1 per AZ) per pod
Maximum number of running VMs per usable node in the workload cluster building block	Use at most 165 small VMs, 150 medium VMs, or 8 large VMs per usable node
Maximum number of VMs per workload cluster building block	Use at most 2,475 small VMs per workload cluster building block
Workload cluster building block node redundancy	Use n + 1 for redundancy
Maximum usable nodes per maximum workload cluster building block for this design	Configure at most 15 usable nodes per maximum workload cluster building block
Workload cluster building blocks in one rack or split across multiple racks	Use one rack per workload cluster building block
Cluster replication factor	Use replication factor 2
Cluster high availability configuration	Guarantee high availability

Platform Selection

Table 16: Platform Selection

Cluster	Management	Workload
Node type	NX-3060-G7	NX-3155G-G8
Node count	4 (increments of 1)	6–16 per building block (increments of 1, up to 16 maximum)
Processor	2 Intel Xeon Gold 5520 18-core 125 W 2.2 GHz	2 Intel Xeon Gold 6354 18-core 205 W 3.0 GHz

Cluster	Management	Workload
RAM	24 × 32 GB 3,200 MHz DDR4 RDIMM (768 GB total)	32 × 32 GB 3,200 MHz DDR4 RDIMM (1 TB total)
SSD	2 × 1.92 TB	6 × 1.9 TB
HDD	4 × 2 TB	N/A
NIC	10 GbE Dual SFP+	25 GbE Dual SFP+
Support	3Y Production	3Y Production

Note: This NVD uses Nutanix NX nodes for scale testing. You can choose any Nutanix-supported hardware vendor with the same or better hardware specifications.

Capacity Management

This NVD sizes the management cluster to host typical workloads as defined in the Management Components section of this document. If the cluster needs more resources, you can expand it one node at a time. Prism Pro can help forecast resource demand.

The main unit of expansion for workload clusters is the building block. In this design, each workload cluster building block has a maximum of 16 nodes, with 15 nodes of usable capacity and 1 node for failure capacity, and a minimum of 4 nodes with 3 usable (following the n + 1 principle). You can expand a workload cluster building block in increments of one node, up to the maximum. Based on the small VM specification, you can have a maximum of 2,475 VMs per workload cluster building block. When a workload cluster building block reaches the maximum number of nodes, the administrator starts a new building block with the four-node minimum, then can expand the new block in increments of one node as needed.

Each pod can support a maximum of four workload cluster building blocks of 16 nodes each. When a pod reaches the maximum of four workload cluster building blocks, the administrator deploys a new pod. This NVD sets the workload cluster building block maximum at 16 nodes to allow you to complete nondisruptive Nutanix software, hardware, firmware, and driver maintenance using Nutanix LCM (Life Cycle Manager) within a 16-hour maintenance window (using Nutanix NX model hardware). You may use a smaller maximum size

per workload building block to shorten maintenance windows and allow more small clusters per pod without changing the maximum number of nodes or VMs each pod supports. For example, an eight-node workload cluster building block reduces maintenance windows by half and allows twice the number of clusters per pod without changing the number of nodes supported. However, the number of usable nodes decreases with the smaller cluster size, as one node per cluster is logically reserved for maintenance and failure.

Note: Nutanix OEM partner hardware platforms may require more or less time depending on the specific OEM partner recommendations.

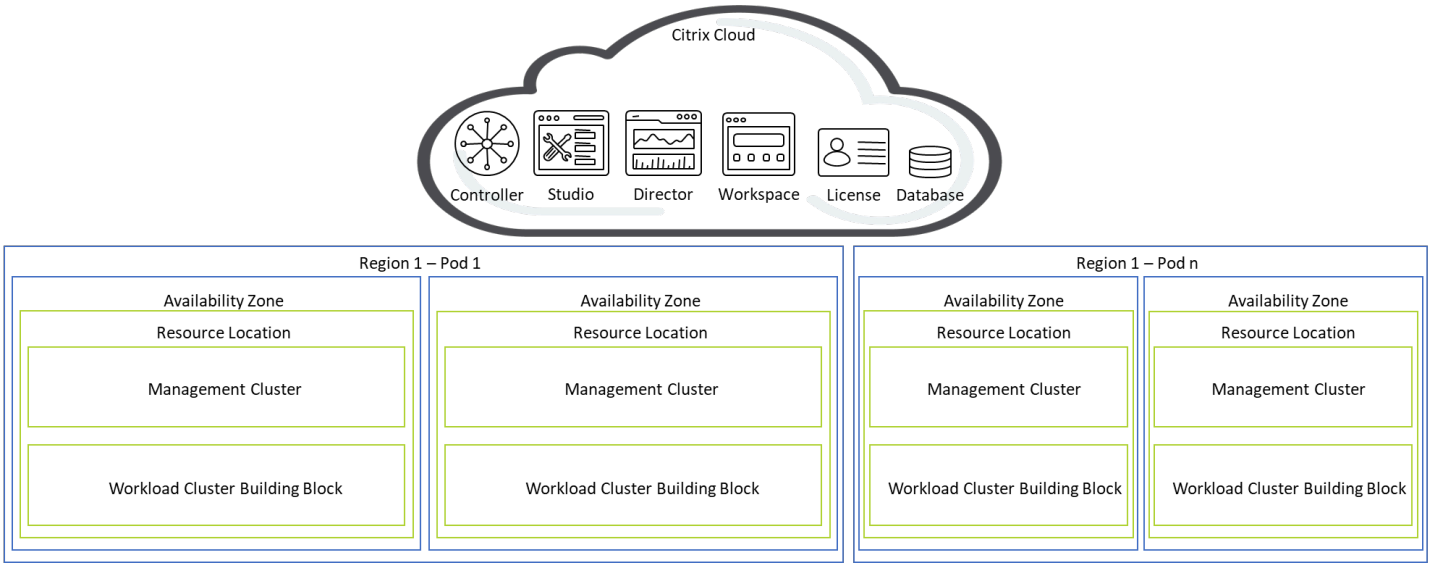

Figure 6: Scaling Beyond a Single Pod

In the previous figure, the first pod (Pod 1) reached capacity, and the administrator started a new pod (Pod n). If existing management clusters have enough capacity for the additional pods, you can reuse them and not implement additional management clusters.

Note: Citrix DaaS has a maximum of 80,000 VDAs, 200 host connections, and 50 resource locations per Citrix Cloud instance. When building out pods and availability zones, do not exceed these configuration limits. If more are required, then additional Citrix Cloud instances for Citrix DaaS will be required.

The following table displays the maximum number of VMs per workload cluster building block and per node.

