v1.2.2 | August 2022 | NVD-2099

NUTANIX VALIDATED DESIGN

Hybrid Cloud: On-Premises 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 with two datacenters for high availability and disaster recovery.

Nutanix can deliver this Nutanix Validated Design (NVD), based on the Nutanix Hybrid Cloud Reference Architecture, as a bundled solution for general server virtualization that includes hardware, software, and services to accelerate and simplify the deployment and implementation process.

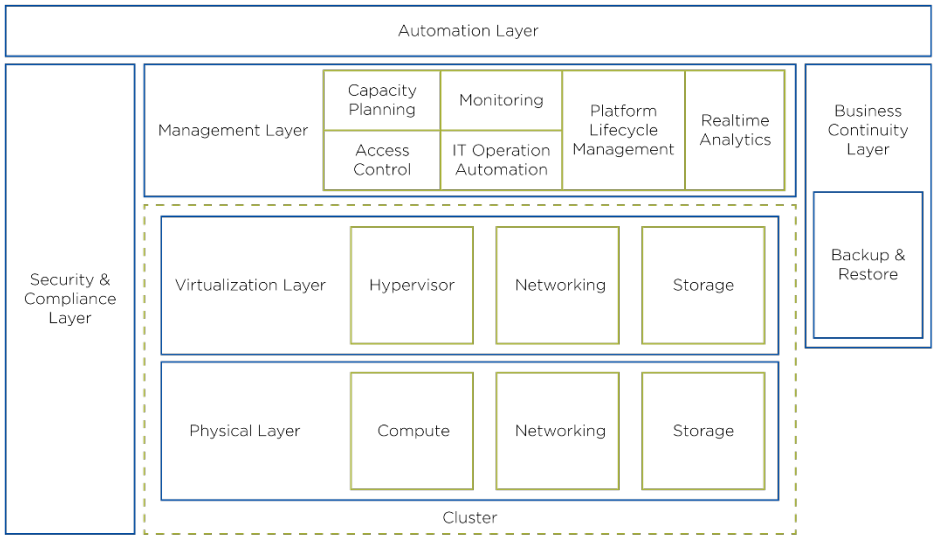

Figure 1: Architectural Layers of the Nutanix Validated Design

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

- Full-stack solution for hybrid cloud deployments that integrates multiple products including AOS, AHV, Prism Pro, Calm, Flow, Nutanix Disaster Recovery, Mine, and HYCU.
- Multidatacenter design built for failure tolerance and 99.999 percent availability.
- Active-active datacenters with two availability zones (AZs) run at 50 percent capacity to allow for full AZ failover in either direction.
- Tested for both planned and unplanned full-site failover with standardized business continuity and disaster recovery (BCDR) service levels.
- Self-service automation through the Calm MarketPlace includes blueprints for Windows, Linux, LAMP, and WISA as well as standardized VM sizes.
- Accelerates customer time-to-value and reduces risk.
- Orderable as a solution with a fully elaborated BOM for hardware, software, and services.
- Total cost of ownership (TCO) is 36 percent less than VxRail and 179 percent less than public cloud over five years.

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

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.
- Backup and disaster recovery for the Nutanix platform and hosted applications.
- Self-service automation with Calm and integration with third-party applications.
- Bill of materials.

Document History

Published	Notes
November 2021	Original publication.
November 2021	Updated the Executive Summary, Virtual Machine Design, Cluster Design, Security and Compliance, and Bill of Materials sections.
December 2021	Updated the Core Infrastructure Design, Backup and Disaster Recovery, and Bill of Materials sections.
March 2022	Updated the Cluster Design, Network Design, Management Components, Backup and Disaster Recovery, Test Plan, and Ordering sections.
May 2022	Updated to align with the Unified Storage and Disaster Recovery to Nutanix Cloud Clusters on AWS Nutanix Validated Designs.
August 2022	Updated the Cluster Conceptual Design and Network Microsegmentation sections and moved the Test Plan section to a separate document.

2. Core Infrastructure Design

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

Core infrastructure design requirements by component:

- Management
  - Deploy a unified management plane at the right scale to manage all clusters and workloads in the environment.
  - Deploy unified management for the dedicated management cluster at each datacenter (dual Prism Central per pod).
  - Configure management to integrate with Active Directory for authentication.
  - Use Active Directory–based groups for access control.
- Virtual Machines
  - Support at least three VM sizes: small, medium, and large.
  - Support Windows Server 2019 and Red Hat Enterprise Linux (RHEL) 8 as VM operating systems.
  - Limit virtual CPU overcommitment to 4:1, or 4 vCPU per physical CPU core.

- Monitoring
  - Enable platform fault monitoring and use email to send alerts.
  - Monitor performance metrics and store historical data for the past 12 months.
  - Keep resource usage under 75 percent; usage over 75 percent generates an email alert.
  - Monitor resources critical to Nutanix AOS operations (for example, CPU, memory, storage, and network resources); resource usage that exceeds configured limits generates an alert.
  - For resources that have high availability reservations, measure the resource utilization threshold against the usable capacity after subtracting the capacity reserved for high availability.
  - Monitor all network links (including host-switch and switch-switch) for bandwidth utilization and store historical data for the past 12 months.
  - Use email as the primary channel for event monitoring alerts.
  - Ensure that event monitoring is resilient. For example, when the management plane is the primary source of alerts, there must be a secondary method for monitoring the management plane itself. Then, if the management plane fails, an alert from the secondary source can trigger the action to recover the management plane.
  - Facilitate automated issue discovery and remote diagnostics.

Core infrastructure design assumptions by component:

- Clusters
  - The maximum number of VMs per workload cluster is 1,860 (124 per usable node).

- Monitoring
  - IT operations teams can continuously staff the mailbox that receives monitoring alerts to address critical issues in a timely manner.
  - IT operations teams can provide email infrastructure with sufficient resilience to send, receive, and access emails even during critical outages.
  - Network security appliances allow the management plane to transmit telemetry data to Nutanix.
- Infrastructure
  - IT operations teams can deploy Active Directory and DNS in a highly available configuration in each management cluster.

Core infrastructure design risks by component:

- Monitoring
  - If Prism Central becomes unavailable for any reason, the platform can no longer send alerts. To mitigate this risk, configure each Prism Element instance to send alerts as well. As this approach results in duplicate alerts during normal operations, send Prism Element alerts to a different mailbox that you can monitor when Prism Central is unavailable.

Core infrastructure design constraints by component:

- Clusters
  - The number of VMs per pod doesn't exceed 7,500 (the limit of Flow policies per Prism Central instance). Monitoring
  - SMTP is an available channel in the environment that can receive event monitoring alerts. Syslog captures logs but doesn't generate alerts on events.

Core Infrastructure Conceptual Design

The conceptual pod design has the following features:

- Two active-active datacenters in separate availability zones (AZs) with less than 5 ms of latency between sites.

- A small management cluster in each AZ that hosts services such as Prism Central and Active Directory.
- An instance of Prism Central hosted in the management cluster of each AZ (dual Prism Central deployment per pod).
- A workload cluster in each AZ that hosts the production workloads.
- A backup cluster in each AZ, replicated between sites for disaster recovery.

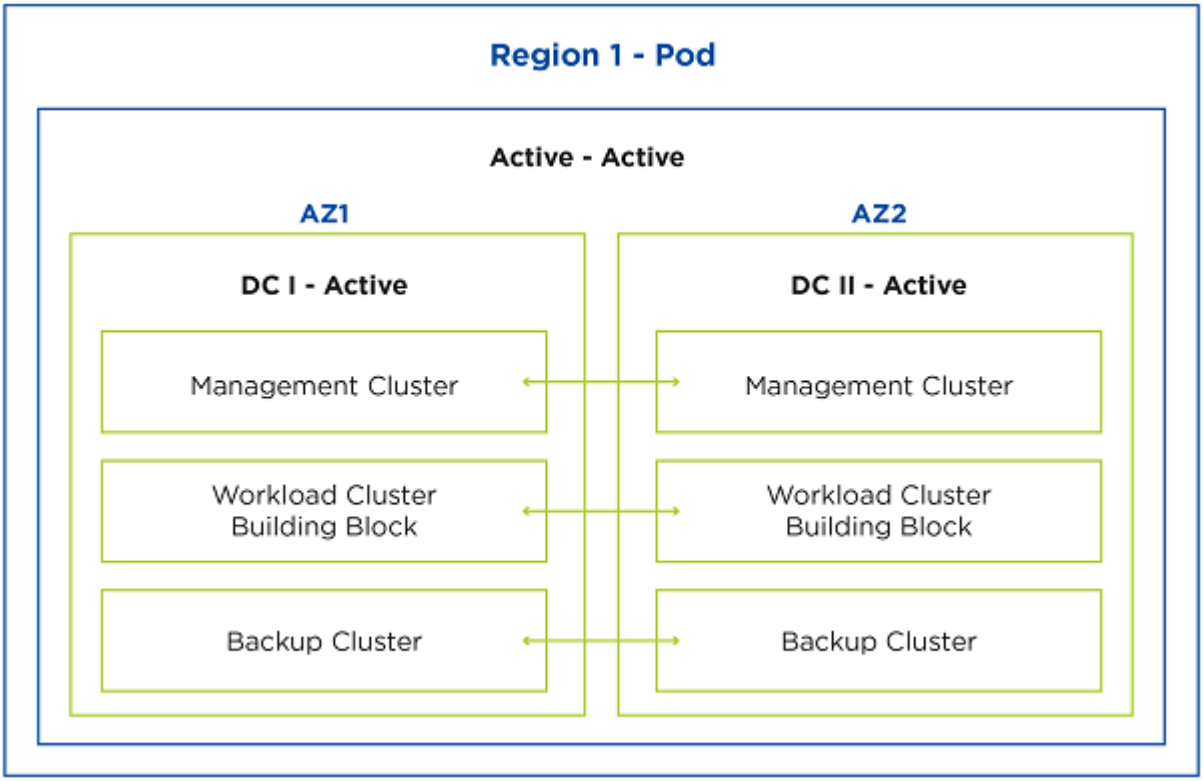

Figure 2: Conceptual Pod Design

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 AZ across two AZs. 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 the infrastructure changes in one AZ, you must upgrade the other AZ accordingly to ensure that a failover can complete successfully.

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.

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 might introduce further constraints; refer to the Datacenter Infrastructure section for more information.

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

Table: Scalability Design Decisions

Decision Name	Decision
Node memory population	Fully populate node memory
Node drive type	Use all-flash drives
Node drive population	Don't fully populate nodes with disk drives
Single rack	Use one rack per AZ
Establish scalability boundaries	Use X-Ray to confirm load per node
Rack availability use	Don't use rack availability

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 Linux LAMP all-in-one VMs could theoretically not hold more than 48 VMs, assuming you could use 100 percent of the available memory for VMs.

Table: Configuration Maximums or Maximum System Values

Entity	Decision
VMs or volume groups	Asynchronous disaster recovery: 200 VMs or volume groups for each protection domain or consistency group
vDisks (including snapshosts)	600,000
Calm node profile	Refer to the Self-Service with Automation Software Versions table
Flow	Refer to the Nutanix Flow Networking Guide.

For more information on Flow, see the Nutanix Flow Networking Guide.

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.

X-Ray test scenarios establish resilience boundaries for various failure scenarios.

Table: Resilience Design Decisions

Decision Name	Decision
Full redundancy of all components	Ensure the full redundancy of all components in the AZ
Established resilience boundaries	Use X-Ray to find resilience constraints

Virtual Machine Design

As the overall objective is to provide a hybrid cloud environment for general server 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 doesn't require any complex configuration.

Virtual Machine Standard Deployment Sizes

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

Table: Supported VM Configurations

VM Size	Small	Medium	Large
Virtual CPU	1	2	4
Virtual memory	8 GB	16 GB	32 GB
Virtual storage	50 GB	100 GB	200 GB
Virtual NIC	1	1	1
Virtual CD-ROM	1	1	1
Volume groups	No	No	No
Maximum VM instances per node	124	92	46

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

Windows Virtual Machines

All Windows VMs in this NVD are based on Windows Server 2019 Datacenter Edition. Windows VMs use the standard blueprints detailed in the following table when provisioned with Nutanix Calm.

Table: Standard Blueprints for Windows VMs

Template	WISA All-in-One	WISA Distributed	Standard Blueprint
Base template size	Large	Medium	Small

Template	WISA All-in-One	WISA Distributed	Standard Blueprint
Virtual CPU	4 per VM	2 per VM	1 per VM
Virtual memory	32 GB per VM	16 GB per VM	8 GB per VM
Virtual storage	200 GB per VM (VirtIO-SCSI)	100 GB per VM (VirtIO-SCSI)	50 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

Note: Flow policies require the kNormal NIC type to function correctly.

The WISA (Windows Server, Internet Information Services, Microsoft SQL Server, and ASP.NET) all-in-one blueprint installs and configures all necessary web, application, and database components when deployed through Nutanix Calm.

The WISA blueprint includes at least two load-balanced VMs for web servers, two load-balanced application servers, and one database server. Nutanix Calm provisions the individual VMs and installs their specific roles. The WISA distributed blueprint predefines and automatically applies Prism Central categories and Flow policies.

Refer to the appendix for Windows VM performance tuning recommendations.

Linux Virtual Machines

All Linux VMs in this NVD are based on Red Hat Enterprise Linux 8. Linux VMs use the standard blueprints detailed in the following table when provisioned with Nutanix Calm.

Table: Standard Blueprints for Linux VMs

Template	LAMP All-in-One	LAMP Distributed	Standard Blueprint
Base template size	Large	Medium	Small
Virtual CPU	4 per VM	2 per VM	1 per VM
Virtual memory	32 GB per VM	16 GB per VM	8 GB per VM

Template	LAMP All-in-One	LAMP Distributed	Standard Blueprint
Virtual storage	200 GB per VM (VirtIO-SCSI)	100 GB per VM (VirtIO-SCSI)	50 GB per VM (VirtIO-SCSI)
Virtual NIC	1 (VirtIO-Net)	1 (VirtIO-Net)	1 (VirtIO-Net)
Virtual CD-ROM	1	1	1

The LAMP (Linux, Apache, MySQL, and PHP) all-in-one blueprint has all necessary web, application, and database components preinstalled and ready to deploy on demand as a single VM through Nutanix Calm.

The LAMP distributed blueprint includes at least two load-balanced VMs for web servers, two load-balanced application servers, and one database server. Nutanix Calm provisions the individual VMs and installs their specific roles. The LAMP multi-VM blueprint predefines and automatically applies Nutanix Prism Pro categories and Flow policies.

Refer to the appendix for Linux VM performance tuning recommendations.

Cluster Design

This design incorporates three distinct cluster types:

1. Management: critical infrastructure and environment management workloads.
2. Workload: the building block for all general server virtualization workloads.
3. Backup: backup storage for the workload and management components.

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

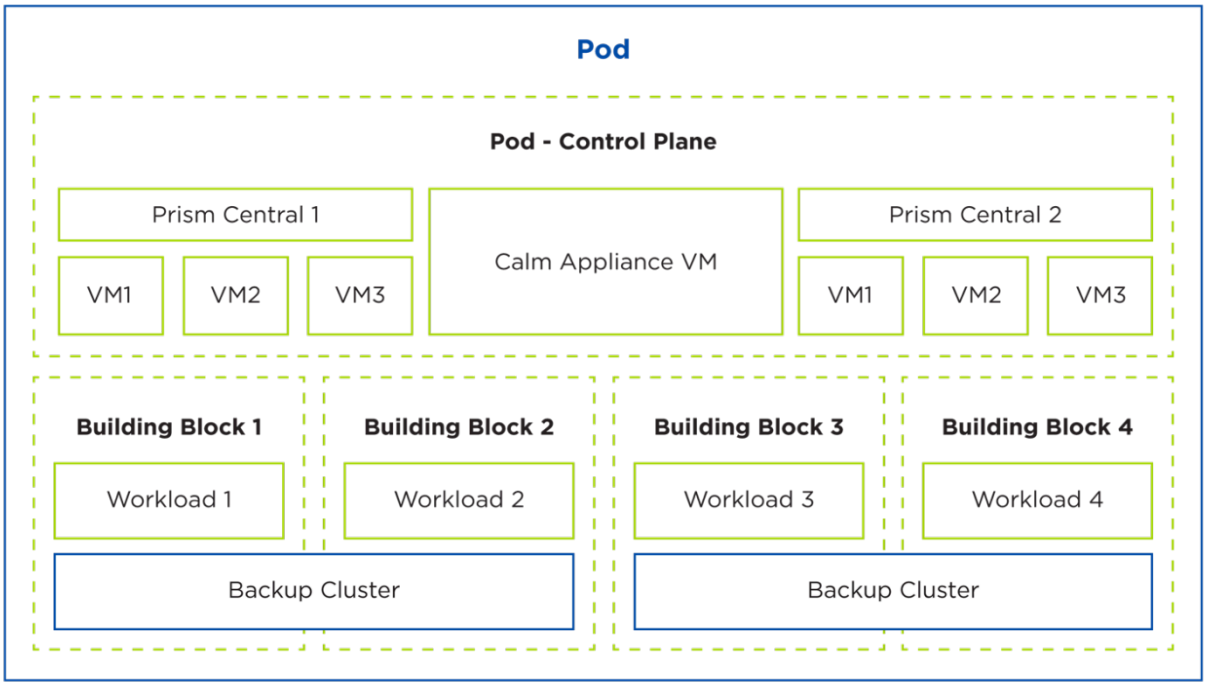

Figure 3: Pod and Building Block Architecture

Cluster Conceptual Design

This NVD solution uses one region with two separate AZs. Both AZs host active workloads, and each AZ provides a replication target for the other's workload cluster building blocks. Cloud-native applications that have built-in redundancy don't require infrastructure-level replication between AZs.

