



FlexPod Datacenter for VMware Horizon 8 with VMware vSphere 8 for up to 3000 Sessions

Deployment Guide for Virtual Desktop Infrastructure
built on Cisco UCS X210c M7 with NVIDIA GPUs,
Cisco Intersight, NetApp Storage for VMware Horizon
8 2212, and VMware vSphere 8.0 U1 Hypervisor

Published: December 2023



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

Cisco Validated Designs (CVDs) consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.

The landscape of desktop and application virtualization is changing constantly. The high-performance Cisco UCS X-Series Compute Node and Cisco UCS Unified Fabric combined as part of the FlexPod Proven Infrastructure with the latest generation NetApp AFF storage result in a more compact, more powerful, more reliable, and more efficient platform.

The solution explains the deployment of a predesigned, best-practice data center architecture with VMware Horizon Remote Desktop Server Hosted (RDSH) sessions and Windows 11 Virtual desktops and VMware vSphere built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches, and NetApp Storage AFF A400 All Flash array supporting iSCSI storage access.

The solution provides outstanding virtual desktop end-user experience as measured by the Login VSI 4.1.40 Knowledge Worker workload running in benchmark mode.

The 2000-3000 seat solution provides a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

When deployed, the architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Desktop Infrastructure (VDI).

Additionally, this FlexPod solution is also delivered as Infrastructure as Code (IaC) to eliminate error-prone manual tasks, allowing quicker and more consistent solution deployments. Cisco Intersight cloud platform delivers monitoring, orchestration, workload optimization and lifecycle management capabilities for the FlexPod solution.

This CVD provides the architecture and design of a virtual desktop infrastructure for up to 3300 end-user compute users. The solution virtualized on Cisco UCS X210c M7 Server, booting VMware vSphere 8.0 U1 through local boot. The virtual desktops are powered using VMware Remote Desktop Server Hosted (RDSH) sessions and VMware Windows 11 Virtual Desktops, with a pooled/non-persistent instant-clones virtual Windows 11 desktops) and persistent full clone virtual Windows 11 desktops.

If you're interested in understanding the FlexPod design and deployment details, including the configuration of various elements of design and associated best practices, refer to Cisco Validated Designs for FlexPod, here: <https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html>

Solution Overview

This chapter contains the following:

- [Introduction](#)
- [Audience](#)
- [Purpose of this Document](#)
- [What's New in this Release?](#)
- [What is FlexPod?](#)
- [FlexPod Cisco Validated Design Advantages for VDI](#)
- [Cisco Desktop Virtualization Solutions: Data Center](#)
- [Physical Topology](#)
- [Configuration Guidelines](#)
- [Solution Summary](#)

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility, and reducing costs. Cisco, NetApp storage, and VMware have partnered to deliver this Cisco Validated Design, which uses best of breed storage, server, and network components to serve for the foundation for desktop virtualization workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

The Cisco UCS X210c M7 Compute Node Server delivers performance, flexibility, and optimization for deployments in data centers, cloud, and remote sites. This enterprise-class server offers market-leading versatility, and density without compromise for workloads, including web infrastructure, distributed databases, Virtual Desktop Infrastructure (VDI), converged infrastructure, and enterprise applications such as SAP HANA and Oracle. The Cisco UCS X210c M7 Compute Node Server can quickly deploy stateless physical and virtual workloads through a programmable, easy-to-use Cisco Intersight and Cisco Intersight and simplified server access through Cisco SingleConnect technology.

Audience

The intended audience for this document includes, but is not limited to IT architects, sales engineers, field consultants, professional services, IT managers, IT engineers, partners, and customers who are interested in learning about and deploying the Virtual Desktop Infrastructure (VDI)

Purpose of this Document

This document provides a step-by-step design, configuration, and implementation guide for the Cisco Validated Design for a large-scale VMware Horizon Remote Desktop Server Hosted (RDSH) sessions and Windows 11 Virtual Desktops with NetApp AFF A400, Cisco UCS X210c M7 Compute Node Servers, and Cisco Nexus 9000 Series Ethernet Switches.

What's New in this Release?

This version of the FlexPod VDI Design is based on the latest Cisco FlexPod Virtual Server Infrastructure and introduces the Cisco UCS M7 Servers featuring the 4th Gen Intel Xeon processors.

Highlights for this design include:

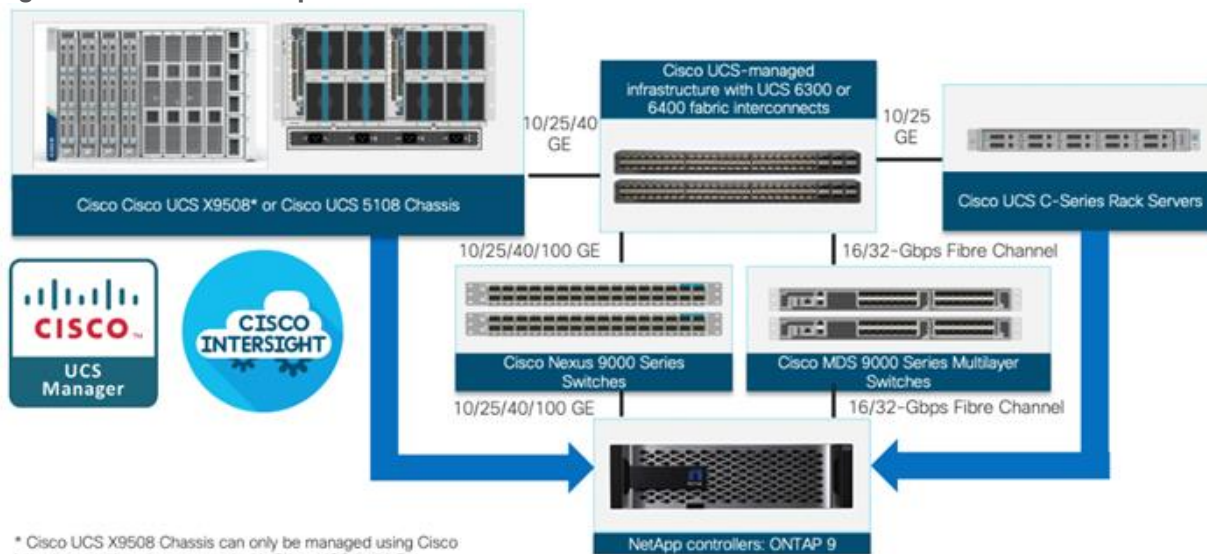
- Deploying and managing Cisco UCS X210c M7 Compute Node Server using Cisco Unified Computing System (Cisco UCS)
- Support for Cisco UCS X210c M7 Compute Node Servers with 4th Gen Intel Xeon Family processors and 3200 MHz memory
- Support for the Cisco Intersight 5.2
- Validation of Cisco Nexus 9000 with NetApp AFF A400 system
- Support for NetApp Storage AFF A400 with ONTAP version 9.13.1P3
- VMware Horizon 8 2212 (ESB)
- Support for VMware vSphere 8.0 U1
- Fully automated solution deployment covering FlexPod infrastructure and vSphere virtualization

What is FlexPod?

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere built on FlexPod includes NetApp FAS, AFF, and ASA storage, Cisco Nexus networking, Cisco MDS storage networking, the Cisco Unified Computing System (Cisco UCS), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to your data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit your requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a Fibre Channel and IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for personalization and investment protection because it truly is a wire-once architecture.

Figure 1. FlexPod Component Families



The following lists the benefits of FlexPod:

- Consistent Performance and Scalability
 - Consistent sub-millisecond latency with 100 percent flash storage
 - Consolidate 100's of enterprise-class applications in a single rack
 - Scales easily, without disruption
 - Continuous growth through multiple FlexPod CI deployments
- Operational Simplicity
 - Fully tested, validated, and documented for rapid deployment
 - Reduced management complexity
 - Auto-aligned 512B architecture removes storage alignment issues
 - No storage tuning or tiers necessary
- Lowest TCO
 - Dramatic savings in power, cooling, and space with 100 percent flash storage
 - Industry leading data reduction
- Enterprise-Grade Resiliency
 - Highly available architecture with no single point of failure
 - Nondisruptive operations with no downtime
 - Upgrade and expand without downtime or performance loss
 - Native data protection: snapshots and replication
 - Suitable for even large resource-intensive workloads such as real-time analytics or heavy transactional databases

FlexPod Cisco Validated Design Advantages for VDI

The data center market segment is shifting toward heavily virtualized private, hybrid and public cloud computing models running on industry-standard systems. These environments require uniform design points that can be repeated for ease of management and scalability.

These factors have led to the need predesigned computing, networking and storage building blocks optimized to lower the initial design cost, simply management, and enable horizontal scalability and high levels of utilization. The use cases include:

- Enterprise Data Center (small failure domains)
- Service Provider Data Center (small failure domains)
- Commercial Data Center
- Remote Office/Branch Office
- SMB Standalone Deployments

Cisco Desktop Virtualization Solutions: Data Center

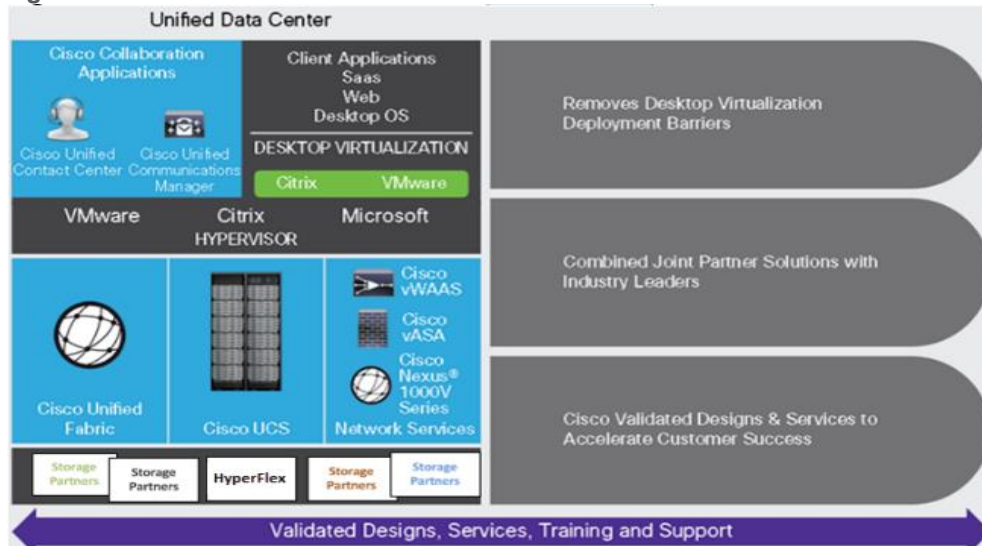
The Evolving Workplace

Today’s IT departments are facing a rapidly evolving workplace environment. The workforce is becoming increasingly diverse and geographically dispersed, including offshore contractors, distributed call center operations, knowledge and task workers, partners, consultants, and executives connecting from locations around the world at all times.

This workforce is also increasingly mobile, conducting business in traditional offices, conference rooms across the enterprise campus, home offices, on the road, in hotels, and at the local coffee shop. This workforce wants to use a growing array of client computing and mobile devices that they can choose based on personal preference. These trends are increasing pressure on IT to ensure protection of corporate data and prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios ([Figure 2](#)).

These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 11 and productivity tools, namely Microsoft Office 2021.

Figure 2. Cisco Data Center Partner Collaboration



Some of the key drivers for desktop virtualization are increased data security, the ability to expand and contract capacity and reduced TCO through increased control and reduced management costs.

Cisco Desktop Virtualization Focus

Cisco focuses on three key elements to deliver the best desktop virtualization data center infrastructure: simplification, security, and scalability. The software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform.

Simplified

Cisco UCS and NetApp provide a radical new approach to industry-standard computing and provides the core of the data center infrastructure for desktop virtualization. Among the many features and benefits of Cisco UCS are the drastic reduction in the number of servers needed, in the number of cables used per server and the capability to rapidly deploy or re-provision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified. Thousands of desktops can be provisioned in minutes with Cisco Intersight server profiles and Cisco storage partners' storage-based cloning. This approach accelerates the time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

Cisco Intersight automates many mundane, error-prone data center operations such as configuration and provisioning of server, network, and storage access infrastructure. In addition, Cisco UCS Servers with large memory footprints enable high desktop density that helps reduce server infrastructure requirements.

Simplification also leads to more successful desktop virtualization implementation. Cisco and its technology partners like VMware have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlexPod. Cisco Desktop Virtualization Solutions have been tested with VMware vSphere.

Secure

Although virtual desktops are inherently more secure than their physical predecessors, they introduce new security challenges. Mission-critical web and application servers using a common infrastructure such as virtual desktops are now at a higher risk for security threats. Inter-virtual machine traffic now poses an important security consideration that IT managers need to address, especially in dynamic environments in which virtual machines, using VMware vMotion, move across the server infrastructure.

Desktop virtualization, therefore, significantly increases the need for virtual machine-level awareness of policy and security, especially given the dynamic and fluid nature of virtual machine mobility across an extended computing infrastructure. The ease with which new virtual desktops can proliferate magnifies the importance of a virtualization-aware network and security infrastructure. Cisco data center infrastructure (Cisco UCS and Cisco Nexus Family solutions) for desktop virtualization provides strong data center, network, and desktop security, with comprehensive security from the desktop to the hypervisor. Security is enhanced with segmentation of virtual desktops, virtual machine-aware policies and administration, and network security across the LAN and WAN infrastructure.

Scalable

The growth of a desktop virtualization solution is all but inevitable, so a solution must be able to scale, and scale predictably, with that growth. The Cisco Desktop Virtualization Solutions built on FlexPod Datacenter infrastructure supports high virtual-desktop density (desktops per server), and additional servers and storage scale with near-linear performance. FlexPod Datacenter provides a flexible platform for growth and improves business agility. Cisco Intersight server profiles allow on-demand desktop provisioning and make it just as easy to deploy dozens of desktops as it is to deploy thousands of desktops.

Cisco UCS servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 16 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 200 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco storage partners NetApp help maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs for End User Computing based on FlexPod solutions have demonstrated scalability and performance, with up to 2000 desktops up and running in less than 15 minutes.

FlexPod Datacenter provides an excellent platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization, data center applications, and cloud computing.

Cisco UCS and Cisco Nexus data center infrastructure provides an excellent platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization, data center applications, and cloud computing.

Savings and Success

The simplified, secure, scalable Cisco data center infrastructure for desktop virtualization solutions saves time and money compared to alternative approaches. Cisco UCS enables faster payback and ongoing savings (better ROI and lower TCO) and provides the industry's greatest virtual desktop density per server, reducing both capital expenditures (CapEx) and operating expenses (OpEx). The Cisco UCS architecture and Cisco Unified Fabric also enables much lower network infrastructure costs, with fewer cables per server and fewer ports required. In addition, storage tiering and deduplication technologies decrease storage costs, reducing desktop storage needs by up to 50 percent.