Note: The maximum deployed VMs per workload cluster is 80 percent of the host CPU usage to allow for high availability capacity.

Table 17: Maximum Number of VMs

	Small VMs	Medium VMs	Large VMs
Maximum running VMs per workload cluster building block	2,475	2,250	120
Maximum running VMs per node	165	150	8

Cluster Resilience

Replication factor 2 protects against the loss of a single component in case of failure or maintenance. During a failure or maintenance scenario, Nutanix rebuilds any data that falls out of compliance much faster than traditional RAID data protection methods. Rebuild performance increases linearly as the cluster grows.

In the Nutanix architecture, rapid recovery in the event of failure is the standard, and there are no single points of failure. You can configure the cluster to maintain three copies of data; however, for general server virtualization, Nutanix recommends that you distribute application and VM components across multiple clusters to provide greater resilience at the application level.

Tip: You can achieve rack-aware resilience when you split clusters evenly across at least three racks, but this NVD doesn't use that approach because it adds configuration and operational complexity. Nutanix cluster replication factor 2 in this design is sufficient to exceed five nines of availability (99.999 percent).

Figure 7: Availability Chart

Storage Design

Nutanix uses a distributed, shared-nothing architecture for storage. For a discussion of Nutanix storage constructs, refer to the Storage Design section in the Nutanix Hybrid Cloud Reference Architecture. For information on node types, counts, and physical configurations, see the Cluster Design section.

Creating a cluster automatically creates the following storage containers:

- NutanixManagementShare: Used for Nutanix features like Files and Objects and other internal storage needs. This storage container does not store workload vDisks.
- SelfServiceContainer: Used by the Nutanix Self-Service Portal and automation services.
- Default-Container-XXXX: Used by VMs to store vDisks for user VMs and applications.

> Note: You can delete the Default-Container and create a new one with your desired naming convention.

In the AZ, the management cluster uses the Default-Container to store VMs and their vDisks. This NVD also enables inline compression on the Default-Container for all management and workload clusters in the AZ. Because these clusters have a fault tolerance level of 1, the replication factor for the containers is 2.

Data Reduction Options

To increase the effective capacity of the cluster, the design enables inline compression with compression delay of zero, as the intended workload is nonpersistent desktop virtualization. We disable deduplication for this container and use the defaults for the other containers.

> Tip: Enabling compression for desktop virtualization workloads is a general best practice; only enable the Elastic Deduplication Engine for full clones (also known as persistent workloads). Erasure coding is not a suitable data reduction technology for desktop virtualization.

The data reduction settings in the following table apply across both the AZs.

Table 18: Data Reduction Settings

Container	Compression	Deduplication	Erasure Coding
Default-Container-XX	On	Off	Off
NutanixManagementShare	On	Off	Off
SelfServiceContainer	On	Off	Off

> Note: This NVD uses a single container for workload VMs. If your workload VMs require different container settings, use multiple containers with required settings.

Table 19: Storage Design Decisions

Decision Name	Decision
Sizing a hybrid cluster	Hybrid storage configuration is suitable for most desktop virtualization deployments, and all flash can be used if required

Decision Name	Decision
Node type vendors	Use all same vendor nodes; don't mix node types from different vendors in the same cluster
Node and disk types	Use similar node types that have similar disks
Sizing for node redundancy for storage and compute	Size all clusters for n + 1 failover capacity
Fault tolerance and replication factor settings	Configure the cluster for fault tolerance 1 and configure the container for replication factor 2
Inline compression	Enable inline compression
Deduplication	Disable deduplication
Erasure coding	Disable erasure coding
Availability domain for workload cluster	Use block awareness
Availability domain for management cluster	Use node awareness

Network Design

A Nutanix cluster can tolerate multiple simultaneous failures because it maintains a set redundancy factor and offers features such as block awareness and rack awareness. However, this level of resilience requires a highly available network connecting a cluster's nodes.

Nutanix clusters send each write to another node in the cluster. As a result, a fully populated cluster sends storage replication traffic in a full mesh, using network bandwidth between all Nutanix nodes. Because storage write latency directly correlates to the network latency between Nutanix nodes, any increase in network latency adds to storage write latency. Protecting the cluster's read and write storage capabilities requires highly available connectivity between nodes. Even with intelligent data placement, if network connectivity between multiple nodes is interrupted or becomes unstable, VMs on the cluster can experience write failures and enter read-only mode.

A Nutanix environment should use datacenter-grade switches designed to handle high-bandwidth server and storage traffic at low latency. Refer to the Nutanix Physical Networking best practice guide for more information.