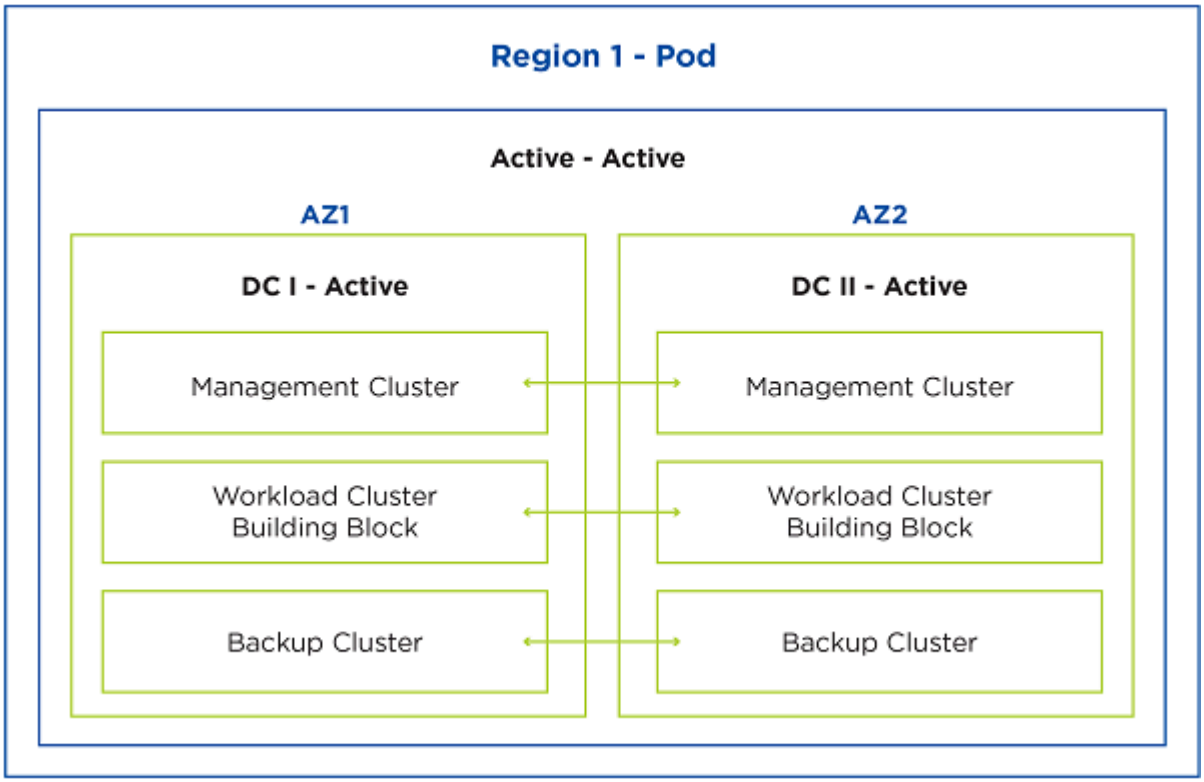

Figure 4: Conceptual Design

Table: Cluster Design Decisions

Decision Name	Decision
Number of regions	Use 1 region
Number of AZs	Use 2 AZs
Number of datacenters	Use 2 datacenters: 1 per AZ
Mixed workloads or dedicated workload per cluster	Mixed workloads per cluster, as this design is for general server virtualization
Minimum workload cluster building block size	Use at least 4 nodes
Workload cluster building block expansion increments	Use 4 nodes
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 8 workload cluster building blocks (4 per AZ) per pod (16 nodes per cluster)
Maximum number of running VMs per usable node in the workload cluster building block	Use at most 124 small VMs, 92 medium VMs, or 46 large VMs per usable node
Maximum number of VMs per workload cluster building block	Use at most 1,860 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 to allow n + 1
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
Percentage of deployed VMs supported during disaster recovery failover	Support 100 percent of deployed VMs
Maximum number of VMs deployed per workload cluster building block to allow for disaster recovery capacity	Deploy at most 930 small VMs per workload cluster building block
Maximum usable resource capacity per workload building block to allow for disaster recovery failover	Use at most 50 percent of the resource capacity

Platform Selection

Table: Platform Selection

Cluster	Management	Workload	Backup (Mine)
Node type	NX-1175S-G7	NX-3170-G8	NX-8155-G7

Cluster	Management	Workload	Backup (Mine)
Node count	4 (increments of 1)	4–16 per building block (increments of 4, up to 16 maximum)	4 (increments of 1)
Processor	1 Intel Xeon Gold 6226 12-core 125 W 2.7 GHz	2 Intel Xeon Gold 5318Y 24-core 165 W 2.1 GHz (Ice Lake)	2 Intel Xeon Silver 4214 12-core 85 W 2.2 GHz
RAM	6 x 64 GB 2,933 MHz DDR4 RDIMM (384 GB total)	12 x 128 GB 2,933 MHz DDR4 RDIMM (1.5 TB total)	4 x 32 GB 2,933 MHz DDR4 RDIMM (128 GB total)
SSD	2 x 1.92 TB	6 x 3.84 TB	2 x 3.84 TB
HDD	N/A	N/A	8 x 18 TB
NIC	10 GbE Dual SFP+	25 GbE Dual SFP+	25 GbE Dual SFP+
Form factor	1RU of single nodes	1RU of single nodes	2RU of single nodes
Support	3Y Production	3Y Production	3Y Production

Capacity Management

This NVD sizes the management and backup (Mine) clusters to host typical workloads as defined in the Management Components and Backup sections of this document. If those clusters need more resources, you can expand them 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 useable 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 4 nodes, up to the maximum. Based on the small VM specification, you can have a maximum of 1,860 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 4-node minimum, then can expand the new block in increments of 4 nodes 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 system starts 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 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 8-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.

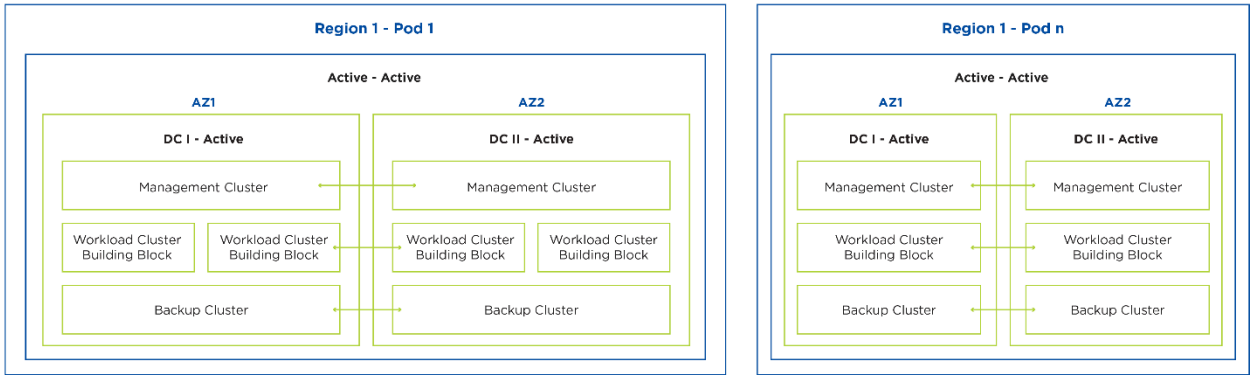

Figure 5: Scaling Beyond a Single Pod

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 50 percent of the running maximum to allow for disaster recovery capacity.

Table: Maximum Number of VMs

Scalability Consideration	Small VMs	Medium VMs	Large VMs
Maximum running VMs per workload cluster building block	1,860	1,380	690
Maximum running VMs per node	124	92	46
Maximum deployed VMs per workload cluster building block to allow for disaster recovery capacity	930	690	345

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.

Note: 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 6: 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 doesn't store workload vDisks.
- SelfServiceContainer: Used by the Nutanix Self-Service Portal and automation services like Calm.
- Default-Container-XXXX: Used by VMs to store vDisks for user VMs and applications.

In both datacenters, the management cluster uses the Default-Container to store VMs and their vDisks. Because this NVD provisions workloads from images with Calm, the SelfServiceContainer serves the workload and backup

clusters here. This NVD enables inline compression and erasure coding on the Default-Container and the SelfServiceContainer for all management, workload, and backup clusters in both datacenters. 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 and erasure coding with the default strip size on the container used for workloads, as the intended workload is general server virtualization.

The data reduction settings in the following table apply across both the primary and disaster recovery clusters.

Table: Data Reduction Settings

Container	Compression	Deduplication	Erasure Coding
Default-Container-XX	On	Off	On
NutanixManagement Share	On	Off	On
SelfService Container	On	Off	On

Table: Storage Design Decisions

Decision Name	Decision
Sizing a cluster	Use an all-flash cluster to provide enough usable SSD capacity to support the application's active data set
Node type vendors	Use all Nutanix NX nodes. Don't mix node types from different vendors in the same cluster
Node and disk types	Use identical node types that have similar disks
Sizing for node redundancy for storage and compute	Size all clusters for n + 1 failover capacity

Decision Name	Decision
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	Don't enable deduplication
Erasure coding	Enable erasure coding
Availability domain for workload cluster	Use block awareness
Availability domain for backup cluster	Use node 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

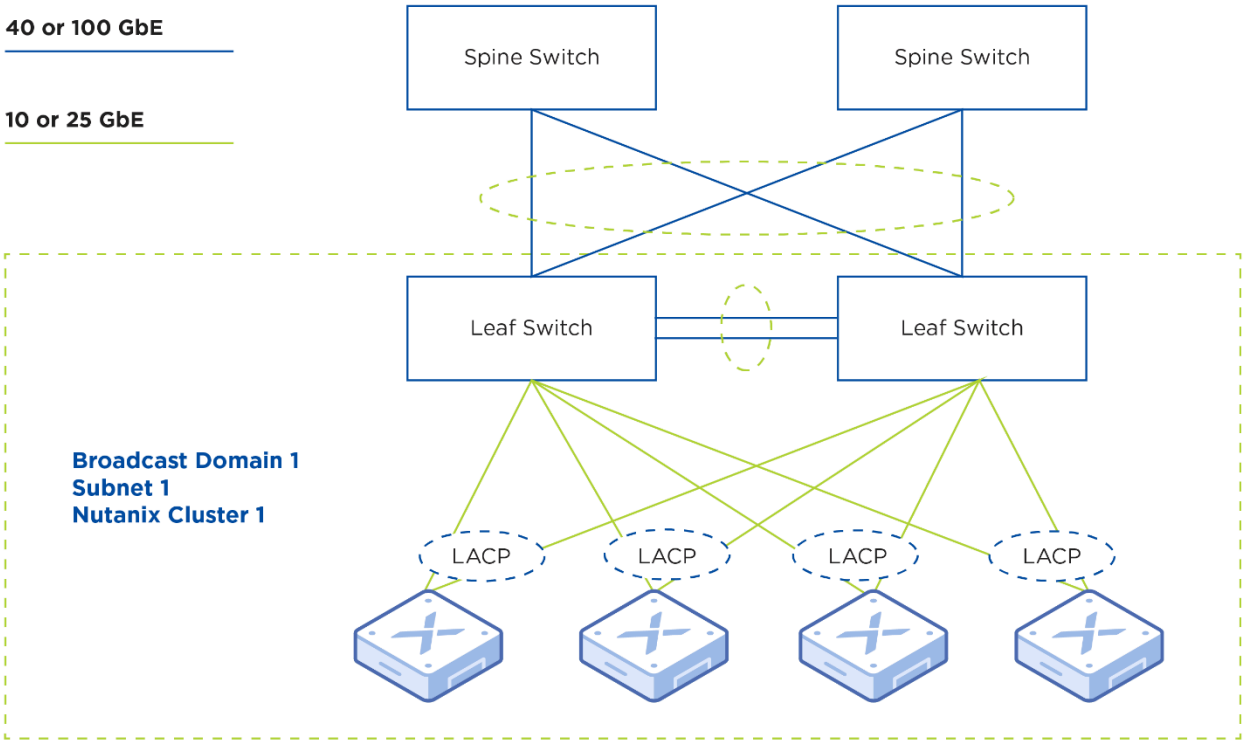

Figure 7: Physical Network Architecture

Table: 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 Gbps or faster top-of-rack switches	25 Gbps switches
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

Decision Name	Decision
Reduce network oversubscription to achieve as close to a 1:1 ratio as possible	1:2
Network design	Layer 2

Table: 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 or LACP load balancing policy	LACP
MTU size	1,500-byte MTU
Terminate L2/L3 networking	Spine

Table: Workload Cluster Networks

Decision Name	Decision
Shared infrastructure network subnet size	/24
VM network subnet size	/22
Number of addresses available per /24 network	254
Number of VM networks	4 per AZ

Decision Name	Decision
Present VM networks to other workload clusters	No
Stretch VM networks to secondary site	No
Number of addresses available per /22 network	1,024

Table: Management Cluster Networks

Decision Name	Decision
Shared infrastructure network subnet size	/24
VM network subnet size	/24
Number of addresses available per /24 network	254
Number of VM networks	1

Network Microsegmentation

Nutanix Flow enables VM- and application-based microsegmentation for traffic visibility and control. This NVD uses Flow to protect the environment from network attacks, create strict traffic controls that segment the network, and gain visibility into application network behavior.

Table: Flow Security Design Decisions

Decision Name	Decision
Prism Central multicluster Flow	Use two Prism Central instances (one per AZ) and replicate security policies between Prism Central instances with a script
VM scale per Prism Central	Limit VM scale to 7,500 per Prism Central instance
Isolation policies	Don't use isolation policies
Application inbound and outbound	Use inbound security policies and allow all outbound traffic

Decision Name	Decision
Category creation	Create a unique AppType category for each application and reuse AppTier categories
Category automation	Create and apply categories using Calm when you deploy applications
Address groups	Create address groups to define the corporate network for easy rule creation
Policy naming convention	Name each Security Policy as follows: <AZ number><policy name><policy number> (for example, AZ01AdProtection01).

Prism Central and Multicluster Design

A Nutanix Flow microsegmentation deployment covers all AHV hosts managed by a single Prism Central instance. Flow applies all categories and security policies uniformly across all clusters and VMs in this single Prism Central instance.

> Note: When two Prism Central instances exist—for example, for disaster recovery or scalability—the categories and policies don't replicate between them, so the designer and administrator need to create a system to either automatically or manually sync categories and policies between sites. This NVD uses a script to sync Flow policies and categories between Prism Central instances in different AZs.

A script runs periodically to sync policies between AZ01 and AZ02. For example, if a policy name starts with AZ01, the script replicates it to AZ02. Security policies for applications such as Active Directory or syslog that are unique per AZ and don't fail over with the AZ don't need to include the AZ01 string.

To enable Flow security synchronization between two Prism Central instances, follow the procedure in KB 12253.

Hypervisor Selection

Nutanix Flow relies on AHV to enforce policies in the hypervisor virtual switch. You can't protect VMs running on the ESXi or Hyper-V hypervisors with Nutanix Flow.

Number of VMs Protected

Nutanix Flow can secure fewer VMs than the maximum number of VMs that Prism Central can manage. Consider these scalability limits when you design clusters and Prism Central deployments. In addition, consider the maximum number of VMs protected in a single policy when you design the individual security policies.

Refer to the Nutanix Flow Microsegmentation Guide for detailed requirements and limitations.

Environment Isolation Requirements

In this NVD, all applications exist inside the same environment, so you don't need isolation policies.

Application Connectivity Requirements

The first step in identifying application connectivity is to define the scope of a single application. This application becomes the center of an application policy, and you can tag all VMs inside the application with the same AppType category (such as AppType: AZ01App1). Using the AZ number in the application name allows for easier identification when replicating across AZs. Next, identify the tiers in the application, such as web and database, and create an AppTier category for each of these tiers. AppTier categories don't need to be unique between AppTypes. Create the smallest number of application types and application tiers required to uniquely identify and group your applications.

Determine whether you need policy hit logs to track allowed and blocked connections. This NVD enables policy hit logs for all policies.

Each AppType category value is associated with a single application policy. For each application policy, determine the inbound traffic required to this application and whether the traffic is from another VM in the Nutanix environment or external. Next, determine the required traffic between tiers of the application and whether you should allow traffic within the same tier.

Finally, decide whether you should allow outbound traffic for this application. You can achieve good application security by strictly controlling the inbound side of the policy and allowing all traffic on the outbound side, but your situation may require stricter outbound traffic regulation. If you don't have a

physical north-south firewall available for this task, the Flow application policy can perform this function.

