

Cisco Compute Hyperconverged with Nutanix for Microsoft SQL Server 2022 Databases

Microsoft SQL Server Databases running on Cisco
Compute Hyperconverged with Nutanix

Published: July 2024



In partnership with:



About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to: <http://www.cisco.com/go/designzone>.

Executive Summary

Digital transformation and data explosion causes the unprecedented growth of innovative applications and services. These applications are highly diversified and distributed, running at various locations such as data centers, edge, and cloud, requiring modern infrastructure and operations to meet the dynamic needs of the business. To continue serving their organization, IT teams need to be able to deploy infrastructure and applications fast using a unified cloud-operating model and enterprise-class systems so they can easily and quickly adapt to the demands of applications, and seamlessly scale from the data center to the edge and the cloud.

Cisco and Nutanix have formed a strategic partnership to introduce complete hyperconverged solutions by integrating and validating Cisco® servers, storage, networking, and SaaS operations with the Nutanix Cloud Platform. Cisco Compute Hyperconverged with Nutanix is built, managed, and supported holistically to deliver a more seamless experience, foster innovation, and accelerate customers' hybrid multicloud journeys.

This document discusses a Cisco Compute Hyperconverged system with Nutanix and provides design and deployment best practices for hosting virtualized Microsoft SQL Server 2022 databases. The hyperconverged system is built with HCI AF240C M7 All-NVMe nodes connected to a pair of Cisco Nexus switches and centrally managed by Cisco Intersight in standalone mode. Various configuration best practices for running Microsoft SQL Server database virtual machines on Nutanix cluster are detailed and some of the validated test cases and their results are discussed in this document.

Solution Overview

This chapter contains the following:

- [Introduction](#)
- [Audience](#)
- [Purpose of this Document](#)
- [What's New in this Release?](#)
- [Solution Summary](#)

Introduction

The new IT business models and innovations have resulted in rapid growth in development of new applications and there is a continuous push to develop and get them to the market as early as possible to take the competitive advantage. On the other hand, technologies like virtualization and containerization have augmented the pace at which new applications are being deployed in DevOps environments. The traditional siloed operating models and tools cannot keep up with such IT business demands. Therefore the IT organizations are looking for datacenter solutions and tools that can address their IT challenges.

Cisco Compute Hyperconverged with Nutanix simplifies and accelerates the delivery of infrastructure and applications, at a global scale, through best-in-class cloud operating models, unparalleled flexibility, and augmented support and resiliency capabilities. Some of the benefits offered by Cisco Compute Hyperconverged with Nutanix cloud platform are listed below.

- **Simplify infrastructure operations** with a cloud operating model that delivers visibility, control, and consistency for hyperconverged systems across highly distributed environments.
- **Effortlessly address modern applications and use cases** with a hyperconverged solution offering flexible deployment options, SaaS innovations, GPU accelerators, and network and drive technologies, plus multicloud integration.
- **Keep systems running and protected** with an augmented joint-solution support model combined with proactive, automated resiliency and security capabilities.

The consolidation of IT applications, particularly databases, has generated considerable interest in the recent years. Being most widely adopted and deployed database platform over several years, Microsoft SQL Server databases have become the victim of a popularly known IT challenge “Database Sprawl.” Some of the challenges of SQL Server sprawl include underutilized Servers, high licensing costs, security, management concerns and so on. Therefore, SQL Server databases would be the right candidates for migrating and consolidating on to a more robust, flexible, and resilient platform. This document discusses a Cisco Compute Hyperconverged with Nutanix reference architecture for deploying and consolidating SQL Server databases.

For more information on Cisco Compute for Hyperconverged with Nutanix, go to: <https://www.cisco.com/go/hci>

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, IT engineers, partners, and customers who want to take advantage of a hyperconverged infrastructure built to deliver IT efficiency and enable IT innovation. It is expected that the reader should have prior knowledge on Cisco UCS, Intersight, Nutanix cloud platform and its components.

Purpose of this Document

The purpose of this document is to provide important implementation best practices, validated use cases and their results for Microsoft SQL Server databases hosted on Cisco Compute Hyperconverged with Nutanix cloud platform.

What's New in this Release?

The following software and hardware products are used in this reference architecture.

- Microsoft SQL Server 2022 database workload validation and performance testing on Cisco Compute Hyperconverged with Nutanix.
- Tested and validated using Cisco Compute Hyperconverged HClAF240C M7 All-NVMe servers and managed in Intersight in Standalone mode.
- Nutanix Acropolis Operating System (AOS) based hyperconverged cluster running on Nutanix Acropolis Hypervisor (AHV).
- Highly performant distributed storage and compute cluster for hosting enterprise database workloads with confidence.
- Cisco Intersight Software as a Service (SaaS) for the UCS infrastructure lifecycle management, Nutanix Foundation central and Prism for deploying, managing and monitoring Nutanix Hyperconverged cluster.

Solution Summary

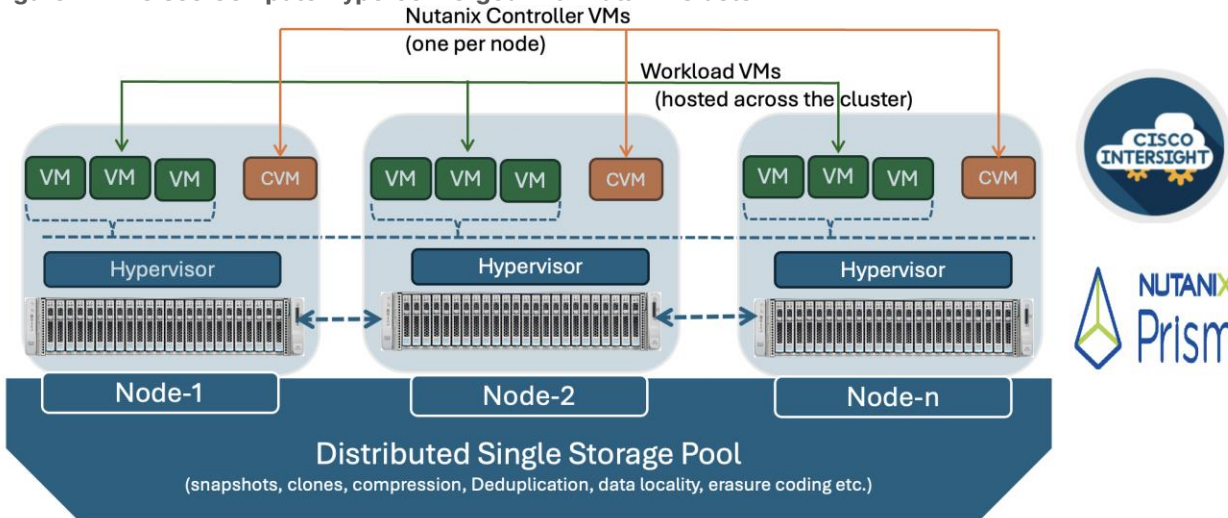
The Cisco Compute Hyperconverged with Nutanix system is built with the following hardware and software components:

- Cisco UCS HClAF240C M7 All-NVMe rack servers
- Cisco Nexus 9000 series switches
- Nutanix Acropolis Hypervisor (AHV) and Nutanix Acropolis Operating System (AOS)
- Cisco Intersight, Nutanix Prism Central and Foundation Central.
- Microsoft Windows Server 2022 and SQL Server 2022

A Nutanix cluster is a group of three or more physical nodes working as a single entity. These servers are connected to a pair of Cisco Nexus switches for internal and external communication; Each node in a Nutanix cluster includes Cisco UCS compute node configured with CPUs, memory, local attached flash or NVMe storage, and runs an industry-standard hypervisor (AHV) and a virtual storage controller called the Controller VM (CVM). The AHV hypervisor virtualizes the computational resources to provide highly available computational cluster (hypervisor cluster) while the CVMs runs Nutanix AOS that pools storage and distributes operating functions across all nodes in the cluster for performance, scalability, and resilience. The underneath UCS servers are life cycle managed using Cisco Intersight while the deployment and management of Nutanix cluster is orchestrated using Nutanix Prism components.

[Figure 1](#) describes high level overview of a standard Nutanix cluster built with Cisco UCS rack servers.

Figure 1. Cisco Compute Hyperconverged with Nutanix Cluster



Cisco Compute Hyperconverged with Nutanix platform components are connected and configured according to both Cisco and Nutanix best practices and provide robust hyperconverged platform for running a variety of enterprise workloads including databases with confidence. Nutanix clusters can be scaled out to the max cluster server limit [documented by Nutanix](#).

Cisco and Nutanix have also built a robust and experienced support team focused on Cisco Compute Hyperconverged with Nutanix system, from customer accounts and technical sales representatives to professional services and technical support engineers. The support alliance between Nutanix and Cisco gives customers and channel services partners direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

For details and specifications of the individual components, go to [Appendix B - References used in this guide](#) where all the necessary links are provided.

Technology Overview

This chapter contains the following:

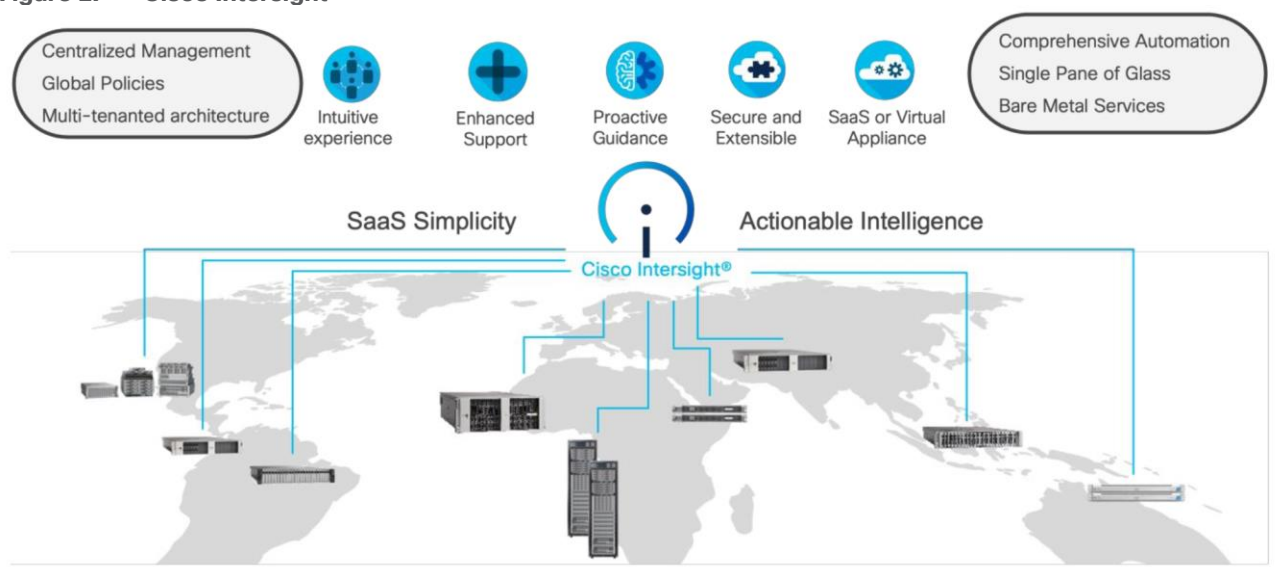
- [Cisco Intersight](#)
- [Cisco Compute Hyperconverged HClAF240C M7 All-NVMe/All-Flash Server](#)
- [Cisco Nexus Switching Fabric](#)
- [Nutanix](#)
- [Nutanix Prism](#)
- [Nutanix Foundation Central](#)
- [Nutanix AHV and AOS](#)
- [Microsoft Windows Server 2022](#)
- [Microsoft SQL Server 2022](#)

The components deployed in this solution are configured using the best practices from both Cisco and Nutanix to deliver an enterprise-class data protection solution deployed on Cisco UCS C-Series rack servers. The following sections provide a summary of the key features and capabilities available in these components.

Cisco Intersight

As applications and data become more distributed from core data center and edge locations to public clouds, a centralized management platform is essential. IT agility will be a struggle without a consolidated view of the infrastructure resources and centralized operations. Cisco Intersight provides a cloud-hosted, management and analytics platform for all Cisco Compute for Hyperconverged, Cisco UCS, and other supported third-party infrastructure deployed across the globe. It provides an efficient way of deploying, managing, and upgrading infrastructure in the data center, ROBO, edge, and co-location environments.

Figure 2. Cisco Intersight



Cisco Intersight provides:

- No Impact Transition: Embedded connector will allow you to start consuming benefits without forklift upgrade.
- SaaS/Subscription Model: SaaS model provides for centralized, cloud-scale management and operations across hundreds of sites around the globe without the administrative overhead of managing the platform.
- Enhanced Support Experience: A hosted platform allows Cisco to address issues platform-wide with the experience extending into TAC supported platforms.
- Unified Management: Single pane of glass, consistent operations model and experience for managing all systems and solutions.
- Programmability: End to end programmability with native API, SDK's and popular DevOps toolsets will enable you to deploy and manage the infrastructure quickly and easily.
- Single point of automation: Automation using Ansible, Terraform, and other tools can be done through Intersight for all systems it manages.
- Recommendation Engine: Our approach of visibility, insight and action powered by machine intelligence and analytics provide real-time recommendations with agility and scale. Embedded recommendation platform with insights sourced from across Cisco install base and tailored to each customer.

For more information, go to the Cisco Intersight product page on cisco.com.

License Requirements

The Cisco Intersight platform uses a new subscription-based license model now with two tiers. You can purchase a subscription duration of one, three, or five years and choose the required Cisco UCS server volume tier for the selected subscription duration. All Cisco UCS M7 servers require either an Essentials or Advantage license listed below. You can purchase any of the following Cisco Intersight licenses using the Cisco ordering tool:

- **Cisco Intersight Essentials:** The Essentials includes Lifecycle Operations features, including Cisco UCS Central and Cisco UCS-Manager entitlements, policy-based configuration with server profiles (IMM), firmware management, Global Monitoring and Inventory, Custom Dashboards, and evaluation of compatibility with the Cisco Hardware Compatibility List (HCL). Also, Essentials includes Proactive Support features, including Proactive RMA, Connected TAC, Advisories, and Sustainability.
- **Cisco Intersight Advantage:** Advantage offers all the features of the Essentials tier plus In-Platform Automation features such as Tunneled KVM, Operating System Install Automation, Storage/Virtualization/Network Automation, and Workflow Designer. It also includes Ecosystem Integrations for Ecosystem Visibility, Operations, Automation, and ServiceNow Integration.

Servers in the Cisco Intersight Managed Mode require at least the Essentials license. For more information about the features provided in the various licensing tiers, see

https://intersight.com/help/saas/getting_started/licensing_requirements/lic_infra.

In this solution, using Cisco Intersight Advantage License Tier enables the following:

- Configuration of Server Profiles for the Nutanix on Cisco UCS C-Series Rack Servers
- Integration of Cisco Intersight with Foundation Central for Day 0 to Day N operations

Cisco Compute Hyperconverged HClAF240c-M7 All-NVMe/All-Flash Servers

The Cisco Compute Hyperconverged HClAF240C M7 All-NVMe/All-Flash Servers extends the capabilities of Cisco's Compute Hyperconverged portfolio in a 2U form factor with the addition of the 4th Gen Intel® Xeon®

Scalable Processors (codenamed Sapphire Rapids), 16 DIMM slots per CPU for DDR5-4800 DIMMs with DIMM capacity points up to 256GB.