The simplified deployment of Cisco NetApp FlexPod solution for desktop virtualization accelerates the time to productivity and enhances business agility. IT staff and end users are more productive, and the business can respond to new opportunities quickly by deploying virtual desktops whenever and wherever they are needed. The high-performance Cisco Systems and network deliver a near-native end-user experience, allowing users to be productive anytime and anywhere.

The key measure of desktop virtualization for any organization is its efficiency and effectiveness in both the near term and the long term. The Cisco Desktop Virtualization Solutions are very efficient, allowing rapid deployment, requiring fewer devices and cables, and reducing costs. The solutions are also extremely effective, providing the services that end users need on their devices of choice while improving IT operations, control, and data security. Success is bolstered through Cisco's best-in-class partnerships with leaders in virtualization and through tested and validated designs and services to help you throughout the solution lifecycle. Long-term success is enabled through the use of Cisco's scalable, flexible, and secure architecture as the platform for desktop virtualization.

The ultimate measure of desktop virtualization for any end-user is a great experience. Cisco NetApp delivers class-leading performance with sub-second base line response times and index average response times at full load of just under one second.

Use Cases

The following are some typical use cases:


- Healthcare: Mobility between desktops and terminals, compliance, and cost

- Federal government: Teleworking initiatives, business continuance, continuity of operations (COOP), and training centers
- Financial: Retail banks reducing IT costs, insurance agents, compliance, and privacy
- Education: K-12 student access, higher education, and remote learning
- State and local governments: IT and service consolidation across agencies and interagency security
- Retail: Branch-office IT cost reduction and remote vendors
- Manufacturing: Task and knowledge workers and offshore contractors
- Microsoft Windows 11 migration
- Graphic intense applications
- Security and compliance initiatives
- Opening of remote and branch offices or offshore facilities
- Mergers and acquisitions

Physical Topology

Figure 3 illustrates the physical architecture.

Figure 3. Physical Architecture



The reference hardware configuration includes:

- Two Cisco Nexus N9K-C93180YC-FX3S switches
- Two Cisco UCS 6536 Fabric Interconnects
- Eight Cisco UCS X210c M7 Compute Node Servers (for VDI workload)
- Infrastructure VMs for VDI were housed on an external cluster
- One NetApp AFF A400 Storage (high-availability pair) System

For desktop virtualization, the deployment includes VMware Horizon 8 Remote Desktop Session Hosts (RDSH) Sessions and Windows 11 virtual desktops running on VMware vSphere 8.0 U1.

The design is intended to provide a large-scale building block for VMware Horizon Remote Desktop Session Hosted (RDSH) Sessions workloads consisting of Remote Desktops Server Hosted (RDSH) sessions with Windows Server 2019 hosted shared desktop sessions and Windows 11 non-persistent and persistent hosted desktops in the following:

- 3360 Floating Hosted Shared (RDSH) Server 2019 user sessions with Microsoft Office 2021 (Instant clones)
- 2240 Floating Windows 11 Desktops with Microsoft Office 2021 (instant clones)
- 2240 Dedicated Windows 11 Desktops with Microsoft Office 2021 (full clones)

The data provided in this document will allow you to adjust the mix of Remote Desktop Server Hosted (RDSH) Sessions and Windows 11 Virtual Desktops to suit their environment. For example, additional Compute Node servers and chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the detailed steps for deploying the base architecture. This procedure explains everything from physical cabling to network, compute, and storage device configurations.

Configuration Guidelines

This Cisco Validated Design provides details for deploying a fully redundant, highly available 2000 seats workload virtual sessions /desktop solution with VMware on a FlexPod Datacenter architecture. Configuration guidelines are provided that refer to the reader to which redundant component is being configured with each step. For example, storage controller 01 and storage controller 02 are used to identify the two AFF A400 storage controllers that are provisioned with this document, Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured.

The Cisco UCS 6536 Fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, VM-Host-RDSH-01, VM-Host-VDI-01 and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure.

Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both Microsoft Windows 11 virtual desktops and RDS server desktop sessions based on Microsoft Server 2019. The solution includes Cisco UCS hardware and Data Platform software, Cisco Nexus switches, the Cisco Unified Computing System (Cisco UCS), VMware Horizon and VMware vSphere software in a single package. The design is efficient such that the networking, computing, and storage components occupy 25 rack units footprint in an industry standard 42U rack. Port density on the Cisco Nexus switches and Cisco

UCS Fabric Interconnects enables the networking components to accommodate multiple UCS clusters in a single Cisco UCS domain.

A key benefit of the Cisco Validated Design architecture is the ability to customize the environment to suit your requirements. A Cisco Validated Design scales easily as requirements and demand change. The unit can be scaled both up (adding resources to a Cisco Validated Design unit) and out (adding more Cisco Validated Design units).

The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a hyper-converged desktop virtualization solution. A solution capable of consuming multiple protocols across a single interface allows for choice and investment protection because it truly is a wire-once architecture.

The combination of technologies from Cisco Systems, Inc., NetApp Inc, and VMware Inc. produced a highly efficient, robust, and affordable desktop virtualization solution for a virtual desktop, hosted shared desktop or mixed deployment supporting different use cases. Key components of the solution include the following:

- **More power, same size.** Cisco UCS X210c M7 Compute Node Servers with Up to 2x 4th Gen Intel® Xeon® Scalable Processors (codenamed Sapphire Rapids) with up to 60 cores per processor with up to 8TB with 32 x 256GB DDR5-4800MT/s DIMMs, in a 2-sockets configuration with VMware Horizon support more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel 32-core processors used in this study provided a balance between increased per-server capacity and cost.
- **Fault-tolerance with high availability built into the design.** The various designs are based on multiple Cisco UCS X210c M7 Compute Node Servers for virtual desktop and infrastructure workloads. The design provides N+1 server fault tolerance for every payload type tested.
- **Stress-tested to the limits during aggressive boot scenario.** The 3360 user Remote Desktop Server Hosted (RDSH) sessions and 2240 Windows 11 Virtual Desktops environment booted and registered with the VMware Connection server in under 20 minutes, providing you with an extremely fast, reliable cold-start desktop virtualization system.
- **Stress-tested to the limits during simulated login storms.** The 3360 user RDSH sessions and 2240 Windows 11 Virtual Desktops environment ready state in 48-minutes without overwhelming the processors, exhausting memory, or exhausting the storage subsystems, providing you with a desktop virtualization system that can easily handle the most demanding login and startup storms.
- **Ultra-condensed computing for the datacenter.** The rack space required to support the initial 2000 user system is 3 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects. Incremental seat Cisco converged solutions clusters can be added one at a time to a total of 32 nodes.
- **100 percent virtualized** This CVD presents a validated design that is 100 percent virtualized on VMware ESXi 8.0 U1 All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, VMware Horizon Connection Server components, VMware Horizon virtual desktops and RDSH servers were hosted as virtual machines.
- **Cisco data center management:** Cisco maintains industry leadership with the new Cisco Intersight 5.2(0.230041) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco Intersight, Cisco UCS Central, and Cisco UCS Director ensure that your environments are consistent locally, across Cisco UCS Domains and across the globe. Cisco UCS software suite offers increasingly simplified operational and deployment management, and it

continues to widen the span of control for your organizations' subject matter experts in compute, storage, and network.

- **Cisco 100G Fabric:** Our 100G unified fabric story gets additional validation on 6500 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- **NetApp AFF A400** array provides industry-leading storage solutions that efficiently handles the most demanding I/O bursts (for example, login storms), profile management, and user data management, deliver simple and flexible business continuance, and help reduce storage cost per desktop. Along with a simple to understand storage architecture for hosting all user data components (VMs, profiles, user data) on the same storage array.
- **NetApp clustered Data ONTAP software** enables to seamlessly add, upgrade, or remove storage from the infrastructure to meet the needs of the virtual desktops.
- **VMware Horizon advantage:** VMware Horizon follows a new unified product architecture that supports both Virtual Desktops and Remote Desktop Server Hosted server sessions. This new VMware Horizon release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of user increase.
- **Optimized for performance and scale.** For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the RDS virtual machines did not exceed the number of hyper-threaded (logical) cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for the virtual machines running virtualized RDS systems.
- **Provisioning desktop machines made easy:** VMware Horizon 8 offers basis of the centralized management to create and provisions Remote Desktop Hosted Sessions (RDS) virtual desktops as well as persistent and non-persistent desktops for this solution. The Instant Clone method greatly reduces the amount of life-cycle spend and the maintenance windows for the guest OS.

Technology Overview

This chapter contains the following:

- [Cisco Unified Computing System](#)

Cisco Unified Computing System

Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10-100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform with a unified management domain for managing all resources.

Cisco Unified Computing System consists of the following subsystems:

- **Compute** - The compute piece of the system incorporates servers based on the fourth-generation Intel Xeon processors. Servers are available in Compute Node and rack form factor, managed by Cisco Intersight.
- **Network** - The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lower costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.
- **Virtualization** - The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtual environments to support evolving business needs.
- **Storage access** - Cisco UCS provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides you with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.
- **Management:** The system uniquely integrates compute, network, and storage access subsystems, enabling it to be managed as a single entity through Cisco Intersight software. Cisco Intersight increases IT staff productivity by enabling storage, network, and server administrators to collaborate on Service Profiles that define the desired physical configurations and infrastructure policies for applications. Service Profiles increase business agility by enabling IT to automate and provision resources in minutes instead of days.

Cisco UCS Differentiators

Cisco Unified Computing System is revolutionizing the way servers are managed in the datacenter. The following are the unique differentiators of Cisco Unified Computing System and Cisco Intersight:

- **Embedded Management** – In Cisco UCS, the servers are managed by the embedded firmware in the fabric interconnects, eliminating the need for any external physical or virtual devices to manage the servers.
- **Unified Fabric** – In Cisco UCS, from Compute Node server chassis or rack servers to FI, there is a single Ethernet cable used for LAN, SAN, and management traffic. This converged I/O results in reduced cables, SFPs and adapters - reducing capital and operational expenses of the overall solution.

-
- **Auto Discovery** – By simply inserting the Compute Node server in the chassis or connecting the rack server to the fabric interconnect, discovery and inventory of compute resources occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.
 - **Policy Based Resource Classification** – When Cisco Intersight discovers a compute resource, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD highlights the policy-based resource classification of Cisco Intersight.
 - **Combined Rack and Compute Node Server Management** – Cisco Intersight can manage Cisco UCS X-Series Compute Node Servers and Cisco UCS C-Series Rack Servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.
 - **Model based Management Architecture** – The Cisco Intersight architecture and management database is model based, and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco Intersight with other management systems.
 - **Policies, Pools, Templates** – The management approach in Cisco Intersight is based on defining policies, pools, and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network, and storage resources.
 - **Loose Referential Integrity** – In Cisco Intersight, a server profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. Profiles allow subject matter experts to work independently from one another, providing more flexibility for specialized skillsets in network, storage, security, server, and virtualization to collaborate and accomplish a complex task.
 - **Policy Resolution** – In Cisco Intersight, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organizational hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then the special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.
 - **Server Profiles and Stateless Computing** – A server profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
 - **Built-in Multi-Tenancy Support** – The combination of policies, pools and templates, loose referential integrity, policy resolution in the organizational hierarchy and a service profiles-based approach to compute resources makes Cisco Intersight inherently friendly to multi-tenant environments typically observed in private and public clouds.
 - **Simplified QoS** – Even though Fibre Channel and Ethernet are converged in the Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco Intersight by representing all system classes in one GUI panel.

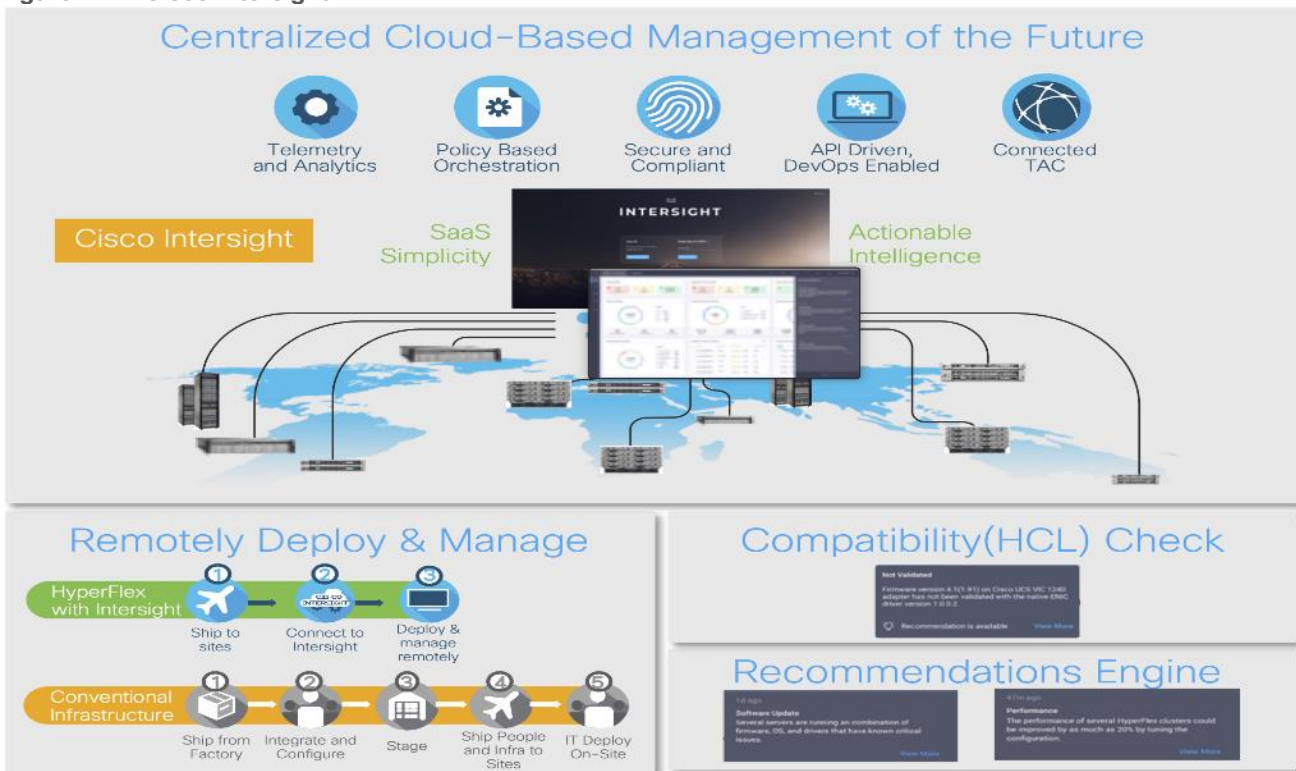
Cisco Intersight

Cisco Intersight is a lifecycle management platform for your infrastructure, regardless of where it resides. In your enterprise data center, at the edge, in remote and branch offices, at retail and industrial sites—all these locations present unique management challenges and have typically required separate tools. Cisco Intersight Software as a Service (SaaS) unifies and simplifies your experience of the Cisco Unified Computing System (Cisco UCS).