For all inbound traffic rules concerning VMs that exist in the Nutanix environment but aren't part of an existing AppType, create new top-level categories that you can use to add the relevant VMs to the policy as sources. For sources and destinations that don't exist in any Nutanix cluster, create an Addresses entity to group these networks and IP addresses for easy policy management.

This NVD creates address groups with the following addresses of corporate servers and clients to allow differentiated access for devices that don't run as AHV VMs. Replace these placeholders with addresses specific to your deployment.

Table: Address Groups

Name	Addresses	Purpose
AddrCorpAll	10.0.0.0/8	Identify all corporate IP addresses
AddrCorpClient	10.50.0.0/16	Identify all IP addresses that belong to corporate client devices
AddrCorpServer	10.38.0.0/16	Identify all IP addresses that belong to corporate server devices
AddrCalmAZ01	10.38.100.10, 10.38.100.11, 10.38.100.12	Identify all Calm IP addresses in AZ01
AddrCalmAZ02	10.38.200.10, 10.38.200.11, 10.38.200.12	Identify all Calm IP addresses in AZ02

The following application security policies protect infrastructure VMs that run on AHV. This infrastructure is unique for each site, so you don't need to replicate the policy between Prism Central instances. Create these infrastructure policies in each Prism Central instance.

For more information on Active Directory, see Microsoft's Active Directory and Active Directory Domain Services Port Requirements article.

Table: Active Directory Application Security Policy InfraAD-001

Purpose	Source	Destination	Port / Protocol
Allow all corp to Active Directory	AddrCorpAll	AppType: ActiveDirectory	See Microsoft documentation
Allow Active Directory out	AppType: ActiveDirectory	Allow All	All

Global Policy Settings: Enable Policy Hit Log

Table: Syslog Application Security Policy InfraSyslog-001

Purpose	Source	Destination	Port / Protocol
Allow corp servers to syslog	AddrCorpServer	AppType: Syslog	UDP 6514, TCP 6514
Allow corp clients to syslog	AddrCorpClient	AppType: Syslog	TCP 9000, UDP 514, TCP 514
Allow syslog out	AppType: Syslog	Allow All	All

Global Policy Settings: Enable Policy Hit Logs

This NVD creates individual applications with a unique policy for each application. The following table provides an example security policy for a single application. Modify the name and specific addresses or categories based on the application you're protecting. The prefix AZ01 indicates that this policy protects VMs that run primarily in AZ01. Use the prefix AZ02 for any policy that protects an app that runs primarily in AZ02.

Table: Example Application Security Policy: AZ01-Example-001

Purpose	Source	Destination	Port / Protocol
Allow corp clients to web	AddrCorpClient	AppType: AZ01-Example-001; AppTier: Web	TCP 80, 443
Allow web to app	AppType: AZ01-Example-001; AppTier: Web	AppType: AZ01-Example-001; AppTier: App	TCP 8080

Purpose	Source	Destination	Port / Protocol
Allow app to DB	AppType: AZ01-Example-001; AppTier: App	AppType: AZ01-Example-001; AppTier: DB	TCP 3306
Allow example app out	AppType: AZ01-Example-001	Allow All	All
Allow Calm to manage app	AddrCalmAZ01	AppType: AZ01-Example-001, all tiers	TCP 22, 5985–5986

Global Policy Settings: Enable Policy Hit Logs

This NVD modifies the forensic quarantine policy to allow quarantine of specific VMs while also allowing access from the security operations team. VMs owned by the security operations team for the explicit purpose of digital forensics and incident response have the category Security: DFIR. Update this policy in both Prism Central instances.

Table: Quarantine Security Policy

Purpose	Source	Destination	Port / Protocol
Allow security VMs to investigate	Security: DFIR	Forensic: Quarantine	All
Block all quarantine outbound	Forensic: Quarantine	None	None

Global Policy Settings: Enable Policy Hit Logs

Category Automation

Applying categories to VMs is a critical component of Flow security and automating this task is a great way to ensure a secure-by-default design. Tools such as Calm can automatically create VMs with the desired categories based on a blueprint, but you can also use external automation with our APIs.

This NVD creates application VMs through Calm with the appropriate AppType and AppTier categories assigned.

When disaster recovery replicates VMs to another site, the categories replicate as well.

Management Components

Management components such as Prism Central, Active Directory, DNS, and NTP are critical services that must be highly available. Prism Central is the global control plane for Nutanix, responsible for VM management, replication, application orchestration (through Calm), microsegmentation (through Flow), and other monitoring and analytics functions. You can deploy Prism Central in either in a single-VM or scale-out (three-VM) configuration.

When you design your management components, decide how many Prism Central instances you need. This NVD uses a scale-out Prism Central instance in each AZ, for a total of two Prism Central instances. This setup provides better scalability and increased disaster recovery functionality when you use additional Nutanix portfolio products such as Flow, Calm, and Objects.

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 and include scale-out Prism Central instances in both AZs. The management clusters only run core infrastructure management components, not general user VM workloads.

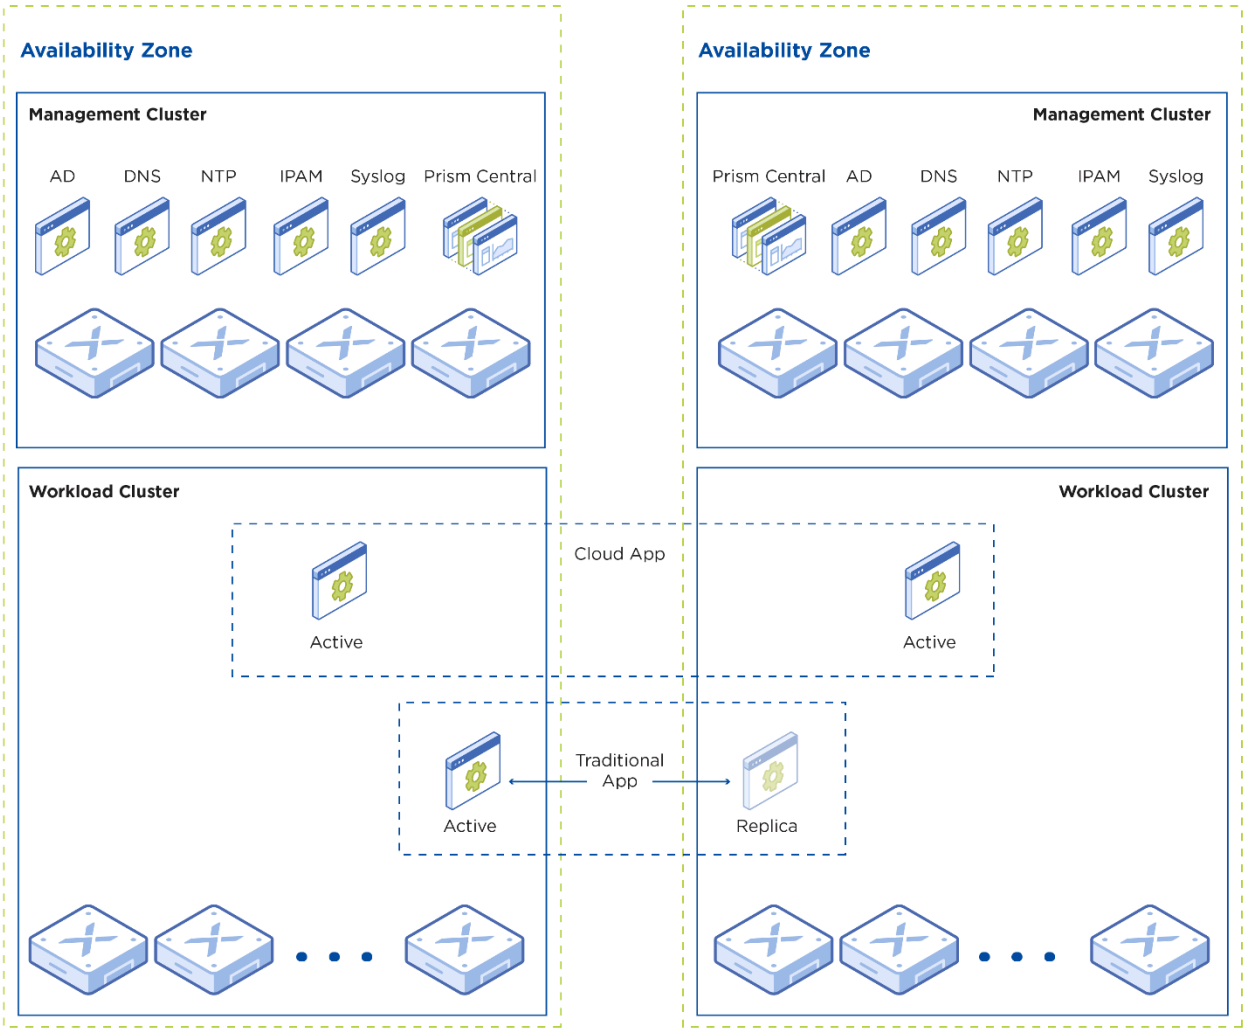

Figure 8: Management Plane

Management Detailed Design

In this NVD, management clusters run AOS 6.0 and workload clusters run AOS 5.20. AOS 6.0 has several Prism Central disaster recovery enhancements, specifically when paired with Prism Central pc.2021.7. Although mixing short-term service (STS) and long-term service (LTS) versions of AOS in a production environment can add some operational complexity, the significant feature gains for Prism Central disaster recovery outweigh the potential downsides. In future

iterations of this validated design, Nutanix plans to harmonize these versions where possible.

Table: Nutanix Management Component Software Versions

Component	Software Version
Prism Central	pc.2021.7
AOS	6.0 (STS)

Table: Management Component Design Decisions

Decision Name	Decision
Management cluster architecture	One management cluster in each AZ
Management cluster size	Four nodes (n + 1)
Management cluster node specifications	See the Platform Selection section
Deploy scale-out Prism Central for enhanced cluster management	Scale-out Prism Central (3 VMs)
Deploy Prism Central in each region or AZ using runbook disaster recovery automation	Deploy a scale-out Prism Central instance at both datacenters
Prism Central deployment size	Large: 3 VMs (each with 10 vCPU, 44 GB of RAM, and 2,500 GiB of storage)
Prism Central deployment locations	One in each AZ, deployed in each management cluster
Prism Central container name	Default container
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 different 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.

Monitoring Conceptual Design

In this NVD, Prism Central performs most of the event monitoring. Prism Central picks up events from the Nutanix clusters that it manages and forwards alerts, as defined by Nutanix Cluster Check (NCC). SMTP-based email alerts serve 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 Central go to a primary email alert recipient that's always monitored.

To cover situations where Prism Central might be unavailable, each Nutanix cluster in this NVD sends out notifications using SMTP as well. The individual Nutanix clusters send alerts to a different receiving mailbox that's only monitored when Prism Central isn't available.

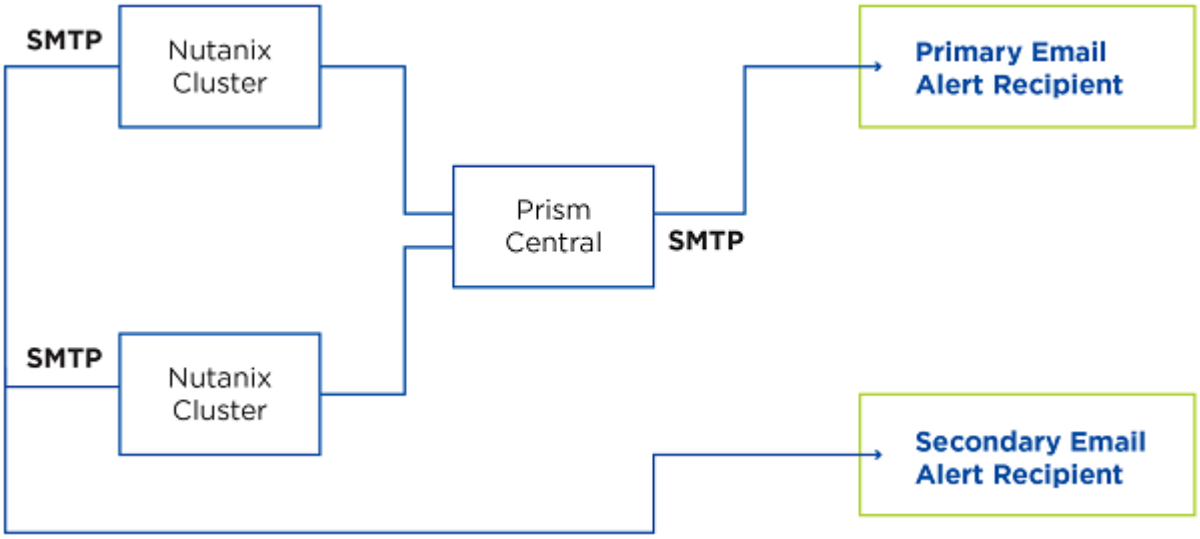

Figure 9: SMTP for Email Alerts from Prism Element and Prism Central

Prism Central monitors cluster performance in key areas such as CPU, memory, network, and storage utilization. Prism Central captures these metrics by default, so you don't need to do much configuration. When a Prism Central instance manages a cluster, Prism Central transmits all Pulse data, so it doesn't originate from individual clusters. When you enable Pulse, it detects known issues affecting cluster stability and automatically opens support cases.

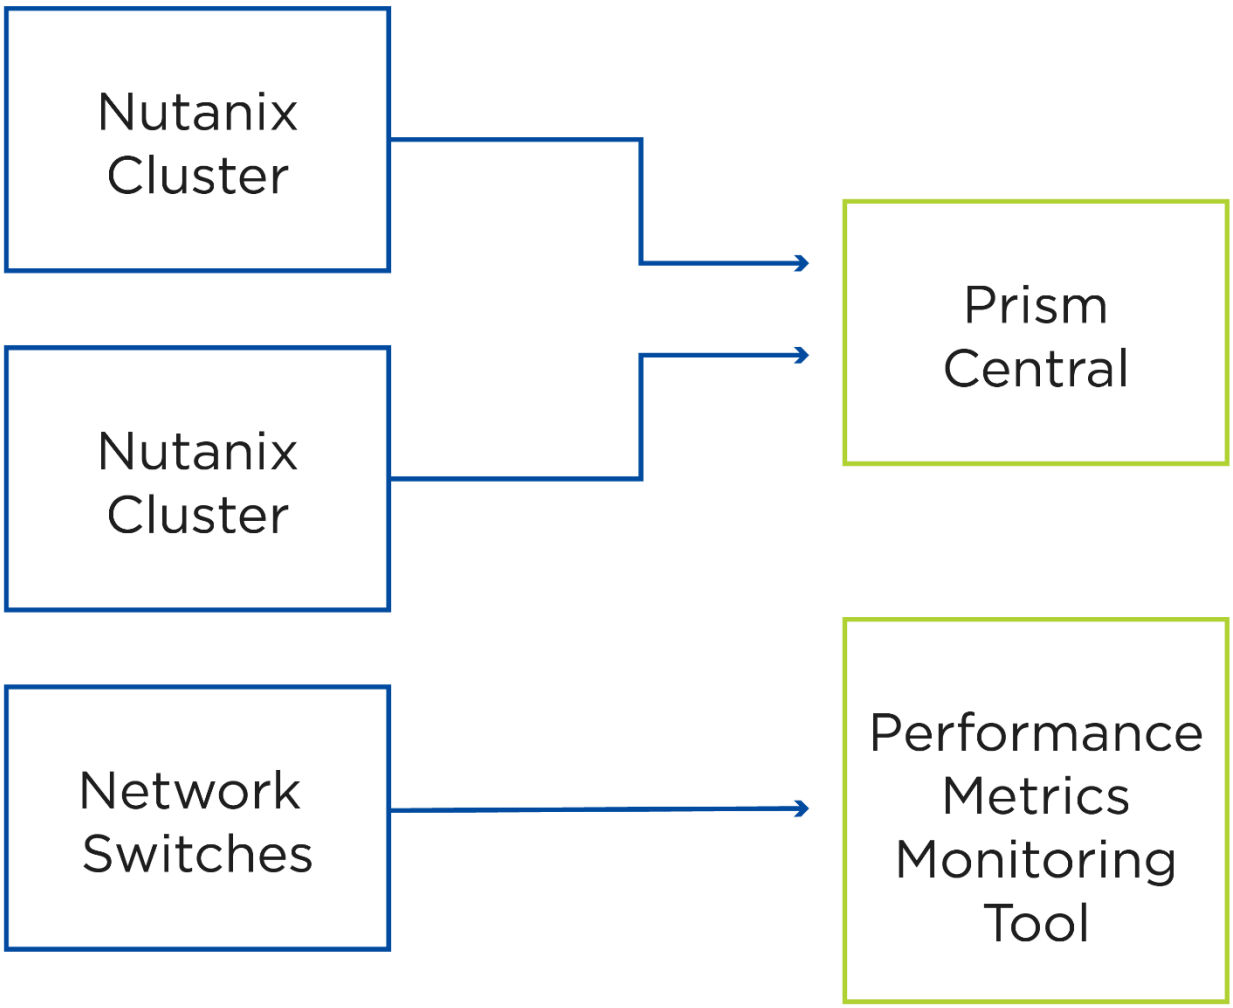

Figure 10: Systems Used to Capture Performance Metrics

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: Monitoring Design Decisions

Decision Name	Decision
Platform performance monitoring	Prism Central monitors Nutanix platform performance
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 5.20.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 and Prism Central
SMTP alerting source email address	Configure the source email address to be `clustername@nutanix.com` to uniquely identify the source of emails. For Prism Central, use the Prism Central host name in place of `clustername`
SMTP alerting Prism Central recipient email address	Configure the Prism Central recipient email address to be `primaryalerts@nutanix.com`
SMTP alerting Prism Element recipient email address	Configure the Prism Element recipient email address to be `secondaryalerts@nutanix.com`
NCC reports	Configure daily NCC reports to run at 6:00 AM local time and send them by email to the primary alerting mailbox
Nutanix Pulse	Configure Nutanix Pulse to send telemetry data back to Nutanix

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 for more information on the network-based security of hosted VMs using microsegmentation policies, and 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 clusters in both AZs. 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 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.