The All-NVMe/all-Flash Server supports 2x 4th Gen Intel® Xeon® Scalable Processors (codenamed Sapphire Rapids) with up to 60 cores per processor. With memory up to 8TB with 32 x 256GB DDR5-4800 DIMMs, in a 2-socket configuration. There are two servers to choose from:

- HCIAF240C-M7SN with up to 24 front facing SFF NVME SSDs (drives are direct-attach to PCIe Gen4 X2)
- HCIAF240C-M7X with up to 24 front facing SFF SAS/SATA SSDs

For more details, go to: [HCIAF240C M7 All-NVMe/All-Flash Server Specification sheet](#)

Figure 3. Front view: HCIAF240C M7 All-NVMe/All-Flash Servers




Cisco VIC 15238 mLOM

The Cisco UCS VIC 15237 is a dual-port quad small-form-factor pluggable (QSFP/QSFP28/QSFP56) mLOM cards designed for Cisco UCS C-series M6/M7 rack servers. The card supports 40/100/200-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

When the UCS rack server with VIC 15238 is connected to a ToR switch such as Cisco Nexus 9000 Series, the VIC adapter is provisioned through the Cisco IMC or Intersight policies for a UCS standalone server. [Figure 4](#) shows Cisco VIC 15238 MLOM network card used for this solution.

Figure 4. Cisco VIC 15238



In this solution, each Cisco HCIAF240C-M7SN node is connected to a pair of Cisco Nexus 9000 series switches using two 100Gbps physical links providing 200Gbps aggregated network bandwidth.

Cisco Nexus Switching Fabric

The Cisco Nexus 9000 Series Switches offer both modular and fixed 1/10/25/40/100 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of nonblocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.

The Cisco Nexus 93180YC-FX3 Switch is a 1RU switch that supports 3.6 Tbps of bandwidth and 1.2 Bpps. The 48 downlink ports on the 93180YC-FX3 can support 1/10/25-Gbps Ethernet, offering deployment flexibility and investment protection. The 6 uplink ports can be configured as 40 or 100-Gbps Ethernet, offering flexible migration options. The Cisco Nexus 93180YC-FX3 switch supports standard PTP telecom profiles with SyncE and PTP boundary clock functionality for telco datacenter edge environments.

The two Nexus 93180YC-FX3 switches are configured with Virtual Port Channel (vPC) feature there by two physical switches acting as single logical switch providing many benefits to the design such as improved network bandwidth, redundancy with independent control planes.

Figure 5. Cisco Nexus 93180YC-FX3 Switch



Tech tip

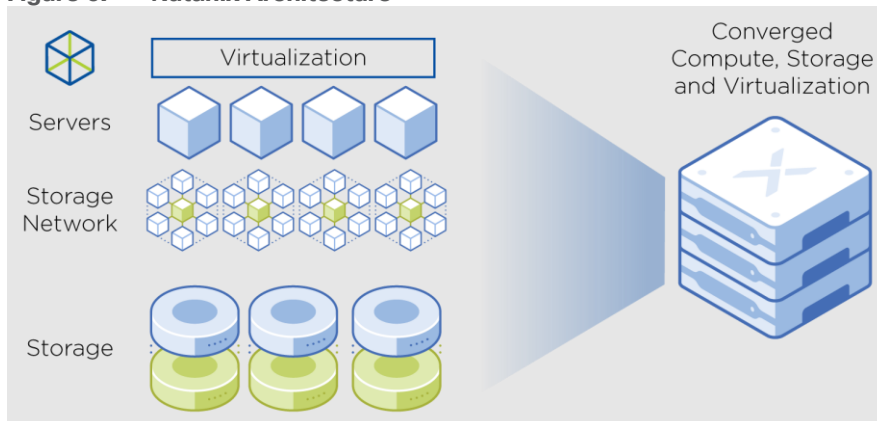
Cisco Nexus 93180YC-FX3 only supports six 100Gbps ports. If you want to expand the Nutanix cluster in the future, you need to use Cisco Nexus switches that support more than 100Gbps ports for cluster expansion. For example, Cisco Nexus 9300 series switches can be used instead. For an overview of the Cisco Nexus 9300 series switches see:

<https://www.cisco.com/c/dam/en/us/products/switches/nexus-9000-series-switches/nexus-9300-40GE-switches-comparison.html>

Nutanix

Nutanix HCI converges the datacenter stack including compute, storage, storage networking, and virtualization, replacing the separate servers, storage systems, and storage area networks (SANs) found in conventional datacenter architectures and reducing complexity. Each node in a Nutanix cluster includes compute, memory, and storage, and nodes are pooled into a cluster. The Nutanix Acropolis Operating System (AOS) software running on each node pools storage across nodes and distributes operating functions across all nodes in the cluster for performance, scalability, and resilience.

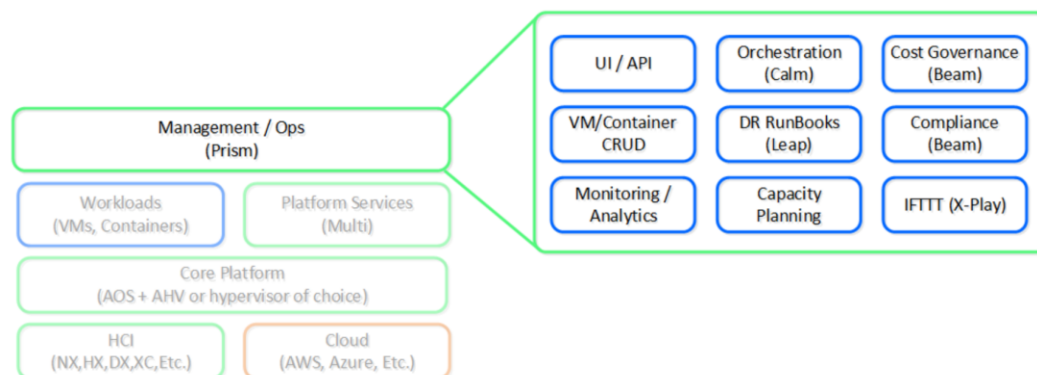
Figure 6. Nutanix Architecture



Nutanix Prism

Nutanix Prism management layer that provides central access to configure, monitor, and manage virtual environments. It uses machine learning to mine large volumes of system data easily and quickly, generating actionable insights for optimizing all aspects of virtual infrastructure management.

Figure 7. Nutanix Prism



Nutanix Prism has two core components:

- Prism Element

Prism Element is a service built into the platform for every deployed Nutanix cluster. Prism Element fully configures, manages, and monitors Nutanix clusters running any supported hypervisor.

- Prism Central

Prism Element is a service built into the platform for every deployed Nutanix cluster. Prism Element fully configures, manages, and monitors Nutanix clusters running any supported hypervisor. Because Prism Element manages only the cluster that it's part of, each deployed Nutanix cluster has a unique Prism Element instance for management. Prism Central allows you to manage different clusters across separate physical locations on one screen and gain an organizational view into a distributed Nutanix environment.

Further information about Prism Central can be found using the [Prism Central Guide](#).

For an overview, see the [Nutanix Prism Tech Note](#).

Nutanix Foundation Central

Nutanix Foundation Central allows the creation of clusters from factory-imaged nodes and the reimage of existing nodes that are already registered with Foundation Central or remotely from Prism Central.

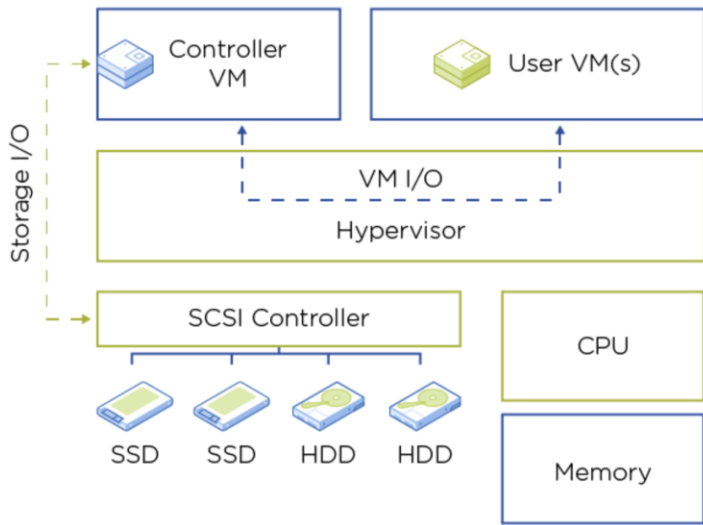
For more information about Foundation Central, see the [Foundation Central Guide](#)

Nutanix AHV and AOS

AHV

AHV is the Nutanix-native hypervisor that natively converges compute and storage into a single application. It offers powerful virtualization capabilities—including core virtual machine (VM) operations, live migration, VM high availability, and virtual network management—as fully integrated features of the infrastructure stack rather than standalone products that require separate deployment and management. For further information, go to [Nutanix AHV Virtualization](#).

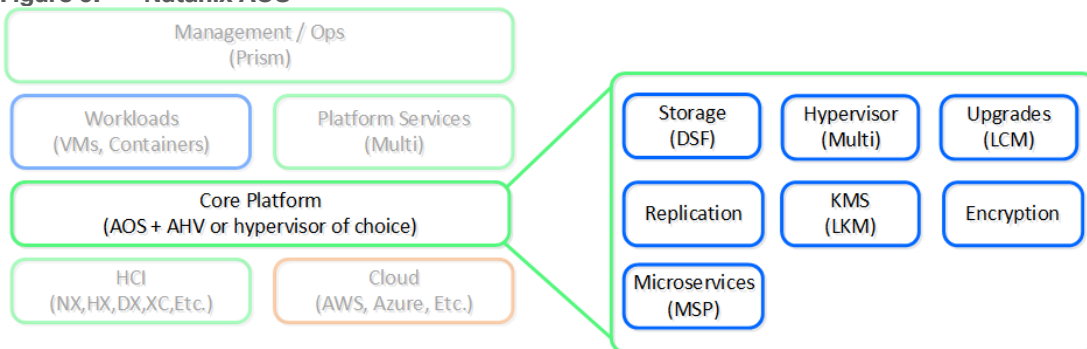
Figure 8. Nutanix AHV Node Architecture



AOS

The Acropolis Operating System (AOS) provides the core functionality leveraged by workloads and services running on the platform. It utilizes a distributed approach that combines the storage resources of all nodes in a Nutanix cluster to deliver the capabilities and performance that you expect from SAN storage while eliminating much of the cost, management overhead, and hassle that comes with managing traditional storage. Intelligent software enables AOS Storage to appear on a hypervisor—such as VMware ESXi or Nutanix AHV—as a single, uniform storage pool. Over the years, Nutanix has expanded its features and capabilities, making AOS Storage a leader in software-defined distributed storage. Each node in a Nutanix cluster runs a VM called the Controller Virtual Machine (CVM) that runs the distributed storage services as well as other services necessary for a cluster environment. Because storage and other Nutanix services are distributed across the nodes in the cluster, no one entity is a single point of failure. Any node can assume leadership of any service. To learn more about the underlying components of AOS Storage, go to the [Nutanix Bible](#).

Figure 9. Nutanix AOS



Microsoft Windows Server 2022

Windows Server 2022 is the latest OS platform release from Microsoft. Windows Server 2022 is an excellent platform for running Microsoft SQL Server 2022 databases. It offers new features and enhancements related to security, patching, domains, clusters, storage, and support for various new hardware features, and so on. It

enables Windows Server to provide best-in-class performance and a highly scalable platform for deploying SQL Server databases.

Microsoft SQL Server 2022

SQL Server 2022 (16.x) is the latest relational database from Microsoft and builds on previous releases to grow SQL Server as a platform that gives you choices of development languages, data types, on-premises or cloud environments, and operating systems. It offers various enhancements and new features that enables SQL Server deployments to be more reliable, highly available, performant, and secured than ever. SQL Server 2022 can leverage new hardware capabilities from partners like Intel to provide extended capabilities. For example, now it can leverage Intel Quick Assist Technology (QAT) for offloading backup compression thereby improving backup and restore performance. For more details about the new capabilities of SQL Server 2022, go to: [Microsoft SQL Server 2022](#)

Solution Design

This chapter contains the following:

- [Physical Topology](#)
- [Software Components](#)

The Cisco Compute Hyperconverged with Nutanix system in this CVD was designed to address the following key goals:

- Resilient design across all the layers of the infrastructure with no single point of failure.
- Highly performing and scalable design with ability to independently scale the compute, storage, and network bandwidth as needed.
- Enterprise-grade storage features and optimizations.
- Best-practices based design, incorporating design, technology, and product best practices.

Physical Topology

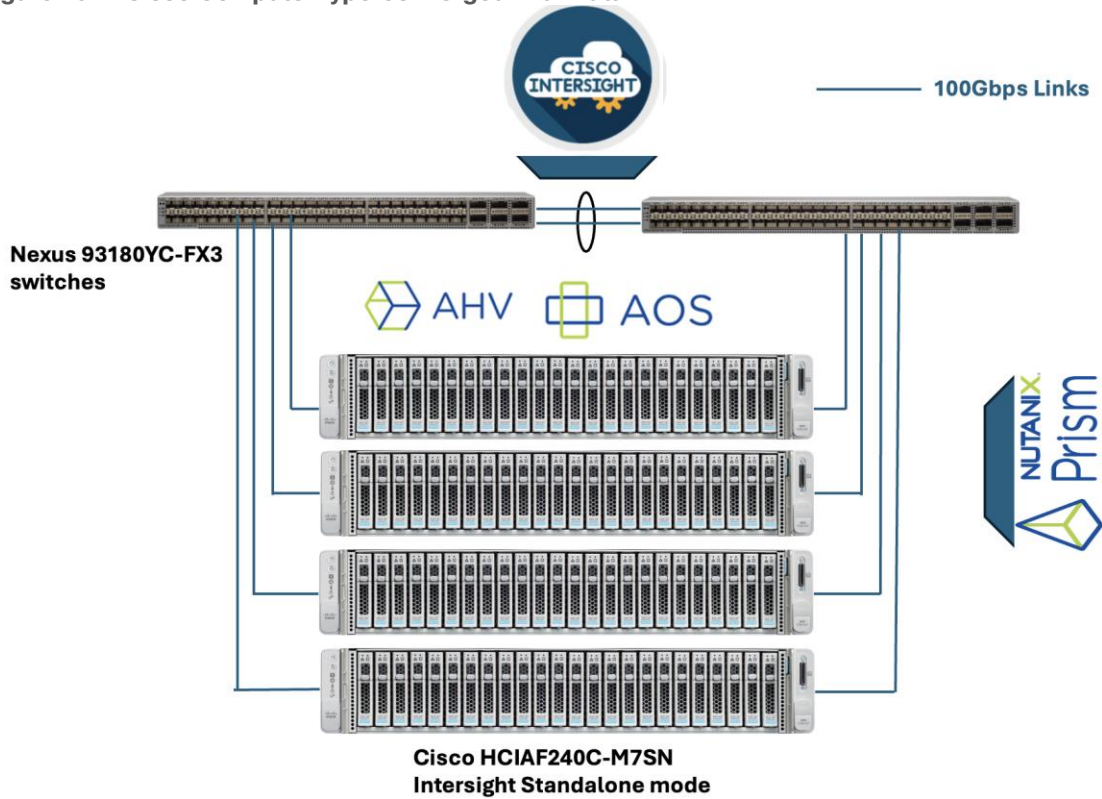
The physical topology of the Cisco Compute Hyperconverged with Nutanix in Intersight standalone mode (ISM) is detailed in below figure. The entire Day0 deployment is managed through Cisco Intersight and Nutanix Foundational Central enabled through Prism Central.

Each HCIAF240C-M7 SN node is configured with following hardware components.