Cisco Intersight software delivers a new level of cloud-powered intelligence that supports lifecycle management with continuous improvement. It is tightly integrated with the Cisco Technical Assistance Center (TAC).

Expertise and information flow seamlessly between Cisco Intersight and IT teams, providing global management of Cisco infrastructure, anywhere. Remediation and problem resolution are supported with automated upload of error logs for rapid root-cause analysis.

Figure 4. Cisco Intersight



- Automate your infrastructure

Cisco has a strong track record for management solutions that deliver policy-based automation to daily operations. Intersight SaaS is a natural evolution of our strategies. Cisco designed Cisco UCS to be 100 percent programmable. Cisco Intersight simply moves the control plane from the network into the cloud. Now you can manage your Cisco UCS infrastructure wherever it resides through a single interface.

- Deploy your way

If you need to control how your management data is handled, comply with data locality regulations, or consolidate the number of outbound connections from servers, you can use the Cisco Intersight Virtual Appliance for an on-premises experience. Cisco Intersight Virtual Appliance is continuously updated just like the SaaS version, so regardless of which approach you implement, you never have to worry about whether your management software is up to date.

- DevOps ready

If you are implementing DevOps practices, you can use the Cisco Intersight API with either the cloud-based or virtual appliance offering. Through the API you can configure and manage infrastructure as code—you are not merely configuring an abstraction layer; you are managing the real thing. Through the API and support [cloud-based RESTful API](#), Terraform providers, Microsoft PowerShell scripts, or Python software, you can automate the deployment of settings and software for both physical and virtual layers. Using the API, you can simplify infrastructure lifecycle operations and increase the speed of continuous application delivery.

- Pervasive simplicity
 - Simplify the user experience by managing your infrastructure regardless of where it is installed.
 - Automate updates to Cisco UCS Data Platform software, reducing complexity and manual efforts.
- Actionable intelligence
 - Use best practices to enable faster, proactive IT operations.
 - Gain actionable insight for ongoing improvement and problem avoidance.
- Manage anywhere
 - Deploy in the data center and at the edge with massive scale.
 - Get visibility into the health and inventory detail for your Intersight Managed environment on-the-go with the Cisco Intersight Mobile App.

For more information about Cisco Intersight and the different deployment options, go to: [Cisco Intersight – Manage your systems anywhere.](#)

Solution Components

This chapter contains the following:

- [Cisco UCS Fabric Interconnect](#)
- [Cisco Unified Computing System X-Series](#)
- [Cisco Switches](#)
- [Cisco Intersight](#)
- [VMware Horizon](#)
- [NetApp A-Series All Flash FAS](#)
- [NetApp ONTAP 9.13.1P3](#)
- [VMware vSphere 8.0 U1](#)
- [Cisco Intersight Assist Device Connector for VMware vCenter and NetApp ONTAP](#)
- [NetApp ONTAP Tools for VMware vSphere](#)
- [NetApp NFS Plug-in for VMware VAAI](#)
- [NetApp XCP File Analytics](#)

Cisco UCS Fabric Interconnect

The Cisco UCS Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. Depending on the model chosen, the Cisco UCS Fabric Interconnect offers line-rate, low-latency, lossless 10 Gigabit, 25 Gigabit, 40 Gigabit, or 100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel connectivity. Cisco UCS Fabric Interconnects provide the management and communication backbone for the Cisco UCS C-Series, X-Series and Cisco UCS B-Series Blade Servers, and Cisco UCS 5100 Series Compute Node Server Chassis. All servers and chassis, and therefore all Compute Nodes, attached to the Cisco UCS Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabrics, the Cisco UCS Fabric Interconnects provide both the LAN and SAN connectivity for all servers within its domain.

For networking performance, the Cisco UCS 6536 uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10/25/40/100 Gigabit Ethernet ports, a switching capacity of 7.42 Tbps per FI and 14.84 Tbps per unified fabric domain, independent of packet size and enabled services. It enables 1600Gbps bandwidth per Cisco UCS X9508 chassis with X9108-IFM-100G in addition to enabling end-to-end 100G ethernet and 200G aggregate bandwidth per X210c compute node. With the X9108-IFM-25G and the IOM 2408, it enables 400Gbps bandwidth per chassis per FI domain. The product family supports Cisco low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric capabilities, which increases the reliability, efficiency, and scalability of Ethernet networks. The Cisco UCS 6536 Fabric Interconnect supports multiple traffic classes over a lossless ethernet fabric from the server through the fabric interconnect. Significant TCO savings come from the Unified Fabric optimized server design in which Network Interface Cards (NICs), Host Bus Adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6536 Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide a single point for connectivity and management for the entire Cisco Unified Computing System. Typically deployed as an active/active pair, the system's FIs integrate all components into a single, highly available management domain controlled by Cisco Intersight. Cisco UCS FIs

provide a single unified fabric for the system, with low-latency, lossless, cut-through switching that supports LAN, SAN, and management traffic using a single set of cables.

Figure 5. Cisco UCS 6536 Fabric Interconnect




The Cisco UCS 6536 36-Port Fabric Interconnect (Figure 5) is a One-Rack-Unit (1RU) 1/10/25/40/100 Gigabit Ethernet, FCoE, and Fibre Channel switch offering up to 7.42 Tbps throughput and up to 36 ports. The switch has 32 40/100-Gbps Ethernet ports and 4 unified ports that can support 40/100-Gbps Ethernet ports or 16 Fiber Channel ports after breakout at 8/16/32-Gbps FC speeds.

Cisco Unified Computing System X-Series

The Cisco UCS X-Series Modular System is designed to take the current generation of the Cisco UCS platform to the next level with its future-ready design and cloud-based management. Decoupling and moving the platform management to the cloud allows Cisco UCS to respond to your feature and scalability requirements in a much faster and efficient manner. Cisco UCS X-Series state of the art hardware simplifies the data-center design by providing flexible server options. A single server type, supporting a broader range of workloads, results in fewer different data center products to manage and maintain. The Cisco Intersight cloud-management platform manages Cisco UCS X-Series as well as integrates with third-party devices, including VMware vCenter and NetApp storage, to provide visibility, optimization, and orchestration from a single platform, thereby driving agility and deployment consistency.

Figure 6. Cisco UCS X9508 Chassis



Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As seen in Figure 7, the Cisco UCS X9508 chassis has only a power-distribution midplane. This midplane-free design provides fewer obstructions for better airflow. For I/O connectivity, vertically oriented compute nodes intersect with horizontally oriented fabric modules, allowing the chassis to support future fabric innovations. Cisco UCS X9508 Chassis' superior packaging enables larger compute nodes, thereby providing more space for actual compute components, such as memory, GPU, drives, and accelerators. Improved airflow through the chassis enables support for higher

power components, and more space allows for future thermal solutions (such as liquid cooling) without limitations.

Figure 7. Cisco UCS X9508 Chassis - Midplane Free Design



The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes and a pool of current and future I/O resources that includes GPU accelerators, disk storage, and nonvolatile memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) that connect the chassis to upstream Cisco UCS 6400 or 6500 Series Fabric Interconnects. At the bottom rear of the chassis are slots to house X-Fabric modules that can flexibly connect the compute nodes with I/O devices. Six 2800W Power Supply Units (PSUs) provide 54V power to the chassis with N, N+1, and N+N redundancy. A higher voltage allows efficient power delivery with less copper and reduced power loss. Efficient, 100mm, dual counter-rotating fans deliver industry-leading airflow and power efficiency, and optimized thermal algorithms enable different cooling modes to best support your environment.

Cisco UCSX-I-9108-100G Intelligent Fabric Modules

In the end-to-end 100Gbps Ethernet design, for the Cisco UCS X9508 Chassis, the network connectivity is provided by a pair of Cisco UCSX-I-9108-100G Intelligent Fabric Modules (IFMs). Like the fabric extenders used in the Cisco UCS 5108 Blade Server Chassis, these modules carry all network traffic to a pair of Cisco UCS 6536 Fabric Interconnects (FIs). IFMs also host the Chassis Management Controller (CMC) for chassis management. In contrast to systems with fixed networking components, Cisco UCS X9508's midplane-free design enables easy upgrades to new networking technologies as they emerge making it straightforward to accommodate new network speeds or technologies in the future.

Figure 8. Cisco UCSX-I-9108-100G Intelligent Fabric Module




Each IFM supports eight 100Gb uplink ports for connecting the Cisco UCS X9508 Chassis to the FIs and 8 100Gb or 32 25Gb server ports for the eight compute nodes. IFM server ports can provide up to 200 Gbps of unified fabric connectivity per compute node across the two IFMs. The uplink ports connect the chassis to the Cisco UCS FIs, providing up to 1600Gbps connectivity across the two IFMs. The unified fabric carries management, VM, and Fibre Channel over Ethernet (FCoE) traffic to the FIs, where server management traffic is routed to the Cisco Intersight cloud operations platform, FCoE traffic is forwarded to either native Fibre Channel interfaces through unified ports on the FI (to Cisco MDS switches) or to FCoE uplinks (to Cisco Nexus switches supporting SAN switching), and data Ethernet traffic is forwarded upstream to the data center network (using Cisco Nexus switches).

Cisco UCS X210c M7 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to 8 Cisco UCS X210c M7 or X210c M6 Compute Nodes. The hardware details of the Cisco UCS X210c M7 Compute Nodes are shown in [Figure 9](#):

Figure 9. Cisco UCS X210c M7 Compute Node

UCS X210c M7 Compute Node – Key features



The Cisco UCS X210c M7 features:

- **CPU:** Up to 2x 4th Gen Intel Xeon Scalable Processors with up to 60 cores per processor and 2.625 MB Level 3 cache per core and up to 112.5 MB per CPU.
- **Memory:** Up to 32 x 256 GB DDR5-4800 DIMMs for a maximum of 8 TB of main memory.
- **Disk storage:** Up to 6 SAS or SATA drives or NVMe drives can be configured with the choice of an internal RAID controller or passthrough controllers. 2 M.2 memory cards can be added to the Compute Node with optional hardware RAID.
- **GPUs:** The optional front mezzanine GPU module allows support for up to two HHL GPUs. Adding a mezzanine card and a Cisco UCS X440p PCIe Node allows up to four more GPUs to be supported with an X210c M7.
- **Virtual Interface Card (VIC):** Up to 2 VICs including an mLOM Cisco UCS VIC 15231 or an mLOM Cisco UCS VIC 15420 and a mezzanine Cisco UCS VIC card 15422 can be installed in a Compute Node.
- **Security:** The server supports an optional Trusted Platform Module (TPM). Additional security features include a secure boot FPGA and ACT2 anticounterfeit provisions.


Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS X210c M7 Compute Nodes supports the Cisco UCS VIC 15231.

Cisco UCS VIC 15231

Cisco UCS VIC 15231 fits the mLOM slot in the Cisco UCS X210c Compute Node and enables up to 100 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 200 Gbps of connectivity per server. Cisco UCS VIC 15231 connectivity to the IFM and up to the fabric interconnects is delivered through 100Gbps. Cisco UCS VIC 15231 supports 512 virtual interfaces (both FCoE and Ethernet) along with the latest networking innovations such as NVMeoF over FC or TCP, VxLAN/NVGRE offload, and so forth.

Figure 10. Cisco UCS VIC 15231 in Cisco UCS X210c M7



Cisco Switches


Cisco Nexus N9K-C93180YC-FX3S Switches

The N9K-C93180YC-FX3S Switch provides a flexible line-rate Layer 2 and Layer 3 feature set in a compact form factor. Designed with Cisco Cloud Scale technology, it supports highly scalable cloud architectures. With the option to operate in Cisco NX-OS or Application Centric Infrastructure (ACI) mode, it can be deployed across enterprise, service provider, and Web 2.0 data centers.

- Architectural Flexibility
 - Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
 - Leaf node support for Cisco ACI architecture is provided in the roadmap
 - Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support
- Feature Rich
 - Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
 - ACI-ready infrastructure helps users take advantage of automated policy-based systems management
 - Virtual Extensible LAN (VXLAN) routing provides network services
 - Rich traffic flow telemetry with line-rate data collection
 - Real-time buffer utilization per port and per queue, for monitoring traffic micro-bursts and application traffic patterns
- Highly Available and Efficient Design
 - High-density, non-blocking architecture
 - Easily deployed into either a hot-aisle and cold-aisle configuration
 - Redundant, hot-swappable power supplies and fan trays

- Simplified Operations
 - Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
 - An Intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
 - Python Scripting for programmatic access to the switch command-line interface (CLI)
 - Hot and cold patching, and online diagnostics

Figure 11. Cisco Nexus N9K-C93180YC-FX3S Switch



Cisco Intersight

The Cisco Intersight platform is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support. The Cisco Intersight platform is designed to be modular, so you can adopt services based on their individual requirements. The platform significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses a unified Open API design that natively integrates with third-party platforms and tools.

Figure 12. Cisco Intersight Overview



The main benefits of Cisco Intersight infrastructure services are as follows:

- Simplify daily operations by automating many daily manual tasks
- Combine the convenience of a SaaS platform with the capability to connect from anywhere and manage infrastructure through a browser or mobile app

-
- Stay ahead of problems and accelerate trouble resolution through advanced support capabilities
 - Gain global visibility of infrastructure health and status along with advanced management and support capabilities
 - Upgrade to add workload optimization and Kubernetes services when needed

Cisco Intersight Virtual Appliance and Private Virtual Appliance

In addition to the SaaS deployment model running on Intersight.com, on-premises options can be purchased separately. The Cisco Intersight Virtual Appliance and Cisco Intersight Private Virtual Appliance are available for organizations that have additional data locality or security requirements for managing systems. The Cisco Intersight Virtual Appliance delivers the management features of the Cisco Intersight platform in an easy-to-deploy VMware Open Virtualization Appliance (OVA) or Microsoft Hyper-V Server virtual machine that allows you to control the system details that leave your premises. The Cisco Intersight Private Virtual Appliance is provided in a form factor specifically designed for users who operate in disconnected (air gap) environments. The Private Virtual Appliance requires no connection to public networks or back to Cisco to operate.

Cisco Intersight Assist

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. A data center could have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight, but does not connect directly with it, will need a connection mechanism. Cisco Intersight Assist provides that connection mechanism. In FlexPod, VMware vCenter and NetApp Active IQ Unified Manager connect to Intersight with the help of Intersight Assist VM.

Cisco Intersight Assist is available within the Cisco Intersight Virtual Appliance, which is distributed as a deployable virtual machine contained within an Open Virtual Appliance (OVA) file format. More details about the Cisco Intersight Assist VM deployment configuration is covered in later sections.

Licensing Requirements

Cisco Intersight offers services that allow you to manage, automate, optimize, and support your physical and virtual infrastructure.

With Cisco Intersight, you can activate licenses for the Infrastructure Service and Cloud Orchestrator. Use these services to manage Cisco endpoints such as the Cisco UCS server and Cisco HyperFlex system.