Physical Network Architecture

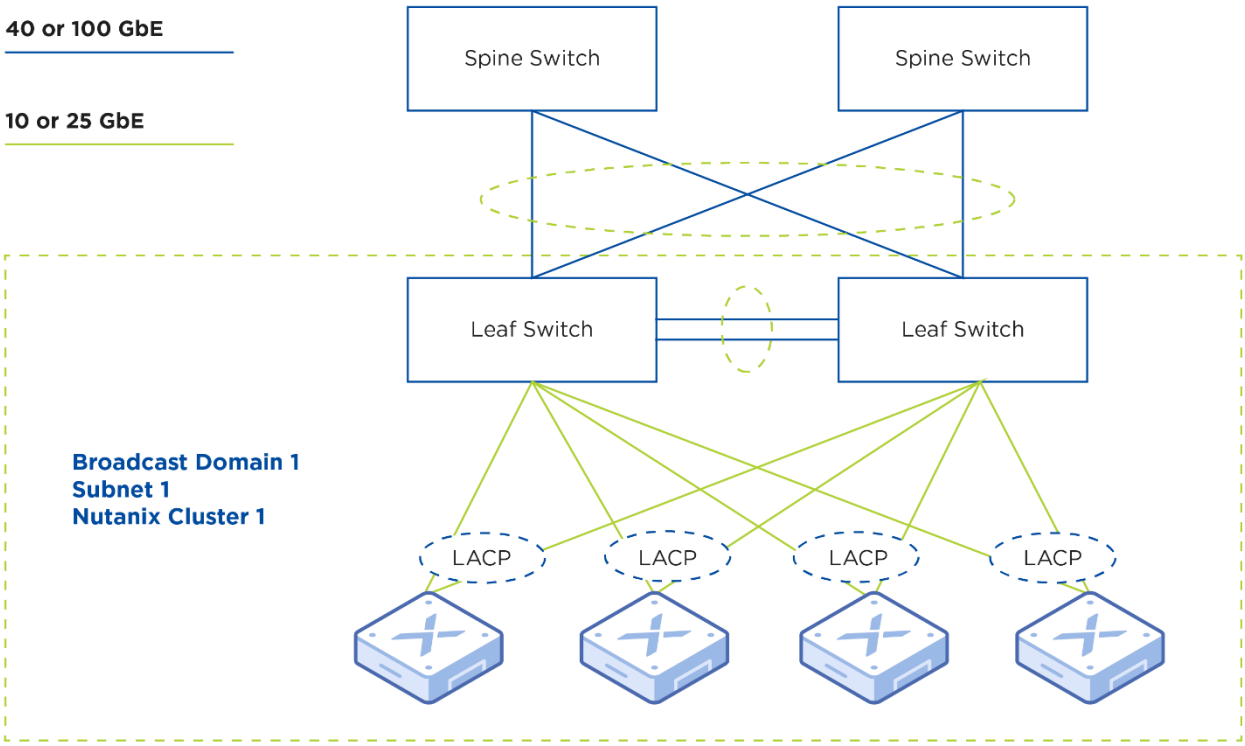

Figure 8: Physical Network Architecture

Table 20: Physical Network Design Decisions

Decision Name	Decision
Use a large-buffer datacenter switch at 10 Gbps or faster	25 Gbps switches
Network topology for new environments	Leaf-spine network topology
Populate each rack with two 10 GbE or faster top-of-rack switches	25 Gbps switches

Decision Name	Decision
Avoid switch stacking to ensure network availability during individual device failure	MLAG configuration to avoid stacking
Switches between nodes	Ensure that there are at most three switches between any two Nutanix nodes in the same cluster
Reduce network oversubscription to achieve as close to a 1:1 ratio as possible	1:2
Network design	Layer 2

Table 21: Node Connectivity Network Design Decisions

Decision Name	Decision
CVM and hypervisor VLAN	Configure the CVM and hypervisor VLAN as native, or untagged, on server-facing switch ports
Switch ports for guest workloads	Use tagged VLANs on the switch ports for all guest workloads
Connect at least one 10 GbE or faster NIC to each Top-of-rack switch	25 GbE NICs
Virtual switch	Use a single vs0 virtual switch with at least two of the fastest uplinks of the same speed
NICs	Use NICs from the same vendor within a bond
Logical network separation	Use VLANs to separate logical networks
Use active-backup uplink load balancing	LACP (Link Aggregation Control Protocol)
MTU size	1,500-byte MTU
Terminate L2/L3 networking	Spine

Table 22: Workload Cluster Networks

Decision Name	Decision
Shared infrastructure network subnet size	/24
VM network subnet size	/20
Number of addresses available per /21 network	4,083
Number of VM networks	1
Present VM networks to other workload clusters	No
Stretch VM networks to secondary site	No

Note: Instead of having one single large network for workload VMs, consider using multiple smaller networks. For example, 11 /24 networks would give you 2,695 usable addresses total with 245 usable addresses per /24 network, or 3 /22 networks would give you 3,033 usable address total with 1,011 usable address per /22 network.

Table 23: Management Cluster Networks

Decision Name	Decision
Shared infrastructure network subnet size	/24
VM network subnet size	/24
Number of addresses available per /24 network	245
Number of VM networks	1

Management Components

Management components such as Active Directory, DNS, and NTP are critical services that must be highly available. Nutanix Files is a software-defined, scale-out file storage solution that provides a repository for unstructured data, such as home directories, user profiles, and departmental shares. Files is a fully integrated core component of Nutanix, designed to recover from a range of service disruptions. You can deploy Files starting with a minimum of three FSVMs and then scale up the FSVMs with more resources or scale out more

FSVMs on the cluster. In a Files deployment, the number of FSVMs the file server has must be fewer than or equal to the number of nodes in the Nutanix cluster; however, you can create multiple Files file server deployments if needed.

Citrix Cloud Connectors allow resources in each resource location to connect to Citrix Cloud. Resources such as Active Directory, AHV, Citrix Virtual Delivery Agents (VDAs), StoreFront, Citrix Gateway, and Provisioning will communicate through Cloud Connectors to connect with Citrix Cloud. Nutanix recommends having at least two Cloud Connectors in each resource location for high availability and redundancy. Cloud Connectors are updated automatically by Citrix in a rolling upgrade process. Citrix Cloud has a Service Continuity feature that allows Citrix DaaS to be resilient in the unlikely event of a Citrix Cloud service failure or outage. Service Continuity allows users to continue to access their resources when Citrix Cloud resources are not reachable.

> Note: To maintain high availability and load management while updating Cloud Connectors, consider deploying at least three Cloud Connectors in each resource location. Citrix periodically updates the Cloud Connector software, and Cloud Connectors are updated one at a time.

Citrix Provisioning allows nonpersistent images to be delivered to VMs through network streaming. When using Provisioning with Citrix DaaS, you must deploy SQL and Citrix Licensing. Provisioning is an optional Citrix DaaS component that you only need if you plan to deploy nonpersistent machines using Provisioning. Nutanix recommends having at least two Provisioning servers per Provisioning site in a farm for redundancy and high availability. Plan for high availability and redundancy so that a single Provisioning server failure does not reduce the amount of target devices supported in a site. Configure the Provisioning boot files with multiple Provisioning servers in a site for high availability (the Provisioning server boot file can have up to four Provisioning servers). Each Provisioning server should have itself listed highest in the boot file order for local boot file configuration. Enable Provisioning server load balancing on the vDisk for load distribution across Provisioning servers in the site.

> Note: Configure 8 vCPU per Provisioning server in larger environments, with multiple Provisioning servers in the site for high availability, load balancing, and redundancy.

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 keeps Provisioning operational for target devices when a database outage occurs. When the Provisioning database is offline, management functions and the management console become unavailable. Offline database support is not enabled by default; if you want to use it, enable the feature after setting up and configuring the Provisioning farm.

> Note: 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. Enable Provisioning offline database support after configuring the farm.

When you design your management components, decide how many Nutanix Files FSVMs, Citrix Cloud Connectors, and Citrix Provisioning servers you need. This NVD uses a Nutanix Files deployment with three FSVMs, a pair of Citrix Cloud Connectors, and a pair of Citrix Provisioning servers (for Provisioning workloads only) in the AZ. This setup provides better scalability and increased resilience for Citrix DaaS deployments.

> Note: You can use MCS or PVS, and both image provisioning options are in this NVD. Citrix Provisioning requires additional management components, while MCS is built into the Citrix DaaS control plane.

Management Conceptual Design

Nutanix recommends that you have a dedicated management cluster in the datacenter AZ for both Nutanix and non-Nutanix environment management and control plane instances. For this validated design, the management clusters contain at least four nodes. The management clusters run only core infrastructure management components, not general user VM workloads.