- 2x Intel® Xeon® Gold 8462Y CPUs. Each CPU has 32 cores running at 2.8GHz base frequency
- 1024GB DDR5 memory
- 2x 240GB M.2 card managed through M.2 RAID controller
- 6x 3.8TB 2.5-inch U.2 NVMe SSD disks (scalable up to 24)
- 1x Cisco VIC 15238 2x 40/100Gbps PCIe mLOM card with Secure Boot option

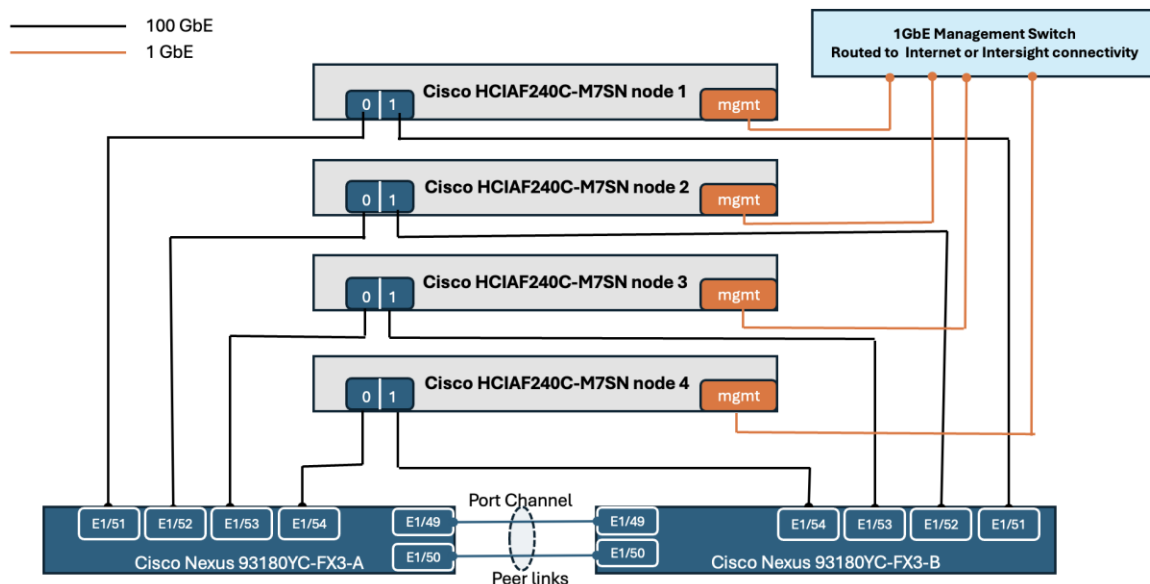
Note: This document illustrates the Cisco HCIAF240C M7 All-NVMe Servers specifications as validated in this document. You have several options to configure CPU, Memory, Network cards, GPUs and storage as detailed in this spec sheet: [Cisco Compute Hyperconverged with Nutanix](#)

Figure 10. Cisco Compute Hyperconverged with Nutanix



In this CVD, the Nutanix cluster is deployed on the HClAF240C-M7SN servers which are managed by Intersight in Intersight Standalone Mode (ISM). The Intersight Standalone Mode requires the Cisco UCS C-Series Rack Servers to be directly connected to ethernet switches and the servers are claimed through Cisco Intersight. Once the servers are claimed into Cisco Intersight, the Nutanix cluster will be installed on these nodes using Nutanix Prism and Foundation Central software components. [Figure 11](#) shows the physical cabling connectivity of the four-node Cisco Compute Hyperconverged with Nutanix.

Figure 11. Cabling Diagram



Software Components

[Table 1](#) lists the software components and the versions validated for in this CVD.

Table 1. Software Components of Cisco Compute Hyperconverged with Nutanix

Component	Hardware
Foundation Central	1.6
Prism Central Deployed on an external ESXi cluster	pc.2022.6.0.10
AOS and AHV bundled	nutanix_installer_package-release-fraser-6.5.5.6
Cisco C240 M7 All NVMe server	4.3(3.240043)
VirtIO Driver	1.2.3-x64

Install and Configure

This chapter contains the following:

- [Prerequisites](#)
- [Cisco IMC Configuration](#)
- [Cisco Intersight Configuration and Keys](#)
- [Claim Servers on Cisco Intersight](#)
- [Prism Central Installation and Configuration](#)
- [Configure Foundation Central](#)
- [Onboard Servers and Create Nutanix Cluster](#)
- [Post Cluster Installation Tasks](#)
- [Create Virtual Machine for SQL Server Instances](#)
- [Install and Configure SQL Server](#)
- [Monitor SQL Server VMs using Prism Element](#)

This chapter describes the solution deployment for Nutanix on Cisco UCS C-Series Rack Servers in the Intersight Standalone Mode (ISM), with step-by-step procedures for implementing and managing the deployment.

Note: This CVD focuses on the important deployment and configuration steps which are more relevant from point of view of hosting and running Microsoft SQL Server database workloads on Nutanix cluster. It is not intended to provide detailed steps for deploying Nutanix on Cisco UCS servers.

Detailed step-by-step procedures for deploying Nutanix on Cisco UCS C-Series Rack Servers are provided in the base infrastructure CVD: [Cisco Compute Hyperconverged with Nutanix in Intersight Standalone Mode Deployment Guide](#).

Prerequisites

Prior to beginning the installation of Nutanix Cluster on Cisco UCS C-Series servers in Intersight Standalone Mode, ensure Nutanix Prism Central is deployed and enable Nutanix Foundation Central through Nutanix marketplace available through Prism Central. Foundation Central can create clusters from factory-imaged nodes and reimage existing nodes that are already registered with Foundation Central from Prism Central. This provides benefits such as creating and deploying several clusters on remote sites, such as ROBO, without requiring onsite visits.

At a high level, to continue with the deployment of Nutanix on Cisco UCS C-Series servers in Intersight standalone mode (ISM), ensure the following:

- Cisco Intersight account must be created, and Intersight Advantage license must be configured
- Prism Central is deployed on either an external Nutanix Cluster or ESXi cluster/node
- Foundation Central 1.6 or later is enabled on Prism Central

- A local webserver or http file share must be available for hosting the Nutanix AOS image and must be reachable from the Cisco IMC network. Download the AOS image and store into the file share:
<https://portal.nutanix.com/page/downloads?product=nos>

Cisco IMC Configuration

Cisco Integrated Management Controller (CIMC) is used for management/monitoring of C-Series Rack servers. CIMC provides options like WebGUI, CLI and IPMI for management/monitoring tasks. Detailed steps to configure an IP Address on CIMC management port are provided here: <https://community.cisco.com/t5/data-center-and-cloud-knowledge-base/configure-or-change-cimc-ip-address-on-ucs-c200-series-servers/ta-p/3141563>. Use this link to configure the CIMC IP address on all four servers.

Cisco Intersight Configuration and Keys

Follow the procedures in this section to enable the software download option for downloading required firmware from cisco.com. You also need to create Intersight API Keys which will be used by Nutanix Foundation Central to make API calls to Cisco Intersight to create the server profiles and associate them to the servers.

Procedure 1. Activate Software Download option and create API keys

Step 1. Log into Intersight and go to System > Settings > Cisco ID and from the Cisco software download option click Activate. Once activated a login window will be displayed. Login with your Intersight credentials and click Generate. This will activate the software download option.

Step 2. To create API keys, Logon to Intersight and go to System > Settings > API Keys and click Generate API Key. On the Generate API Key window, provide a description, select API Key version 3 and provide an expiry date for the key. Click Generate.

Generate API Key [X]

i Generate API key to authenticate and authorize API requests with Intersight account.

Description * ⓘ
API Key for FC

API Key Purpose ⓘ

API key for OpenAPI schema version 3 ⓘ
 API key for OpenAPI schema version 2 (deprecated) ⓘ

Expiration Time * ⓘ
Dec 31, 2024 01:23 AM [Calendar Icon]

Close Generate

Step 3. Once API key generated, save the API key and Secret key at a secure place. These keys will be used in the Foundation Central.

Claim Servers on Cisco Intersight

The following high-level steps describe the process to claim servers on Cisco Intersight. Ensure CIMC of all servers have been configured with proper DNS for Cisco Intersight reachability. All the nodes that part of the

Nutanix cluster must be claimed into Intersight. Perform the below steps on all the servers that are going to be the part of Nutanix cluster.

Procedure 1. Claim UCS C-Series Servers on Intersight

Step 1. Log into the server CIMC session using its CIMC IP and credentials which are set during the CIMC IP configuration.

Step 2. Go to Admin > Device Connector. On the Device Connector page, click Settings and provide the DNS, Proxy details for the server to be able reach Intersight through internet. After that, collect the Device ID and Claim code. This will be used to claim the server on Cisco Intersight.

Step 3. Log into the Cisco Intersight and navigate to System > Targets.

Step 4. Click All and select Cisco UCS Server (Standalone) option to claim C-Series rack server in standalone mode.

Step 5. Provide Device ID and Claim Code collected from the previous step. Click Claim.

Step 6. Repeat steps 1 - 5 on all the four servers. Once claimed, all the servers should be discovered in the Intersight as shown below. You can view all the claimed devices from System > Targets

<input type="checkbox"/>	Name	Health	Status	Type	Claimed Time	Claimed By
<input type="checkbox"/>	C240-WZP27510W0V	Healthy	Connected	Standalone M7 Server	May 10, 2024 12:12 AM	rgopunar@cisco.com
<input type="checkbox"/>	C240-WZP28030WVD	Healthy	Connected	Standalone M7 Server	May 10, 2024 12:16 AM	rgopunar@cisco.com
<input type="checkbox"/>	C240-WZP28030WVP	Healthy	Connected	Standalone M7 Server	May 10, 2024 12:19 AM	rgopunar@cisco.com
<input type="checkbox"/>	C240-WZP28030WVQ	Healthy	Connected	Standalone M7 Server	May 10, 2024 12:22 AM	rgopunar@cisco.com

Prism Central Installation and Configuration

This section provides the procedures to deploy Prism Central on an external ESXi cluster.

Procedure 1. Download and deploy Prism Central

Step 1. Download Prism Central (PS) for ESXi from <https://portal.nutanix.com/page/downloads?product=prism>

Step 2. Identify an ESXi host and deploy the OVA file downloaded in Step 1. Provide the required inputs for creating and running the Prism Central virtual machine on the identified ESXi node as shown below.

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Select networks
- 7 Ready to complete

Ready to complete

Review your selections before finishing the wizard

- ▼ Select a name and folder

Name	pc.2022.6.0.10
Template name	pc.2022.6.0.10
Folder	RTP-B4-Infra
- ▼ Select a compute resource

Resource	AA06
----------	------
- ▼ Review details

Download size	5.5 GB
---------------	--------
- ▼ Select storage

Size on disk	570.0 GB
Storage mapping	1
All disks	Datastore: infra_datastore_02; Format: Thick provision lazy zeroed
- ▼ Select networks

Network mapping	1
VM Network	Tenant6_10_106_1_NET

IP allocation settings

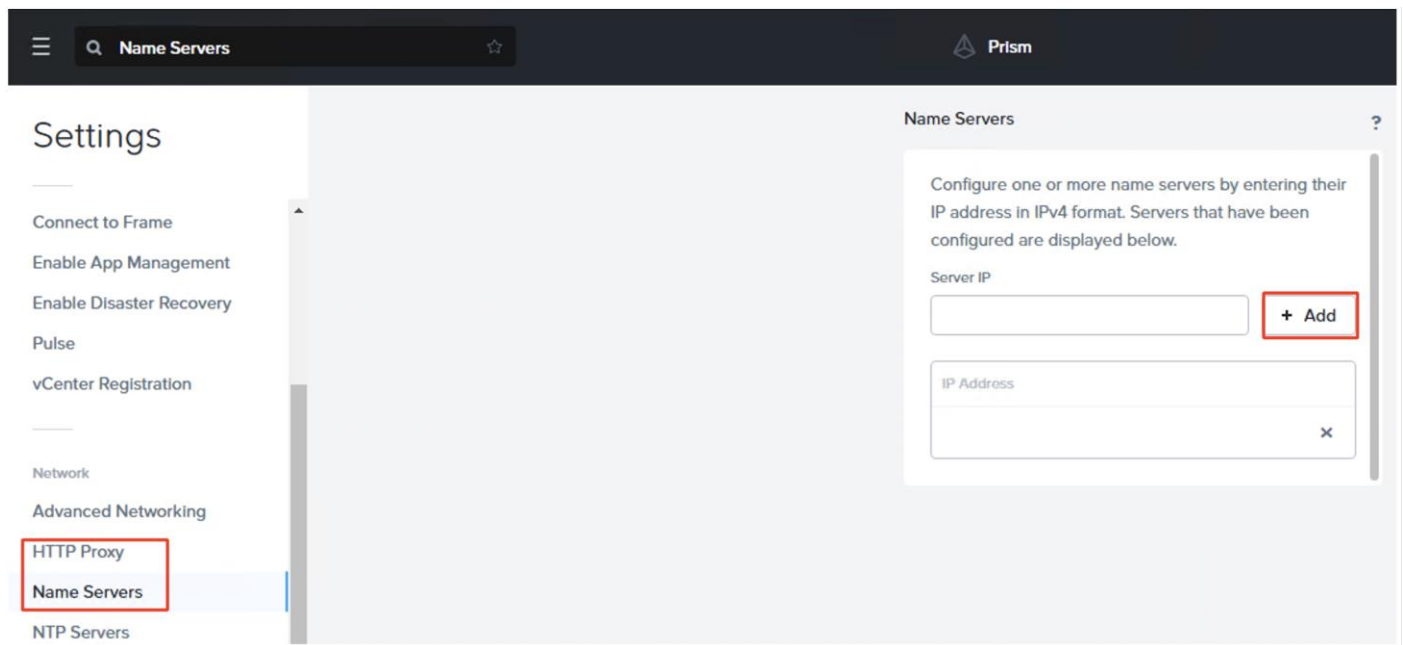
IP protocol	IPv4
IP allocation	Static - Manual

CANCEL BACK FINISH

Step 3. Once OVA is deployed, power on the VM. Perform the Post installation steps as detailed here: https://portal.nutanix.com/page/documents/details?targetId=Acropolis-Upgrade-Guide-v6_5:upg-vm-install-wc-t.html. Follow the instructions to set the IP address and create the prism cluster.

Step 4. Once completed, log into the Prism Central by using the default password and changing the default password options.

Step 5. The Prism Central and Foundation Central should be able to reach Intersight. If Prism Central is deployed behind a proxy, you need to set the DNS and Proxy. To do so, go to Prism Central Settings > Name Servers and HTTP Proxy, provide these two details, and click Add.



Procedure 2. Enabling and upgrading Foundation Central (FC) using Prism Central

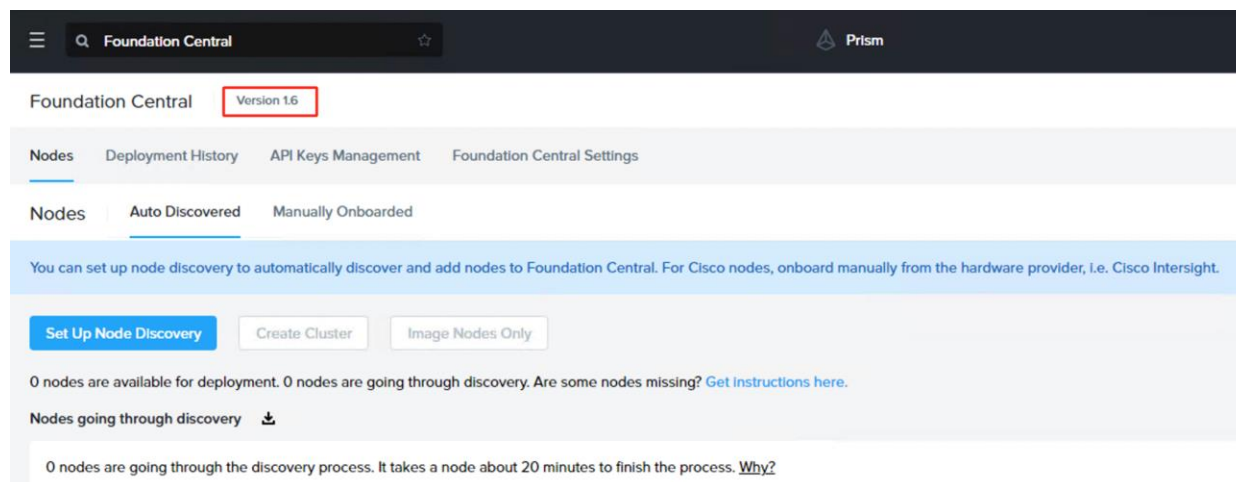
Step 1. On the Prism Central portal, go to Services > Foundation Central and click Enable Foundation Central.