For more information on the endpoints and the features available in each tier, see the [Cisco Infrastructure Services License](#) section.

Difference between Tech Preview and General Availability on Intersight Features

The Cisco Intersight features delivered as Tech Preview are indicated in the platform UI with the Tech Preview designation. Tech Preview provides a preview of a functionality that is still under development. The Tech Preview features are not intended to be used in production environments. These features, including their GUI and API interfaces, may change between Tech Preview and General Availability. Additionally, the license tier in which the Tech Preview features are delivered may differ from the license tier in which the features become Generally Available. For more information on the difference between Tech Preview and General Availability on Intersight features, see [Licensing FAQ](#).

VMware Horizon

VMware Horizon is a modern platform for running and delivering virtual desktops and apps across the hybrid cloud. It provides administrators with simple, automated, and secure desktop and app management. For users, it provides a consistent experience across devices and locations.

VMware Horizon is a VDI solution that allows users to access their desktops, apps, and data from any device in a secure manner. This end-to-end solution provides complete management, delivery, and security of virtual desktops and applications. The latest version comes with various new features and enhancements, including advanced security measures like certificate pinning for Android and iOS clients, better performance, scalability for cloud-hosted virtual desktops and applications, compatibility with new platforms and operating systems, and more functional remote access.

VMware Horizon 8 also offers a simplified and streamlined deployment process, making it easier for administrators to manage their virtual desktop infrastructure. It includes various tools for monitoring and optimizing virtual desktop and app performance and has advanced automation and customization options.

VMware Horizon 8 is a flexible and powerful VDI solution that can help organizations improve productivity, reduce costs, and ensure better security and compliance.

For more information, go to: [VMware Horizon](#).

VMware Horizon 8 2212

VMware Horizon 8 2212 is an [Extended Service Branch \(ESB\)](#). VMware provides periodic service packs (SP) updates for ESB releases, which only include cumulative critical bug fixes and security fixes without any addition features. This allows customers to deploy a stable Horizon platform for their critical deployments.

[VMware Horizon version 2012](#) provides the following new features and enhancements. This information is grouped by installable component.

- Virtual Desktops and Applications

VMware Horizon 8 version 2212 in conjunction with App Volumes 4 version 2212 introduces Horizon Published Apps on Demand. With this new feature, administrators can use App Volumes applications directly in their instant-clone RDS farms. Now applications can be delivered dynamically to a generic Windows OS as users launch them. This greatly simplifies static image management and gives administrators the ability to reduce their application specific farms. This also brings the Horizon and App Volumes administration consoles closer together, allowing Horizon administrators to add App Volumes Manager servers and entitle applications to users without the need for duplicate entitlements in App Volumes. This feature creates an opportunity to reduce the time-consuming management of application installations on RDS Farms and enables scenarios such as multiple users being able to use different versions of the same application while logged in to the same RDS Server.

- Microsoft MAK licenses are now supported with Instant Clones

For vTPM-enabled Instant Clone desktop pools, Horizon previously always used Mode A provisioning (Instant Clones with Parent VM) due to a bug in older ESXi versions. With the resolution of this bug, Horizon now also supports Mode B provisioning (Instant Clones without Parent VM) for vTPM-enabled desktop pools if all hosts in the cluster are running ESXi 7.0 Update 3f or later with Horizon 8 version 2212 or later.

When you create an automated pool of full clone desktops, you can now specify an active directory OU in which computer accounts can be created. Previously, computer accounts would get created in the default

OU and administrators would manually move them after pool creation. This feature, which already exists for Instant Clone desktop pools, addresses this pain point for administrators.

Improved GPU performance on Physical Machines running Windows Server 2019 with Horizon Indirect Display Driver based setup.

The network settings for a create Instant Clone pool or farm workflow are now set to the network settings of a golden image instead of a snapshot. This simplifies management for administrators as they only have to keep track of network settings of a golden image rather than many of its snapshots.

- Horizon Connection Server

Horizon 8 now supports a maximum of 500 Virtual Machines per ESXi host when using non-vSAN storage. The achievable maximum depends on the workload and specifics of the hardware. See VMware Configuration Maximums for all Horizon Configuration Maximums.

Cloud Pod Architecture is supported with IPv6 environments for more security and added address spaces.

Administrators can now generate a CSR configuration file, import a CA-signed certificate to Connection Server, and monitor health of the certificate from Horizon Console.

Hybrid Azure Active Directory for SSO is now supported on instant clone desktop pools.

- Horizon Agent for Windows

The Horizon Agent for Windows has been migrated from Azul OpenJDK to BellSoft OpenJDK.

- Horizon Agent for Linux

This release adds support for the following Linux distributions:

- Debian 10.13 and 11.5
- Red Hat Enterprise Linux (RHEL) Workstation 8.7 and 9.1
- Red Hat Enterprise Linux (RHEL) Server 8.7 and 9.1
- SUSE Linux Enterprise Desktop (SLED) 15 SP4
- SUSE Linux Enterprise Server (SLES) 15 SP4

- This release supports the MATE desktop environment on desktops running RHEL 7.9.

- Beginning with this release, the following Linux distributions are no longer supported:

- RHEL Workstation 7.8 and earlier, 8.5, 8.3 and earlier
- RHEL Server 7.8
- CentOS 7.8 and earlier
- SLED/SLES 12 SP3 and 15 SP2
- The Horizon Agent for Linux has been migrated from Azul OpenJDK to BellSoft OpenJDK.

- Horizon Client

For information about new features in a Horizon Client, including HTML Access, see the release notes for that client.

- General

- You can now enable or disable TrueSSO Trigger Mode in the add or edit SAML Authenticator workflow on the Horizon Console.

- All Horizon Console grids can persist hide/unhide column preferences.
- Horizon Console login username and domain can be persisted in browser storage.
- Horizon administrators with Smartcard bypass privilege can authenticate and consume APIs even if connection server mandates Smartcard authentication.
- Horizon 8 has been tested to work with Microsoft Defender Endpoint.
- Horizon RESTful APIs
 - New RESTful APIs and new versions of existing RESTful APIs have been added to help in automation for customer deployments. To get the latest documentation for the Horizon RESTful APIs:
 - a. Install or Upgrade to the latest released version of Connection Server.
 - b. Navigate to <https://<CS-IP//FQDN>rest/swagger-ui.html> from any browser.
 - c. Click Select a spec from the top right of the browser. Select Latest to see the latest version of APIs. Select Default to view all versions of all APIs.

NetApp A-Series All Flash FAS

Powered by [NetApp ONTAP data management software](#), [NetApp AFF A-Series systems](#) (NetApp AFF) deliver the industry’s highest performance, superior flexibility, and best-in-class data services and cloud integration to help you accelerate, manage, and protect business-critical data across the hybrid cloud. It is a robust scale-out platform built for virtualized environments, combining low-latency performance with best-in-class data management, built-in efficiencies, integrated data protection, multiprotocol support, and nondisruptive operations. Deploy as a stand-alone system or as a high-performance tier in a NetApp ONTAP configuration.

A wide range of organizations, from enterprise to midsize businesses, rely on NetApp AFF A-Series to:

- Simplify operations with seamless data management, on the premises and in the cloud
- Accelerate traditional and new-generation applications
- Keep business-critical data available, protected, and secure
- Accelerates applications and future-proofs your infrastructure

In the modern data center, IT is charged with driving maximum performance for business-critical workloads, scaling without disruption as the business grows, and enabling the business to take on new data-driven initiatives. NetApp AFF A-Series systems handle all of it with ease.

The NetApp AFF A-Series lineup includes the A150, A250, A400, A700, A800 and A900. These controllers and their technical specifications are listed in [Table 1](#). For more information about the NetApp A-Series AFF controllers, see:

<http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>

<https://hww.netapp.com/Controller/Index?platformTypeld=5265148>

Table 1. NetApp AFF Technical Specifications

Specifications	AFF A150	AFF A250	AFF A400	AFF A800	AFF A900
Maximum scale-out	2-24 nodes (12 HA pairs)	2-24 nodes (12 HA pair)	2-24 nodes (12 HA pair)	2-24 nodes (12 HA pair)	2-24 nodes (12 HA pair)
Maximum SSDs	864	576	5760	2880	5760

Specifications	AFF A150	AFF A250	AFF A400	AFF A800	AFF A900
Max effective capacity	26PB	35PB	702.7PB	316.3PB	702.7PB
Controller form factor	2U; 24 internal SSD slots	2U	4U	4U with 48 SSD slots	8U
PCIe expansion slots	n/a	4	10	8	20
FC target ports (32Gb autoranging)	n/a	16	24	32	64
FC target ports (16Gb autoranging)	n/a	n/a	32(with FC mezzanine card)	32	64
FCoE target ports, UTA2	8	n/a	n/a	n/a	64
100GbE ports (40GbE autoranging)	n/a	4	16	20	32
25GbE ports (10GbE autoranging)	n/a	20	16	16	64
10GbE ports	4	n/a	32	32	64
12Gb/6Gb SAS ports	4	8	32	n/a	64
Storage networking supported	NVMe/TCP, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon S3	NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3	NFSv4/RDMA, NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3	NFSv4/RDMA, NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3	NVMe/TCP, NVMe/FC, FC, iSCSI, NFS, pNFS, CIFS/SMB, Amazon s3
OS version	ONTAP 9.12 1P1 or later	ONTAP 9.8 RC1 or later	ONTAP 9.7 RC1 or later	ONTAP 9.7 RC1 or later	ONTAP 9.10.1 RC2 or later

The following are few advantages of NetApp AFF A-Series:

- Maximum performance for your most demanding applications:
 - NetApp AFF A-Series systems deliver industry-leading performance proven by SPC-1 and SPEC SFS industry benchmarks, making them ideal for demanding, highly transactional applications such as Oracle, Microsoft SQL Server, MongoDB databases, VDI, and server virtualization.
 - With the power of front-end NVMe/FC and NVMe/TCP host connectivity and back-end NVMe-attached SSDs, our high-end AFF A900 systems deliver latency as low as 100µs. Based on a high-resiliency design, the A900 also delivers high RAS and enables non-disruptive in-chassis upgrade from its

predecessor A700. The A800 delivers high performance in a compact form factor and is especially suited for EDA and Media and Entertainment workloads. The midrange, most versatile NetApp AFF A400 system features hardware acceleration technology that significantly enhances performance and storage efficiency. And our entry-level, budget-friendly NetApp AFF A400, provides 40% more performance and 33% more efficiency at no extra cost compared with its predecessor.

- NetApp AFF A-Series also lets you:
 - Drive mission-critical SAN workloads with symmetric active-active host connectivity for continuous availability and instant failover.
 - Consolidate workloads to deliver up to 14.4 million IOPS at 1ms latency in a cluster with a truly unified scale-out architecture. Built-in adaptive quality of service (QoS) safeguards SLAs in multi-workload and multitenant environments.
 - Manage massively scalable NAS containers of up to 20PB and 400 billion files with a single namespace.
 - Improve the speed and productivity of collaboration across multiple locations and increase data throughput for read-intensive applications with NetApp FlexCache software.
 - Modernize with advanced connectivity

NetApp AFF A-Series all-flash systems deliver industry-leading performance, density, scalability, security, and network connectivity. As the first enterprise-grade storage systems to support both NVMe/TCP and NVMe/FC, NetApp AFF A-Series systems boost performance with modern network connectivity. With NVMe/TCP, which uses the commonly available Ethernet infrastructure, you don't have to invest in new hardware to take advantage of the faster host connectivity. With NVMe/FC, you can get twice the IOPS and cut application response time in half compared with traditional FC. These systems support a range of ecosystems, including VMware, Microsoft Windows 11, and Linux, with storage path failover. For most, integrating NVMe/FC into an existing SAN is a simple, nondisruptive software upgrade.

- Scale without disruption

With NetApp AFF A-Series, you can integrate new technologies and private or public cloud into your infrastructure nondisruptively. NetApp AFF A-Series is the only all-flash array that enables you to combine different controllers, SSD sizes, and new technologies so that your investment is protected. The NVMe-based AFF systems also support SAS SSDs, maximizing the flexibility and cost effectiveness of your upgrade:

- Best balance between price, technology, features, and performance.
- Increase operational efficiency

IT departments are striving to make budgets go further and to allow IT staff to focus on new value-added projects rather than on day-to-day IT management. NetApp AFF systems simplify IT operations, which therefore reduces data center cost. In particular, our entry-level system, the NetApp AFF A400, delivers best-in-class performance and efficiency to mid-size businesses so they can consolidate more workloads and eliminate silos.

- Provision storage in minutes

NetApp AFF systems offer broad application ecosystem support and deep integration for enterprise applications, virtual desktop infrastructure (VDI), database, and server virtualization, supporting Oracle, Microsoft SQL Server, VMware, SAP, MySQL, and more. You can quickly provision storage in less than 10 minutes with NetApp ONTAP System Manager. In addition, infrastructure management tools simplify and automate common storage tasks so you can:

- Easily provision and rebalance workloads by monitoring clusters and nodes.
- Use one-click automation and self-service for provisioning and data protection.
- Upgrade OS and firmware with a single-click
- Import LUNs from third-party storage arrays directly into an AFF system to seamlessly migrate data.

Additionally, the [NetApp Active IQ Digital Advisor](#) engine enables you to optimize your NetApp systems with predictive analytics and proactive support. Fueled by the massive NetApp user base, AI and machine learning create actionable insights that help you prevent problems, optimize your configuration, save time, and make smarter decisions.

- Achieve outstanding storage savings

NetApp employs various capabilities to promote optimal capacity savings and to drive down your TCO. AFF A-Series system's support for solid-state drives (SSDs) with multistream write technology, combined with advanced SSD partitioning, provides maximum usable capacity, regardless of the type of data that you store. Thin provisioning; NetApp Snapshot copies; and inline data reduction features, such as deduplication, compression, and compaction, provide substantial additional space savings—without affecting performance—enabling you to purchase the least amount of storage capacity possible.

- Build your hybrid cloud with ease

Your data fabric built by NetApp helps you simplify and integrate data management across cloud and on-premises environments to meet business demands and gain a competitive edge. With AFF A-Series, you can connect to more clouds for more data services, data tiering, caching, and disaster recovery. You can also:

- Maximize performance and reduce overall storage costs by automatically tiering cold data to the cloud with FabricPool
- Instantly deliver data to support efficient collaboration across your hybrid cloud
- Protect your data by taking advantage of Amazon Simple Storage Service (Amazon S3) cloud resources—on premises and in the public cloud
- Accelerate read performance for data that is shared widely throughout your organization and across hybrid cloud deployments
- Keep data available, protected, and secure

As organizations become more data driven, the business impact of data loss can be increasingly dramatic—and costly. IT must protect data from both internal and external threats, ensure data availability, eliminate maintenance disruptions, and quickly recover from failures.