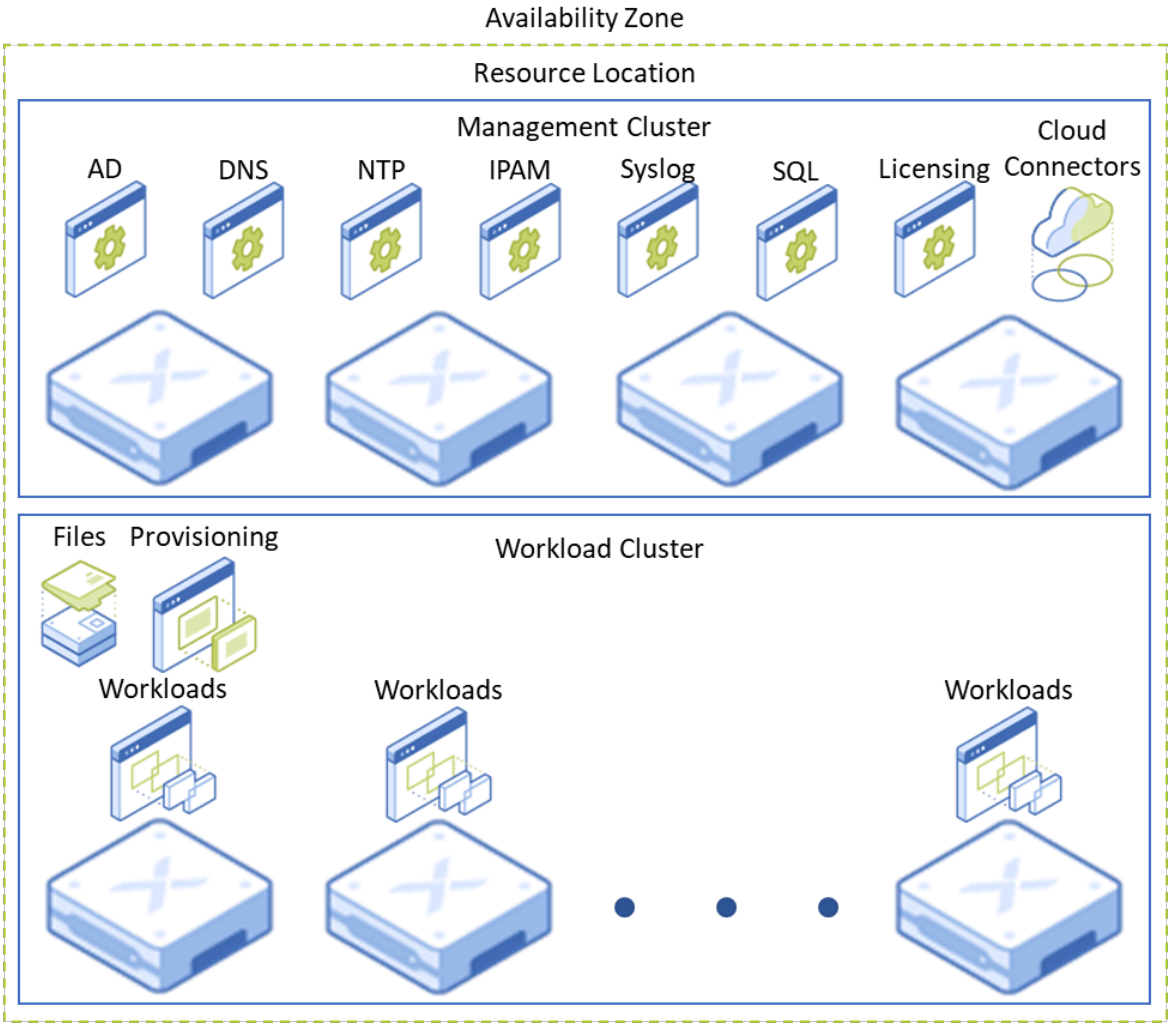

Figure 9: Management Plane

Management Detailed Design

In this NVD, management clusters run AOS 5.20.3 and workload clusters run AOS 5.20.3.

Table 24: Nutanix Management Component Software Versions

Component	Software Version
AOS	5.20.3 (LTS)

The following table lists the design decisions for the Nutanix management components.

Table 25: Management Component Design Decisions

Decision Name	Decision
Management cluster architecture	One management cluster at each datacenter
Management cluster size	Four nodes (n + 1)
Management cluster node specifications	See the Platform Selection section
Active Directory authentication	Use Active Directory authentication
Connection to Active Directory	Use SSL or TLS for Active Directory

Monitoring

Monitoring in the NVD falls into two categories: event monitoring and performance monitoring. Each category addresses unique needs and different issues.

In a highly available environment, you must monitor events to maintain high service levels. When faults occur, the system must raise alerts in a timely manner so that administrators can take remediation actions as soon as possible. This NVD configures the Nutanix platform's built-in capability to generate alerts in case of failure.

In addition to keeping the platform healthy, maintaining a healthy level of resource usage is also essential to the delivery of a high-performing environment. Performance monitoring continuously captures and stores metrics that are essential when you need to troubleshoot application performance. A comprehensive monitoring approach should track the following areas:

- Application and database metrics.
- Operating system metrics.
- Hyperconverged platform metrics.

- Network environment metrics.
- Physical environment metrics.

By tracking a variety of metrics in these areas, the Nutanix platform can also provide capacity monitoring across the stack. Most enterprise environments inevitably grow, so you need to understand resource utilization and the rate of expansion to anticipate changing capacity demands and avoid any business impact caused by lack of resources.

This NVD only includes Prism Element. If you would like to use Prism Central for monitoring, refer to the Nutanix Hybrid Cloud Validated Design.

> Note: While this NVD focuses on Prism Element, Nutanix Cloud Platform VDI licensing requires a Prism Central instance, which you can deploy in the management cluster. See the Nutanix Hybrid Cloud Validated Design for Prism Central considerations.

Monitoring Conceptual Design

In this NVD, Prism Element performs most of the event monitoring. We use SMTP-based email alerts as the channel for notifications in this design.

> Note: This NVD uses syslog for log collection; for more information, refer to the Security and Compliance section. All alerts from Prism Element go to a primary email alert recipient that's always monitored.

Each Nutanix cluster in this NVD sends out notifications using SMTP. The individual Nutanix clusters send alerts to a primary email recipient mailbox that is always monitored.