For more information on self-service and hosted VM access, refer to the Self-Service with Automation section.

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. Syslog coverage also extends to microsegmentation event logging from Prism Central with Flow.

> Note: This NVD assumes that the centralized syslog servers in each AZ can replicate log messages between sites, allowing for inspection in case the primary log system is unavailable.

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 nondisruptvely 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.

The 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: Security Design Decisions

Decision Name	Decision
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)
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: Security Configuration References

Configuration Target	Key:Value
Active Directory	AD-admin-group:ntnx-ctrl-admins
Syslog Server	infra-az-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.

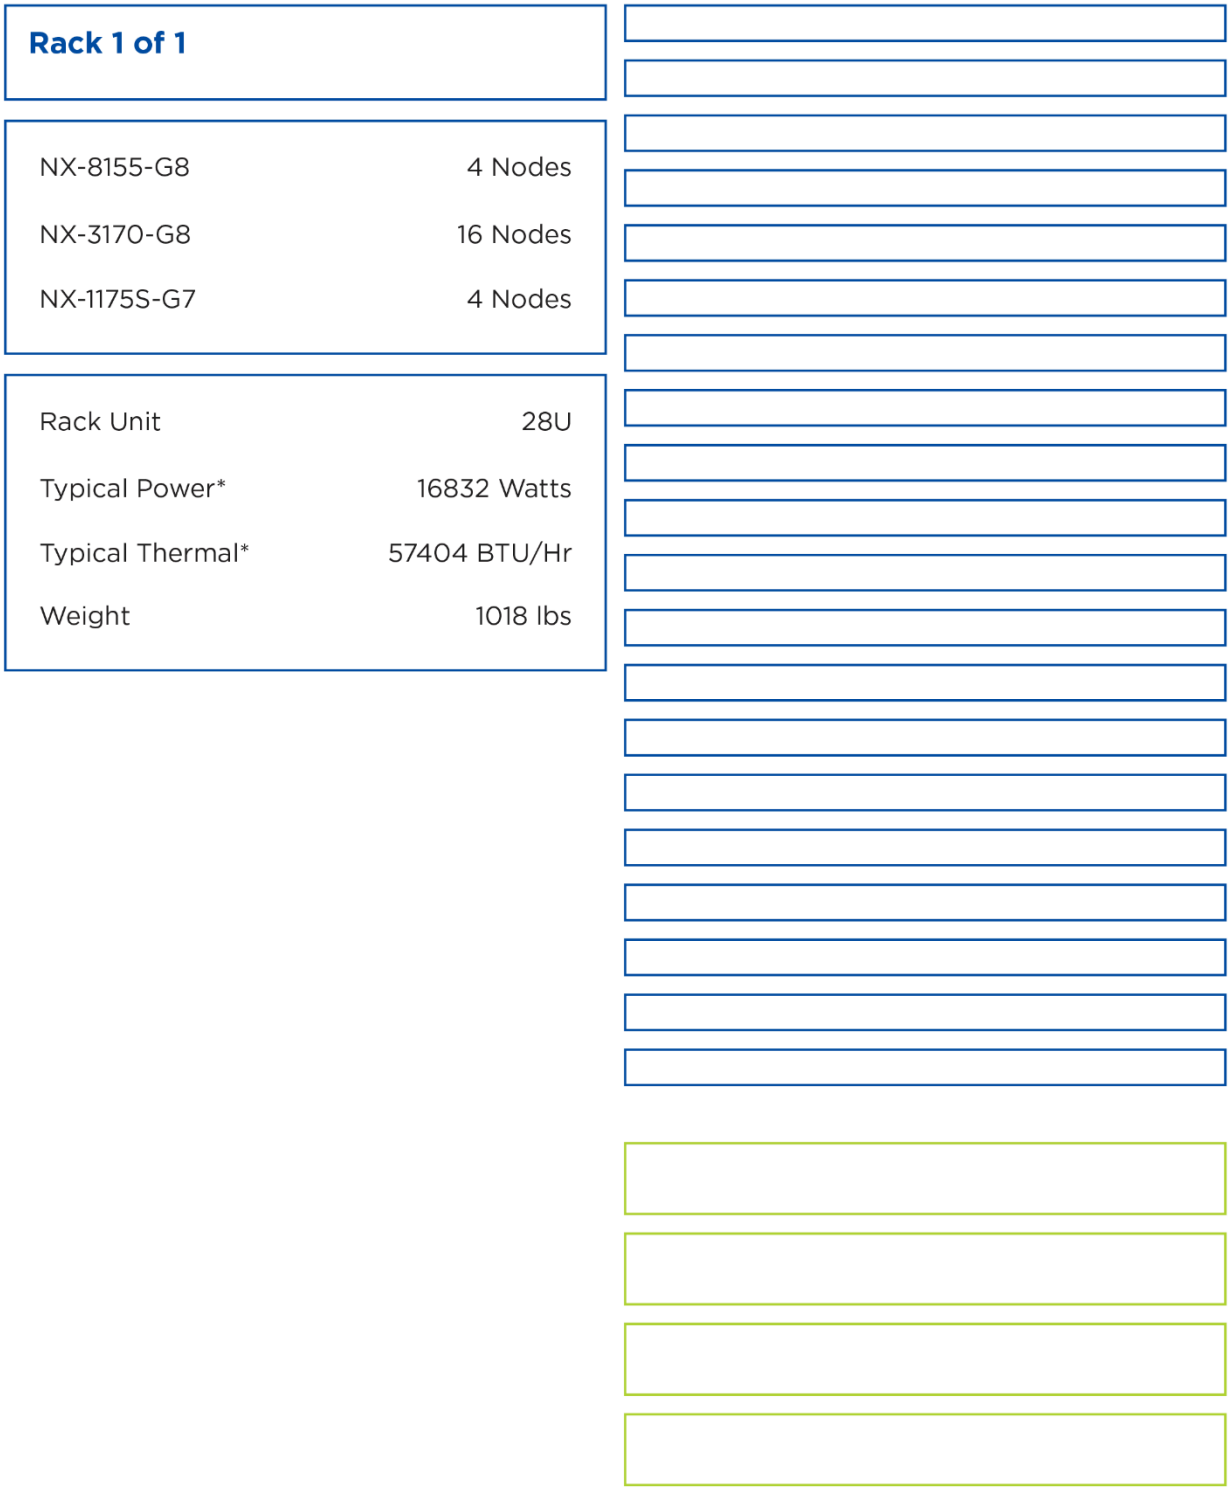

Figure 11: Rack Layout

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 28RU of systems with 3RU reserved for two data switches and one out-of-band switch, leaving 11RU of space available.

For network ports, the 24 nodes in this NVD consume 24 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 20 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. Backup and Disaster Recovery

This NVD uses Nutanix Disaster Recovery (on-premises) and Nutanix Mine to provide a BCDR solution to protect against different types of events. This section defines the overall high-level disaster recovery, backup, and backup storage designs.

For applications with native BCDR capabilities (for example, Microsoft SQL Always On availability groups), use the native disaster recovery resilience. For applications that lack this capability, use infrastructure BCDR. This design provides four levels of recovery point objective (RPO) for data protection:

- Gold Tier RPO: 0 minutes
- Silver Tier RPO: 15 minutes
- Bronze Tier RPO: 1 hour
- Recovery from backup RPO: 24 hours

The solution provides the following recovery time objective (RTO) levels:

- Gold Tier RTO: 2 hours
- Silver Tier RTO: 3 hours
- Bronze Tier RTO: 4 hours

To protect workloads against security threats like ransomware attacks, this NVD also provides protection to an external backup system.

NVD BCDR requirements:

- Use crash-consistent snapshots.
- Place workloads from different protection tiers into separate protection policies.
- Configure Nutanix snapshot schedules to retain the lowest number of snapshots while still meeting the retention policy.

- Provide an RPO between 0 min and 15 min for the application.
- Application requires an RPO of at least 1 hour.
- Subset of all applications require an RTO of 2 hours.
- Support full failover (including networking).
- Support automatic re-IP on workloads after failover.
- Provide maximum automation and orchestration for failover and failback.
- Provide VM-centric disaster recovery capabilities.
- Support disaster recovery testing without affecting production workloads.
- Simplify disaster recovery exercise, reducing human interaction to minimum during disaster recovery.
- Support the following disaster recovery events:
  - Datacenter outage.
  - Single cluster outage.
  - Ransomware attack.
  - Top-of-rack switch outage.
  - Single VLAN outage.
  - Human error.
  - Software bug.
  - Performance degradation caused by infrastructure (Nutanix cluster or network) or hardware components.
- Provide disaster recovery avoidance.
- Choose a backup vendor to use with Nutanix Mine.
- Use Nutanix Objects as an archival tier for backups.
- Choose a backup solution with native Nutanix API integration.

- Choose a backup solution that supports Nutanix Files backup and restore using API.
- Choose a backup solution that supports Nutanix Files file-level backup and restore.
- Choose a backup solution that supports S3-compatible storage as a backup target.
- Choose a target backup storage system that supports ransomware protection.
- Choose a target backup storage system that supports write once, read many (WORM).
- Choose a target backup storage system that supports file immutability.
- Choose a backup solution that supports replication to a secondary location.
- Choose a backup solution that supports archiving to S3-compatible storage.

Note: The customer must confirm every assumption in the following list.

NVD BCDR Assumptions:

- Disaster recovery avoidance causes minimal application and VM downtime.
- Customer provides redundant WAN connectivity between AZs.
- Customer provides WAN connectivity with sufficient bandwidth and latency (round-trip time (RTT) below 5 ms) to meet RPO requirements.
- Supporting infrastructure elements like DNS, Active Directory, and IPAM are available in both AZs.
- Solution doesn't provide partial failover capabilities.

Table: NVD BCDR Risks

Risk Description	Impact	Likelihood	Mitigation
Full outage of active AZ	Large	Unlikely	Fail over to remote AZ.

Risk Description	Impact	Likelihood	Mitigation
Full outage of management cluster	Medium	Unlikely	Fail management cluster over to remote AZ.
WAN link outage	Large	Unlikely	Provide redundant WAN connection.
Ransomware attack	Large	Likely	Implement backup solution with immutability. Replicate backup data between AZs.
Top-of-rack switch outage or misconfiguration	Large	Unlikely	Use two top-of-rack switches for redundancy.
Single Nutanix cluster outage	Medium	Unlikely	Replicate data and fail over to remote AZ.
Single VLAN outage or misconfiguration	Medium	Unlikely	Replicate data and fail over to remote AZ.
Human error	Large	Likely	Introduce automation. Replicate data and fail over to remote AZ.
Performance degradation caused by infrastructure or hardware components (Nutanix clusters, network)	Large	Unlikely	Replicate data and fail over to remote AZ.
Latency spikes above 5 ms in WAN cause application performance degradation	Large	Unlikely	Implement WAN monitoring to check latency on the link. Create SEV1 ticket for WAN latency spike events.

Table: NVD BCDR Constraints

Constraint Description	Comment
Use Nutanix Mine for backup	Currently limited to HYCU, Veeam, and Commvault.
Use Nutanix Disaster Recovery for disaster recovery orchestration	Nutanix Disaster Recovery is the solution of choice to provide disaster recovery orchestration.

Table: NVD BCDR Design Decisions

Decision Name	Decision
Define boundaries for recovery plans	Recovery plans don't span multiple protection policies
Use categories or VM names in recovery plans	Use categories in recovery plans so you can cover more VMs in each recovery plan (the maximum number of VMs in a recovery plan is 500 when you use categories versus 275 when you use VM names)
Use separate categories for different products	Disaster recovery and backup have separate categories
Choose disaster recovery orchestration product	Use Nutanix Disaster Recovery for disaster recovery orchestration, automation, and testing
Automate and orchestrate disaster recovery failover and disaster recovery testing	Use Nutanix Disaster Recovery to orchestrate disaster recovery
Provide solution to support an RPO of 15 min and an RPO of 1 hour	Use Nutanix Disaster Recovery with synchronous, NearSync, and asynchronous replication
Simplify disaster recovery management and VM placement	Use Nutanix categories to simplify VM disaster recovery and backup manageability
Determine the maximum number of entities for the protection policy (PP)	Asynchronous: 500 VMs; NearSync: 500 VMs; Synchronous: 200 VMs

Decision Name	Decision
Determine the maximum number of entities for the recovery plan (RP)	Asynchronous: 500 VMs; NearSync: 500 VMs; Synchronous: 200 VMs
Determine the PP-to-RP ratio	Keep a ratio of one PP to one RP (1:1)
Determine the Nutanix local and remote snapshot retention policies	Keep a maximum of 12 hours of snapshot history on Nutanix for both local and remote sites
Use dedicated network for failover	To simplify network management, use dedicated failover networks to accommodate VMs after failover
Protect 7,500 VMs in three protection tiers	Bronze: 4,000 VMs; Silver: 2,500 VMs; Gold: 1,000 VMs
Determine which backup product to use	Use Nutanix Mine with HYCU for applications
Determine the maximum number of VMs assigned to a single HYCU backup controller (BC)	Up to 1,500 VMs (for backup and restore) per HYCU VM (based on HYCU recommendations)
Back up workloads within AZs or across AZs	To optimize the backup window and save WAN bandwidth, Mine clusters back up workloads that are in the local AZ
Determine the RPO to set on backup policies	Set 24-hour RPO on backup policies
Determine how many backup policies to configure per HYCU BC	Single HYCU policy with up to 1,500 VMs
Determine which storage solution to use as a backup repository	Use Nutanix Objects as backup target
Determine how many S3 buckets to use as the backup repository	Use one object store with one bucket as the backup repository
Determine which advanced features to enable on S3 storage	Enable WORM and set it for 365 days
Determine which method to use to replicate backups between AZs	Use HYCU to manage backup replication

Backup and Disaster Recovery Conceptual Design

Nutanix Prism Central is the management and control plane for disaster recovery capabilities. Both disaster recovery and backup use categories to sort VMs into logical groups to automate their association with a protection policy, a recovery plan, and a backup policy.

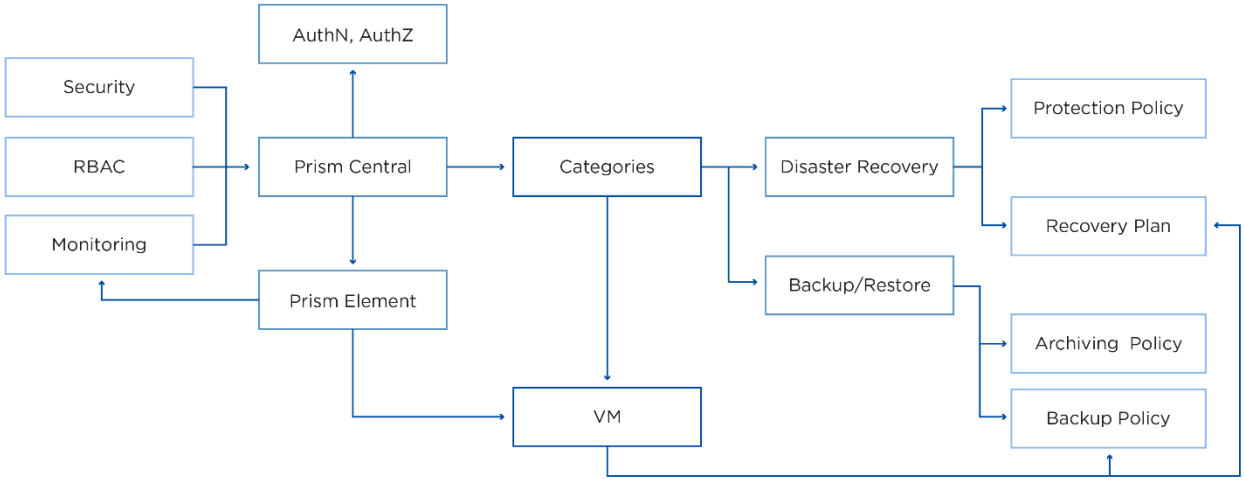

Figure 12: NVD BCDR Conceptual Design

Disaster Recovery

Disaster Recovery Logical Design

This NVD provides comprehensive disaster recovery protection for applications across both AZs in a single region. Applications can take advantage of underlying infrastructure to provide disaster recovery resilience based on three protection levels with bidirectional replication between AZs. The design provides granular disaster recovery to the single VM, IP address, or IP subnet level.