- Integrated data protection

NetApp AFF A-Series systems come with a full suite of acclaimed NetApp integrated and application-consistent data protection software. Key capabilities include:

- Native space efficiency with cloning and NetApp Snapshot copies reduce storage costs and minimize performance impact. Up to 1,023 copies are supported.
- [NetApp SnapCenter](#) software provides application-consistent data protection and clone management to simplify application management.
- [NetApp SnapMirror](#) technology replicates to any NetApp FAS or AFF system on the premises or in the cloud, reducing overall system costs.

- Business continuity and fast disaster recovery

With NetApp AFF, you can maintain constant data availability with zero data loss and zero downtime. NetApp MetroCluster software provides synchronous replication to protect your entire system, and NetApp SnapMirror Business Continuity provides a more flexible, cost-effective business continuity to even with more granular replication of selected critical data.

- Security everywhere

Flexible encryption and key management help guard your sensitive data on the premises, in the cloud, and in transit. The market-leading anti-ransomware protection for both preemption and post-attack recovery safeguards your critical data from ransomware attacks and can prevent catastrophic financial consequences. With the simple and efficient security solutions, you can:

- Achieve FIPS 140-2 compliance (Level 1 and Level 2) with self-encrypting drives and use any type of drives with software-based encryption
- Meet governance, risk, and compliance requirements with security features such as secure purge; logging and auditing monitors; and write once, read many (WORM) file locking
- Protect against threats with multifactor authentication, role-based access control, secure multitenancy, and storage-level file security

NetApp AFF A400

The NetApp AFF A400 provides full end-to-end NVMe support. The frontend FC-NVMe connectivity makes it possible to achieve optimal performance from an all-flash array for workloads that include artificial Intelligence, machine learning, and real-time analytics as well as business-critical databases. The frontend NVMe-TCP connectivity enables customers to take advantage of NVMe technology over existing ethernet infrastructure for faster host connectivity. On the back end, the A400 supports both serial-attached SCSI (SAS) and NVMe-attached SSDs, offering the versatility for current customers to move up from their legacy A-Series systems and satisfying the increasing interest that all customers have in NVMe-based storage.

The NetApp AFF A400 has greater port availability, network connectivity, and expandability. The NetApp AFF A400 has 10 PCIe Gen3 slots per high availability pair. The NetApp AFF A400 is available with 10GbE, 25GbE or 100GbE, as well as 16/32Gb/FC and NVMe/FC network connectivity. This model was created to keep up with changing business needs and performance and workload requirements by merging the latest technology for data acceleration and ultra-low latency in an end-to-end NVMe storage system.

Figure 13. NetApp AFF A400 Front View



Figure 14. NetApp AFF A400 Rear View





Note: We used 4 port 32Gb FC HBA on slot 1 (1a,1b, other two ports unused) for front-end FC SAN connection, 4x25Gb Ethernet NICs on slot 0 (e0e, e0f, e0g, e0h) for NAS connectivity, 2x100Gb ethernet ports on slot 3 (e3a, e3b) used for cluster interconnect, 2x25Gb ethernet on slot 0 (e0a, e0b) used for Node HA interconnect, 2x100Gb ethernet on slot 0 (e0c, e0d) and 2x100Gb ethernet on slot 5 (e5a, e5b) are used for backend NVMe storage connectivity.

NetApp AFF C-Series Storage

The NetApp AFF C-Series all-flash systems are ideal for tier 1 and tier 2 workloads while also delivering better performance than disks. These systems are suited for large-capacity deployment with a small footprint as an affordable way to modernize data center to all flash and also connect to the cloud. Powered by NetApp ONTAP data management software, NetApp AFF C-Series systems deliver industry-leading efficiency, superior flexibility, and best-in-class data services and cloud integration to help scale IT infrastructure, simplify data management, reduce storage cost, rack space usage, power consumption, and improve sustainability significantly. NetApp offers several AFF C-series controllers to meet varying demands of the field. The high-end NetApp AFF C800 systems offer superior performance. The midrange NetApp AFF C400 delivers high performance and good expansion capability. The entry-level NetApp AFF C250 systems provide balanced performance, connectivity, and expansion options for a small footprint deployment.

For more information about the NetApp AFF C-series controllers, see the NetApp AFF C-Series product page: <https://www.netapp.com/data-storage/aff-c-series/>

You can view or download more technical specifications of the NetApp AFF C-Series controllers here: <https://www.netapp.com/media/81583-da-4240-aff-c-series.pdf>

You can look up the detailed NetApp storage product configurations and limits here: <https://www.netapp.com/>

FlexPod CVDs provide reference configurations and there are many more supported IMT configurations that can be used for FlexPod deployments, including NetApp hybrid storage arrays.

NetApp ASA (All-flash SAN Array)

NetApp ASA, a family of modern all-flash block storage that's designed for customers who need resilient, high-throughput, low-latency solutions for their mission-critical workloads. Many businesses see the benefit of SAN solutions. Especially when every minute of downtime can cost hundreds of thousands of dollars, or when poor performance prevents you from fulfilling your mission. Unified storage is often a convenient consolidated solution for file and block workloads, but customers might prefer a dedicated SAN system to isolate these workloads from others. NetApp ASA is block optimized and supports NVMe/TCP and NVMe/FC as well as standard FC and iSCSI protocols. Building upon the foundation of well-architected SAN, ASA offers your organization the following benefits:

- Six nines (99.9999%) availability that's backed by an industry-leading 6 Nines Data Availability Guarantee
- Massive scalability with the NetApp ONTAP cluster capability, which enables you to scale out ASA storage to more than 350PB of effective capacity
- Industry-leading storage efficiency that's built-in and supported by a simple, straightforward [Storage Efficiency Guarantee](#)
- The most comprehensive cloud connectivity available

- Cost-effective integrated data protection

For more information about NetApp ASA, see the NetApp ASA product page: <https://www.netapp.com/data-storage/all-flash-san-storage-array/>

You can view or download more technical specifications of the NetApp ASA controllers here: https://www.netapp.com/media/87298-NA-1043-0523_ASA_AFF_Tech_Specs_HR.pdf

NetApp ONTAP 9

NetApp storage systems harness the power of ONTAP to simplify the data infrastructure from edge, core, and cloud with a common set of data services and 99.9999 percent availability. NetApp ONTAP 9 data management software from NetApp enables you to modernize your infrastructure and transition to a cloud-ready data center. NetApp ONTAP 9 has a host of features to simplify deployment and data management, accelerate and protect critical data, and make infrastructure future-ready across hybrid-cloud architectures.

NetApp ONTAP 9 is the data management software that is used with the NetApp AFF A400 all-flash storage system in this solution design. NetApp ONTAP software offers secure unified storage for applications that read and write data over block- or file-access protocol storage configurations. These storage configurations range from high-speed flash to lower-priced spinning media or cloud-based object storage. NetApp ONTAP implementations can run on NetApp engineered AFF, FAS or ASA series arrays and in private, public, or hybrid clouds (NetApp Private Storage and NetApp Cloud Volumes ONTAP). Specialized implementations offer best-in-class converged infrastructure, featured here as part of the FlexPod Datacenter solution or with access to third-party storage arrays (NetApp FlexArray virtualization). Together these implementations form the basic framework of the NetApp Data Fabric, with a common software-defined approach to data management, and fast efficient replication across systems. FlexPod and ONTAP architectures can serve as the foundation for both hybrid cloud and private cloud designs.

Read more about all the capabilities of ONTAP data management software here: <https://www.netapp.com/us/products/data-management-software/ontap.aspx>.

NetApp ONTAP 9.13.1P3

NetApp ONTAP Features for VDI

The following are the NetApp ONTAP features for VDI:

- Secure Multi-Tenancy—Tenants can be in overlapping subnet or can use identical IP subnet range.
- Multi-Protocol—Same storage system can be used for Block/File/Object storage demands.
- FlexGroup Volumes—High performance and massive capacity (~20PB and ~40 billion files) for file shares and for hosting VDI pools.
- FlexCache—Enables Single Global Namespace can be consumed around the clouds or multi-site.
- File System Analytics—Fast query to file metadata on the SMB file share.
- Ease of management with vCenter Plugins—Best practices are validated and implemented while provisioning. Supports VAAI and VASA for fast provisioning and storage capability awareness.
- SnapCenter integration with vCenter—Space efficient data protection with snapshots and FlexClones.
- Automation support—Supports RESTapi, has modules for Ansible, PowerShell, and so on.
- Storage Efficiency—Supports inline dedupe, compression, thin provisioning, and so on. Guaranteed dedupe of 8:1 for VDI.

- Adaptive QoS—Adjusts QoS setting based on space consumption.
- ActiveIQ Unified Manager—Application based storage provisioning, Performance Monitoring, End-End storage visibility diagrams.

Storage Efficiency

Storage efficiency has always been a primary architectural design point of ONTAP. A wide array of features allows businesses to store more data using less space. In addition to deduplication and compression, businesses can store their data more efficiently by using features such as unified storage, multi-tenancy, thin provisioning, and NetApp Snapshot technology.

Starting with ONTAP 9, NetApp guarantees that the use of NetApp storage efficiency technologies on AFF systems reduce the total logical capacity used to store your data by 75 percent, a data reduction ratio of 4:1. This space reduction is a combination of several different technologies, such as deduplication, compression, and compaction, which provide additional reduction to the basic features provided by ONTAP.

Compaction, which is introduced in ONTAP 9, is the latest patented storage efficiency technology released by NetApp. In the NetApp WAFL file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on the disk-to-save space. This process is illustrated in [Figure 16](#).

Storage Efficiency Features

The storage efficiency features are as follows:

- Deduplication

Deduplication reduces the amount of physical storage required for a volume (or all the volumes in an AFF aggregate) by discarding duplicate blocks and replacing them with references to a single shared block. Reads of deduplicated data typically incur no performance charge. Writes incur a negligible charge except on overloaded nodes.

As data is written during normal use, WAFL uses a batch process to create a catalog of block signatures. After deduplication starts, ONTAP compares the signatures in the catalog to identify duplicate blocks. If a match exists, a byte-by-byte comparison is done to verify that the candidate blocks have not changed since the catalog was created. Only if all the bytes match is the duplicate block discarded and its disk space reclaimed.

- Compression

Compression reduces the amount of physical storage required for a volume by combining data blocks in compression groups, each of which is stored as a single block. Reads of compressed data are faster than in traditional compression methods because ONTAP decompresses only the compression groups that contain the requested data, not an entire file or LUN.

You can perform inline or postprocess compression, separately or in combination:

- Inline compression compresses data in memory before it is written to disk, significantly reducing the amount of write I/O to a volume, but potentially degrading write performance. Performance-intensive operations are deferred until the next postprocess compression operation, if any.
- Postprocess compression compresses data after it is written to disk, on the same schedule as deduplication.
- Compaction

- Small files or I/O padded with zeros are stored in a 4 KB block whether or not they require 4 KB of physical storage. Inline data compaction combines data chunks that would ordinarily consume multiple 4 KB blocks into a single 4 KB block on disk. Compaction takes place while data is still in memory, so it is best suited to faster controllers

Figure 15. Storage Efficiency Deduplication & Compression Features

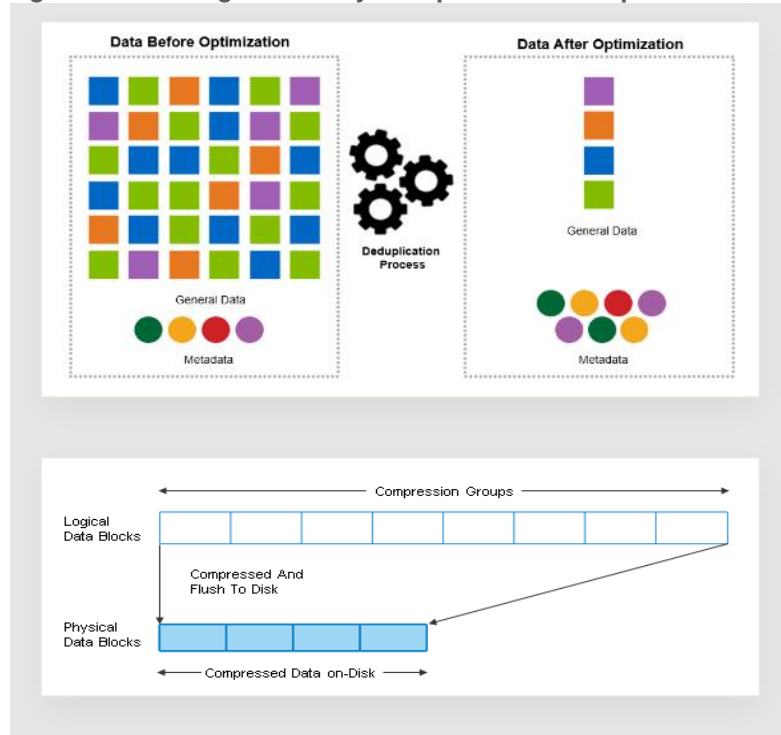
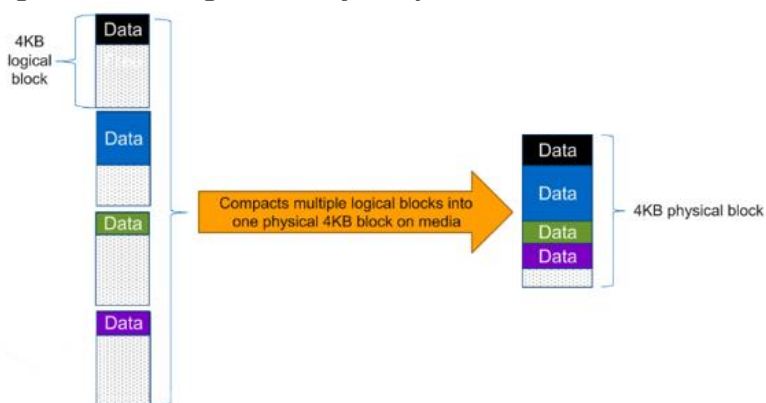


Figure 16. Storage Efficiency Compaction features



Note: Some applications such as Oracle and SQL have unique headers in each of their data blocks that prevent the blocks to be identified as duplicates. So, for such applications, enabling deduplication does not result in significant savings. So, deduplication is not recommended to be enabled for databases. However, NetApp data compression works very well with databases, and we strongly recommend enabling compression for databases. [Table 2](#) lists some guidelines where compression, deduplication and/or inline Zero block deduplication can be used. These are guidelines, not rules; environment may have different performance requirements and specific use cases.