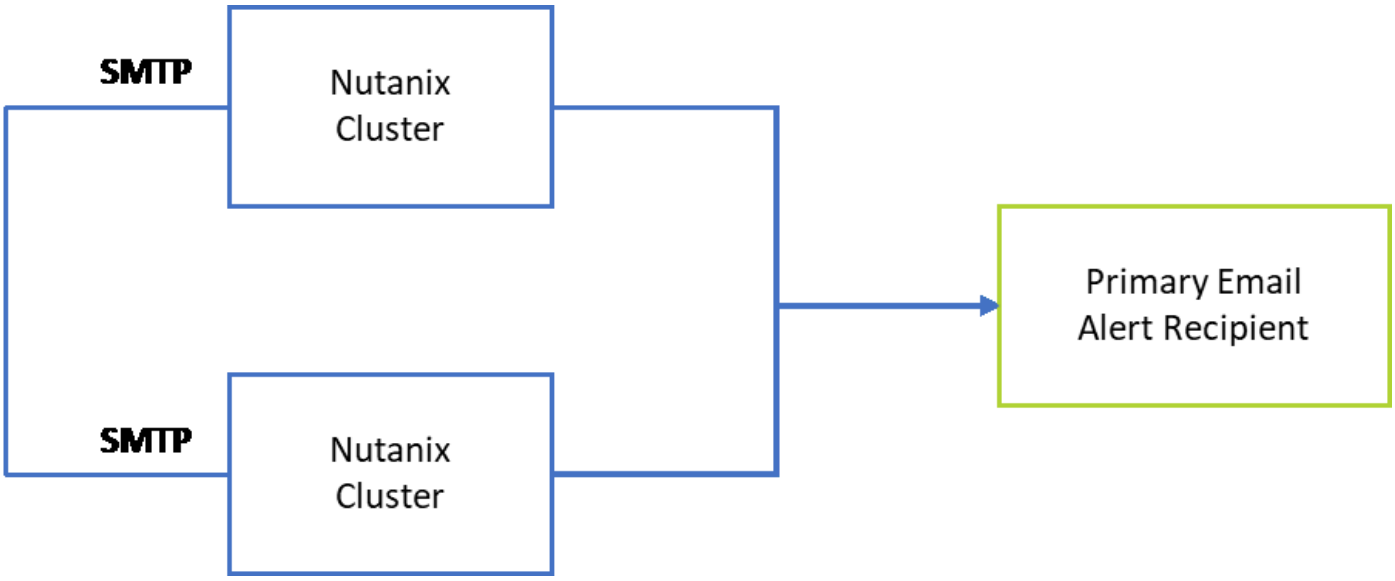

Figure 10: SMTP for Email Alerts from Prism Element and Prism Central

Prism Element monitors cluster performance in key areas such as CPU, memory, network, and storage utilization. Prism Element captures these metrics by default.

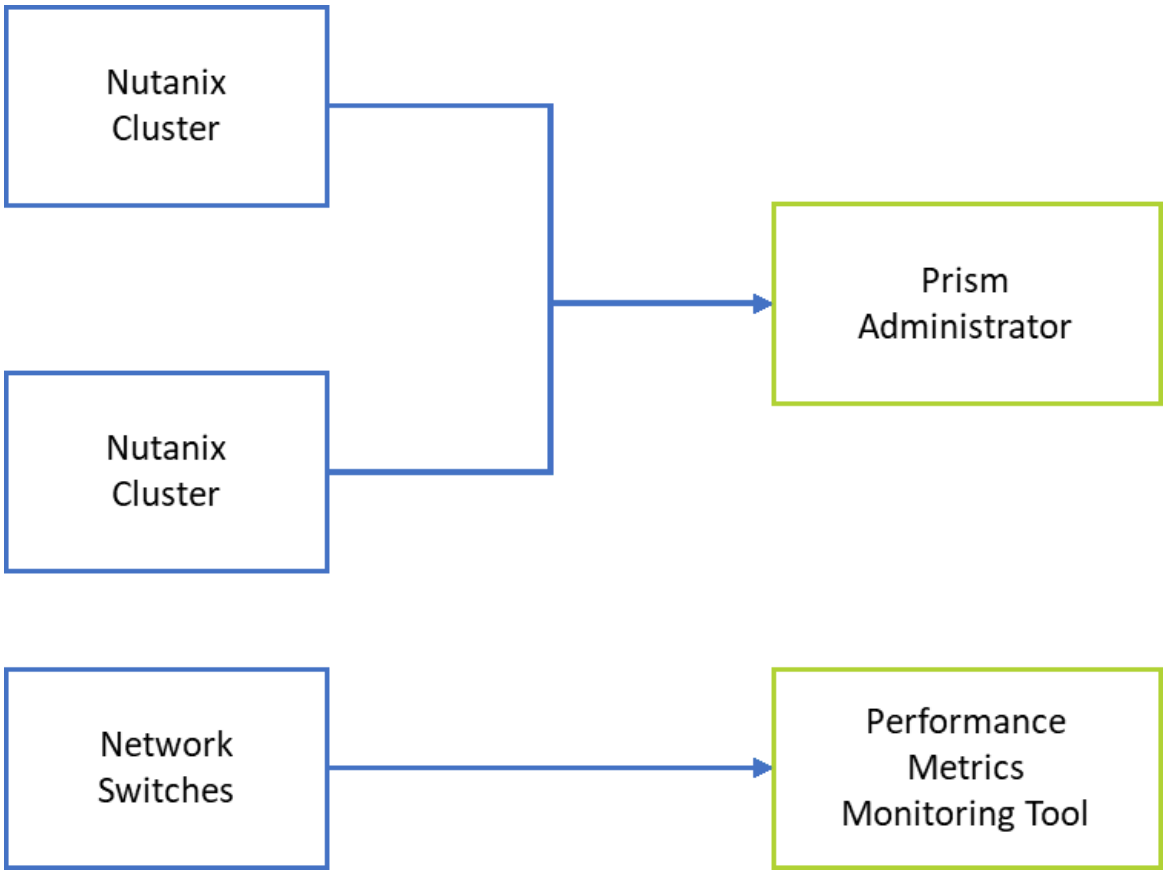

Figure 11: Systems Used to Capture Performance Metrics

The network switches that connect the cluster also play an important role in cluster performance. A separate monitoring tool that's compatible with the deployed switches can capture switch performance metrics. For example, an SNMP-based tool can regularly poll counters from the switches.

The following table provides descriptions of the monitoring design decisions.