Step 2. Download the latest FC bundles here:
<https://portal.nutanix.com/page/downloads?product=foundationcentral>.

Note: FC v1.6 is the version we used when documenting this deployment guide.

Step 3. Upgrade the FC to the version v1.6 by following the detailed steps provided here:
https://portal.nutanix.com/page/documents/details?targetId=Foundation-Central-v1_6:v1-upgrade-fc-cli-t.html

Once FC is successfully upgraded to the latest version, the deployment options will be displayed as shown below:



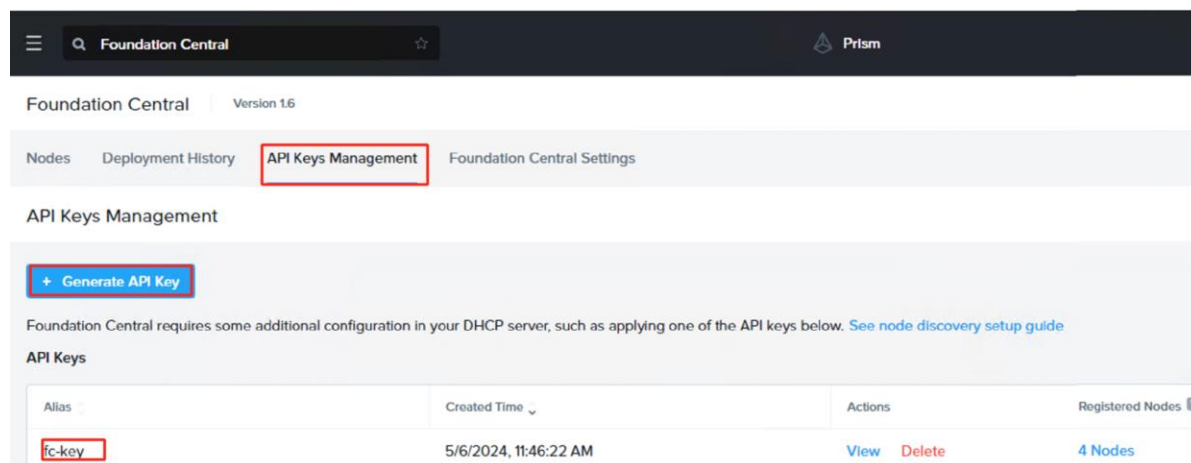
Configure Foundation Central

This procedure describes the Foundation Central configuration required for a successful Nutanix cluster creation with Cisco UCS C-Series nodes in ISM mode.

Note: API key authenticates API communication between the Nutanix node and Foundation Central. It is recommended that you create a unique key for each remote site.

Procedure 1. Generate FC API Keys

Step 1. Log into Prism Central and navigate to Foundation Central > API Keys Management. Click on Generate API Key button. The API Keys will be displayed. It will be added during the Nutanix cluster creation through FC.

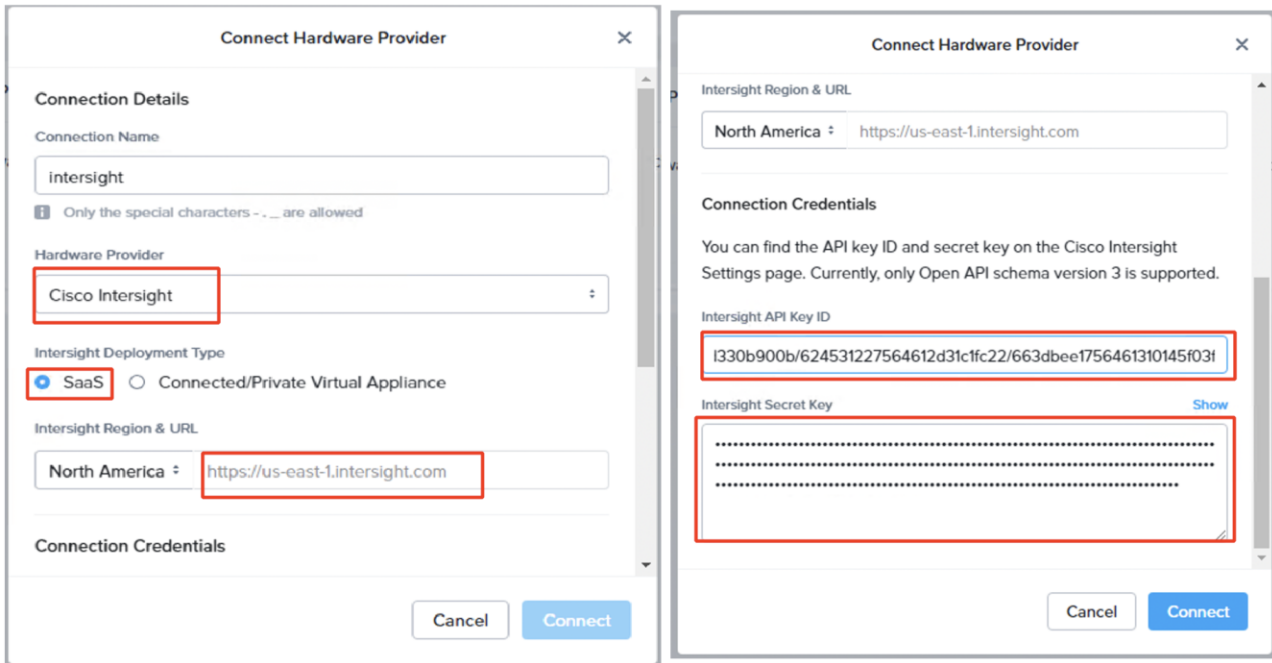


Procedure 2. Connect Foundation Central to the Intersight

Follow these steps to onboard Intersight into the FC using the Intersight API keys. This allows the FC to interact with Intersight and discover the nodes claimed on the Intersight.

Step 1. Log into Prism Central and navigate to Foundation Central > Foundation Central Settings. Click Connect Hardware Provider.

Step 2. Provide a friendly name under the Connection Name text box, select Cisco Intersight for the hardware provider, and select SaaS for the deployment type. The Intersight URL will be automatically populated when the SaaS option is selected. Provide the API Key and Secret Keys you gathered from Intersight and click Connect.



Once the authentication succeeds, the connection details will be displayed under the Foundation Central Settings tab.

Procedure 3. Onboard Servers and Create Nutanix Cluster

This procedure describes the process to onboard the nodes on Foundation Central and thereafter create the cluster for Cisco UCS C-Series nodes managed in Intersight Standalone Mode (ISM).

Step 1. Go to Foundation Central, select the Nodes tab and click the Manually Onboard tab. Click Onboard Nodes. A screen will be displayed with Cisco Intersight connection details configured in the previous step. Select Intersight and click Next.

Step 2. FC will connect to Cisco Intersight and fetches all the unconfigured nodes claimed in the Intersight. Select the nodes provisioned for Nutanix and click Onboard. Once the servers are onboarded into the FC, select the nodes and click Create Cluster.

Congratulations! 90 days of Prism Ultimate is on us, enjoy! [See what's included here.](#)

Foundation Central
Prism

Foundation Central Version 1.6

Nodes Deployment History API Keys Management Foundation Central Settings

Nodes Auto Discovered Manually Onboarded

Onboard Nodes
Create Cluster
Actions ▾
View by : Group by :

Type in a query Modify Filters

4 selected out of 4 Cisco Nodes 1 - 4 of 4 | 20 rows

Node Serial	Name	Node Status	Model	Organizations	Tags	Classification	Onboarded On	
<input checked="" type="checkbox"/>	WZP28030WVD	C240-WZP28030WVD	Onboarded	UCSC-C240-M7SN	default	-	Nutanix	May 9, 2024, 11:36 PM
<input checked="" type="checkbox"/>	WZP28030WVP	C240-WZP28030WVP	Onboarded	UCSC-C240-M7SN	default	-	Nutanix	May 9, 2024, 11:36 PM
<input checked="" type="checkbox"/>	WZP27510W0V	C240-WZP27510W0V	Onboarded	UCSC-C240-M7SN	default	-	Nutanix	May 9, 2024, 11:36 PM
<input checked="" type="checkbox"/>	WZP28030WVQ	C240-WZP28030WVQ	Onboarded	UCSC-C240-M7SN	default	-	Nutanix	May 9, 2024, 11:36 PM

Step 3. On the Cluster Details tab, enter the name of the Nutanix cluster, Cluster Replication Factor (RF) will be set to 2 as minimum of five nodes are required to select RF between 2 and 3. Set Intersight Organization to default and click Next.

Step 4. From the Hypervisor/AOS tab, select first radio button and provide the http file share location where the AOS 6.5.5.6 image downloaded and stored. Select AHV for the Hypervisor and check the box stating AOS and AHV are bundled together into a single image as shown below. Click Next.

1 Cluster Details
2 Hypervisor / AOS
3 Network Settings
4 CVM Settings
5 Configure Nodes
6 Security

● I want to image all nodes with custom AOS and hypervisor installers.

AOS Download URL

http://10.106.1.2/images/Nutanix/nutanix_installer_package-release-frasei

AOS package will be downloaded from this location during the imaging process of each node.

Pre-computed SHA256 Checksum of AOS Installer (Optional)

Provide the checksum if you want Foundation Central to run a checksum test of the downloaded installer before starting the imaging process. This checksum also helps Foundation Central skip a re-download if a file with the same checksum already exists.

Hypervisor

AHV

Choose the hypervisor type that you want to install on each node.

Instead of downloading AHV from a URL, I want to use the AHV bundled inside the AOS package specified above (please ensure your AOS package bundles an AHV installer)

Hypervisor Download URL

Hypervisor installer will be downloaded from this location during the imaging process of each node.

Step 5. Under the Network Settings tab, provide gateway, subnet, and Cluster IP details. The network ports on the servers can be in access or trunk mode. Use trunk mode to allow multiple VLANs for the different traffics (management, guest VMs, and so on) Providing a VLAN here will configure the server port as trunk port and allow multiple traffics. For this deployment, VLAN 1061 is used for all management and guest traffic for the keeping the deployment simple.

You have a choice to enable LACP with AHV. Default mode is active-backup. Go to: <https://portal.nutanix.com/page/documents/solutions/details?targetId=BP-2071-AHV-Networking:bp-ahv-networking-best-practices.html>. Enable LACP after cluster configuration and is supported only during re-imaging of nodes.

In the event of LACP configuration failure wherein cluster creation fails with error as “Failed to receive the first heartbeat of the node.” For resolution, go to: <https://portal.nutanix.com/page/documents/kbs/details?targetId=kA0VO0000001w0L0AQ>

LACP is not enabled for this cluster and the default CL91 or AUTO (RS-FEC) is selected which is applicable for 10/40/50/100 GbE cables and optics.

The screenshot shows a configuration page with a breadcrumb trail: 1 Cluster Details, 2 Hypervisor / AOS, 3 Network Settings, 4 CVM Settings, 5 Configure Nodes, 6 Security. The 'Network Settings' section is active.

Host and CVM Network

Nutanix requires all hosts and CVMs of a cluster to have static IP addresses in the same subnet.

Gateway of Every Host and CVM Reuse Existing
10.106.1.254

Netmask of Every Host and CVM
255.255.255.0

Cluster Virtual IP (Optional)
10.106.1.30

This IP will always point to an online node, even in case of a node failure. Must be in the host-CVM subnet. Your subnet range is: 10.106.1.0 - 10.106.1.255

i If you plan to deploy Nutanix Objects, click here to learn about important network requirements.

Host and CVM VLAN

If your host-CVM subnet has a VLAN configuration, enter the tag below. All packets leaving the hosts and the CVMs will be wrapped with this VLAN tag.

VLAN Tag of Every Host and CVM (Optional)
1061

Minimum 1, maximum 4094, if left blank, VLAN 0 will be used.

Hypervisor LACP Configuration

Enable LACP

FEC Mode for VIC Adapter

The FEC mode on the VIC adapter must match what is configured on the ports on the uplink switch. Setting the FEC mode to cl91 is suitable for most cases, but a different value may be required for some switches and transceivers/cable combinations. Check the configuration of your specific networking equipment to determine what mode to use.

cl91

< Back Next >

As mentioned in the cabling diagram, each server is connected to a pair of Cisco Nexus switches which are configured with vPC. The port configuration on the Nexus switches where four servers are connected to both the switches is shown below:

```

AA06-C93180YC-FX3-1#
AA06-C93180YC-FX3-1# show running-config interface ethernet 1/51-54
!Command: show running-config interface Ethernet1/51-54
!Running configuration last done at: Tue Jul 9 11:52:54 2024
!Time: Tue Jul 9 11:56:58 2024

version 9.3(7) Bios:version 01.08

interface Ethernet1/51
description NX-C240M7-A-N1
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

interface Ethernet1/52
description NX-C240M7-A-N2
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

interface Ethernet1/53
description NX-C240M7-A-N3
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

interface Ethernet1/54
description NX-C240M7-A-N4
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

AA06-C93180YC-FX3-2#
AA06-C93180YC-FX3-2# show running-config interface ethernet 1/51-54
!Command: show running-config interface Ethernet1/51-54
!Running configuration last done at: Tue Jul 9 11:56:04 2024
!Time: Tue Jul 9 11:57:43 2024

version 9.3(7) Bios:version 01.08

interface Ethernet1/51
description NX-C240M7-B-N1
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

interface Ethernet1/52
description NX-C240M7-B-N2
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

interface Ethernet1/53
description NX-C240M7-B-N3
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

interface Ethernet1/54
description NX-C240M7-B-N4
switchport mode trunk
spanning-tree port type edge trunk
mtu 9216
no shutdown

AA06-C93180YC-FX3-2#

```

Step 6. On the CVM settings tab, set 64Gb memory for the controller VMs. 64GB memory is recommended for controller VMs for the Nutanix cluster that hosts critical enterprise workloads like databases. Provide time zone, NTP server, and DNS servers details. Click Next

Step 7. From the Configure Nodes tab, click Set Range for Host IP, CVM IP, and Hostname and provide values along with a number as a suffix for the first node. The IPs and host names for the other nodes will be automatically populated, as shown below. Click Next.

The screenshot shows the 'Create Deployment' wizard in Foundation Central, specifically the 'Set up Hostname Range' step. The wizard progress is: 1 Cluster Details, 2 Hypervisor / AOS, 3 Network Settings, 4 CVM Settings, 5 Set up Hostname Range. Below the progress bar, there is a text input field for 'Enter the IP/hostnames you want each node to have.' Below that is a table with the following data:

Node Serial	Node Name	Host IP Set Range	CVM IP Set Range	Set Range
WZP27510W0V	C240-WZP27510W0V	10.106.1.31	10.106.1.35	NXSQ1
WZP28030WVP	C240-WZP28030WVP	10.106.1.32	10.106.1.36	NXSQ2
WZP28030WVQ	C240-WZP28030WVQ	10.106.1.33	10.106.1.37	NXSQ3
WZP28030WVD	C240-WZP28030WVD	10.106.1.34	10.106.1.38	NXSQ4

A modal dialog titled 'Set up Hostname Range' is open, showing the instruction: 'Set up an incremental hostname for the entire column in the input below.' The input field contains 'NXSQ1' and a 'Clear' button is visible to the right. At the bottom of the wizard, there are 'Back' and 'Next' buttons.

Step 8. From the Security tab, select the Foundation Central API key created in the previous steps and click on Submit to start the deployment of Nutanix. Once the cluster deployment completes successfully, log into the Prism Element by clicking the Open Prism Element link and complete post-installation tasks such as cluster configuration, storage container creation, VM network configuration, and so on.