Table 2. Compression and Deduplication Guidelines

Workload	Storage Efficiency Guidelines		
	All Flash FAS (AFF)	Flash Pool (Sized as per Flash Pool Best Practice)	Hard Disk Drives
Database (Oracle, SQL)	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Inline zero-block deduplication Inline deduplication (Data ONTAP 8.3.2 and above) 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Inline zero-block deduplication Inline deduplication (Data ONTAP 8.3.2 and above) 	For primary workloads, use: <ul style="list-style-type: none"> Inline zero-block deduplication For secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Adaptive background compression Inline zero-block deduplication
VDI and SVI	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication Inline deduplication (Data ONTAP 8.3.2 and above) 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication Inline deduplication (Data ONTAP 8.3.2 and above) 	For primary workloads, use: <ul style="list-style-type: none"> Deduplication Inline zero-block deduplication For secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Adaptive background compression Deduplication Inline zero-block deduplication
Exchange	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Set schedule to off peak hours Inline zero-block 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Inline secondary compression Background secondary compression Deduplication

Workload	Storage Efficiency Guidelines		
		deduplication	<ul style="list-style-type: none"> Inline zero-block deduplication
File Services	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Adaptive background compression Deduplication Inline zero-block deduplication
Mixed Workload	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary and secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Deduplication Inline zero-block deduplication 	For primary workloads, use: <ul style="list-style-type: none"> Deduplication Inline zero-block deduplication For secondary workloads, use: <ul style="list-style-type: none"> Adaptive inline compression Adaptive background compression Deduplication Inline zero-block deduplication

Space Savings

[Table 3](#) lists the storage efficiency data reduction ratio ranges for different applications. A combination of synthetic datasets and real-world datasets has been used to determine the typical savings ratio range. The savings ratio range mentioned is only indicative.

Table 3. Typical Savings Ratios with ONTAP 9—Sample Savings Achieved with Internal and Customer Testing


Typical Savings Ratios with ONTAP 9	
Workload (with deduplication, data compaction, adaptive compression and FlexClone volumes, where applicable, technologies)	Ratio Range
Home directories	1.5:1-.2:1
Software development	2:1 - 10:1
VDI VMware Horizon full clone desktops (persistent)	6:1 - 10:1
VDI VMware Horizon instant clone desktops (nonpersistent)	5:1 - 7:1
VDI VMware Horizon Remote Desktop Server Hosted (RDSH) sessions Instant clone desktops (nonpersistent)	6:1 - 10:1
Virtual Servers (OS and Applications)	2:1-.4:1
Oracle databases (with no database compression)	2.1 - 4:1
SQL 2014 databases (with no database compression)	2.1 - 4:1
Microsoft Exchange	1.6:1
Mongo DB	1.3:1 - 1.5:1
Recompressed data (such as video and image files, audio files, pdfs, and so on)	No Savings

NetApp Storage Virtual Machine (SVM)

An SVM is a logical abstraction that represents the set of physical resources of the cluster. This adds extra security and peace of mind to your VDI environment, giving you another place besides vCenter to apply HA, High Availability. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and may reside on any node in the cluster to which the SVM has been given access. An SVM may own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node to another. For example, a flexible volume can be non-disruptively moved to a new node and aggregate, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware, and it is not tied to any specific physical hardware.

An SVM can support multiple data protocols concurrently. Volumes within the SVM can be joined together to form a single NAS namespace, which makes all of an SVM's data available through a single share or mount point to create a VMware NFS datastore for your VDI desktop folders. SVMs also support block-based protocols, and LUNs can be created and exported by using iSCSI, FC, or FCoE. Any or all of these data protocols can be configured for use within a given SVM to support your VDI needs.

Figure 17. NetApp Storage Virtual Machine



Service providers use SVMs in multitenant environments to isolate tenant data and simplify chargeback.

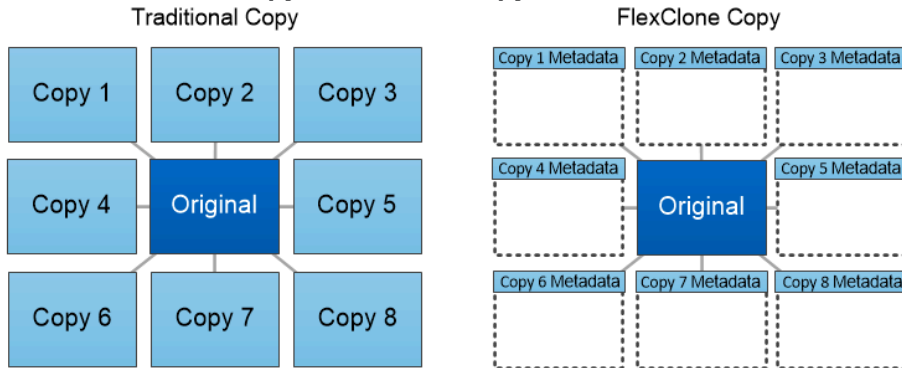
FlexClones

FlexClone technology references Snapshot metadata to create writable, point-in-time copies of a volume. Copies share data blocks with their parents, consuming no storage except what is required for metadata until changes are written to the copy. FlexClone files and FlexClone LUNs use identical technology, except that a backing Snapshot copy is not required.

Where traditional copies can take minutes or even hours to create, FlexClone software lets you copy even the largest datasets almost instantaneously. That makes it ideal for situations in which you need multiple copies of identical datasets (a virtual desktop deployment, for example) or temporary copies of a dataset (testing an application against a production dataset).

You can clone an existing FlexClone volume, clone a volume containing LUN clones, or clone mirror and vault data. You can split a FlexClone volume from its parent, in which case the copy is allocated its own storage.

Figure 18. Traditional Copy vs. FlexClone Copy




FlexClone copies share data blocks with their parents, consuming no storage except what is required for metadata.

SAN Boot

NetApp recommends implementing SAN boot for Cisco UCS servers in the FlexPod Datacenter solution. Doing so enables the ESXI host to be safely secured by the NetApp All Flash FAS storage system, providing better performance. In this design, the FC SAN boot is validated.

In FC SAN boot, each Cisco UCS server boots by connecting the NetApp All Flash FAS storage to the Cisco MDS switch. The 32G FC storage ports, in this example 1a and 1b, are connected to Cisco MDS switch. The FC LIFs are created on the physical ports and each FC LIF is uniquely identified by its target WWPN. The storage system target WWPNs can be zoned with the server initiator WWPNs in the Cisco MDS switches. The FC boot LUN is exposed to the servers through the FC LIF using the MDS switch; this enables only the authorized server to have access to the boot LUN.

Figure 19. FC - SVM ports and LIF layout




Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead if a network interface becomes unavailable, the ESXI host chooses a new optimized path to an available network interface. ALUA is a standard supported by NetApp used to provide information about SCSI targets, which allows a host to identify the best path to the storage.

FlexGroups

ONTAP 9.3 brought an innovation in scale-out NAS file systems: NetApp FlexGroup volumes, which plays a major role to give ONTAP the ability to be scaled nondisruptively out to 24 storage nodes while not degrading the performance of the VDI infrastructure.

With FlexGroup volumes, a storage administrator can easily provision a massive single namespace in a matter of seconds. FlexGroup volumes have virtually no capacity or file count constraints outside of the physical limits of hardware or the total volume limits of ONTAP. Limits are determined by the overall number of constituent member volumes that work in collaboration to dynamically balance load and space allocation evenly across all members. There is no required maintenance or management overhead with a FlexGroup volume. You simply create the FlexGroup volume and share it with your NAS clients. NetApp ONTAP does the rest.

Figure 20. NetApp FlexGroups



Storage QoS

Storage QoS (Quality of Service) can help you manage risks around meeting your performance objectives. You use Storage QoS to limit the throughput to workloads and to monitor workload performance. You can reactively limit workloads to address performance problems and you can pro-actively limit workloads to prevent performance problems.


A workload represents the input/output (I/O) operations to one of the following kinds of storage objects:

- FlexVol volumes
- LUNs

You assign a storage object to a policy group to control and monitor a workload. You can monitor workloads without controlling them.

[Figure 21](#) illustrates an example environment before and after using Storage QoS. On the left, workloads compete for cluster resources to transmit I/O. These workloads get "best effort" performance, which means you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right are the same workloads assigned to policy groups. The policy groups enforce a maximum throughput limit.

Figure 21. Before and After using Storage QoS



NetApp storage quality of service (QoS) works with both SAN and NAS storage, and it runs across the entire NetApp product line from entry to enterprise. Storage QoS offers significant benefits for all types of VDI environments. It lets you:

- Achieve greater levels of consolidation
- Set maximum and minimum limits on multiple VDI workloads that require separate service level agreements (SLAs)
- Add additional workloads with less risk of interference
- Make sure your customers get what they pay for, but not more

Adaptive QoS

Adaptive QoS automatically scales the policy group (A *policy group* defines the throughput ceiling for one or more workloads) value to workload (A *workload* represents the I/O operations for a storage object: a volume, file, qtree or LUN, or all the volumes, files, qtrees, or LUNs in an SVM) size, for the size of the workload changes. That is a significant advantage when you are managing hundreds or thousands of workloads in a VDI deployment. With Adaptive QoS, Ceiling and Floor limit can be set using allocated or used space. The QoS also address HA and Scaling as it will assist in both efforts to produce a non-disruptive change during VDI growth by

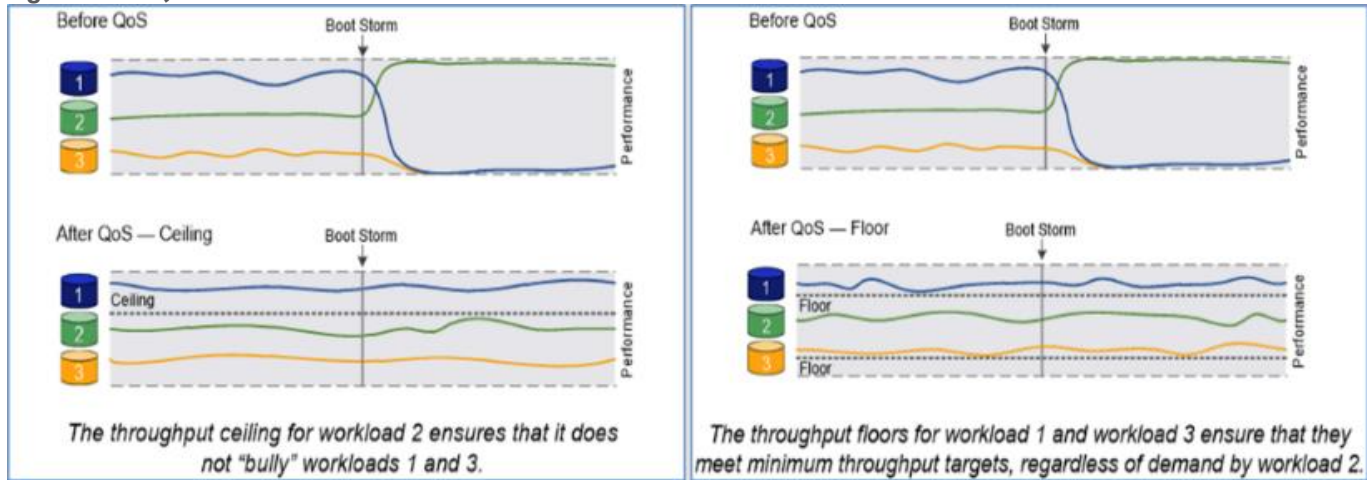
maintaining the ratio of IOPS to TBs/GBs. To assist in managing your QoS, Active IQ unified manager will provide QoS suggestions based on historical performance and usage.

Three default adaptive QoS policy groups are available, as listed in [Table 4](#). You can apply these policy groups directly to a volume.

Table 4. Available Default Adaptive QoS Policy Groups

Default Policy Group	Expected IOPS/TB	Peak IOPS/TB	Absolute Min IOPS
extreme	6,144	12,288	500
performance	2,048	4,096	500
Value	128	512	75

Figure 22. QOS boot storm illustration



Security and Data Protection

Vscan

With Vscan you can use integrated antivirus functionality on NetApp storage systems to protect data from being compromised by viruses or other malicious code. NetApp virus scanning, called Vscan, combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when.


Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors. The ONTAP Antivirus Connector, provided by NetApp and installed on the external server, handles communication between the storage system and the antivirus software.

You can use *on-access scanning* to check for viruses when clients open, read, rename, or close files over CIFS. File operation is suspended until the external server reports the scan status of the file. If the file has already been scanned, ONTAP allows the file operation. Otherwise, it requests a scan from the server.

You can use *on-demand scanning* to check files for viruses immediately or on a schedule. You might want to run scans only in off-peak hours, for example. The external server updates the scan status of the checked files, so that file-access latency for those files (assuming they have not been modified) is typically reduced when they are next accessed over CIFS. You can use on-demand scanning for any path in the SVM namespace, even for volumes that are exported only through NFS.

Typically, you enable both scanning modes on an SVM. In either mode, the antivirus software takes remedial action on infected files based on your settings in the software.

Figure 23. Scanning with Vscan




NetApp Volume Encryption(NVE) and NetApp Aggregate Encryption (NAE)

NetApp Volume Encryption is a software-based, data-at-rest encryption solution that is FIPS 140-2 compliant. NVE allows ONTAP to encrypt data for each volume for granularity. NAE, is an outgrowth of NVE; it allows ONTAP to encrypt data for each volume, and the volumes can share keys across the aggregate. NVE and NAE enable you to use storage efficiency features that would be lost with encryption at the application layer. For greater storage efficiency, you can use aggregate deduplication with NAE.

Here’s how the process works: The data leaves the disk encrypted, is sent to RAID, is decrypted by the CryptoMod, and is then sent up the rest of the stack. This process is illustrated in [Figure 24](#).

Figure 24. NVE and NAE Process



To view the latest security features for ONTAP 9, go to: [Security Features in ONTAP 9 | NetApp](#).

ONTAP Rest API

ONTAP Rest API enables you to automate the deployment and administration of your ONTAP storage systems using one of several available options. The ONTAP REST API provides the foundation for all the various ONTAP automation technologies.

Beginning with ONTAP 9.6, ONTAP includes an expansive workflow-driven REST API that you can use to automate deployment and management of your storage. In addition, NetApp provides a Python client library, which makes it easier to write robust code, as well as support for ONTAP automation based on Ansible.

AutoSupport and Active IQ Digital Advisor

ONTAP offers artificial Intelligence-driven system monitoring and reporting through a web portal and through a mobile app. The AutoSupport component of ONTAP sends telemetry that is analyzed by Active IQ Digital Advisor. Active IQ enables you to optimize your data infrastructure across your global hybrid cloud by delivering actionable predictive analytics and proactive support through a cloud-based portal and mobile app. Data-driven insights and recommendations from Active IQ are available to all NetApp customers with an active SupportEdge contract (features vary by product and support tier).

The following are some things you can do with Active IQ:

- Plan upgrades. Active IQ identifies issues in your environment that can be resolved by upgrading to a newer version of ONTAP and the Upgrade Advisor component helps you plan for a successful upgrade.
- View system wellness. Your Active IQ dashboard reports any issues with wellness and helps you correct those issues. Monitor system capacity to make sure you never run out of storage space.
- Manage performance. Active IQ shows system performance over a longer period than you can see in ONTAP System Manager. Identify configuration and system issues that are impacting your performance.
- Maximize efficiency. View storage efficiency metrics and identify ways to store more data in less space.
- View inventory and configuration. Active IQ displays complete inventory and software and hardware configuration information. View when service contracts are expiring to ensure you remain covered.