Table 26: Monitoring Design Decisions

Decision Name	Description
Platform performance monitoring	Prism Element monitors Nutanix platform performance

Decision Name	Description
Network switch performance monitoring	A separate tool that performs SNMP polling to the switches monitors network switch performance
Management cluster storage utilization warning threshold	On a management cluster with AOS 6.0.x, leave the Prism Element storage utilization warning threshold at 75 percent (the default value)
Workload cluster storage utilization warning	On a workload cluster with AOS 6.0.x, leave the Prism Element storage utilization warning threshold at 75 percent (the default value)
Prism Element health check CPU utilization warning threshold	For the Prism Element health check, leave the host CPU utilization warning threshold at 75 percent (the default value)
SMTP alerting	Use SMTP alerting; use enterprise SMTP service as the primary SMTP gateway for Prism Element
SMTP alerting source email address	Configure the source email address to be clustername@nutanix.com to uniquely identify the source of emails
SMTP alerting Prism Element recipient email address	Configure the Prism Element recipient email address to be primaryalerts@nutanix.com
NCC reports	Configure daily NCC checks to run at 6:00 AM local time and send them by email to the primary alerting mailbox

Security and Compliance

Nutanix recommends a defense-in-depth strategy for layering security throughout any enterprise datacenter solution. This design section focuses on validating the layers that Nutanix can directly oversee at the control and data plane levels. Refer to the Network Design section of the Nutanix Hybrid Cloud Validated Design for more information on the network-based security of hosted VMs using microsegmentation policies, as this NVD doesn't include Flow Network Security or Prism Central. Read the Security and Compliance Layer section of the Nutanix Hybrid Cloud Reference Architecture for additional details.

Authentication and Authorization

All Nutanix control plane endpoints use Active Directory–hosted LDAPS. Active Directory itself is redundant across the management cluster in the AZ. Only administrative accounts are mapped to admin roles, which are controlled through a named Active Directory group.

This NVD rotates all default passwords for all accounts that aren't integrated with the Active Directory, such as emergency accounts or local accounts for out-of-band interfaces. Because clusters don't have lockdown mode enabled, password SSH is enabled by default.

AOS Hardening

In each AOS cluster, this NVD enables additional nondefault hardening options:

- Advanced Intrusion Detection Environment (AIDE).
- Hourly security configuration management automation (SCMA).

Both features are trivial to enable, introduce little to no discernible system overhead, and help detect and prevent internal system configuration changes that may otherwise compromise service availability. These features add to the intrinsic hardening built into AOS.

Syslog

For each control plane endpoint, system-level internal logging goes to a centralized third-party syslog server that runs in the local management cluster in each AZ. The system is configured to send logs for all available modules when they reach the syslog Error severity level. TCP transport via TLS is preferred where available.

Certificates

SSL endpoints serve all Nutanix control plane web pages. This NVD replaces the default self-signed certificates with certificates signed by an internal certificate authority from a Microsoft public key infrastructure (PKI). Any client endpoints that interact with the control plane should have the trusted certificate authority chain preloaded, preventing browser security errors.

Note: Certificate management is an ongoing activity, and certificates need to be rotated periodically. The NVD signs all certificates for one year of validity.

Data-at-Rest Encryption

Nutanix AOS can perform data-at-rest encryption (DaRE) at the cluster level; however, as the NVD doesn't have a stated requirement that warrants enabling it, this design doesn't use it. If requirements change, you can enable DaRE nondisruptively after cluster creation and data population. Once you enable DaRE, existing data is encrypted in place and all new data is written in an encrypted format.

Note: To enable DaRE, you must also deploy an encryption key management solution.

Our decision to not use DaRE doesn't preclude the use of in-guest encryption techniques such as system-level encryption, database encryption (for example, Microsoft SQL Transparent Data Encryption (TDE)), or the storage of encrypted files; however, in-guest encrypted data can't be compressed in most cases. As this design enables compression, but in-guest encrypted data isn't likely to be compressible, using in-guest encryption might affect the amount of available storage.

Table 27: Security Design Decisions

Decision Name	Description
DaRE	Disable DaRE; don't deploy a key management server
SSL endpoints	Sign control plane SSL endpoints with an internal certificate authority (Microsoft PKI)
Certificates	Provision certificates with a yearly expiration date and rotate accordingly
Authentication	Use Active Directory LDAPS authentication (port 636)
Control plane endpoint administration	Use a common administrative Active Directory group for all control plane endpoints
Cluster lockdown mode	Don't enable cluster lockdown mode (allow password-driven SSH)

Decision Name	Description
Nondefault hardening options	Enable AIDE and hourly SCMA
System-level internal logging	Enable error-level logging to external syslog server for all available modules
Syslog delivery	Use TCP transport for syslog delivery

Table 28: Security Configuration References

Configuration Target	Key:Value
Active Directory	AD-admin-group:ntnx-ctrl-admins
Syslog Server	infra-AZ[1..2]-syslog:6514 (tcp)

Datacenter Infrastructure

This design assumes that datacenters in the hosting region can sustain two AZs without intraregional fate-sharing—in other words, that failures in one datacenter's physical plant or supporting utilities don't affect the other datacenter. This NVD addresses points where the Nutanix gear touches the datacenter equipment to make sure all your needs are met.

Rack Design

Each cluster is confined to a single rack. You can add more racks as needed, depending on top-of-rack network switch density as well as the datacenter's power, weight, and cooling density capabilities per square foot. Refer to the Platform Selection section for the specific node models selected for this NVD. The following figure shows the initial density for this design, with the designated requirements, assumptions, and constraints.