Foundation Central | Prism

Nodes | **Deployment History** | API Keys Management | Foundation Central Settings

Deployment History

It may take up to 15 minutes for a deployment progress to start being reported. [Why?](#) Viewing 15 deployments Show Only Non-archived

NTNX-SQL | Deployment complete | Start Date and Time: 5/31/2024, 11:36 AM Open Prism Element Archive

Phase 1A: Node Preparation 100% 4 nodes prepared | Phase 1B: Node Imaging 100% 4 nodes finished | **Phase 2: Cluster Formation 100% All operations completed successfully**

Cluster Details

Redundancy Factor 2	Host-CVM Subnet 10.106.1.254 / 255.255.255.0	CVM NTP Servers 172.20.10.11, 172.20.10.12	AOS Installer URL http://10.106.1.2/images/Nutanix/nutanix_installer_package-release-fraser-6.5.5.6-stable-e36389d5f36fa8e7cacfdccdc2049f85206c6d7c-x86_64.tar.gz
Cluster External IP 10.106.1.30	Intersight Organization default	CVM DNS Servers 10.106.1.10	Hypervisor Installer URL Not provided
CVM VLAN Tag 1061	Deployment UUID c1fc2c71-a559-4170-551a-451d0542ec56	LACP No	

Step 9. Log into Intersight and look for the UCS C-Series servers, policies, server profiles, and so on, created as part of the Day 0 deployment and associated with the C240CM7 servers.

Intersight | Infrastructure Service

Servers

8 items found | 18 per page | 1 of 1

Health 8 Critical 4 Healthy 4 | **Power** Off 1 On 7 | **HCL Status** Incomplete 5 Not Listed 3 | **Bundle Version** 5.2(0.230041) 4 4.3(3.240043) 4 | **Firmware Version** 5.2(0.230041) 4 4.3(3.240043) 4 | **Models** 8

Name	Health	Manag...	Model	U.	Server Profile
C240-WZP27510W0V	Healthy	10.106.0.13	UCSC-C240-M7SN		ntnx_WZP27510W0V_c1fc2c71-a559-4170-551a-451d...
C240-WZP28030WVD	Healthy	10.106.0.15	UCSC-C240-M7SN		ntnx_WZP28030WVD_c1fc2c71-a559-4170-551a-451d...
C240-WZP28030WVP	Healthy	10.106.0.14	UCSC-C240-M7SN		ntnx_WZP28030WVP_c1fc2c71-a559-4170-551a-451d...
C240-WZP28030WVQ	Healthy	10.106.0.11	UCSC-C240-M7SN		ntnx_WZP28030WVQ_c1fc2c71-a559-4170-551a-451d...

Post Cluster Installation Tasks

The following steps provide the list of recommended settings used for this solution.

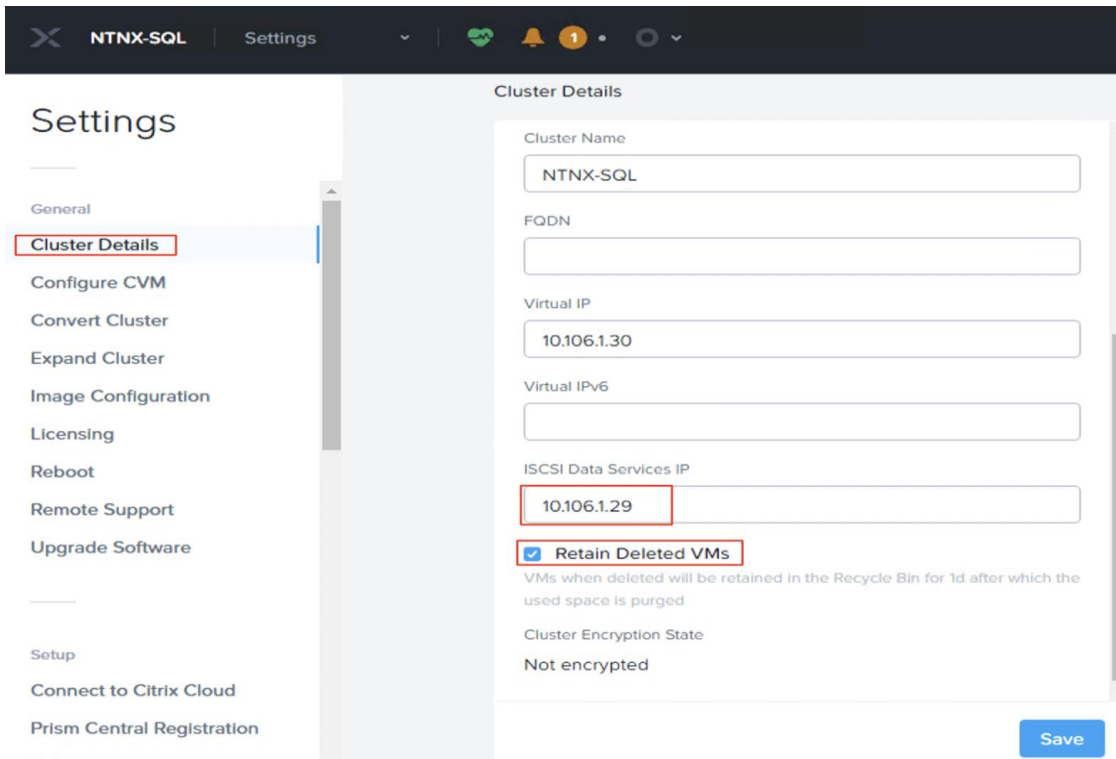
Procedure 1. Post Cluster Creation Tasks

Step 1. Log into Prism Element with 'admin' user, default password nutanix/4u and change the default password.

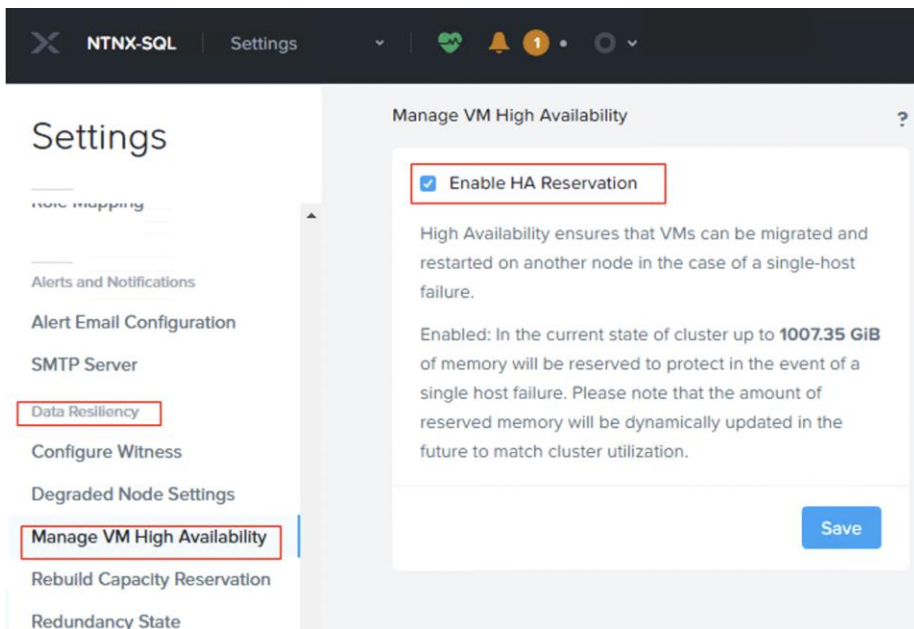
Step 2. After logging into the Prism Element, select Storage from the drop-down list and create a storage container as shown below.

Step 3. Enable the Rebuild Capacity Reservation option as shown below.

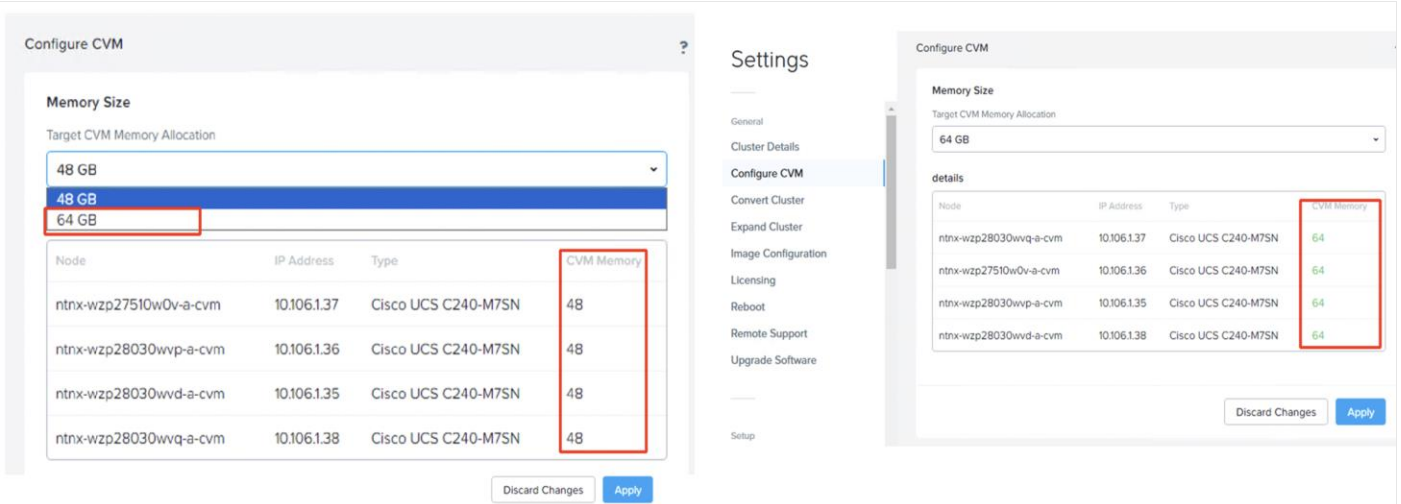
Step 4. Go to Cluster details and enter iSCSI data services IP and enable Retain Deleted VMs for 1 day. Click Save.



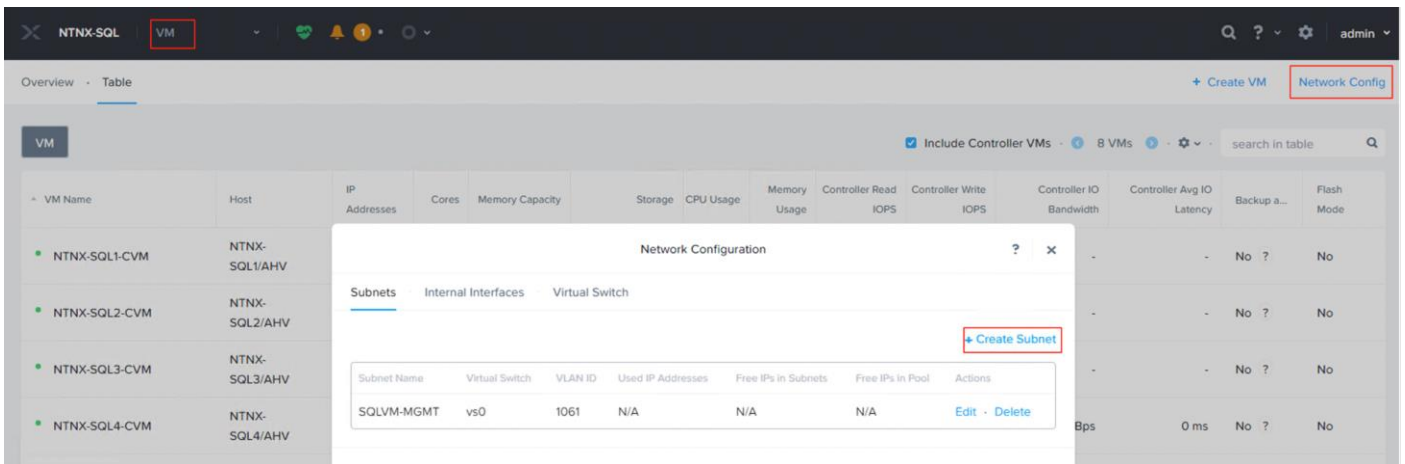
Step 5. Go to Settings > Manage VM High Availability, Enable HA Reservation and click Save.



Step 6. If the cluster is hosting enterprise-critical workloads, it is recommended that the Controller VM memory be changed from the default 48GB to 64GB. If 64GB memory is already selected during the Nutanix installation, this step can be ignored. Otherwise, go to Settings > Configure CVM, select 64GB from the drop-down list, and click Apply. Wait for memory changes to be applied to all the CVMs.




Step 7. Create Subnet under default Virtual Switch vs0 for virtual machines to connect to the network. Go to VM > Network Config > Create Subnet. Set the VLAN as 1061 and select default Virtual Switch vs0.



Step 8. Newer storage media such as NVMe have user space libraries such as SPDK (Storage Performance Development Kit) to manage the device I/O directly eliminating the need to make system calls which avoid context switches there by improving overall IO performance. Blockstore enables AOS to leverage Intel SPDK for direct access of NVMe backed disks. SPDK is automatically enabled on a cluster when Blockstore is active on the cluster with NVMe. To ensure that Blockstore with SPDK is enabled, ssh into any controller VM and run the below command. If you see any NVMe devices in the /dev/spdk/ table, then the Blockstore with SPDK is enabled.

```
nutanix@cvm$ allssh ls /dev/spdk/*
```

The following screenshot shows that each node has 6 NVMe disks supporting Blockstore with SPDK:



Step 9. Optionally, the Prism Element can be registered with Prism Central which enables us to monitor and manage multiple Nutanix clusters from Prism Central. Log into the Prism Element, go to Settings > Prism Central

Registration and click Register. Select the second option I already have Prism Central Instance Deployed and provide the Prism Central IP, port and credentials and click Connect.

Step 10. Run the NCC check and resolve warnings such as changing default passwords of CVMs, AHV nodes and so on.

For more information about changing the default passwords, go to:

<https://portal.nutanix.com/page/documents/kbs/details?targetId=kA00e000000LkXcCAO>

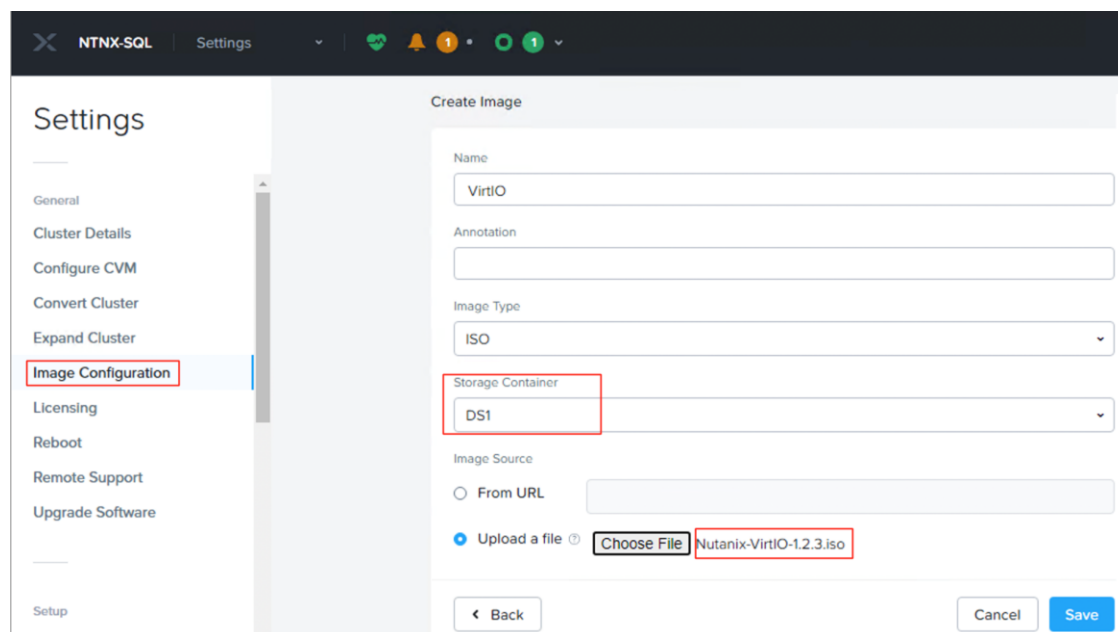
Create Virtual Machine for SQL Server Instances

The following steps provides the list of best practices and settings used for this CVD validation with SQL Server databases.