Disaster recovery testing, failover, and failback are fully orchestrated and require only minimal human involvement.

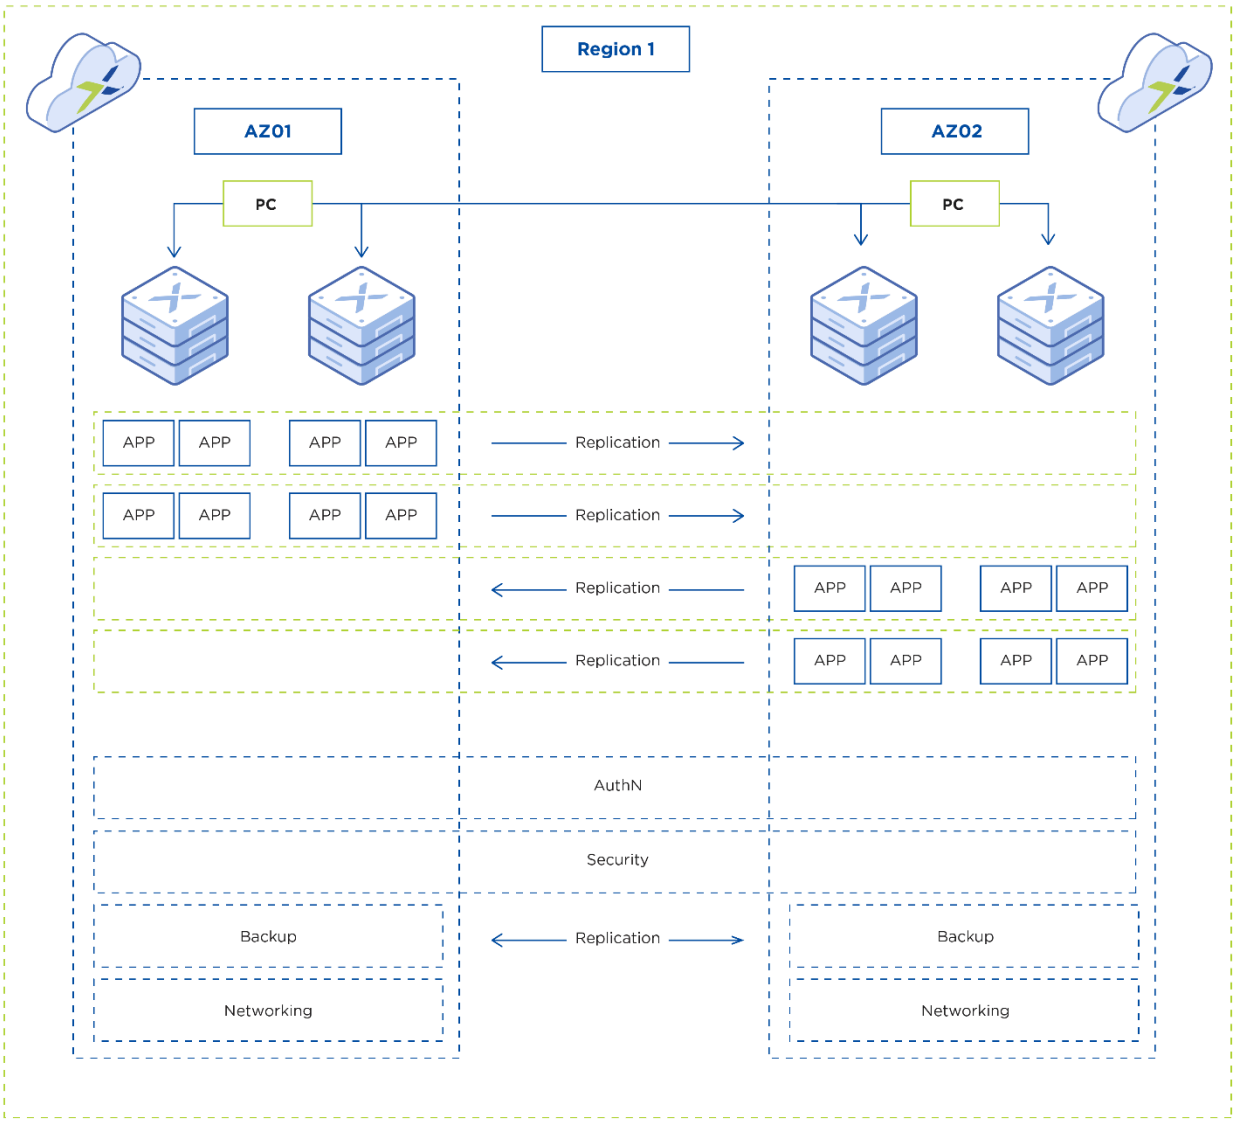

Figure 13: NVD BCDR Logical Diagram

Disaster Recovery Detailed Design

This NVD provides three protection tiers.

Table: NVD BCDR Design Decisions

Tier	RPO	RTO
Gold	0	2 hours

Tier	RPO	RTO
Silver	15 minutes	3 hours
Bronze	1 hour	4 hours

The BCDR section of this NVD uses the following software versions.

Table: Software Versions for Disaster Recovery

Component	Software Version
Prism Central	pc.2021.7
AOS	5.20.1.1 (LTS)

This NVD uses categories in Prism Central to automate VM placement in the target protection policy. To simplify failover and failback, the design assigns VMs to a local category (for example, it assigns VMs that run on AZ02 to a category with the prefix `AZ02`). Nutanix Disaster Recovery categories present three data protection levels:

- Bronze: RPO = 1 hour
- Silver: RPO = 15 minutes
- Gold: RPO = 0 minutes

The following table provides guidance on how to design Nutanix Disaster Recovery categories for 7,500 VMs.

Table: Nutanix Disaster Recovery Categories

Tier	Category Name	Value	Max # of VMs
Asynchronous	AZ01-DR-Bronze-01	RPO1h	500
Asynchronous	AZ01-DR-Bronze-02	RPO1h	500
Asynchronous	AZ01-DR-Bronze-03	RPO1h	500
Asynchronous	AZ01-DR-Bronze-04	RPO1h	500
Asynchronous	AZ02-DR-Bronze-01	RPO1h	500
Asynchronous	AZ02-DR-Bronze-02	RPO1h	500

Tier	Category Name	Value	Max # of VMs
Asynchronous	AZ02-DR-Bronze-03	RPO1h	500
Asynchronous	AZ02-DR-Bronze-04	RPO1h	500
NearSync	AZ01-DR-Silver-01	RPO15m	500
NearSync	AZ01-DR-Silver-02	RPO15m	500
NearSync	AZ01-DR-Silver-03	RPO15m	500
NearSync	AZ02-DR-Silver-01	RPO15m	500
NearSync	AZ02-DR-Silver-02	RPO15m	500
NearSync	AZ02-DR-Silver-03	RPO15m	500
Synchronous	AZ01-DR-Gold-01	RPOZero	200
Synchronous	AZ01-DR-Gold-02	RPOZero	200
Synchronous	AZ01-DR-Gold-03	RPOZero	100
Synchronous	AZ02-DR-Gold-01	RPOZero	200
Synchronous	AZ02-DR-Gold-02	RPOZero	200
Synchronous	AZ02-DR-Gold-03	RPOZero	100

The following two tables provide details on protection policy configuration for 7,500 VMs. Each protection policy has VMs located on a single AZ.

Table: Protection Policy Configuration for AZ01

Policy Name	Category	# of VMs	Source Cluster	Target Cluster	RPO
AZ01-AZ02-Bronze-01	AZ01-DR-Bronze-01	500	AZ01-CLS-0X	AZ02-CLS-0X	1 hour
AZ01-AZ02-Bronze-02	AZ01-DR-Bronze-02	500	AZ01-CLS-0X	AZ02-CLS-0X	1 hour
AZ01-AZ02-Bronze-03	AZ01-DR-Bronze-03	500	AZ01-CLS-0X	AZ02-CLS-0X	1 hour
AZ01-AZ02-Bronze-04	AZ01-DR-Bronze-04	500	AZ01-CLS-0X	AZ02-CLS-0X	1 hour

Policy Name	Category	# of VMs	Source Cluster	Target Cluster	RPO
AZ01-AZ02-Silver-01	AZ01-DR-Silver-01	500	AZ01-CLS-0X	AZ02-CLS-0X	15 minutes
AZ01-AZ02-Silver-02	AZ01-DR-Silver-02	500	AZ01-CLS-0X	AZ02-CLS-0X	15 minutes
AZ01-AZ02-Silver-03	AZ01-DR-Silver-03	500	AZ01-CLS-0X	AZ02-CLS-0X	15 minutes
AZ01-AZ02-Gold-01	AZ01-DR-Gold-01	200	AZ01-CLS-0X	AZ02-CLS-0X	0 minutes
AZ01-AZ02-Gold-02	AZ01-DR-Gold-02	200	AZ01-CLS-0X	AZ02-CLS-0X	0 minutes
AZ01-AZ02-Gold-03	AZ01-DR-Gold-03	100	AZ01-CLS-0X	AZ02-CLS-0X	0 minutes

Table: Protection Policy Configuration for AZ02

Policy Name	Category	# of VMs	Source Cluster	Target Cluster	RPO
AZ02-AZ01-Bronze-01	AZ02-DR-Bronze-01	500	AZ02-CLS-0X	AZ01-CLS-0X	1 hour
AZ02-AZ01-Bronze-02	AZ02-DR-Bronze-02	500	AZ02-CLS-0X	AZ01-CLS-0X	1 hour
AZ02-AZ01-Bronze-03	AZ02-DR-Bronze-03	500	AZ02-CLS-0X	AZ01-CLS-0X	1 hour
AZ02-AZ01-Bronze-04	AZ02-DR-Bronze-04	500	AZ02-CLS-0X	AZ01-CLS-0X	1 hour
AZ02-AZ01-Silver-01	AZ02-DR-Silver-01	500	AZ02-CLS-0X	AZ01-CLS-0X	15 minutes
AZ02-AZ01-Silver-02	AZ02-DR-Silver-02	500	AZ02-CLS-0X	AZ01-CLS-0X	15 minutes
AZ02-AZ01-Silver-03	AZ02-DR-Silver-03	500	AZ02-CLS-0X	AZ01-CLS-0X	15 minutes
AZ02-AZ01-Gold-01	AZ02-DR-Gold-01	200	AZ02-CLS-0X	AZ01-CLS-0X	0 minutes

Policy Name	Category	# of VMs	Source Cluster	Target Cluster	RPO
AZ02-AZ01-Gold-02	AZ02-DR-Gold-02	200	AZ02-CLS-0X	AZ01-CLS-0X	0 minutes
AZ02-AZ01-Gold-03	AZ02-DR-Gold-03	100	AZ02-CLS-0X	AZ01-CLS-0X	0 minutes

The following two tables provide detailed information about recovery plans. To simplify failover and failback, the design assigns VMs to a recovery plan from the AZ. For example, VMs located in AZ01 are assigned to the recovery plan for AZ01.

Table: Details of Recovery Plans for AZ01 VMs

Name	Stage	VM Category	Delay	Source Network	Failover Networks	Test Failover Network	# of VMs
AZ01-AZ02-Bronze-01	Stage1	AZ01-DR-Bronze-01	0	Source-PG	Failover-PG	Test-PG	500
AZ01-AZ02-Bronze-02	Stage1	AZ01-DR-Bronze-02	0	Source-PG	Failover-PG	Test-PG	500
AZ01-AZ02-Bronze-03	Stage1	AZ01-DR-Bronze-03	0	Source-PG	Failover-PG	Test-PG	500
AZ01-AZ02-Bronze-04	Stage1	AZ01-DR-Bronze-04	0	Source-PG	Failover-PG	Test-PG	500
AZ01-AZ02-Silver-01	Stage1	AZ01-DR-Silver-01	0	Source-PG	Failover-PG	Test-PG	500
AZ01-AZ02-Silver-02	Stage1	AZ01-DR-Silver-02	0	Source-PG	Failover-PG	Test-PG	500

Name	Stage	VM Category	Delay	Source Network	Failover Networks	Test Failover Network	# of VMs
AZ01-AZ02-Silver-03	Stage1	AZ01-DR-Silver-03	0	Source-PG	Failover-PG	Test-PG	250
AZ01-AZ02-Gold-01	Stage1	AZ01-DR-Gold-01	0	Source-PG	Failover-PG	Test-PG	200
AZ01-AZ02-Gold-02	Stage1	AZ01-DR-Gold-02	0	Source-PG	Failover-PG	Test-PG	200
AZ01-AZ02-Gold-03	Stage1	AZ01-DR-Gold-03	0	Source-PG	Failover-PG	Test-PG	100

Table: Details of Recovery Plans for AZ02 VMs

Name	Stage	VM Category	Delay	Source Network	Failover Networks	Test Failover Network	# of VMs
AZ02-AZ01-Bronze-01	Stage1	AZ02-DR-Bronze-01	0	Source-PG	Failover-PG	Test-PG	500
AZ02-AZ01-Bronze-02	Stage1	AZ02-DR-Bronze-02	0	Source-PG	Failover-PG	Test-PG	500
AZ02-AZ01-Bronze-03	Stage1	AZ02-DR-Bronze-03	0	Source-PG	Failover-PG	Test-PG	500
AZ02-AZ01-Bronze-04	Stage1	AZ02-DR-Bronze-04	0	Source-PG	Failover-PG	Test-PG	500
AZ02-AZ01-Silver-01	Stage1	AZ02-DR-Silver-01	0	Source-PG	Failover-PG	Test-PG	500

Name	Stage	VM Category	Delay	Source Network	Failover Networks	Test Failover Network	# of VMs
AZ02-AZ01-Silver-02	Stage1	AZ02-DR-Silver-02	0	Source-PG	Failover-PG	Test-PG	500
AZ02-AZ01-Silver-03	Stage1	AZ02-DR-Silver-03	0	Source-PG	Failover-PG	Test-PG	250
AZ02-AZ01-Gold-01	Stage1	AZ02-DR-Gold-01	0	Source-PG	Failover-PG	Test-PG	200
AZ02-AZ01-Gold-02	Stage1	AZ02-DR-Gold-02	0	Source-PG	Failover-PG	Test-PG	200
AZ02-AZ01-Gold-03	Stage1	AZ02-DR-Gold-03	0	Source-PG	Failover-PG	Test-PG	100

The following two tables provide details about mapping between protection policies, recovery plans, and categories for 7,500 VMs.

Table: Protection Policy to Recovery Plan Mapping for AZ01

Policy Name	RP Name	Category Name	RPO	RTO	# of VMs
AZ01-AZ02-Bronze-01	AZ01-RP-Bronze-01	AZ01-DR-Bronze-01	1 hour	4 hours	500
AZ01-AZ02-Bronze-02	AZ01-RP-Bronze-02	AZ01-DR-Bronze-02	1 hour	4 hours	500
AZ01-AZ02-Bronze-03	AZ01-RP-Bronze-03	AZ01-DR-Bronze-03	1 hour	4 hours	500
AZ01-AZ02-Bronze-04	AZ01-RP-Bronze-04	AZ01-DR-Bronze-04	1 hour	4 hours	500
AZ01-AZ02-Silver-01	AZ01-DR-Silver-01	AZ01-DR-Silver-01	15 minutes	3 hours	500

Policy Name	RP Name	Category Name	RPO	RTO	# of VMs
AZ01-AZ02-Silver-02	AZ01-DR-Silver-02	AZ01-DR-Silver-02	15 minutes	3 hours	500
AZ01-AZ02-Silver-03	AZ01-DR-Silver-03	AZ01-DR-Silver-03	15 minutes	3 hours	250
AZ01-AZ02-Gold-01	AZ01-DR-Gold-01	AZ02-DR-Gold-01	0 minutes	2 hours	200
AZ01-AZ02-Gold-02	AZ01-DR-Gold-02	AZ02-DR-Gold-02	0 minutes	2 hours	200
AZ01-AZ02-Gold-03	AZ01-DR-Gold-03	AZ02-DR-Gold-03	0 minutes	2 hours	100

Table: Protection Policy to Recovery Plan Mapping for AZ02

Policy Name	RP Name	Category Name	RPO	RTO	# of VMs
AZ02-AZ01-Bronze-01	AZ02-RP-Bronze-01	AZ02-DR-Bronze-01	1 hour	4 hours	500
AZ02-AZ01-Bronze-02	AZ02-RP-Bronze-02	AZ02-DR-Bronze-02	1 hour	4 hours	500
AZ02-AZ01-Bronze-03	AZ02-RP-Bronze-03	AZ02-DR-Bronze-03	1 hour	4 hours	500
AZ02-AZ01-Bronze-04	AZ02-RP-Bronze-04	AZ02-DR-Bronze-04	1 hour	4 hours	500
AZ02-AZ01-Silver-01	AZ02-DR-Silver-01	AZ02-DR-Silver-01	15 minutes	3 hours	500
AZ02-AZ01-Silver-02	AZ02-DR-Silver-02	AZ02-DR-Silver-02	15 minutes	3 hours	500
AZ02-AZ01-Silver-03	AZ02-DR-Silver-03	AZ02-DR-Silver-03	15 minutes	3 hours	250
AZ02-AZ01-Gold-01	AZ02-DR-Gold-01	AZ02-DR-Gold-01	0 minutes	2 hours	200
AZ02-AZ01-Gold-02	AZ02-DR-Gold-02	AZ02-DR-Gold-02	0 minutes	2 hours	200

Policy Name	RP Name	Category Name	RPO	RTO	# of VMs
AZ02-AZ01-Gold-03	AZ02-DR-Gold-03	AZ02-DR-Gold-03	0 minutes	2 hours	100

Backup

Backup Logical Design

This NVD provides a backup option for workloads running in both AZs. To protect backup data against cluster failure and datacenter failure, data replicates bidirectionally between two backup instances across both AZs in one region. This design optimizes the backup solution to back up workloads that run locally to the backup cluster. Using categories helps organize VMs and ensures that the Nutanix Mine instance that's local to the AZ can back them up.