Security and Ransomware Protection

Ransomware is a type of malware that threatens to publish the victim's personal data or permanently block access to it unless a ransom is paid off. A ransomware attack can have direct and indirect impacts so it's important to detect it as early as possible so that you can prevent its spread and avoid costly downtime.

NetApp storage administrators use local or remote login accounts to authenticate themselves to the cluster and storage VM. Role-Based Access Control (RBAC) determines the commands to which an administrator has access. In addition to RBAC, NetApp ONTAP supports multi-factor authentication (MFA) and multi-admin verification (MAV) to enhance the security of the storage system.

With NetApp ONTAP, you can use the security login create command to enhance security by requiring that administrators log in to an admin or data SVM with both an SSH public key and a user password. Beginning with ONTAP 9.12.1, you can use Yubikey hardware authentication devices for SSH client MFA using the FIDO2 (Fast Identity Online) or Personal Identity Verification (PIV) authentication standards.

With ONTAP 9.11.1, you can use multi-admin verification (MAV) to ensure that certain operations, such as deleting volumes or Snapshot copies, can be executed only after approvals from designated administrators. This prevents compromised, malicious, or inexperienced administrators from making undesirable changes or deleting data.

Also, with ONTAP 9.10.1, the Autonomous Ransomware Protection (ARP) feature uses workload analysis in NAS (NFS and SMB) environments to proactively detect and warn about abnormal activity that might indicate a ransomware attack. While NetApp ONTAP includes features like FPolicy, Snapshot copies, SnapLock, and Active IQ Digital Advisor to help protect from ransomware, ARP utilizes machine-learning and simplifies the detection of and the recovery from a ransomware attack. ARP can detect the spread of most ransomware

attacks after only a small number of files are encrypted, take action automatically to protect data, and alert users that a suspected attack is happening. When an attack is suspected, the system takes a volume Snapshot copy at that point in time and locks that copy. If the attack is confirmed later, the volume can be restored to this proactively taken snapshot to minimize the data loss.

For more information on MFA, MAV, and ransomware protection features, refer to the following:

- <https://docs.netapp.com/us-en/ontap/authentication/setup-ssh-multifactor-authentication-task.html>
- <https://www.netapp.com/pdf.html?item=/media/17055-tr4647pdf.pdf>
- <https://docs.netapp.com/us-en/ontap/multi-admin-verify/>
- <https://docs.netapp.com/us-en/ontap/anti-ransomware/>

VMware vSphere 8.0 U1

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructures (resources including CPUs, storage, and networking) as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

VMware vSphere 8.0 U1 has several improvements and simplifications including, but not limited to:

- Fully featured vSphere Client (HTML5) client. The flash-based vSphere Web Client has been deprecated and is no longer available.
- Improved Distributed Resource Scheduler (DRS) – a very different approach that results in a much more granular optimization of resources.
- Assignable hardware – a new framework that was developed to extend support for vSphere features when you utilize hardware accelerators.
- vSphere Lifecycle Manager – a replacement for VMware Update Manager, bringing a suite of capabilities to make lifecycle operations better.
- Refactored vMotion – improved to support today’s workloads

For more information about VMware vSphere and its components, see:

<https://www.vmware.com/products/vsphere.html>.

VMware vSphere vCenter

VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest OS, and other critical components of a virtual infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

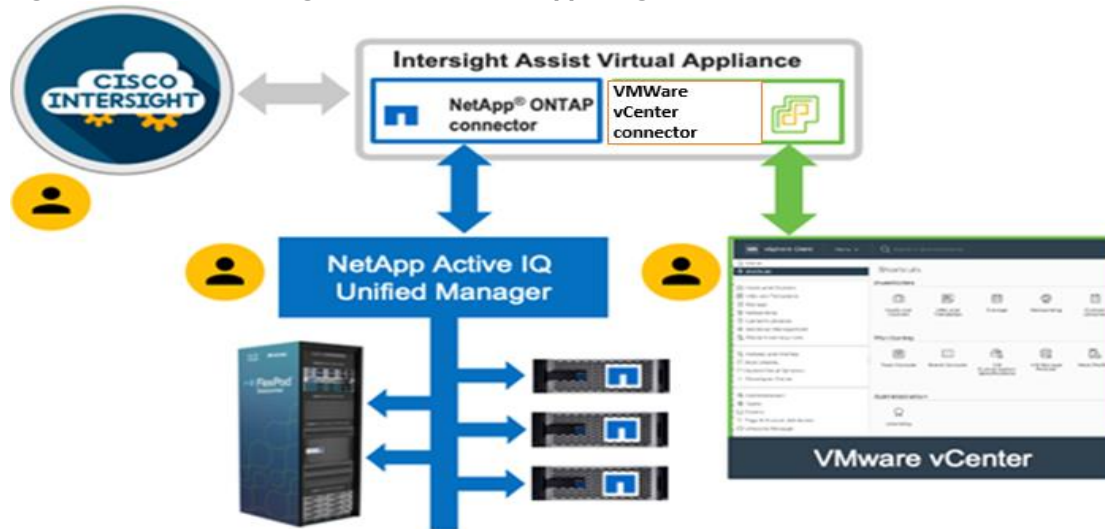
Cisco Intersight Assist Device Connector for VMware vCenter and NetApp ONTAP

Cisco Intersight integrates with VMware vCenter and NetApp storage as follows:

- Cisco Intersight uses the device connector running within Cisco Intersight Assist virtual appliance to communicate with the VMware vCenter.

- Cisco Intersight uses the device connector running within a Cisco Intersight Assist virtual appliance to integrate with NetApp Active IQ Unified Manager. The NetApp AFF A400 should be added to NetApp Active IQ Unified Manager.

Figure 25. Cisco Intersight and vCenter/NetApp Integration



The device connector provides a secure way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware hypervisor and ONTAP data storage environments.

Enterprise SAN and NAS workloads can benefit equally from the integrated management solution. The integration architecture enables FlexPod customers to use new management capabilities with no compromise in their existing VMware or ONTAP operations. IT users will be able to manage heterogeneous infrastructure from a centralized Cisco Intersight portal. At the same time, the IT staff can continue to use VMware vCenter and NetApp Active IQ Unified Manager for comprehensive analysis, diagnostics, and reporting of virtual and storage environments. The functionality provided through this integration is covered in the upcoming solution design section.

NetApp ONTAP Tools for VMware vSphere

NetApp ONTAP tools for VMware vSphere is a unified appliance that includes vSphere Storage Console (VSC), VASA Provider and SRA Provider. This vCenter web client plug-in that provides Context sensitive menu to provision traditional datastores and Virtual Volume (vVol) datastore.

ONTAP tools provides visibility into the NetApp storage environment from within the vSphere web client. VMware administrators can easily perform tasks that improve both server and storage efficiency while still using role-based access control to define the operations that administrators can perform. It includes enhanced REST APIs that provide vVols metrics for SAN storage systems using ONTAP 9.7 and later. So, NetApp OnCommand API Services is no longer required to get metrics for ONTAP systems 9.7 and later.

To download ONTAP tools for VMware vSphere, go to:

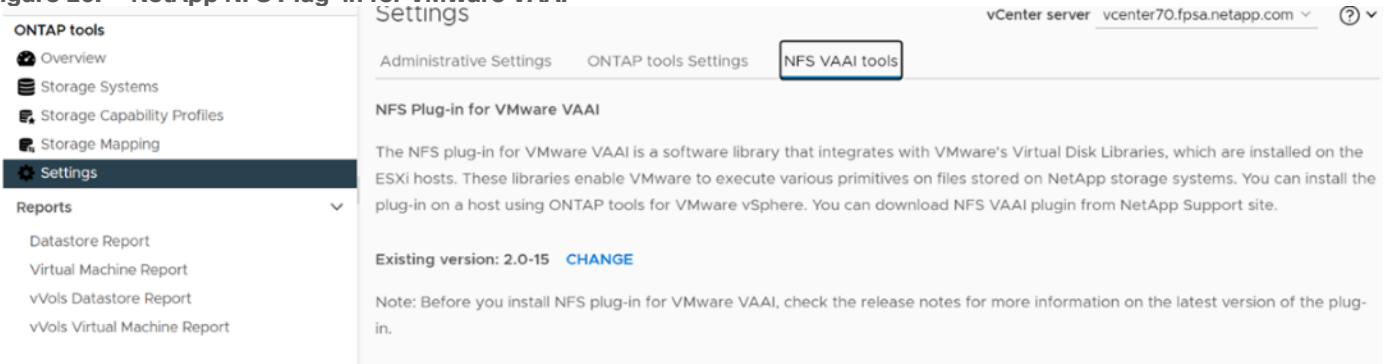
<https://mysupport.netapp.com/site/products/all/details/otv/downloads-tab>.

NetApp NFS Plug-in for VMware VAAI

The NetApp NFS Plug-in for VMware vStorage APIs - Array Integration (VAAI) is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. The VMware VAAI package

enables the offloading of certain tasks from the physical hosts to the storage array. Performing those tasks at the array level can reduce the workload on the ESXi hosts.

Figure 26. NetApp NFS Plug-in for VMware VAAI



The copy offload feature and space reservation feature improve the performance of VSC operations. The NetApp NFS Plug-in for VAAI is not shipped with VSC, but you can install it by using VSC. You can download the plug-in installation package and obtain the instructions for installing the plug-in from the NetApp Support Site.

For more information about the NetApp VSC for VMware vSphere, see the [NetApp Virtual Infrastructure Management Product Page](#).

NetApp XCP File Analytics

NetApp XCP file analytics is host-based software to scan the file shares, collect and analyzes the data and provide insights into the file system. NetApp XCP file analytics works for both NetApp and non-NetApp systems and runs on Linux or Windows host. For more information, go to: <http://docs.netapp.com/us-en/xcp/index.html>

Architecture and Design Considerations for Desktop Virtualization

This chapter contains the following:

- [Understanding Applications and Data](#)
- [Project Planning and Solution Sizing Sample Questions](#)
- [Hypervisor Selection](#)
- [Storage Considerations](#)
- [VMware Horizon Design Fundamentals](#)
- [Desktop Virtualization Design Fundamentals](#)
- [VMware Horizon Design Fundamentals](#)

There are many reasons to consider a virtual desktop solution such as an ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Device (BYOD) to work programs. The first step in designing a virtual desktop solution is to understand the user community and the type of tasks that are required to successfully execute their role. The following user classifications are provided:

- Knowledge Workers today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- External Contractors are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.
- Task Workers perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.
- Mobile Workers need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- Shared Workstation users are often found in state-of-the-art University and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications for the needs of the organization change, tops the list.

After the user classifications have been identified and the business requirements for each user classification have been defined, it becomes essential to evaluate the types of virtual desktops that are needed based on user requirements. There are essentially five potential desktops environments for each user:

- Traditional PC: A traditional PC is what typically constitutes a desktop environment: a physical device with a locally installed operating system.
- Hosted Shared Desktop: A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2019, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Remoted Desktop Server Hosted

Server sessions: A hosted virtual desktop is a virtual desktop running on a virtualization layer (ESX). The user does not work with and sit in front of the desktop, but instead, the user interacts through a delivery protocol.

- **Published Applications:** Published applications run entirely on the VMware Horizon Remote Desktop Server Hosted (RDSH) sessions, RDS server virtual machines, and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft Office, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.
- **Streamed Applications:** Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly, but the resources may only be available while they are connected to the network.
- **Local Virtual Desktop:** A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a primary type and is synchronized with the data center when the device is connected to the network.

For the purposes of the validation represented in this document, both VMware Horizon Remote Desktop Server Hosted (RDSH) sessions and Windows 11 Virtual Desktops sessions were validated. Each section provides some fundamental design decisions for this environment.

Understanding Applications and Data

When the desktop user groups and sub-groups have been identified, the next task is to catalog group application and data requirements. This can be one of the most time-consuming processes in the VDI planning exercise but is essential for the VDI project's success. If the applications and data are not identified and co-located, performance will be negatively affected.

The process of analyzing the variety of application and data pairs for an organization will likely be complicated by the inclusion of cloud applications, for example, Salesforce.com. This application and data analysis is beyond the scope of this Cisco Validated Design but should not be omitted from the planning process. There are a variety of third-party tools available to assist organizations with this crucial exercise.

Project Planning and Solution Sizing Sample Questions

Now that user groups, their applications, and their data requirements are understood, some key project and solution sizing questions may be considered.

General project questions should be addressed at the outset, including:

- Has a VDI pilot plan been created based on the business analysis of the desktop groups, applications, and data?
- Is there infrastructure and budget in place to run the pilot program?
- Are the required skill sets to execute the VDI project available? Can we hire or contract for them?
- Do we have end user experience performance metrics identified for each desktop sub-group?
- How will we measure success or failure?
- What is the future implication of success or failure?

Below is a short, non-exhaustive list of sizing questions that should be addressed for each user sub-group:

- What is the desktop OS planned? Windows 10 or Windows 11?
- 32 bit or 64 bit desktop OS?

- How many virtual desktops will be deployed in the pilot? In production? All Windows 11?
- How much memory per target desktop group?
- Are there any rich media or graphics-intensive workloads?
- Are there any applications installed? What application delivery methods will be used, Installed, Streamed, Layered, Hosted, or Local?
- What is the OS planned for RDS Server Roles? Windows Server 2019 or Server 2022?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- What is the SQL server version for the database? SQL server 2017 or 2019?
- Is user profile management (for example, non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Hypervisor Selection

VMware vSphere has been identified for the hypervisor for both VMware Horizon Remoted Server Desktop Hosted (RDSH) sessions and Windows 11 Virtual Desktops.

VMware vSphere comprises the management infrastructure or virtual center server software and the hypervisor software that virtualizes the hardware resources on the servers. It offers features like Distributed Resource Scheduler, vMotion, high availability, Storage vMotion, VMFS, and a multi-pathing storage layer. More information on VMware vSphere can be found here:

<http://www.vmware.com/products/datacentervirtualization/vsphere/overview.html>.

Note: For this CVD, the hypervisor used was VMware ESXi 8.0 U1.

Storage Considerations

Boot from SAN

When utilizing Cisco UCS Server technology, it is recommended to configure Boot from SAN and store the boot partitions on remote storage, this enabled architects and administrators to take full advantage of the stateless nature of service profiles for hardware flexibility across lifecycle management of server hardware generational changes, Operating Systems/Hypervisors, and overall portability of server identity. Boot from SAN also removes the need to populate local server storage creating more administrative overhead.

NetApp AFF Storage Considerations

Note: Make sure each NetApp AFF Controller is connected to BOTH storage fabrics (A/B).

Within NetApp, the best practice is to map Hosts to iGroups and then iGroups to LUNs, this ensures same LUN ID to all hosts and allows for simplified management of ESXi Clusters across multiple nodes.