Procedure 1. Create Images for Windows Server and VirtIO drivers

Note: Before creating a VM and installing Windows Server Guest Operating System (OS), it is required to create images for Windows Server and VirtIO drivers. VirtIO drivers are required for detecting Nutanix vDisks for deploying the Operating System.

Step 1. Download Windows Server 2022 ISO image from Microsoft website, VirtIO from Nutanix website, and create images for them by uploading the files into the Storage Container created in the previous step. Go to Settings > Image Configuration > Upload Image. The following screenshot shows creating an image for VirtIO file. Image file for Windows Server 2022 should also be created.



Procedure 2. Create Virtual Machine for SQL Server Instance

This procedure provides recommendations to create a virtual machine for running SQL Server database workloads.

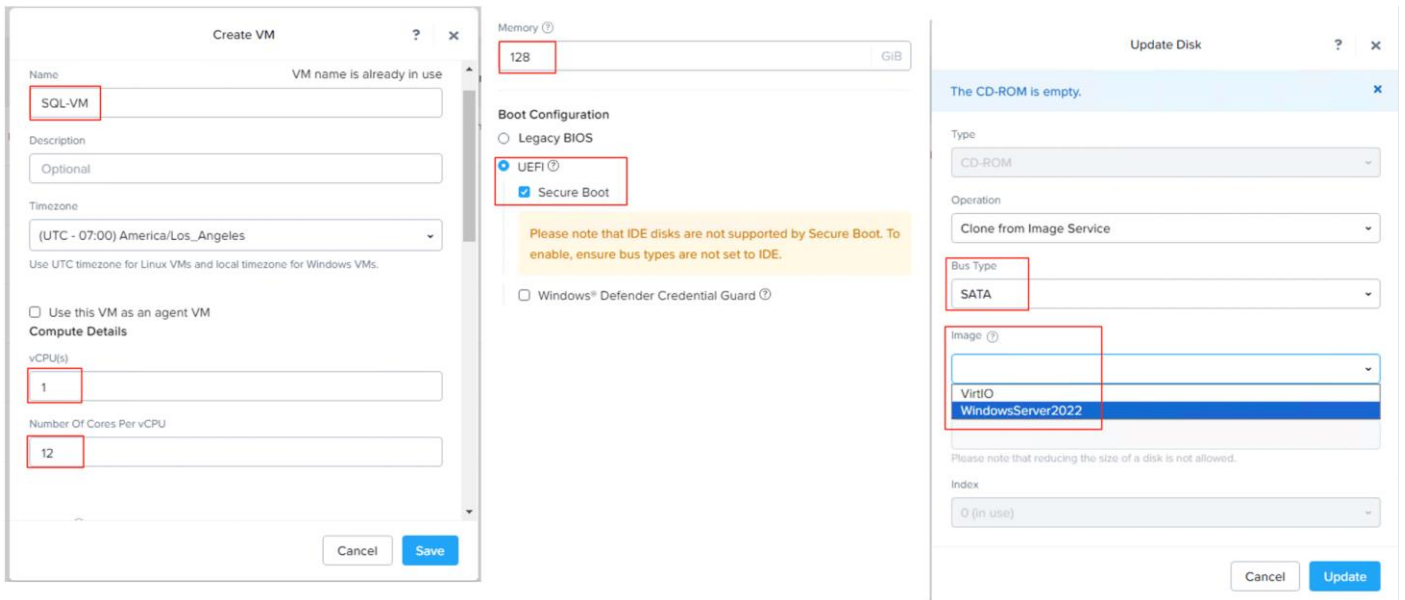
[Table 2](#) provides more details on the CPU and vDisk configuration used for SQL Server Virtual Machine hosted on the Nutanix cluster.

Table 2. CPU and Storage Configuration used for SQL Server Virtual Machine hosting 500G Database

Component	Hardware
vCPU(s)	1
Number Of Cores Per vCPU	12
Memory	128GB
Disk Layout	1x 120G vDisk for Widows OS + SQL Binaries + System Databases Following disks are used for storing 500G user/test database (5000 Warehouse IDs) (created with 8x data files and 1x T-Log files) 4x 400G vDisks for user Database data files 2x 300G vDisks for TempDB data files 1x 600G vDisk for user database and TempDB T-Log files 4x 500G vDisks for storing database backup (optional)
Disk Bus Type	scsi

Step 1. Spread the SQL Server database data (.mdf and .ndf) files across multiple vDisks. This distribution maximizes parallel I/O processing and provides better IO performance for the SQL Server VMs. For more information about databases on Nutanix and the best practices, see [Databases on Nutanix](#).

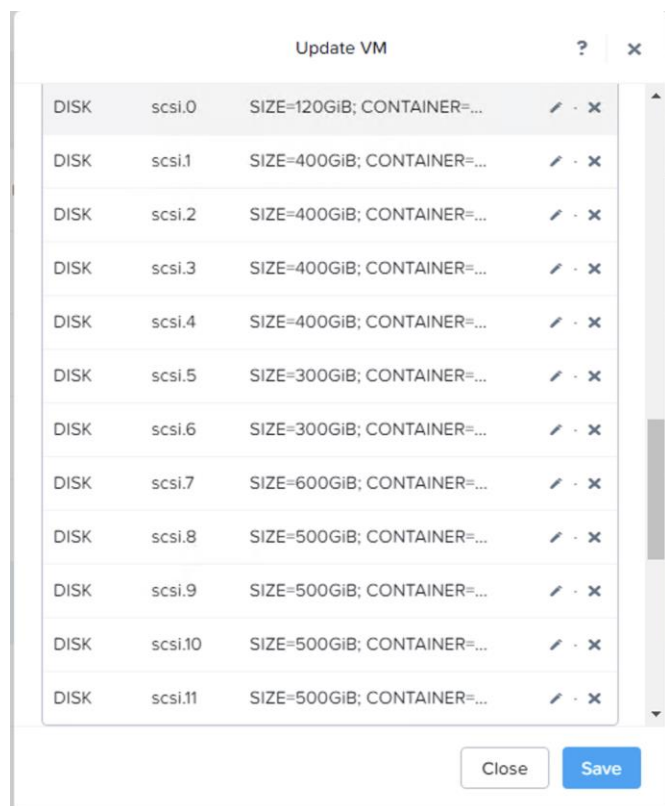
The following screenshot shows the VM configuration used for running SQL Server database.



Step 2. Select the Windows Image that was created in the previous step to install Windows Guest OS in the VM.

Note: the UEFI Secure boot option can be enabled only after changing the CD-ROM bus-type from IDE (Default) to SATA.

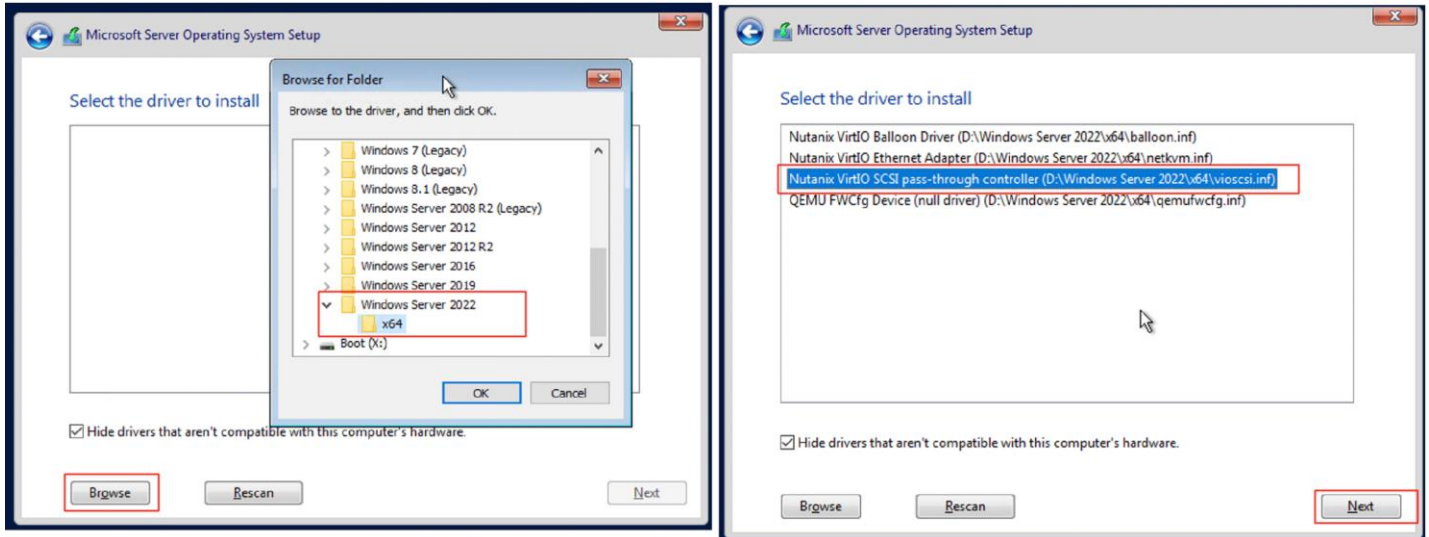
The following screenshot shows the vDisks created for the SQL Server VM. All the vDisks are stored in a single storage container DS1 created in the previous steps.



Step 3. Add one Network Adapter from the Subnet with VLAN created in the previous step. Click Save to save the VM configuration and power on the VM.

Procedure 3. Install and Configure Windows Guest Operating System

- Step 1.** Once the VM is powered on, launch VM console and hit enter to start Windows server installation.
- Step 2.** You need to mount and load the VirtIO scsi device drivers for the windows installation wizard to detect the vDisk that was added as boot disk. Click the Mount ISO option located on right corner of the console and select the VirtIO image file and click Mount.
- Step 3.** Click Browse and go to D:\ Windows Server 2022\X64\ and select the vioscsi.inf file and click Next.



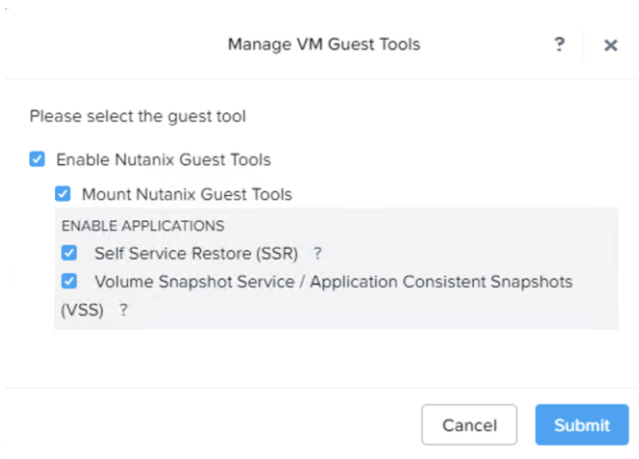
The SCSI driver is loaded. It scans for all the disks attached to the VM and displays the disk list.

Step 4. Mount back the Windows Server 2022 image and click Refresh to reload the Windows server image. Select boot disk from the disk list and click Next. Complete the windows installation.

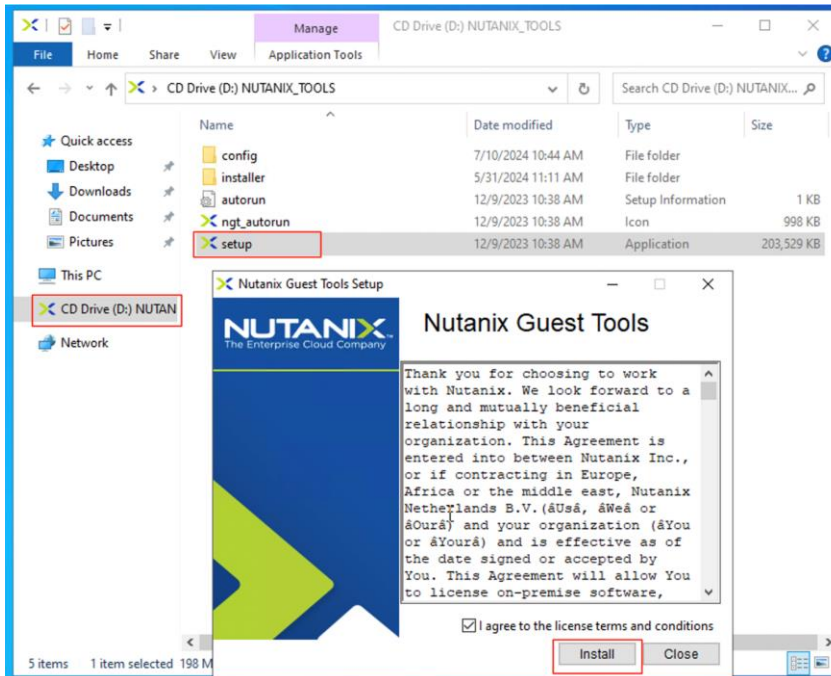
Procedure 4. Install Nutanix Guest Tools for Windows

Step 1. Before installing Guest Tools, set the CD-ROM as empty. Select the VM, right-click it and select update. Go to the Disks section and edit the CD-ROM disk and set the operation as Empty CD-ROM. Click Save.

Step 2. To mount the Guest Tools to the VM, select the VM, right-click it and click Manage Guest Tools. Enable all the options and Submit.



Step 3. Connect to the VM console and complete the Guest Tools installation. After the guest tools installation, restart the VM. Log in again with your local administrator account, assign an IP address, if required rename the hostname, and join the VM to a domain.



Step 4. Change the power plan to High performance as shown below.

```

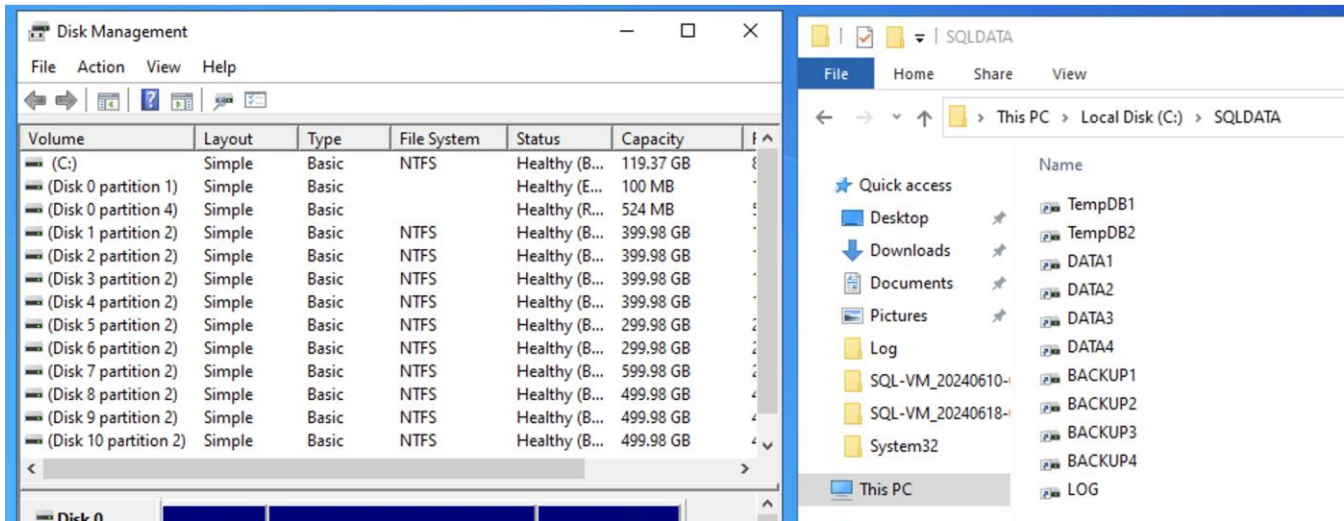
PS C:\Users\Administrator> powercfg.cpl -l
PS C:\Users\Administrator> powercfg -l

Existing Power Schemes (* Active)
-----
Power Scheme GUID: 381b4222-f694-41f0-9685-ff5bb260df2e (Balanced) *
Power Scheme GUID: 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c (High performance)
Power Scheme GUID: a1841308-3541-4fab-bc81-f71556f20b4a (Power saver)
PS C:\Users\Administrator> powercfg -s 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c
PS C:\Users\Administrator> powercfg -l

Existing Power Schemes (* Active)
-----
Power Scheme GUID: 381b4222-f694-41f0-9685-ff5bb260df2e (Balanced)
Power Scheme GUID: 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c (High performance) *
Power Scheme GUID: a1841308-3541-4fab-bc81-f71556f20b4a (Power saver)
PS C:\Users\Administrator>

```

Step 5. Open Disk management tool and initialize, partition, and format all the data and log disks using 64K allocation unit and use 1MB for backup disks. Optionally, these disks can be mounted to folders to ease the management of the disks instead of assigning drive letters to each disk.