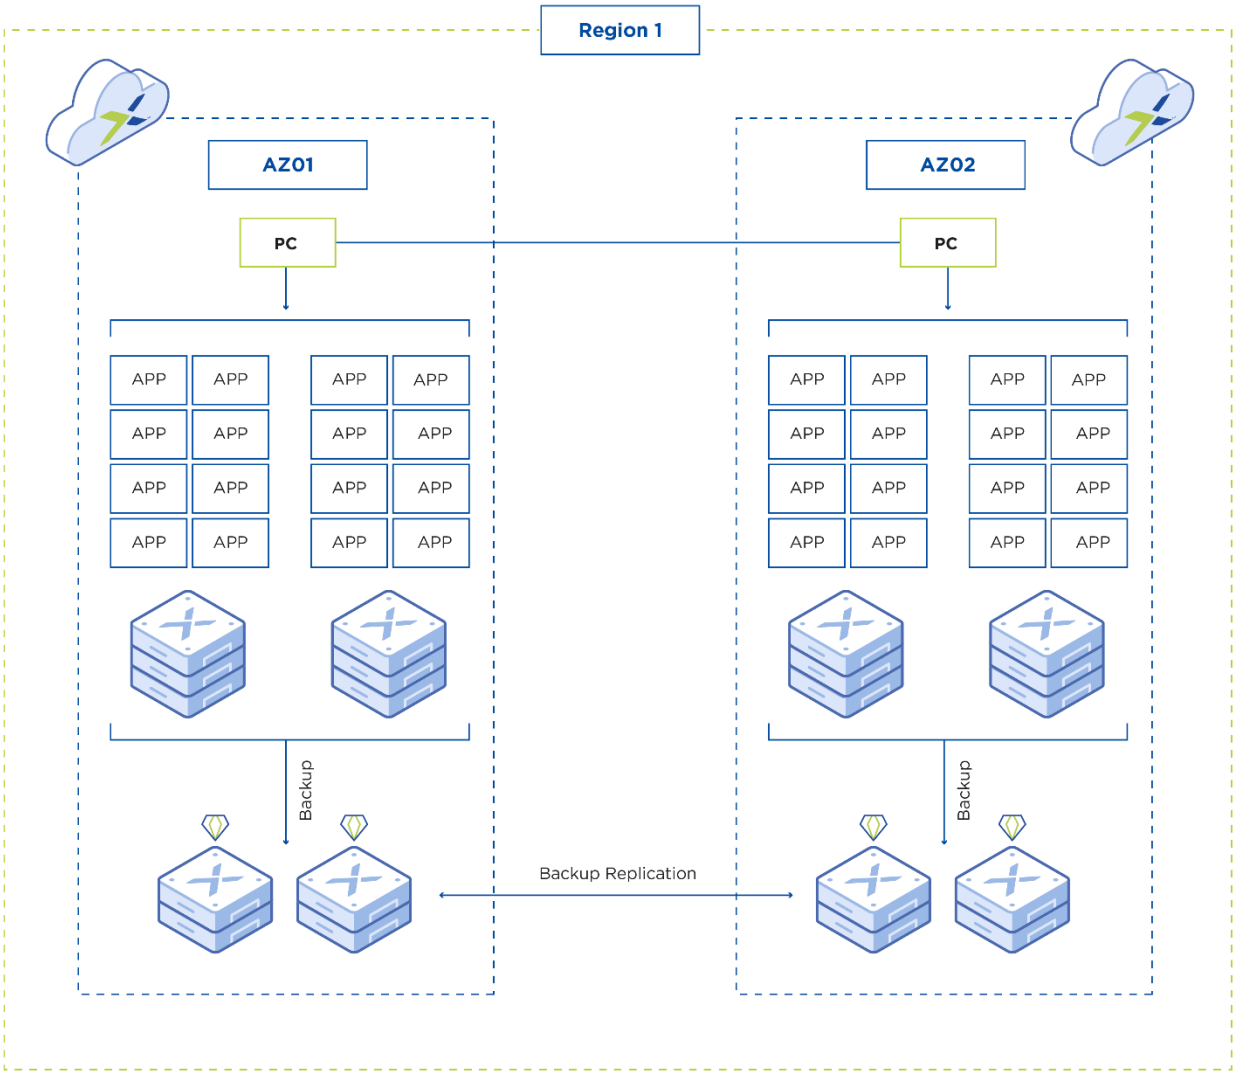

Figure 14: NVD Backup Architecture Logical Design

Nutanix Mine Logical Design

To provide additional protection against data loss, this NVD has an external backup system: Nutanix Mine. Each datacenter contains an instance of the backup system, local to the workloads you want to back up and restore.

To provide maximum performance and the desired RPO and RTO across the environment, each Nutanix Mine setup has multiple backup proxies. To simplify backup policy management, there is a 1:1 mapping between the backup policy

and the backup proxy. This approach helps scale the solution linearly as it grows.

For maximum performance, all backup components use the same network subnet:

- CVM
- AHV
- Object networks (storage and client)
- HYCU backup VMs

Nutanix Mine is a high-performance backup target that's compatible with S3 storage. S3-compatible storage provides advanced security features to help protect data against common security threats.

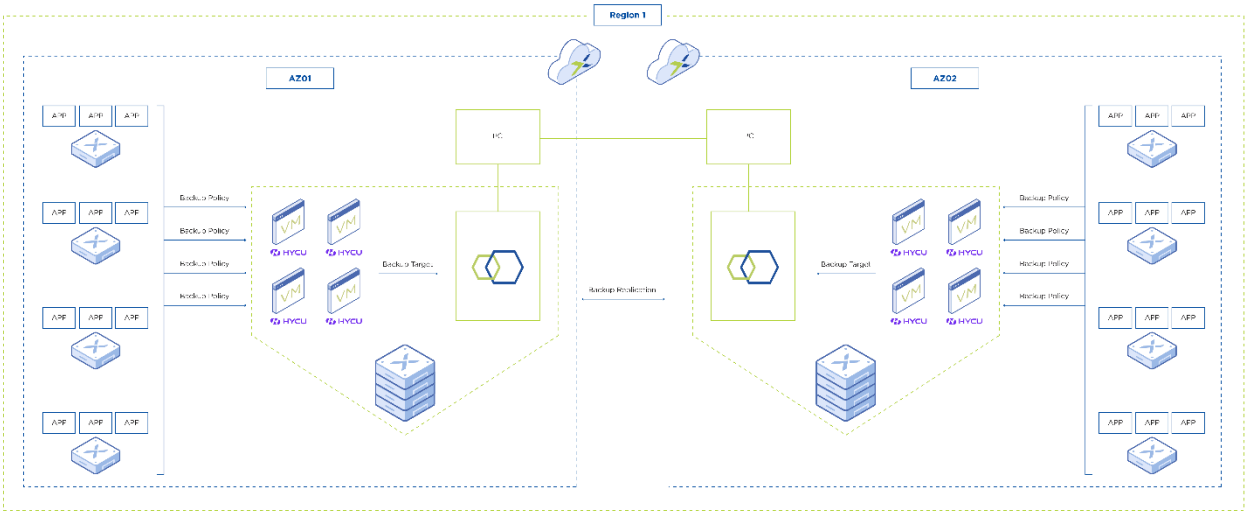

Figure 15: Nutanix Mine Logical Design

Backup Detailed Design

The following sections describe in detail the physical components of BCDR with Nutanix Mine as the external backup system.

Table: Backup Software Versions

Component	Software Version
Prism Central	pc.2021.7
AOS	5.20.x (LTS)
Mine	3
HYCU	4.2.1
Objects	3.2.x
Object manager	3.2.x

Categories in Prism Central automate VM placement in the target backup policy. The backup system uses categories to identify VM location (local AZ versus remote AZ) and RPO tier. A single category can have at most 750 VMs, and each category has an RPO of 24 hours.

Table: Backup Categories

Location	Category Name	Value	Max # of VMs
AZ01	AZ01-Backup-01	RPO24h	750
AZ01	AZ01-Backup-02	RPO24h	750
AZ01	AZ01-Backup-03	RPO24h	750
AZ01	AZ01-Backup-04	RPO24h	750
AZ01	AZ01-Backup-05	RPO24h	750
AZ02	AZ02-Backup-01	RPO24h	750
AZ02	AZ02-Backup-02	RPO24h	750
AZ02	AZ02-Backup-03	RPO24h	750
AZ02	AZ02-Backup-04	RPO24h	750
AZ02	AZ02-Backup-05	RPO24h	750

One HYCU backup server can have up to two backup policies and a total of 1,500 VMs.

Table: Backup Policies

Backup VM Name	Policy Name	RPO	Category	Max # of VMs
AZ01HycuBP01	AZ01-Backup-01	24 hours	AZ01-Backup-01	750
AZ01HycuBP01	AZ01-Backup-01	24 hours	AZ01-Backup-02	750
AZ01HycuBP02	AZ01-Backup-02	24 hours	AZ01-Backup-03	750
AZ01HycuBP02	AZ01-Backup-02	24 hours	AZ01-Backup-04	750
AZ01HycuBP03	AZ01-Backup-03	24 hours	AZ01-Backup-05	750
AZ02HycuBP01	AZ02-Backup-01	24 hours	AZ02-Backup-01	750
AZ02HycuBP01	AZ02-Backup-01	24 hours	AZ02-Backup-02	750
AZ02HycuBP02	AZ02-Backup-02	24 hours	AZ02-Backup-03	750
AZ02HycuBP02	AZ02-Backup-02	24 hours	AZ02-Backup-04	750
AZ02HycuBP03	AZ02-Backup-03	24 hours	AZ02-Backup-05	750

Backup proxies transfer backup data from source clusters to target storage. Each backup proxy has the resource configuration shown in the following tables.

Table: HYCU Backup Proxy Resources

Location	Host Name	vCPU	RAM	Storage (GB)	OS
AZ01	AZ01HycuBP01	16	32	200	Appliance
AZ01	AZ01HycuBP02	16	32	200	Appliance
AZ01	AZ01HycuBP03	16	32	200	Appliance
AZ02	AZ02HycuBP01	16	32	200	Appliance

Location	Host Name	vCPU	RAM	Storage (GB)	OS
AZ02	AZ02HycuBP02	16	32	200	Appliance
AZ02	AZ02HycuBP03	16	32	200	Appliance

Table: HYCU Backup Proxy Virtual Hardware Configuration

Virtual Hardware	Value	Type
Virtual CPU	16	vCPU
Virtual memory	32 GB	RAM
Virtual storage	200 GB	VirtIO-SCSI
Virtual NIC	1	VirtIO-Net
Virtual CD-ROM	1	IDE

Nutanix Objects

The NVD uses Nutanix Objects as backup target storage for all backup data. For maximum performance, deploy three worker nodes and two load balancers.

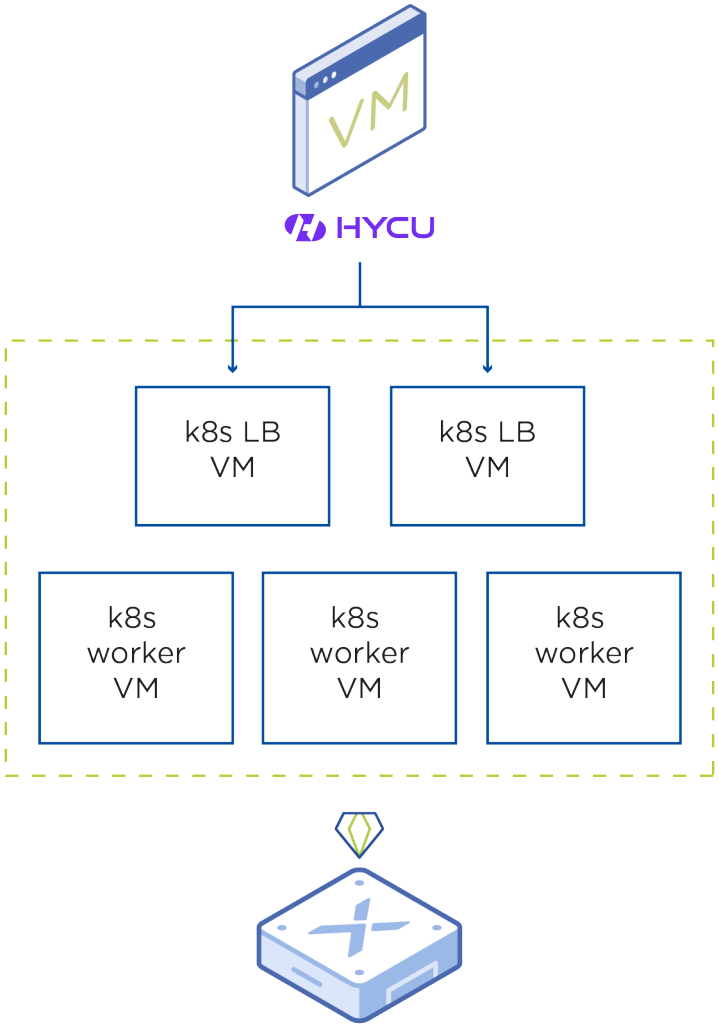

Figure 16: Nutanix Objects Logical Design

The object store that hosts backup data has the resource configuration shown in the following tables.

Table: Compute Resources for Objects VMs

Function	Number of Instances	vCPU	RAM
Load balancer VM	2	2	4

Function	Number of Instances	vCPU	RAM
Worker VM	3	10	32

Table: Object Store Details

Location	Nutanix Cluster	Name	Storage Size	# of LB VMs	# of Worker VMs
AZ01	<MineCluster>	AZ01Backup01.domain.local	Allocate maximum available storage	2	3
AZ02	<MineCluster>	AZ02Backup01.domain.local	Allocate maximum available storage	2	3

The Nutanix S3 bucket has the configuration shown in the following table. The number of S3 buckets depends on the number of backup VMs. There is 1:1 matching between the backup VM, the S3 bucket for primary backup data, and the S3 bucket for the backup data copy on the remote site.

Table: Nutanix Bucket Configuration Details

Location	Object Store Name	Bucket Name	Versioning	WORM
AZ01	AZ01Backup01.domain.local	<AZ01 Backup Server Name>	No	Yes: 365 days
AZ01	AZ01Backup01.domain.local	<AZ01 Backup Server Name>	No	Yes: 365 days
AZ01	AZ01Backup01.domain.local	<AZ01 Backup Server Name>	No	Yes: 365 days
AZ01	AZ02Backup01.domain.local	<AZ02 Backup Server Name>-copy	No	Yes: 365 days
AZ01	AZ02Backup01.domain.local	<AZ02 Backup Server Name>-copy	No	Yes: 365 days

Location	Object Store Name	Bucket Name	Versioning	WORM
AZ01	AZ02Backup01.domain.local	<AZ02 Backup Server Name>-copy	No	Yes: 365 days
AZ02	AZ02Backup01.domain.local	<AZ02 Backup Server Name>	No	Yes: 365 days
AZ02	AZ02Backup01.domain.local	<AZ02 Backup Server Name>	No	Yes: 365 days
AZ02	AZ02Backup01.domain.local	<AZ02 Backup Server Name>	No	Yes: 365 days
AZ02	AZ01Backup01.domain.local	<AZ01 Backup Server Name>-copy	No	Yes: 365 days
AZ02	AZ01Backup01.domain.local	<AZ01 Backup Server Name>-copy	No	Yes: 365 days
AZ02	AZ01Backup01.domain.local	<AZ01 Backup Server Name>-copy	No	Yes: 365 days

See the earlier HYCU Backup Proxy Resources table for the backup server host name to use as the bucket name.

4. Self-Service with Automation

Self-Service with Automation Introduction

This design incorporates Nutanix Calm to provide self-service with automation to IT users. With a marketplace experience, users can deploy VMs and applications in a secure and consistent manner.

In a common enterprise scenario, you must configure every application deployment with IP addresses that can come from an IPAM system or DNS, join directory services for authentication, or involve getting a virtual IP (VIP) address from a load balancer. The blueprints in this design include integrations with these foundational services.

Use Prism Central categories in Calm blueprints to mitigate the risk of not applying a category, something that's likely to happen in a manual deployment.