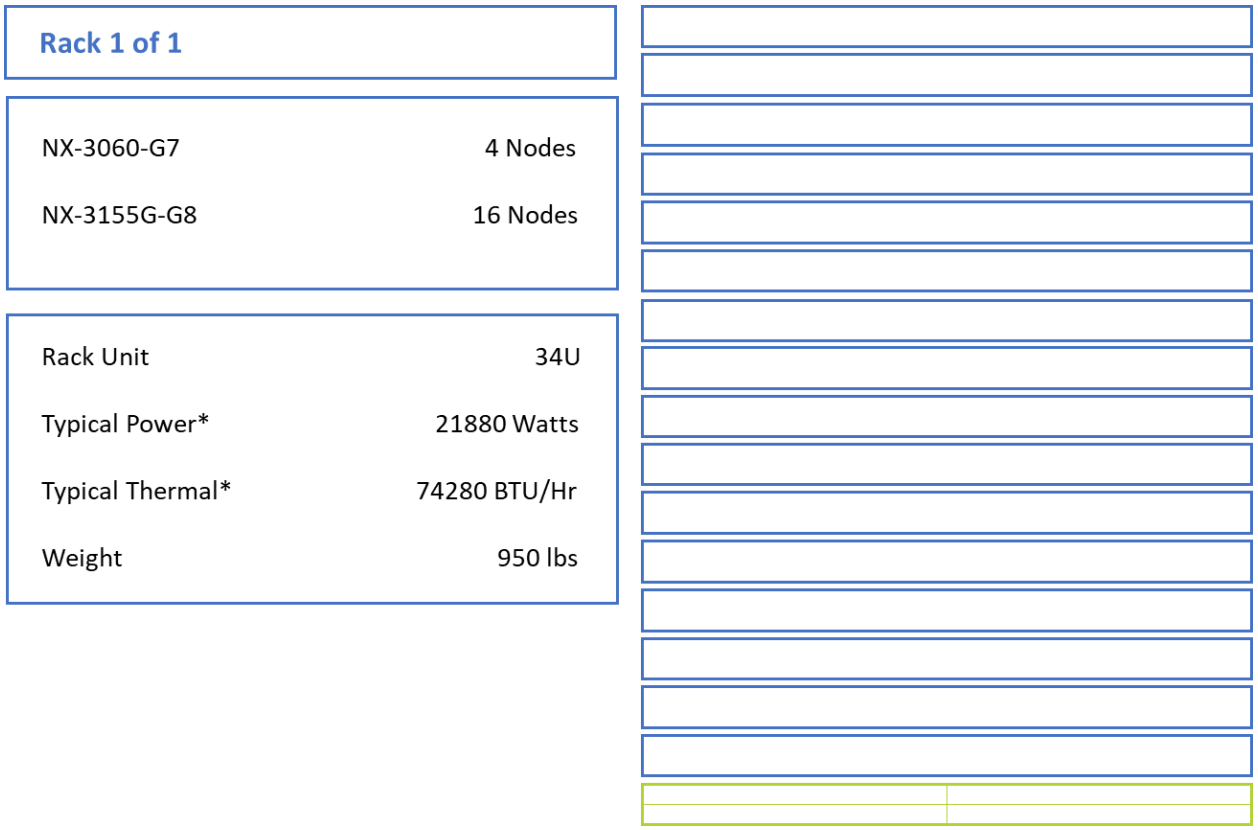

Figure 12: Rack Layout

When you scale the environment, consider physical rack space, network port availability, and the datacenter's power and cooling capacity. In most environments the workload clusters are the most likely to grow, followed by the backup clusters.

In this design's physical rack space, one generic 42RU rack contains 34RU of systems with 3RU reserved for two data switches and one out-of-band switch, leaving 5RU of space available.

For network ports, the 20 nodes in this NVD consume 20 ports on each of the two data switches. Assuming that there are two Inter-Switch Links (ISLs) and two uplinks to the upstream network, this configuration leaves 24 ports available per data switch.

For power, cooling, and weight, you need the minimums specified in the previous figure and should assume at least double these values for a fully

loaded rack including network switches. Datacenter selection is beyond the scope of this design; however, you should have a conversation about fully loaded racks with datacenter management prior to initial deployment, as planning to properly support the environment's long-term growth may change where in the facility you want to set up the equipment.

3. Test Plan

An NVD provides a framework of components and certifies the operational functionality of the integrated modular design. This detailed test plan describes what Nutanix tested, along with the expected results, to certify that Nutanix has tested the solution with the configuration specified and confirmed that it functions as designed.

Core Infrastructure Detailed Test Plan

Table 29: Core Infrastructure Tests

Summary	Validation Tasks	Expected Result
Verify app integration.	Launch Prism Element.	Interfaces open without error.
Verify storage container creation.	Create a new storage container.	Storage container created with inline compression enabled by default.
Verify Files file server deployment.	Create a new Files file server.	File server created with proper capacity configuration, network configuration, and joined to domain.
Verify Files file share creation.	Create a new Files file share.	File share created for SMB as a distributed share with compression and ABE (Access Based Enumeration) enabled.
Verify Citrix DaaS connection to workload cluster.	Add a hosting connection for the workload cluster from Connections and Resources from Citrix DaaS management.	Hosting connection adds without error.

Summary	Validation Tasks	Expected Result
Verify Citrix DaaS connection to workload cluster using Test Connection.	Run Test Connection from Citrix DaaS management.	All tests pass.
Verify Citrix DaaS MCS machine catalog deployment.	Create an MCS single-session or multisession machine catalog using the previously created hosting connection from Citrix DaaS management.	Able to deploy MCS VDAs using Citrix DaaS.
Verify Citrix DaaS MCS delivery group deployment.	Create a delivery group using the previously created machine catalog and test power operations.	Able to power operate and register MCS VDAs using Citrix DaaS.
Verify Citrix DaaS MCS image update.	Update the previously created MCS machine catalog.	Machine catalog uses the updated image.
Verify Citrix DaaS MCS image rollback.	Roll back the previously updated MCS machine catalog.	Machine catalog uses the previous image.
Verify Citrix DaaS PVS streaming.	Create an initial PVS single-session or multisession image and test PVS streaming.	Able to create an image and stream from PVS.
Verify Citrix DaaS PVS machine catalog deployment.	Create a PVS single-session or multisession machine catalog using a previously created hosting connection from PVS management.	Able to deploy PVS VDAs using PVS.
Verify Citrix DaaS PVS delivery group deployment.	Create a delivery group using a previously created machine catalog and test power operations.	Able to power operate and register PVS VDAs using Citrix DaaS.

Summary	Validation Tasks	Expected Result
Verify Citrix DaaS PVS image update.	Update a previously created PVS machine catalog.	Machine catalog uses the updated image.
Verify Citrix DaaS PVS image rollback.	Roll back a previously updated PVS machine catalog.	Machine catalog uses the previous image.
Verify Citrix DaaS resource launch.	Launch a Citrix DaaS resource from Workspace.	Session launches from Workspace to previously created MCS or PVS machine.
Verify Files file share profile creation.	Review the Files file share for a user folder with FSLogix Profile Container.	FSLogix Profile Container is inside the user's directory in the Files file share.
Verify email receipt in primary mailbox.	Wait up to 24 hours for the first emails from Prism Element to arrive in the primary mailbox.	The primary mailbox receives alerts from Prism Element.
Verify email receipt in secondary mailbox.	Wait up to 24 hours for the first emails from Prism Element to arrive in the secondary mailbox.	The secondary mailbox receives alerts from Nutanix clusters (Prism Element).
Verify NCC report receipt in secondary mailbox.	Verify that the correct email address receives reports from daily NCC runs.	The secondary mailbox receives NCC reports.

4. Ordering

This bill of materials (BoM) reflects the validated and tested hardware, software, and services that Nutanix recommends to achieve the outcomes described here. Consider the following points when you build your orders:

- All software is based on core licensing whenever possible.
- Nutanix Xpert Services or an affiliated partner selected by Nutanix provides all services.
- Nutanix based the functional testing described in this document on NX series models with similar configurations to validate the interoperability of software and services.