Install and Configure SQL Server

This section discusses a few important SQL Server installation and configuration best practices used for SQL Server validation on Cisco Compute Hyperconverged with a Nutanix system. SQL Server installation on Windows guest OS is a standard practice and well documented by Microsoft here: [SQL Server 2022 Installation Guide](#).

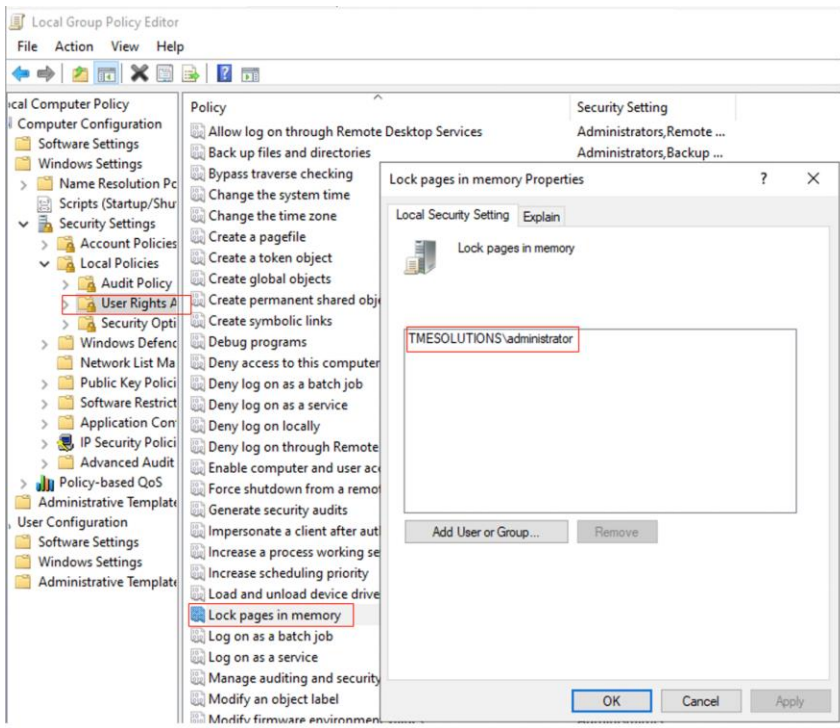
Procedure 1. Microsoft SQL Server Installation and Configuration

Step 1. In the Server Configuration window of the SQL Server 2022 Setup wizard, make sure that instant file initialization is enabled by selecting the checkbox Grant Perform Volume Maintenance Task Privilege to SQL Server Database Engine. With this setting enabled, SQL Server data files are instantly initialized, avoiding zeroing operations.

Step 2. In the Database Engine Configuration window on the TempDB tab, make sure that the number of TempDB data files is equal to 8 when the number of virtual CPUs (vCPUs) or logical processors of the SQL Server virtual machine is less than or equal to 8. If the number of logical processors is more than 8, start with 8 data files and try adding data files in multiples of 4 when you notice contention on the TempDB resources.

Step 3. After SQL Server is installed successfully, use the Windows Group Policy Editor to add a SQL Server service account (used for SQL Server database service) to the Lock pages in memory policy. Granting the Lock pages in memory user the right to the SQL Server service account prevents the Windows server from paging out SQL Server buffer pool pages.

The following screenshot shows how to enable this option. Also, if a domain account is used as a SQL Server service account that is not a member of the local administrator group, then add a SQL Server service account to the Perform volume maintenance tasks policy using the Local Security Policy Editor.



Step 4. SQL Server can consume all the memory allocated to the virtual machine. Setting the maximum server memory allows you to reserve sufficient memory for the operating system and other processes running on the virtual machine. Ideally, you should monitor the overall memory consumption of SQL Server and determine the memory requirements. To start, allow SQL Server to consume about 80 percent of the total memory, or leave at least 2 to 4 GB of memory for the operating system. The Maximum Server Memory setting can be dynamically adjusted based on your memory requirements.

Step 5. For databases with intensive Data Manipulation Language (DML) operations, you should create multiple data files of the same size to reduce access contention. Use multiple vDisks for the database data files for parallel I/O distribution across the AHV nodes. Refer to Nutanix SQL Server configuration best practices for vDisks here: [Storage Platform for SQL Server](#).

Step 6. SQL Server automatically creates soft-NUMA nodes if a socket has 8 or more CPUs. For this validation, Automatic soft-NUMA is disabled to avoid the soft-NUMA nodes creation. Execute this command **alter server configuration set softNUMA off** on SQL Server management studio to disable soft-NUMA.

Step 7. Ensure that multiple TempDB database data files of equal size are configured and stored across multiple vDisks.

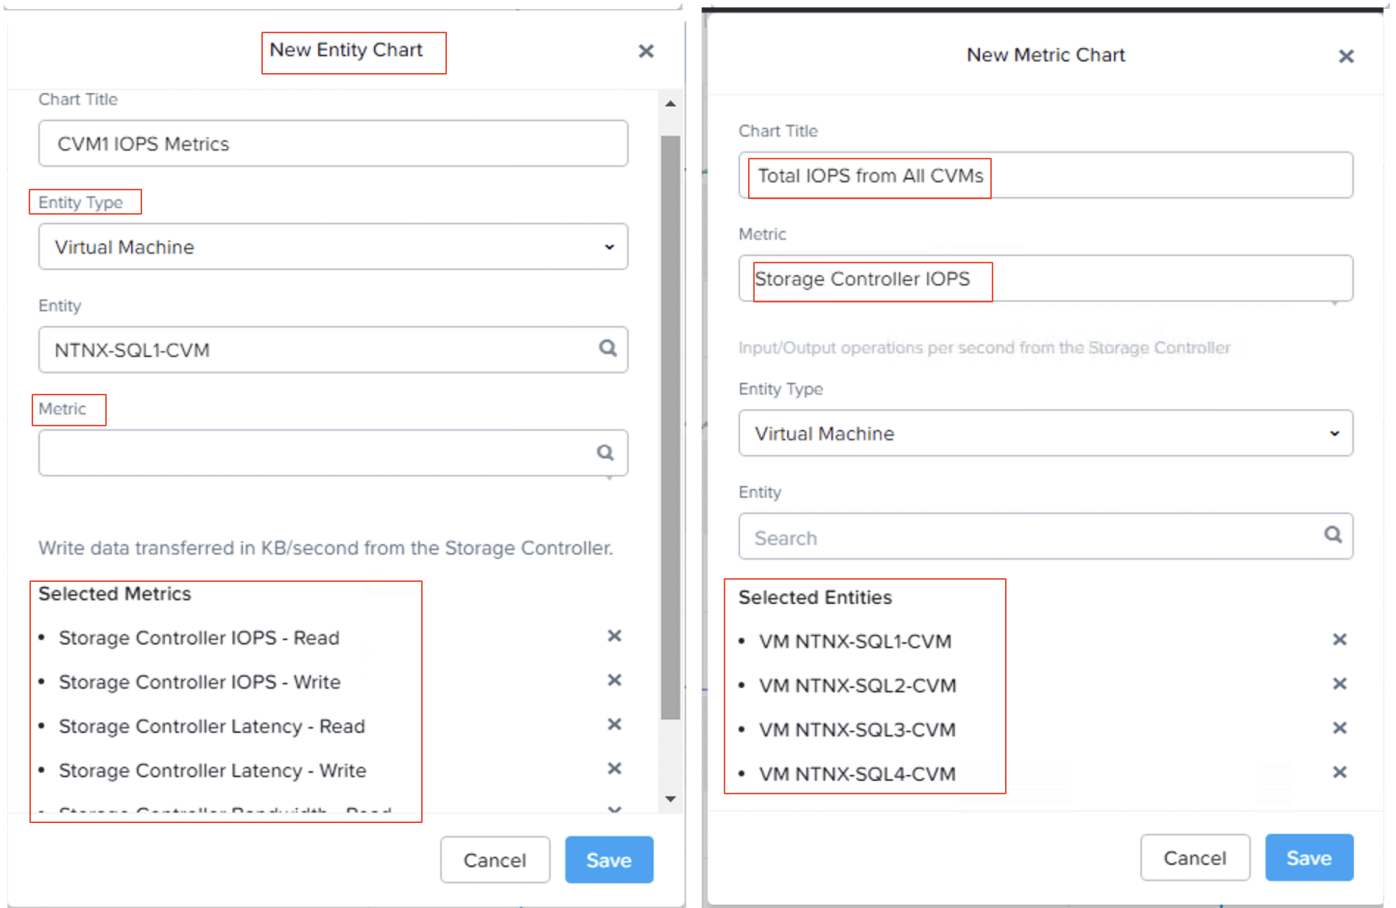
Monitor SQL Server VMs using Prism Element

Nutanix Prism Element comes with two innovative charts to help monitor the cluster or its components, like VM or troubleshoot an issue. They are Entity Charts and Metrics Charts, and they are very easy to build and customize as per our needs.

- **Entity Charts:** These allow you to monitor multiple metrics of a specific object, such as cluster, host, storage container, VM, and so on. For example, monitoring IOPS, latency, and CPU utilization of a SQL VM or Controller VM is useful when troubleshooting performance-related issues.
- **Metric Charts:** It is the opposite of Entity Chart, wherein we monitor a metric of one or more objects. For example, CPU utilization of all controller VMs in a cluster.

Procedure 1. Create Metric Charts

Step 1. Log into Prism Element, go to the Analysis page, click on New, and select Create Entity Chart to create an Entity Chart. Select 'Create Metric Chart' to create a metric chart. In the screenshot below, the entity chart shows IOPS-related metrics of controller VM #1, while the Metric chart shows total IOPS from all the controller VMs.



Step 2. Create entity and metric charts for the required objects (like AHV nodes, storage containers, vDisks, controller VMs and workload VMs, and so on) and monitor the required performance metrics.

Solution Validation

This chapter contains the following:

- [Infrastructure Performance Testing with Nutanix X-Ray Tool](#)
- [Single VM Performance Test with SQL Server](#)
- [Performance Scalability Test with SQL Server VMs](#)
- [SQL Server Always on Availability Group \(AG\) Testing](#)

This chapter provides a high-level summary of the performance testing and validation results for the Cisco Compute Hyperconverged with Nutanix system.

Infrastructure Performance Testing with Nutanix X-Ray Tool

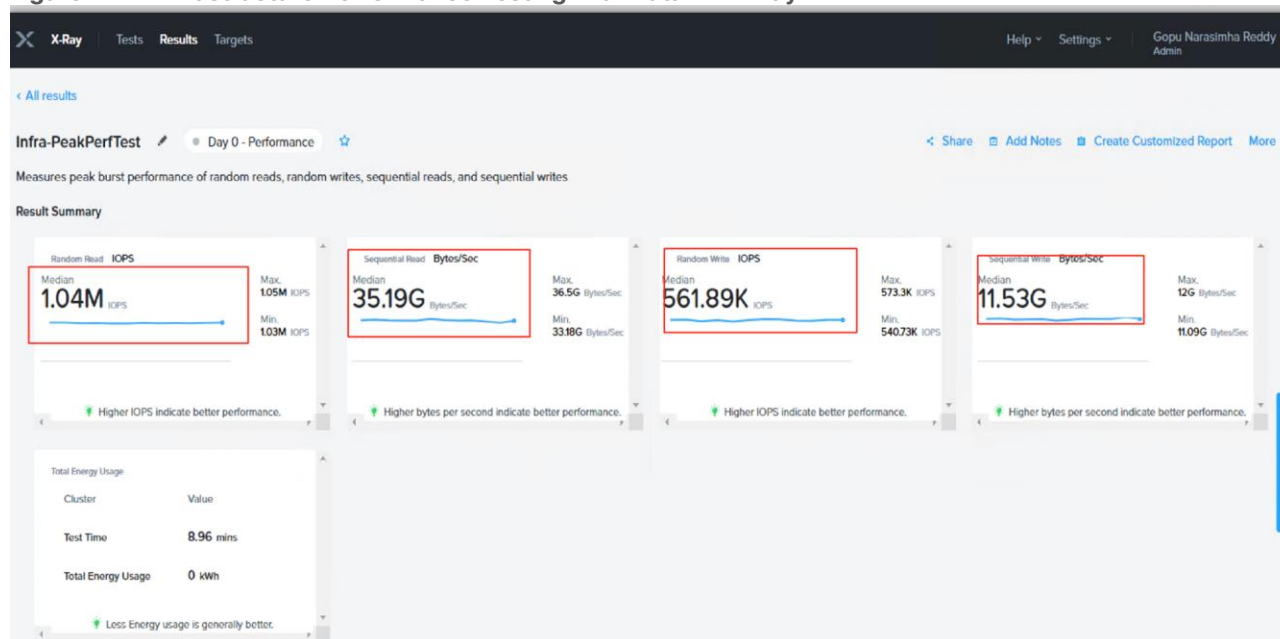
Before deploying and running any workload on the cluster, it is important to validate the system with tools to ensure that the system is optimally configured.

Nutanix X-Ray is a hyperconverged infrastructure (HCI) assessment solution for testing the resiliency, performance, and scalability across combinations of HCI, hypervisor, and hardware platform products. This enables organizations to make informed infrastructure decisions based on how their applications will react under load for these characteristics. For this CVD, the Infrastructure Performance test was run, and the results are shown below. For more details on how to deploy and run the X-Ray test, go to:

https://portal.nutanix.com/page/documents/details?targetId=X-Ray-Guide-v4_4:X-Ray-Guide-v4_4

The results shows that the four-node All-NVMe Nutanix cluster achieved more than 1 million random reads, 561.89K random writes, a sequential read bandwidth of 31GB/s and 11.5GB/s sequential write bandwidth.

Figure 12. Infrastructure Performance Testing with Nutanix X-Ray



Single VM Performance Test with SQL Server

The objective of the single SQL Server VM test is to show the sustained and consistent performance delivered by the Nutanix cluster for a large working set database. [Table 3](#) lists the test configuration details used for all the tests conducted, as detailed below.

Table 3. SQL Server VM Configuration

Component	Details
vCPUs	1
No of Cores	12
Memory	128GB
Storage Layout	1x 120G disk for Widows OS + SQL Binaries + System Databases Following disks are used for storing 500G user/test database 4x 400G disks for user Database data files 2x 300G disks for TempDB data files 1x 600G disk for user database and TempDB T-Log files 4x 500G disks for backup
Database Site and file Layout	500G Created with 8x data files each is 100G and 1x T-Log file of size 300G
SQL Server Settings	Max Memory = 122 Soft-NUMA disabled Enabled Lock Pages in memory and Instant file Initialization
Workload and Testing tool details	SQL Server Operational database workload (OLTP) generated with HammerDB tool (v4.10) Database Size= 500GB Warehouse IDs= 5000 Use All Warehouses= true Ramp Up time and Run time= 5 and 115 Mins Number of HammerDB virtual users= 30

The HammerDB tool is used to simulate and run TPROC-C-like workloads on the SQL Server virtual machines. It is a leading benchmarking and load testing software for the world’s most popular databases like Microsoft SQL Server. It implements a fair usage of TPC specifications for benchmarking the database workloads such as Online Transactional (OLTP) and Decision Support System(DSS). TPC is an industry body most widely recognized for defining benchmarks.