Self-Service with automation requirements by component:

- Calm
  - Provide self-service for Windows, Linux, LAMP, and WISA applications.
  - Be secure by design. Following DevSecOps principles, protect all application networking and data from ransomware attacks.
  - Support more than 5,000 VMs.
  - Let IT users deploy applications in different clusters and locations.
  - Provide cloud governance.
  - Present application costs.
  - Notify IT users when their applications are ready.
  - Provide a seamless hybrid multicloud experience.
  - Standardize the virtual hardware specifications for VMs.

- Integration
  - Integrate with IPAM for configuring VM addresses.
  - Integrate with directory services for authentication.
  - Integrate with backup for VM protection.
  - Integrate with datacenter load balancers for configuring application VIP addresses.

Self-Service with automation assumptions by component:

- Calm
  - Calm can access any third-party system that the blueprint must integrate with.
  - As part of the blueprints, Calm has WinRM (HTTP or HTTPS) or SSH access to the networks where VMs are deployed.
  - Calm can connect to Nutanix Beam in the cloud.
  - VMs deployed by Calm can communicate with email infrastructure to send notifications.
- Integration
  - IPAM infrastructure has sufficient resilience for the system to request, register, and release IP addresses, even during critical outages.
  - Directory services infrastructure has sufficient resilience for adding and removing VMs, even during critical outages.
  - Backup services infrastructure has sufficient resilience for backing up and restoring VMs, even during critical outages.
  - Email infrastructure has sufficient resilience to send, receive, and access emails, even during critical outages.
  - Load balancer infrastructure has sufficient resilience for handling API requests, even during critical outages.
  - Blueprints are also available in a source code management system.

Self-Service with automation risks by component:

- Calm
  - During Calm upgrades, the service is unavailable.
  - During Calm downtime, the service is unavailable.
  - Single-instance Calm is a single point of failure.
  - In the event of a disaster, applications recovered in another Prism Central instance are unavailable in Calm until you run the Prism Central–to–Prism Central sync script.
- Integration
  - During IPAM downtime, new Calm deployments might fail.
  - During directory service downtime, new Calm deployments might fail.
  - During load balancer downtime, new Calm deployments that need load balancing might fail.

Self-Service with automation constraints by component:

- Calm
  - Blueprints must use existing approved VM templates.
  - VM names must adhere to existing naming conventions.
  - Virtual hardware has a maximum of three sizes.
- Integration
  - The IPAM solution is Infoblox.
  - The backup solution is Mine with HYCU.
  - The network security solution is Nutanix Flow.
  - The BCDR solution is Nutanix Disaster Recovery.
  - The directory service is Microsoft Active Directory.
  - The load balancer is F5.

Table: Self-Service with Automation Design Decisions

Component	Description
Calm deployment model	Standalone single virtual appliance
Calm deployment size (small or large)	Large
Define process for Calm recoverability	Calm is protected using a Nutanix Disaster Recovery Protection Policy and a recovery plan as well as using a Mine category for backup and archiving
Align Calm project structure and role-based access control (RBAC) configuration	Don't use the default Calm project; instead, use dedicated Calm projects with RBAC based on your Nutanix Services architecture workshop
Use Active Directory authentication	Use Active Directory for Calm access and project RBAC
Enable the Calm policy engine	Yes: the Calm policy engine is required for functionalities like quotas
Enable Calm showback	Yes: provide showback for Nutanix AHV provider accounts
Enable Show App Protection Status	Yes: provide application tracking in the event of a disaster when recovered in a different location
Choose a method for self-service	Calm marketplace is the self-service portal for IT users
Use SSL or TLS connection to Active Directory	Use WinRM over HTTPS for Windows blueprints
Secure by design with Prism Central categories	Blueprints use Prism Central categories for security (Flow), protection and recovery (Nutanix Disaster Recovery), and backup (Mine) policies to address security earlier in the development process (DevSecOps)
Email notification after application deployment	An in-guest script sends emails from the VMs deployed by Calm

Component	Description
Blueprint development method	Develop blueprints using Calm DSL, which generates multi-VM blueprints even if there is a single service (IaaS)
Blueprint development project	Develop blueprints in a dedicated project and make them available to other projects through the marketplace manager
Standard sizing model	Standardize the virtual hardware on small (1 vCPU, 8 GB of memory), medium (2 vCPU, 16 GB of memory), and large (4 vCPU, 32 GB of memory) sizing models
Integration with IPAM	Blueprints using Escript tasks in the pre-create stage communicate with the Infoblox API for CRUD tasks
Integration with F5	Blueprints using Escript tasks in the create stage communicate with the F5 API for CRUD tasks
Integration with Active Directory	Blueprints using PowerShell or Shell script tasks in the package install stage communicate with Active Directory

Self-Service with Automation Conceptual Design

Nutanix Calm is the automation and orchestration software that runs in Prism Central. From Calm, IT users can request infrastructure and applications and operate them throughout their life cycle.

Calm can deploy workloads in any AZ as part of the marketplace request. In this request, users can specify different aspects of their workloads such as compute type, location, and data protection SLA and preview how much the resources they're asking for cost.

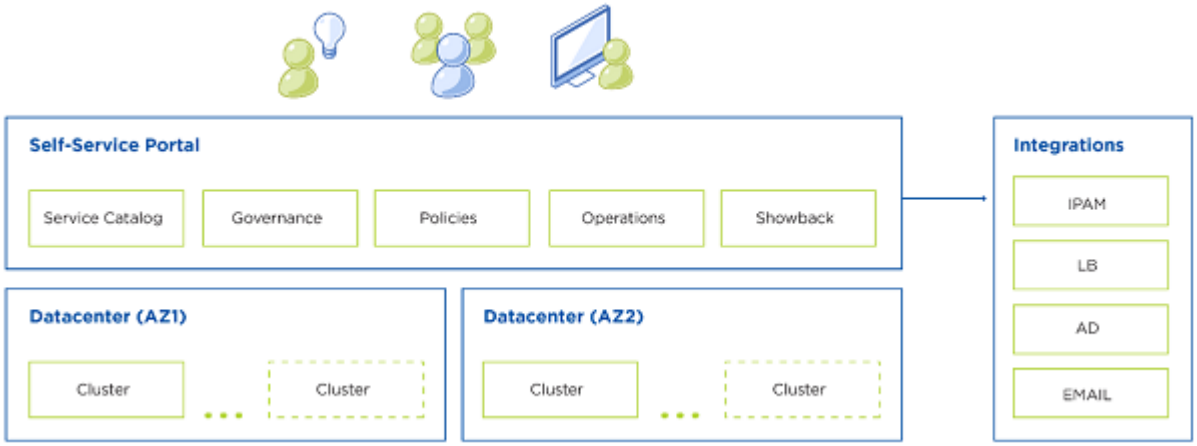

Figure 17: Self-Service with Automation Conceptual Design

Self-Service with Automation Logical Design

Nutanix Calm uses a modular approach for meeting enterprise multitenancy requirements following governance policies.

The following diagram shows the relationships between the components configured in Calm as part of this NVD.

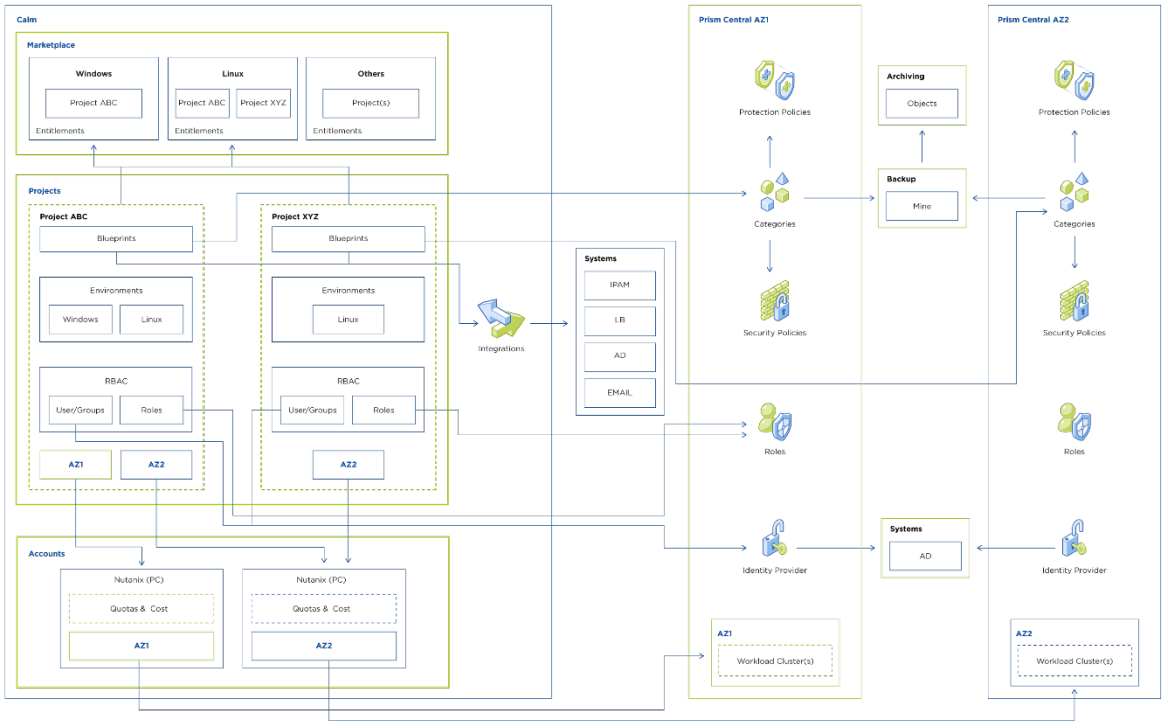

Figure 18: Self-Service with Automation Logical Design

Accounts

Calm needs at least one provider account so projects can deploy workloads. By default, enabling Calm in Prism Central creates the NTNX_LOCAL_AZ account. This account automatically discovers the AHV clusters registered in Prism Central. Because the NVD uses a standalone Calm instance, there are no clusters registered in Prism Central in this case. This NVD adds two Nutanix accounts that connect to the Prism Central instances that manage the clusters in each different AZ.

Projects

Projects are like tenants, delivering governance and multitenancy. Usually, projects are aligned with environments (development or production), operating systems (Windows or Linux), departments (human resources or finance), or applications (Exchange or SAP). A project must have at least one account, RBAC using Prism Central roles and Active Directory, and an environment before project users can request workloads from

the marketplace. This NVD has one project for blueprint design and four projects to validate the security aspects of tenant workloads.

Blueprints

Blueprints are project-specific and define how to automate workload deployment. An important part of this design is the use of Prism Central categories in a blueprint to drive DevSecOps and help prevent ransomware. To make a blueprint available for other projects, you must publish it in the marketplace. This NVD has four blueprints: Windows, Linux, WISA, and LAMP. All the blueprints integrate with IPAM, Active Directory, and email. WISA and LAMP also integrate with the load balancer.

Marketplace

When projects submit blueprints for approval, an administrator reviews, categorizes, and versions them. After they publish a blueprint, an administrator can assign it to projects for consumption through the marketplace page. This NVD uses two projects with different blueprint assignments to validate its security.

Integrations

Integrations are part of the blueprint and occur at different stages of the life cycle. In this NVD, most integrations use Calm Escript (a Python library), with some instances of PowerShell and Shell scripts for integrations where only a CLI is available.

Categories

Security policies with Flow microsegmentation, protection and recovery policies with Nutanix Disaster Recovery, backup policies with Mine, and HYCU archiving to Objects all use Prism Central categories. Using categories in blueprints helps prevent ransomware because every deployment is secure by design.

Self-Service with Automation Detailed Design

Note: Nutanix Services customizes this NVD to meet individual customer requirements following an architecture workshop. In the following tables, items marked TBD represent a value that Nutanix Services and the customer collaboratively determine during the workshop.

Calm

Calm is a standalone instance in this NVD.

Table: Self-Service with Automation Deployment Model

Setting	Value
Deployment	One instance of Calm on AHV
Resources	10 vCPU, 52 GB of memory, 581 GB disk
Network	Management (requires IP address)
Protection	Nutanix Disaster Recovery Protection Policy and Recovery Plan

This NVD uses the following software versions for Calm.

Table: Self-Service with Automation Software Versions

Component	Software Version
Prism Central	pc.2021.7
AOS	6.0 (STS)
Calm	3.3

This NVD uses the following settings to protect the Calm virtual appliance from a disaster scenario.

Table: Nutanix Categories

Category Name	Value	Assigned	Used By
AppType	CalmAppliance	Calm_on_AHV (VM)	Nutanix Disaster Recovery
AZ01-Backup-01	RPO24h	Calm_on_AHV (VM)	Mine (backup)

Table: Protection Policy Configuration

Policy Name	Category	Source Cluster	Target Cluster	RPO
AZ01-AZ02-Calm	AppType: CalmAppliance	AZ01-MGMT-01	AZ02-MGMT-01	1 hour

Table: Recovery Plan

Name	Stage	Category	Delay	Source Network	Failover Networks	Test Failover Network
AZ01-RP-Calm	Stage1	AppType: CalmAppliance	0	Source-PG	Failover-PG	Test-PG

Table: Protection Policy to Recovery Plan Mapping

PP Name	RP Name	Category	Value
AZ01-AZ02-Calm	AZ01-RP-Calm	AppType: CalmAppliance	RPO = 1 h

This NVD uses the following Flow security policy to let the Calm virtual appliance connect to VMs to support automation.

Table: Calm Security Policy

Purpose	Source	Destination	Port / Protocol
Allow Calm management	AppType: CalmAppliance	AppType: AZ01-Example-001; AppTier: Web; AppTier: App; AppTier: DB	TCP 22, 5985–5986

This NVD uses the following software settings for Calm.

Table: Self-Service with Automation Calm Settings

Setting	Value
Default landing page	Yes
Marketplace apps	No

Setting	Value
Showback	Yes
Policy engine	Yes (requires IP address)
Protection status	Yes

Table: Self-Service with Automation Calm Accounts

Account	Provider	Cluster	Cost	Sync Settings	Quotas
Region A: AZ01	Nutanix	Clusters AZ01	TBD	15 min	N/A
Region A: AZ02	Nutanix	Clusters AZ02	TBD	15 min	N/A

Table: Self-Service with Automation Calm Project: Blueprints-design

Project	RBAC	Accounts	Allow List Subnets and Quotas	Environments
Blueprints-design	Blueprint_Designers (Active Directory group)	All	All clusters and subnets, no quotas	Windows and Linux

Table: Other Self-Service with Automation Calm Projects

Cluster	Project	RBAC	Accounts	Allow List Subnets	Quotas	Envir.s
Clusters AZ01	TBD	TBD	All	TBD	TBD	TBD
Clusters AZ02	TBD	TBD	All	TBD	TBD	TBD

From left to right, the following figure shows the workflow of a workload deployment, with the configuration and dependencies for each stage.

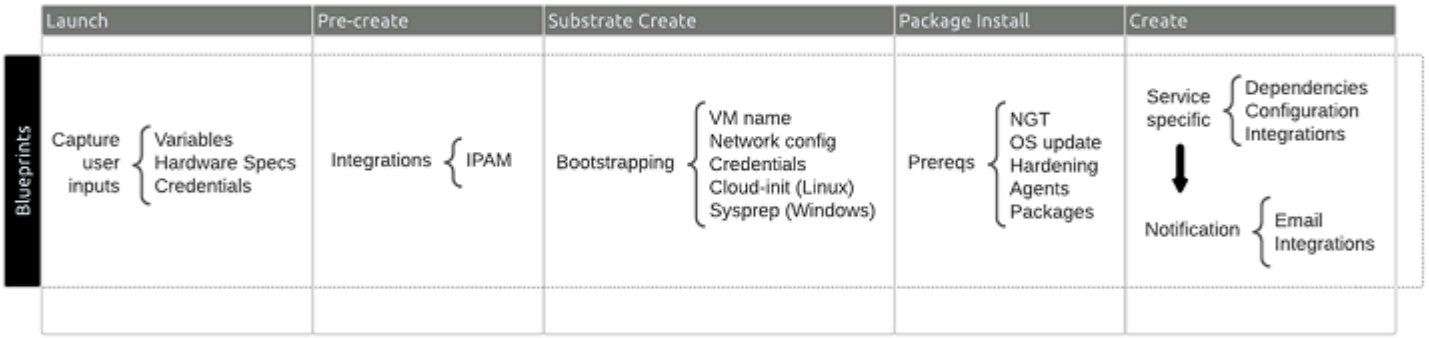

Figure 19: Self-Service with Automation Calm Blueprint Stages

Refer to the Nutanix Validated Designs GitHub repository for the blueprints used in this design.