Substitutions

- Nutanix recommends that you purchase the exact hardware configuration reflected in the BoM whenever possible. If a specific hardware configuration is unavailable, choose a similar option that meets or exceeds the recommended specification.
- You can make hardware substitutions to suit your preferences; however, such changes may result in a solution that doesn't follow the recommended Nutanix configuration.
- Avoid software product code substitutions except when:
  - You need different quantities to maintain software licensing compliance.
  - You prefer a higher license tier or support level for the same software product code.
- Adding any software or workloads that aren't specified in this design to the environment (including additional Nutanix products) may affect the validated density calculations and result in a solution that doesn't follow the recommended Nutanix configuration.

- Professional Services substitutions to accommodate customer preferences aren't possible.

Sizing Considerations

This NVD is based on a block-and-pod architecture. A pod consists of one 16-node workload cluster in a single datacenter. For this NVD, a pod consists of the following components:

- One 4-node management cluster.
- One 16-node workload cluster.

Once the number of nodes, VMs, or clusters exceeds the maximum specified for the solution, create a new pod with a new management cluster.

For smaller environments, you can downsize the workload clusters to 4, 8, or 12 nodes based on your capacity requirements, but don't change the hardware configuration or sizing associated with the management clusters.

Bill of Materials

The following tables show the BoMs for the primary datacenter management clusters and the primary datacenter workload clusters. Keep in mind that your deployment may include the following options:

- Two clusters—one management (4-node) and one workload (16-node)—and a single physical datacenter location.
- Citrix DaaS with PVS installation.
- Nutanix Xpert Services design workshops.
- Nutanix Files with FSLogix deployed on the workload cluster.
- Up to two template images with base optimizations.

Note: Installing Citrix DaaS with MCS rather than PVS does not change the BoM.

Table 30: Primary Datacenter Management Cluster: Hardware, Software, and Services

Product Code	Description	Quantity
	Hardware	
NX-3060-G7	Model: NX-3060-G8, 4-node configuration Type: Hybrid Hardware support: — Support level: Production — NRDK support: No — NR node support: No	1
	Per-Node Hardware Configuration	
	Processor: Intel Xeon-Gold 5520 (2.2 GHz/18-Core)	2
	Memory: 32 GB (3,200 MHz DDR4 RDIMM)	24
	HDD: 2 TB	4
	SSD: 1.92 TB	2
	Network adapter: 10 GbE, 2-port, SFP+ (Intel 82599ES)	1
	Software	
SW-VDI-1U-PRO-PRD	Subscription, Acropolis (AOS) License tier: Pro VDI Support level: Production	Software Enterprise VDI solution
	Cluster Install Services	
Not required	See Services BoM that includes installation for one workload tenant cluster as part of product code CNS-EUC-STR-SML	4

Table 31: Primary Datacenter Workload Cluster: Hardware, Software, and Services

Product Code	Description	Quantity
	Hardware	
NX-3155G-G8	Model: NX-3155G-G8, 1-node configuration Type: All flash Hardware support: — Support level: Production — NRDK support: No — NR node support: No	16
	Per-Node Hardware Configuration	
	Processor: Intel Xeon-Gold 6354 (3.0 GHz/18-Core)	2
	Memory: 32 GB (3,200 MHz DDR4 RDIMM)	32
	HDD: No HDD included	N/A
	SSD: 1.9 TB	6
	Network adapter: 25 GbE, 2-port (NVIDIA MCX512A-ACUT ConnectX-5)	1
	Software	
SW-VDI-1U-PRO-PRD	Subscription, Acropolis (AOS) License tier: Pro VDI Support level: Production	Software Enterprise VDI solution
	Cluster Install Services	
CNS-INF-A-SVC-DEP-STR	Xpert Services, HCI Cluster Deployment Starter	16

Professional Services

The following professional services allow Nutanix to implement this NVD as designed, built, and tested. These services are outcome-based, with fixed prices for the scope described by the services SKUs included in the BoM. See the

Xpert Services information available on Nutanix.com for more details on each of the SKUs included.

Table 32: Professional Services for Platform

Product Code	Description	Quantity
CNS-EUC-STR-SML	Xpert Services, EUC Modernization: Starter Small	1
CNS-EUC-A-SVC-UWS-STD	Xpert Services, EUC Nutanix Files for User Workspace Deployment	1
CNS-EUC-A-SVC-DEP-ENV	Xpert Services, EUC Advanced Environment Management Deployment	1

Appendix

Windows VM Performance Tuning

For performance tuning on Windows VMs:

- Run the Citrix Optimizer.
- Review the Citrix Windows 10 Optimization Guide.
- Review the Citrix Tech Zone Optimizations for Citrix Virtual Apps and Desktops Guide .

Software

Software versions used in this NVD:

AOS 5.20.3

AHV build 20201105.2244

Nutanix Files 3.8.1.3

Citrix DaaS VDA 7.1912 CU5

Citrix Provisioning 7.1912 CU5

Windows Server 2019: Build 17763.2686

Windows 10, version 21H2 build 19044.1466

FSLogix 2.9.7979.62170

References

1. Nutanix Hybrid Cloud Reference Architecture
2. Nutanix Hybrid Cloud Validated Desgin
3. Physical Networking

4. Citrix DaaS on Nutanix G8: MCS and Windows Desktop Reference Architecture
5. Citrix DaaS on Nutanix G8: PVS and Windows Desktop Reference Architecture
6. Citrix DaaS on Nutanix G8: MCS and Windows Server Reference Architecture
7. Citrix DaaS on Nutanix G8: PVS and Windows Server Reference Architecture
8. Nutanix Citrix DaaS Best Practices Guide
9. Nutanix Microsoft SQL Best Practices Guide
10. Nutanix Files Sizing Guide
11. Citrix DaaS Limits

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 social media @nutanix.

Citrix Ready

The Citrix Ready program makes it easy for customers to identify complementary products and solutions that can enhance Citrix environments. To earn its Citrix Ready designation, Nutanix has passed a series of tests established by Citrix to verify that Citrix DaaS runs effectively on Nutanix.

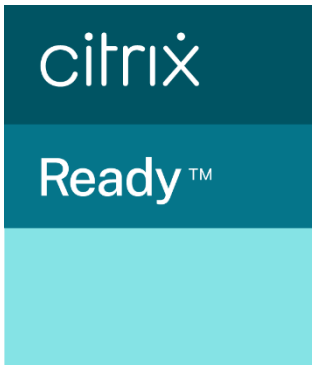

Figure 13: Citrix Ready Logo

List of Figures

List of Tables