For this solution validation with SQL Server database, a 500Gb database is loaded using 5000 warehouse IDs and stored on multiple vDisks as detailed in the previous sections. The test was run for more than two hours and noticed a consistent performance throughout the test duration without any dips in the performance.

Figure 13. Cisco Compute Hyperconverged with Nutanix



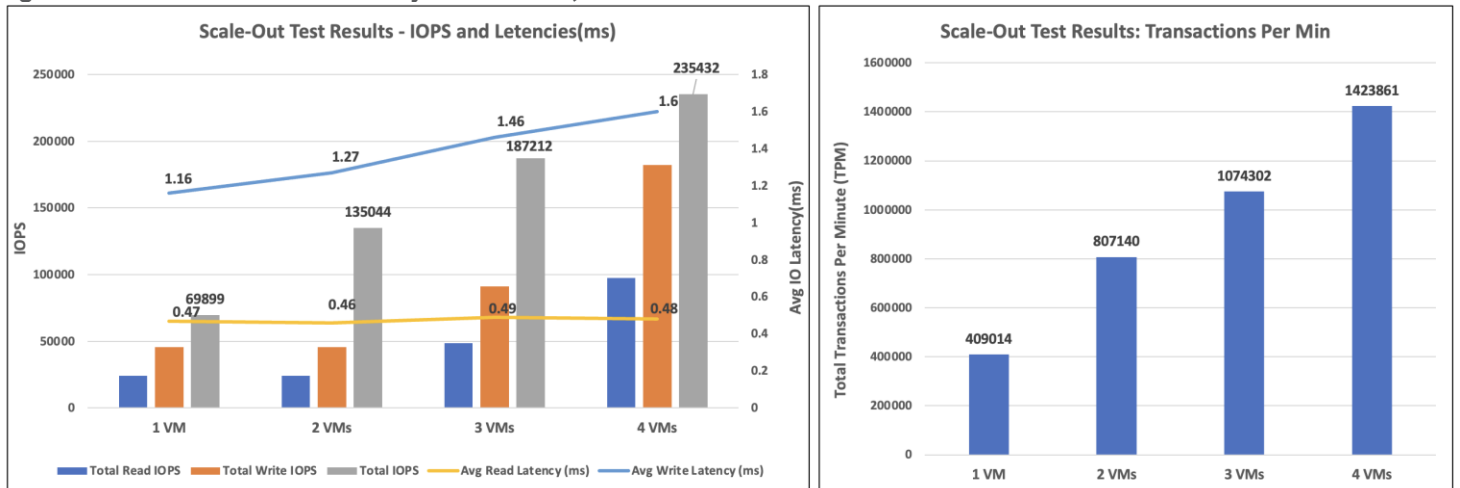
The single SQL VM delivered about 70,000 IOPS under 1.2ms write latency. Nutanix Data locality is one of the major contributors to achieve high consistent performance. Data locality reduces network traffic by avoiding the network fabric bottleneck in both centralized and remotely accessed storage architectures. It also improves performance by taking advantage of local flash and memory capabilities—providing lower latencies, higher throughput, and better use of performance resources.

Performance Scalability Test with SQL Server VMs

The objective of this test is to analyze how database performance scales as more SQL VMs are deployed across the cluster. The test is started with one SQL VM on one AHV node and scaled up to four VMs spread across four AHV nodes by scaling one VM at a time.

The same VM configuration is used as detailed in the previous section. Four different HammerDB clients, hosted on an external server, are used to stress the four SQL VMs hosted on the Nutanix cluster. The following graphs show how database IOPS, along with latencies and database transactions per minute (TPM), scaled from one VM to four VMs.

Figure 14. Performance Scalability Test with SQL Server VMs



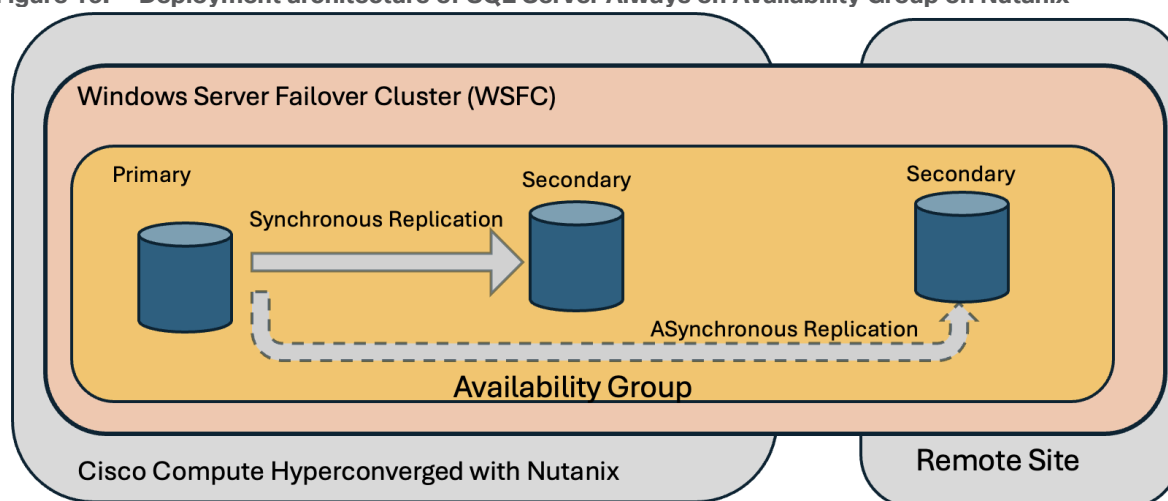
As shown in [Figure 14](#), it demonstrates that both IOPS scaled linearly as SQL VMs are added to the cluster. There is no change in the read latency while the write latency is slightly increased by half a millisecond, staying under 2ms with four VM tests. It is worth noting that with four VMs, nearly 235,000 IOPS is achieved with a read-write mix of nearly 35-65% with varying IO sizes caused by data reads, writes, T-Log writes, and database checkpoints. The TPM is also scaled linearly and in-line with IOPS.

The scale-out test results clearly demonstrate excellent performance scalability capabilities of Nutanix cluster for enterprise database workloads.

SQL Server Always on Availability Group (AG) Testing

The objective of this test is to validate SQL Server Always on Availability Group on Nutanix cluster and analyze the performance with standard Always on Availability Group deployment. [Figure 15](#) illustrates the deployment tested on Cisco Compute Hyperconverged with Nutanix system.

Figure 15. Deployment architecture of SQL Server Always on Availability Group on Nutanix

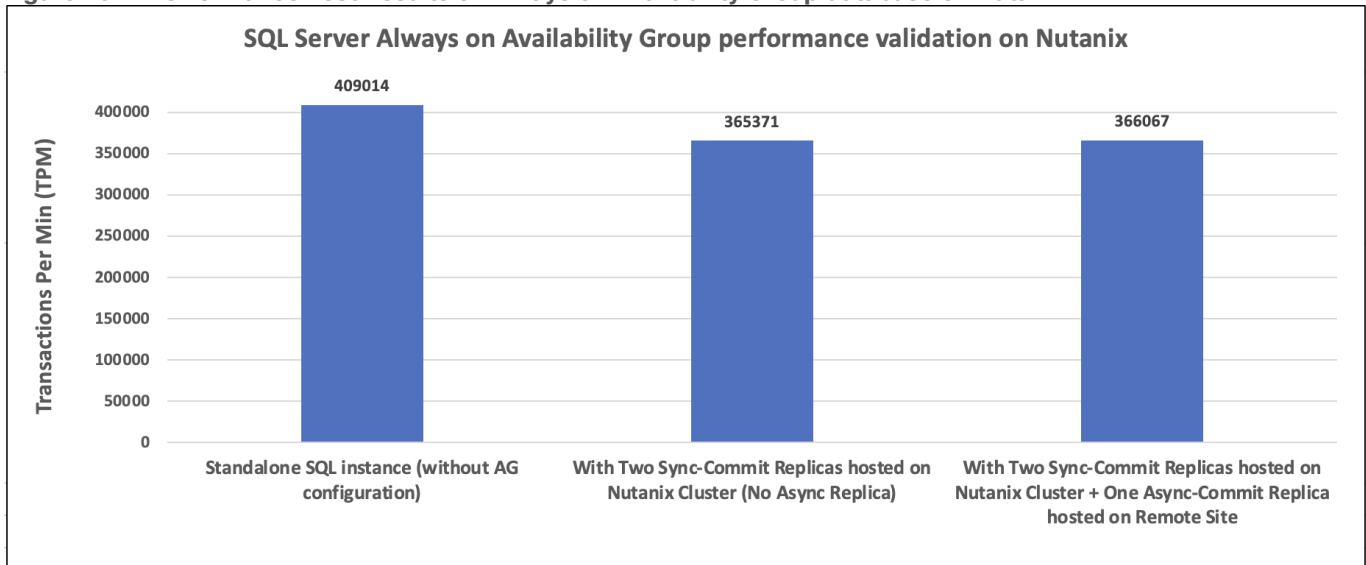


The two SQL Server AG replicas hosted on the Nutanix cluster are configured with synchronous-commit replication with an automatic failover option. Synchronous-commit mode ensures that once a given secondary database is synchronized with the primary database, committed transactions are fully protected. This protection comes at the cost of increased transaction latency. The secondary replica executes the same commands as the primary using the same transaction to be consistent with the primary so that it is able to take over the primary role in case of failure of the original primary replica. Thereby providing high availability to the databases.

The third replica, hosted on a remote site or cloud, is configured with Asynchronous-commit replication with manual failover option. Under asynchronous-commit mode, the primary replica commits transactions without waiting for acknowledgment from asynchronous-commit secondary replicas to harden their transaction logs. Asynchronous-commit mode minimizes transaction latency on the secondary databases but allows them to lag behind the primary databases, making some data loss possible. In case the primary replicas or site hosted in the Nutanix cluster are completely unavailable, 'Forced failover' needs to be performed manually, and this Forced failover is a disaster recovery option.

[Figure 16](#) shows the test results of the Always on Availability Group deployed as explained in the above figure. The SQL Server VMs on Nutanix cluster and Remote site are configured as detailed in [Table 3](#).

Figure 16. Performance Test Results of Always on Availability Group database on Nutanix



As shown in [Figure 16](#), in the ideal case, the standard SQL Server instance on the Nutanix cluster, without AG configuration, achieves about 400,000 transactions per minute (TPM). The second SQL Server AG replicas, hosted on the same Nutanix cluster and configured with Synchronous-commit replication, achieved around 365,000 TPM, a nearly 11% reduction in performance. The reduction in TPM is due to the synchronous replication delays between the two replicas. In the last case, the addition of Asynchronous replication did not have any impact on the performance as the transactions on the primary replica do not need to wait for transactions to be hardened on the third replica configured with Asynchronous commit replication.

Conclusion

Cisco Compute Hyperconverged with Nutanix is built with best-of-breed technologies from Cisco and Nutanix, providing customers with truly hyperconverged infrastructure solutions to consolidate and run a variety of IT workloads. The system is designed to meet the needs of modern applications and improve operational efficiency, agility, and scale through Cisco Intersight and Nutanix AOS storage fabric.

The performance tests demonstrated that the Cisco Compute Hyperconverged with Nutanix delivers the low-latency, consistent, and scalable database performance required by critical enterprise database workloads. Nutanix storage features like data locality, tiering, snapshots and clones etc. help customers achieve greater database consolidation ratios, reduced datacenter footprints, storage efficiencies and higher return on investments (ROI).

About the Authors

Gopu Narasimha Reddy, Technical Marketing Engineer, Cisco Systems, Inc.

Gopu Narasimha Reddy is a Technical Marketing engineer with the UCS Solutions team at Cisco. He is currently focused on validating and developing solutions on various Cisco UCS platforms for enterprise database workloads with different operating environments including Windows, VMware, Linux, and Kubernetes. Gopu is also involved in publishing database benchmarks on Cisco UCS servers. His areas of interest include building and validating reference architectures, development of sizing tools in addition to assisting customers in database deployments.

Database Solutions Engineering Team, Nutanix:

- Pri Abeyratne, Sr Solutions Architect, Nutanix.
- Jisha J, Sr Solutions Architect, Nutanix.
- Krishna Kapa, Staff Solutions Architect, Nutanix.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- John McAbel, Senior Product Manager, Cisco Systems, Inc.
- Anil Dhiman, Technical Marketing Engineer, Cisco Systems, Inc.
- Chris O'Brien, Director, Technical Marketing, Cisco System, Inc.
- Bruno Sousa, Technical Director Database Solutions, Nutanix.

Appendices

This appendix contains the following:

- [Appendix A - Bill of Materials](#)
- [Appendix B - References used in this guide](#)

Appendix A - Bill of Materials

[Table 4](#) provides an example the Bill of Materials used for one (4) node cluster deployed, used in the testing and reference design described in this document.

Table 4. Bill of Materials

Part Number	Description	Quantity
HCI AF240C-M7SN	Cisco Compute Hyperconverged HCI AF240cM7 All Flash NVMe Node	4
CON-L1NCO-HCIAFM7C	CX LEVEL 1 8X7XNCDOS Cisco Compute Hyperconverged HCI AF240cM	4
HCI-NVME4-3840	3.8TB 2.5in U.2 15mm P5520 Hg Perf Med End NVMe	24
HCI-M2-240G	240GB M.2 SATA Micron G2 SSD	8
HCI-M2-HWRAID	Cisco Boot optimized M.2 Raid controller	4
HCI-RAIL-M7	Ball Bearing Rail Kit for C220 & C240 M7 rack servers	4
HCI-TPM-002C	TPM 2.0, TCG, FIPS140-2, CC EAL4+ Certified, for servers	4
UCSC-HSHP-C240M7	UCS C240 M7 Heatsink	8
UCSC-BBLKD-M7	UCS C-Series M7 SFF drive blanking panel	72
UCS-DDR5-BLK	UCS DDR5 DIMM Blanks	64
UCSC-M2EXT-240-D	C240M7 2U M.2 Extender board	4
UCSC-FBRS2-C240-D	C240 M7/M8 2U Riser2 Filler Blank	4
UCSC-FBRS3-C240-D	C240 M7/M8 2U Riser3 Filler Blank	4
HCI-CPU-I8462Y+	Intel I8462Y+ 2.8GHz/300W 32C/60MB DDR5 4800MT/s	8
HCI-MRX64G2RE1	64GB DDR5-4800 RDIMM 2Rx4 (16Gb)	64
HCI-RIS1A-24XM7	C240 M7 Riser1A; (x8;x16x, x8); StBkt; (CPU1)	4
HCI-MLOM	Cisco VIC Connectivity	4
HCI-M-V5D200G	Cisco VIC 15238 2x 40/100/200G mLOM C-Series	4
HCI-PSU1-1200W	1200W Titanium power supply for C-Series Servers	8
NO-POWER-CORD	ECO friendly green option, no power cable will be shipped	8

Appendix B - References used in this guide

Cisco Compute Hyperconverged with Nutanix: <https://www.cisco.com/c/en/us/products/hyperconverged-infrastructure/compute-hyperconverged/index.html>

Cisco Compute Hyperconverged with Nutanix Design and Deployment Guide: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/CCHC_Nutanix_ISM.html

HCIAF240C M7 All-NVMe/All-Flash Server: <https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hc-240m7-specsheet.pdf>

Cisco Intersight: <https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html>

Nutanix Reference Documentation: <https://portal.nutanix.com/>

Reference to information about Nutanix: <https://www.nutanixbible.com>

Database workload on Nutanix: <https://www.nutanix.com/architecture#database-workloads>

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](https://cs.co/en-cvds) at <https://cs.co/en-cvds>.

CVD Program

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS X-Series, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trade-marks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_P3)

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)