Table: Self-Service with Automation Calm Blueprints

Blueprint	Component	Software Version	Component Dependencies	Integrations
Windows	Single service	Windows Server 2019	N/A	IPAM, Active Directory, Email
Linux	Single service	CentOS 8.2	N/A	IPAM, Active Directory, Email

Blueprint	Component	Software Version	Component Dependencies	Integrations
WISA	Load balancer	BIG-IP 16.1.0 Build 0.0.19 Final	Scale-out web	IPAM, Active Directory, Email, LB
WISA	Scale-out web	Server 2019 + IIS 10	Database	IPAM, Active Directory, Email, LB
WISA	Database	Server 2019 + SQL 2019	N/A	IPAM, Active Directory, Email, LB
LAMP	Load balancer	TBD	Scale-out web	IPAM, Active Directory, Email, LB
LAMP	Scale-out web	CentOS 8 + PHP 8	Database	IPAM, Active Directory, Email, LB
LAMP	Database	CentOS 8 + MariaDB 10.6	N/A	IPAM, Active Directory, Email, LB

Table: Self-Service with Automation Calm Marketplace

Blueprint	Available To	Version	Category
Windows	TBD	1.0	TBD
Linux	TBD	1.0	TBD
WISA	TBD	1.0	TBD
LAMP	TBD	1.0	TBD

Directory Services

This NVD adds every workload provisioned using Calm to Active Directory and uses the following software version for Active Directory integration.

Table: Self-Service with Automation Active Directory Software Version

Component	Software Version
Active Directory	Windows Server 2019

Table: Self-Service with Automation Active Directory Connection Details

Connection	Details
Domain	nutanix.nvd
Domain Controller	X.X.X.X or FQDN
Username	svc_calm
Password	xxx

IPAM

Based on user input, Infoblox provides every workload provisioned using Calm with an IP address in the selected network and configured DNS.

This NVD uses the following software version for Infoblox DNS, DHCP, and IP address management (DDI) integration.

Table: Self-Service with Automation IPAM Software Version

Component	Software Version
Infoblox	8.4.4-386831

Table: Self-Service with Automation IPAM Connection Details

Connection	Details
Infoblox API	X.X.X.X or FQDN
Username	svc_calm
Password	xxx
Networks	TBD

Load Balancing

WISA and LAMP workloads integrate with the F5 load balancer for the web server tier.

This NVD uses the following software versions for load balancer integration.

Table: Self-Service with Automation Load Balancer Software Version

Component	Software Version
F5	BIG-IP 16.1.0 Build 0.0.19 Final

Table: Self-Service with Automation Load Balancer Connection Details

Connection	Details
F5 API	X.X.X.X or FQDN
Username	svc_calm
Password	xxx

This NVD uses the following Flow security policy to let the F5 load balancer send HTTP and HTTPS traffic to the application tier.

Table: Address Group

Name	Addresses	Purpose
AddrLoadBalancer	10.38.218.44/32	F5 load balancing for WISA and LAMP

Table: Example Load Balancing Security Policy

Purpose	Source	Destination	Port / Protocol
Allow F5 LB to app	AddrLoadBalancer	AppType: AZ01-Example-0001; AppTier: App	TCP 80,443

Notifications

Every workload provisioned using Calm sends an email to the requester.

This NVD uses the following software versions for notification integration.

Table: Self-Service with Automation Notifications Software Versions

OS	Notification	Library
Windows	Email	Send-MailMessage
Linux	Email	smtplib and email.message

Table: Self-Service with Automation Notifications Connection Details

Connection	Details
SMTP	X.X.X.X or FQDN
Port	465
Sender	`no_reply@nutanix.nvd`
Username	svc_calm
Password	xxx
Recipients	Calm requester, distribution list, or both

5. 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 block consists of 32 nodes, or two 16-node workload clusters—one in each datacenter for BCDR. A pod consists of the following components:

- Two 4-node management clusters.
- Enough 32-node blocks (sets of two 16-node workload clusters) to meet the desired capacity.
- Two Mine backup clusters.

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

For smaller environments, you can downsize the workload clusters to 4, 8, or 12 nodes based on your capacity requirements, but note the following limitations:

- Don't change the hardware configuration or sizing associated with the management clusters.
- Don't change the hardware configuration or sizing associated with the Mine backup clusters.
- You can reduce the number of HYCU licenses in both the primary and secondary datacenter according to the following table.

Table: HYCU Licenses

# of Nodes	# of HYCU Licenses
4	200 VMs per cluster
8	400 VMs per cluster
12	600 VMs per cluster

Bill of Materials

The following sections show the BoMs for the primary and secondary datacenter management clusters, the primary and secondary datacenter workload clusters, and the primary and secondary datacenter backup clusters.

Note: X-1175-G8 can replace NX-1175-G7 when it becomes available.

Primary Datacenter Management Cluster: Hardware, Software, and Services

Hardware

- Product code: NX-1175S-G7
  - Quantity: 4
  - Model: NX-1175S-G7, 1-node configuration
  - Type: All flash
  - Hardware support:
    - Support level: Production
    - NRDK support: No
    - NR node support: No
  - Per-node hardware configuration:
    - Processor: Intel Xeon-Gold 6226R (2.9 GHz / 16-core) x 1
    - Memory: 64 GB (3,200 MHz DDR4 RDIMM) x 12
    - HDD: No HDD Included
    - SSD: 1.92 TB x 4
    - Network adapter: 10 GbE, 2-port, SFP+ (Intel 82599ES) x 1

Software

- Product code: SW-AOS-PRO-PRD
  - Quantity: From hardware
  - Subscription, Acropolis (AOS)
  - License tier: Pro
  - Support level: Production
- Product code: SW-PRS-PRO-CORE
  - Quantity: From hardware
  - Prism Pro software license subscription for 1 CPU core
  - Type: Core-based licensing
  - License type: Prism Pro

Cluster Install Services

- Product code: CNS-INF-A-SVC-DEP-STR
  - Quantity: 4
  - Xpert Services, HCI Cluster Deployment Starter

Secondary Datacenter Management Cluster: Hardware, Software, and Services

Hardware

- Product code: NX-1175S-G7
  - Quantity: 4
  - Model: NX-1175S-G7, 1-node configuration
  - Type: All flash
  - Hardware support:
    - Support level: Production
    - NRDK support: No
    - NR node support: No
  - Per-node hardware configuration:
    - Processor: Intel Xeon-Gold 6226R (2.9 GHz / 16-core) x 1
    - Memory: 64 GB (3,200 MHz DDR4 RDIMM) x 12
    - HDD: No HDD included
    - SSD: 1.92 TB x 4
    - Network adapter: 10 GbE, 2-port, SFP+ (Intel 82599ES) x 1

Software

- Product code: SW-AOS-PRO-PRD
  - Quantity: From hardware
  - Subscription, Acropolis (AOS)
  - License tier: Pro
  - Support level: Production

- Product code: SW-PRS-PRO-CORE
  - Quantity: From hardware
  - Prism Pro software license subscription for 1 CPU core
  - Type: Core-based licensing
  - License type: Prism Pro

Cluster Install Services

- Product code: CNS-INF-A-SVC-DEP-STR
  - Quantity: 4
  - Xpert Services, HCI Cluster Deployment Starter

Primary Datacenter Workload Cluster: Hardware, Software, and Services

Hardware

- Product code: NX-3170-G8
  - Quantity: 16
  - NX-3170-G8, 1-node configuration
  - Type: All flash
  - Hardware support:
    - Support level: Production
    - NRDK support: No
    - NR node support: No
  - Per-node hardware configuration:
    - Processor: Intel Xeon-Gold 5318Y (2.1 GHz / 24-core) x2
    - Memory: 64 GB (3,200 MHz DDR4 RDIMM) x 24
    - HDD: No HDD included
    - SSD: 3.84 TB x 6
    - Network adapter: 25 GbE, 2-port (Mellanox ConnectX-5) x 1

Software

- Product code: SW-AOS-PRO-PRD
  - Quantity: From hardware
  - Subscription, Acropolis (AOS)
  - License tier: Pro
  - Support level: Production

- Product code: SW-AOS-ADVREP-PRD
  - Quantity: From hardware
  - Add on: Advanced replication
- Product code: SW-PRS-PRO-CORE
  - Quantity: From hardware
  - Prism Pro software license subscription for 1 CPU core
  - Type: Core-based licensing
  - License type: Prism Pro
- Product code: SW-CALM-CORE-PRD
  - Quantity: From hardware
  - Support level: Production
- Product code: SW-FLOW-CORE
  - Quantity: From hardware
  - Type: Core-based licensing

Cluster Install Services

Not required. Refer to the Professional Services section for product code CNS-INF-STR-LRG that includes installation for one workload tenant cluster.

Secondary Datacenter Workload Cluster: Hardware, Software, and Services

Hardware

- Product code: NX-3170-G8
  - Quantity: 16
  - NX-3170-G8, 1-node configuration
  - Type: All flash
  - Hardware support:
    - Support level: Production
    - NRDK support: No
    - NR node support: No
  - Per-node hardware configuration:
    - Processor: Intel Xeon-Gold 5318Y (2.1 GHz / 24-core) x 2
    - Memory: 64 GB (3,200 MHz DDR4 RDIMM) x 24
    - HDD: No HDD included
    - SSD: 3.84 TB x 6
    - Network Adapter: 25 GbE, 2-port (Mellanox ConnectX-5) x 1

Software

- Product code: SW-AOS-PRO-PRD
  - Quantity: From hardware
  - Subscription, Acropolis (AOS)
  - License tier: Pro
  - Support level: Production

- Product code: SW-AOS-ADVREP-PRD
  - Quantity: From hardware
  - Add on: Advanced replication
- Product code: SW-PRS-PRO-CORE
  - Quantity: From hardware
  - Prism Pro software license subscription for 1 CPU core
  - Type: Core-based licensing
  - License type: Prism Pro
- Product code: SW-CALM-CORE-PRD
  - Quantity: From hardware
  - Support level: Production
- Product code: SW-FLOW-CORE
  - Quantity: From hardware
  - Type: Core-based licensing

Cluster Install Services

- Product code: CNS-INF-A-SVC-DEP-STR
  - Quantity: 16
  - Xpert Services, HCI Cluster Deployment Starter

Primary Datacenter Backup Cluster: Hardware, Software, and Services

Hardware

- Product code: NX-8155-G8
  - Quantity: 4
  - NX-8155-G8, 1-node configuration
  - Type: Hybrid
  - Hardware support:
    - Support level: Production
    - NRDK support: No
    - NR node support: No
  - Per-node hardware configuration:
    - Processor: Intel Xeon-Gold 6326 (2.9 GHz / 16-core) x 2
    - Memory: 32 GB (3,200 MHz DDR4 RDIMM) x 8
    - HDD: 18 TB, 3.5 inch x 8
    - SSD: 3.84 TB x 2
    - Network adapter: 25 GbE, 2-port (Mellanox ConnectX-5) x 1

Software

- Product code: SW-OBJECTS-DED-PRD
  - Quantity: 1 TiB
  - Objects dedicated
  - Support level: Production

Note: You only need SW-OBJECTS-DED-PRD on the BoM so that you can quote 18 TB HDD. A quantity of 1 TiB is sufficient.

- Product code: SW-PRS-PRO-CORE
  - Quantity: From hardware
  - Prism Pro software license subscription for 1 CPU core
  - Type: Core-based licensing
  - License type: Prism Pro
- Product code: SW-MINE-S3-PRD-3YR
  - Quantity: 264 TiB
  - Mine software
  - Support level: Production
- Product code: SW-H-MINE-RNTXDPVM3yr
  - Quantity: 750 VMs
  - HYCU License Bundle for Nutanix Mine, 1Ct for 3YR

Cluster Install Services

- Product code: CNS-INF-A-SVC-DPD-MIN
  - Quantity: 1
  - Xpert Services, HCI Backup Architecture Deployment

Secondary Datacenter Backup Cluster: Hardware, Software, and Services

Hardware

- Product code: NX-8155-G8
  - Quantity: 4
  - NX-8155-G8, 1-node configuration
  - Type: Hybrid
  - Hardware support:
    - Support level: Production
    - NRDK support: No
    - NR node support: No
  - Per-node hardware configuration:
    - Processor: Intel Xeon-Gold 6326 (2.9 GHz / 16-core) x 2
    - Memory: 32 GB (3,200 MHz DDR4 RDIMM) x 8
    - HDD: 18 TB, 3.5 inch x 8
    - SSD: 3.84 TB x 2
    - Network adapter: 25 GbE, 2-port (Mellanox ConnectX-5) x 1

Software

- Product code: SW-OBJECTS-DED-PRD
  - Quantity: 1 TiB
  - Objects dedicated
  - Support level: Production

Note: You only need SW-OBJECTS-DED-PRD on the BoM so that you can quote 18 TB HDD. A quantity of 1 TiB is sufficient.

- Product code: SW-PRS-PRO-CORE
  - Quantity: From hardware
  - Prism Pro software license subscription for 1 CPU core
  - Type: Core-based licensing
  - License type: Prism Pro
- Product code: SW-MINE-S3-PRD-3YR
  - Quantity: 264 TiB
  - Mine software
  - Support level: Production
- Product code: SW-H-MINE-RNTXDPVM3yr
  - Quantity: 750 VMs
  - HYCU License Bundle for Nutanix Mine, 1Ct for 3YR

Cluster Install Services

- Product code: CNS-INF-A-SVC-DPD-MIN
  - Quantity: 1
  - Xpert Services, HCI Backup Architecture Deployment

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: Professional Services for Platform

Product Code	Description	Quantity
CNS-INF-STR-LRG	Xpert Services Infra Modernization: Starter Large	1
CNS-INF-A-WRK-MCR-STD	Xpert Services, HCI Microsegmentation Workshop	1
CNS-INF-A-SVC-MCR-STD	Xpert Services, HCI MicroSegmentation Deployment Service	1
CNS-INF-A-SVC-DRD-LEAP	Xpert Services, HCI Disaster Recovery Deployment Leap	1
FLEX-CST-CR	Flexible Services Credits (1 credit = $100) (Credits to deliver a HCI Disaster Recovery Leap Workshop)	166
CNS-CAS-PRO-STD	Xpert Services Cloud and IT Automation: Pro	1

6. Appendix

Windows VM Performance Tuning

For Windows VMs, consider the following performance tuning settings:

- In the base OS image, navigate to the `Configure Advanced Settings for Maximum Performance System Properties` page and click the `Advanced` tab. In the `Performance` section, click the `Settings` button and navigate to the `Visual Effects` tab. Select the `Adjust for best performance` option and click `OK`.
- To set the VM graphics adapter hardware acceleration to full, open the Control Panel. In the `Display` section, navigate to the `Settings` tab and click the `Advanced` button. In the `Troubleshooting` tab, set the `Hardware Acceleration` option to `full`.
- Disable screen savers and Windows search indexing.

Linux VM Performance Tuning

For Linux VMs, consider the following performance tuning settings:

- Specifically for Java Virtual Machine (JVM) systems, enable large pages with enough memory to cover the JVM HEAP and any other memory requirements. Reserve the memory required for the JVM HEAP, any other JVM memory, and basic OS functions. Start with 2 vCPU and increase only if necessary. To enable large pages for a Sun JVM, use the parameter `-XX:+UseLargePages`. To enable large pages on an IBM JVM, use the parameter `–Xlp`.
- Edit the grub config and add the following settings to the correct kernel boot line: `transparent_hugepage=never iommu=soft elevator=noop powersaved=off selinux=0 noselinux apm=off`
- Add the following settings to sysctl.conf:

```
vm.overcommit_memory = 1
vm.dirty_background_ratio = 5
```



```
vm.dirty_ratio = 15
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs=100
vm.swappiness = 0
fs.aio-max-nr=3145728
fs.file-max = 6815744
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 262144
net.core.wmem_default = 262144
net.ipv4.tcp_rmem = 4096 262144 16777216
net.ipv4.tcp_wmem = 4096 262144 16777216
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
# these should be on by default, but just to be sure
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
# disable ip forwarding
net.ipv4.ip_forward = 0
# Controls source route verification
net.ipv4.conf.default.rp_filter = 1
# Do not accept source routing
net.ipv4.conf.default.accept_source_route = 0
# Controls the use of TCP syncookies
net.ipv4.tcp_syncookies = 1
# Allow many more connections
net.core.netdev_max_backlog = 5000
net.core.somaxconn = 10000
net.ipv4.tcpkeepalive_intvl = 15
net.ipv4.tcp_fin_timeout = 15
net.ipv4.tcp_keepalive_probes = 5
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_max_syn_backlog = 5000
# Controls the maximum number of shared memory segments, in pages, dependent on
 memory configured on server, try the defaults first before using these
#kernel.shmall = 4294967296 # 90% memory, in pages
#kernel.shmmni = 4096
kernel.shmmax = 5368709120
kernel.sem = 250 256000 128 1024
# Large pages, replace the nr pages with the amount of memory to reserve divided
 by 2M,
# which is the page size, and replace the group id (gid) with the ID of the group
 id that locks the pages in memory, replace values in <>
vm.nr_hugepages = 2304
vm.hugetlb_shm_group = 1002
```

- Add the following lines to limits.conf:

```
<gid java user> soft nofile 131070 # This ensures there are enough file
 descriptors to handle all the TCP ports and filesystem handles
<gid java user> hard nofile 131070 # Set same as above
@<gid java user> soft memlock 4718592 # This needs to be sufficient to cover the
 number of reserved huge pages
@<gid java user> hard memlock 4718592 # Should be same value as above
```

References

1. Nutanix Hybrid Cloud Reference Architecture

2. Nutanix Calm
3. Flow Microsegmentation Guide
4. Nutanix Disaster Recovery (formerly Leap)
5. Nutanix Mine
6. Nutanix Objects
7. Data Protection and Disaster Recovery
8. Physical Networking

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