Port Connectivity

10/25/40/100 Gbe connectivity support – while both 10 and 25 Gbe is provided through 2 onboard NICs on each AFF controller, if more interfaces are required or if 40Gbe connectivity is also required, then make sure to provision for additional NICs have been included in the original AFF BOM.

16/32Gb Fiber Channel supports the NetApp Storage up to 32Gb FC support on the latest AFF A400 series arrays. Always make sure the correct number of HBAs and the speed of SFPs are included in the original AFFBOM.

Oversubscription

To reduce the impact of an outage or maintenance scheduled downtime it is good practice when designing fabrics to provide oversubscription of bandwidth, this enables a similar performance profile during component failure and protects workloads from being impacted by a reduced number of paths during a component failure or maintenance event. Oversubscription can be achieved by increasing the number of physically cabled connections between storage and compute. These connections can be utilized to deliver performance and reduced latency to the underlying workloads running on the solution.

Topology

When configuring your SAN, it's important to remember that the more hops you have, the more latency you will see. For best performance, the ideal topology is a "Flat Fabric" where the AFF is only one hop away from any applications being hosted on it.

VMware Virtual Volumes Considerations

vCenters that are in Enhanced Linked Mode will each be able to communicate with the same AFF, however vCenters that are not in Enhanced Linked Mode must use CA-Signed Certificates using the same AFF. If multiple vCenters need to use the same AFF for vVols, they should be configured in Enhanced Linked Mode.

There are some AFF limits on Volume Connections per Host, Volume Count, and Snapshot Count. For more information about NetApp AFF limits review the following: <https://hww.NetApp.com/Controller/Index>

When a Storage Policy is applied to a vVol VM, the volumes associated with that VM are added to the designated protection group when applying the policy to the VM. If replication is part of the policy, be mindful of the amount of VMs using that storage policy and replication group. A large amount of VMs with a high change rate could cause replication to miss its schedule due to increased replication bandwidth and time needed to complete the scheduled snapshot. NetApp Storage recommends vVol VMs that have Storage Policies applied be balanced between protection groups.

Desktop Virtualization Design Fundamentals

An ever-growing and diverse base of user devices, complexity in the management of traditional desktops, security, and even Bring Your Own (BYO) device to work programs are prime reasons for moving to a virtual desktop solution.

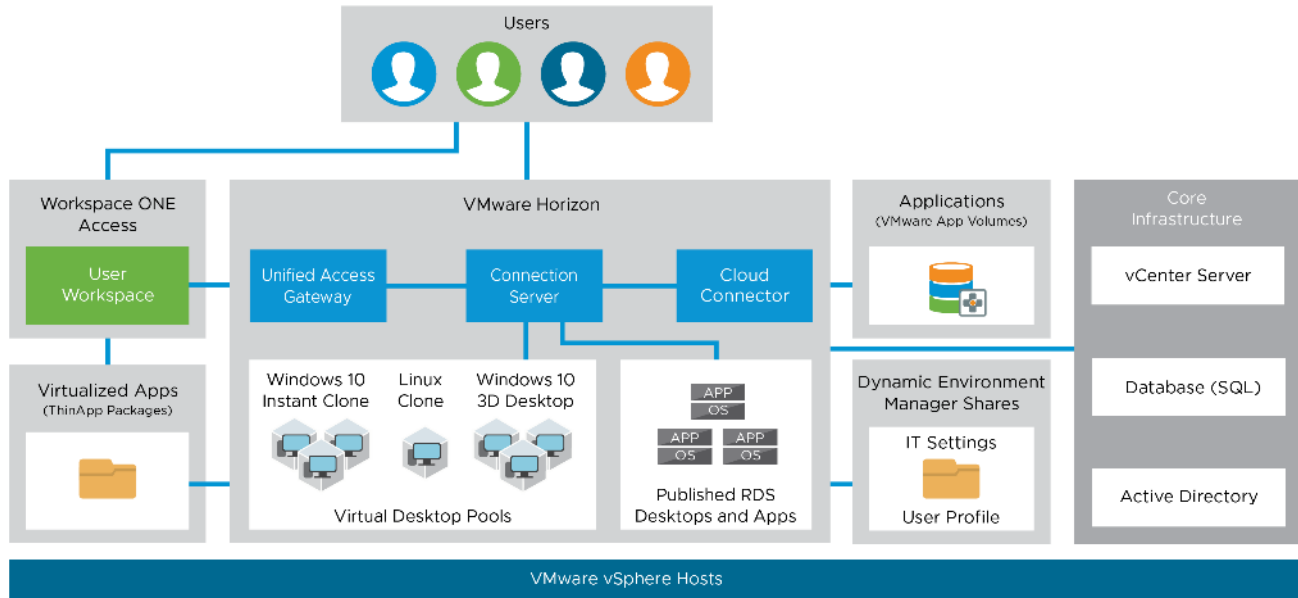
VMware Horizon Design Fundamentals

VMware Horizon 8 integrates VDI desktop virtualization technologies into a unified architecture that enables scalable, simple, efficient, and manageable solutions for delivering Windows applications and desktops a service.

VMware Horizon 8 delivers a native touch-optimized experience via PCoIP or Blast Extreme high-definition performance, even over mobile networks. Users can select applications from an easy-to-use “store” accessible from tablets, smartphones, PCs, Macs, and thin clients.

Several components must be deployed to create a functioning Horizon environment to deliver the VDI. These components refer to as “core infrastructure” and encompass: Domain Controllers, DNS, DHCP, User Profile managers, SQL, vCenters, VMware Horizon View Connection Servers.

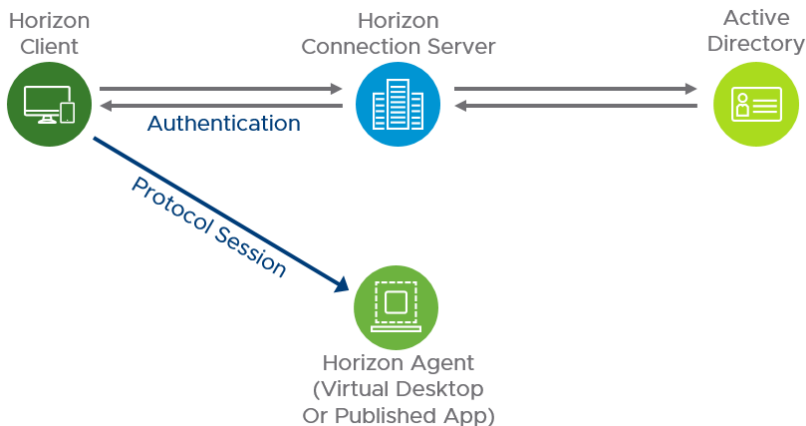
Figure 27. VMware Horizon Design Overview



Horizon VDI Pool and RDSH Servers Pool

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Desktop Pool. The VM provisioning relies on VMware Horizon Connection Server, vCenter Server, and AD components. The Horizon Client then forms a session using PCoIP, Blast, or RDP protocols to a Horizon Agent running on a virtual desktop, RDSH server, or physical computer. In this CVD, virtual machines in the Desktop Pools are configured to run a Windows Server 2019 OS (for RDS Hosted shared sessions using RDP protocol) and a Windows 11 Desktop OS (for pooled VDI desktops using Blast protocol).

Figure 28. Horizon VDI and RDSH Desktop Delivery Based on Display Protocol (PcoIP/Blast/RDP)

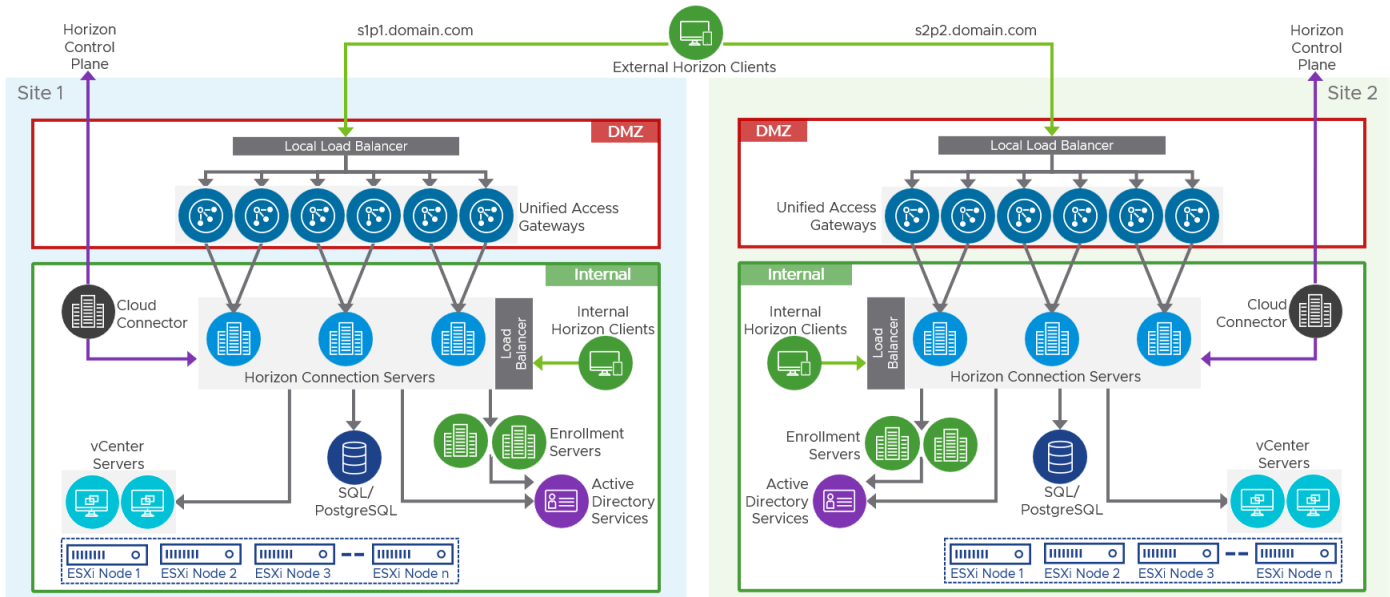


Multiple-Site Configuration

If you have multiple regional sites, you can use any of the Load Balances Tools to direct the user connections to the most appropriate site to deliver the desktops and applications to users.

Figure 29 illustrates the logical architecture of the Horizon multisite deployment. Such architecture eliminates any single point of failure that can cause an outage for desktop users.

Figure 29. Multisite Configuration Overview



Based on the requirement and the number of data centers or remote locations, we can choose any available load-balancing software to increase security and optimize the user experience.

Note: Multisite configuration is shown as the example and was not used in this CVD testing.

Designing a Virtual Desktop Environment for Different Workloads

With VMware Horizon, the method you choose to provide applications or desktops to users depends on the types of applications and desktops you are hosting and available system resources, as well as the types of users and user experience you want to provide.

Table 5. Desktop type and user experience

Desktop Type	User Experience
Server OS machines	<p>You want: Inexpensive server-based delivery to minimize the cost of delivering applications to a large number of users, while providing a secure, high-definition user experience.</p> <p>Your users: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations.</p> <p>Application types: Any application.</p>
Desktop OS machines	<p>You want: A client-based application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition.</p>

Desktop Type	User Experience
	<p>Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications.</p> <p>Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines.</p> <p>Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.</p>
Remote PC Access	<p>You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the data center.</p> <p>Your users: Employees or contractors that have the option to work from home but need access to specific software or data on their corporate desktops to perform their jobs remotely.</p> <p>Host: The same as Desktop OS machines.</p> <p>Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.</p>

For the Cisco Validated Design described in this document, a Remote Desktop Server Hosted sessions (RDSH) using RDS based Server OS and VMware Horizon pooled Instant and Full Clone Virtual Machine Desktops using VDI based desktop OS machines were configured and tested.

Deployment Hardware and Software

This chapter contains the following:

- [Architecture](#)
- [Products Deployed](#)
- [Physical Topology](#)
- [Logical Architecture](#)
- [Configuration Guidelines](#)

Architecture

The architecture deployed is highly modular. While each environment might vary in its exact configuration, the reference architecture contained in this document once built, can easily be scaled as requirements, and demands change. This includes scaling both up (adding additional resources within a Cisco UCS Domain) and out (adding additional Cisco UCS Domains and NetApp storage).

The FlexPod Datacenter solution includes Cisco networking, Cisco UCS and NetApp AFF A400, which efficiently fit into a single data center rack, including the access layer network switches.

Products Deployed

This CVD details the deployment of up to 2500 Multi-session OS, 2000 Single-session OS VDI users featuring the following software:

- VMware vSphere ESXi 8.0 U1 Hypervisor
- Microsoft SQL Server 2019
- Microsoft Windows Server 2019 and Windows 11 64-bit virtual machine Operating Systems
- VMware Horizon 8 2212
- Microsoft Office 2021
- VMware Horizon 8 2212
- FSLogix for User profile management
- Cisco Intersight platform to deploy, maintain, and support the UCS components
- Cisco Intersight Assist virtual appliance to help connect the VMware vCenter with the Cisco Intersight platform
- NetApp ONTAP Tools for VMware vSphere 9.12
- NetApp ONTAP 9.13.1P3

FlexPod with Cisco UCS M7 servers, VMware Horizon Remote Desktop Services Hosted (RDSH) sessions and Windows 11 virtual desktops on vSphere 8.0 U1 delivers a Virtual Desktop Infrastructure that is redundant, using the best practices of Cisco and NetApp Storage. The solution includes VMware vSphere 8.0 U1 hypervisor installed on the Cisco UCS M7 Compute Node Server configured for stateless compute design using boot from SAN. NetApp Storage AFF A400 provides the storage infrastructure required for setting up the VDI workload. Cisco Intersight is utilized to configure and manage the Cisco UCS infrastructure with Cisco Intersight providing lifecycle management capabilities. The solution requirements and design details are covered in this section.

Physical Topology

FlexPod VDI with Cisco UCS X210c M7 servers is an NFS-based storage access design. NetApp Storage AFF and Cisco UCS are connected through Cisco Nexus switches and storage access utilizes the NFS network. For VDI IP based file share storage access NetApp Storage AFF Cisco UCS are connected through Cisco Nexus N9K-C93180YC-FX3S switches. The physical connectivity details are explained below.

Figure 30. FlexPod VDI - Physical Topology




Figure 30 details the physical hardware and cabling deployed to enable this solution:

- 2 Cisco Nexus N9K-C93180YC-FX3S Switches in NX-OS Mode.
- 2 Cisco MDS 9132T Switches for Fiber Channel traffic
- 8 Cisco UCS X210c M7 Compute Nodes with Intel Xeon 6448H 2.40GHz 32-core processors, 2TB 4400MHz RAM, and one Cisco UCS VIC 15231 card, providing N+1 server fault tolerance.
- NetApp AFF A400 Storage System with dual redundant controllers, 2x disk shelves, and 48 x 1.75 TB solid-state NVMe drives providing storage and NVME/NFS/CIFS connectivity.

Note: The common services and LoginVSI Test infrastructure are not a part of the physical topology of this solution.

[Table 6](#) lists the software versions of the primary products installed in the environment.

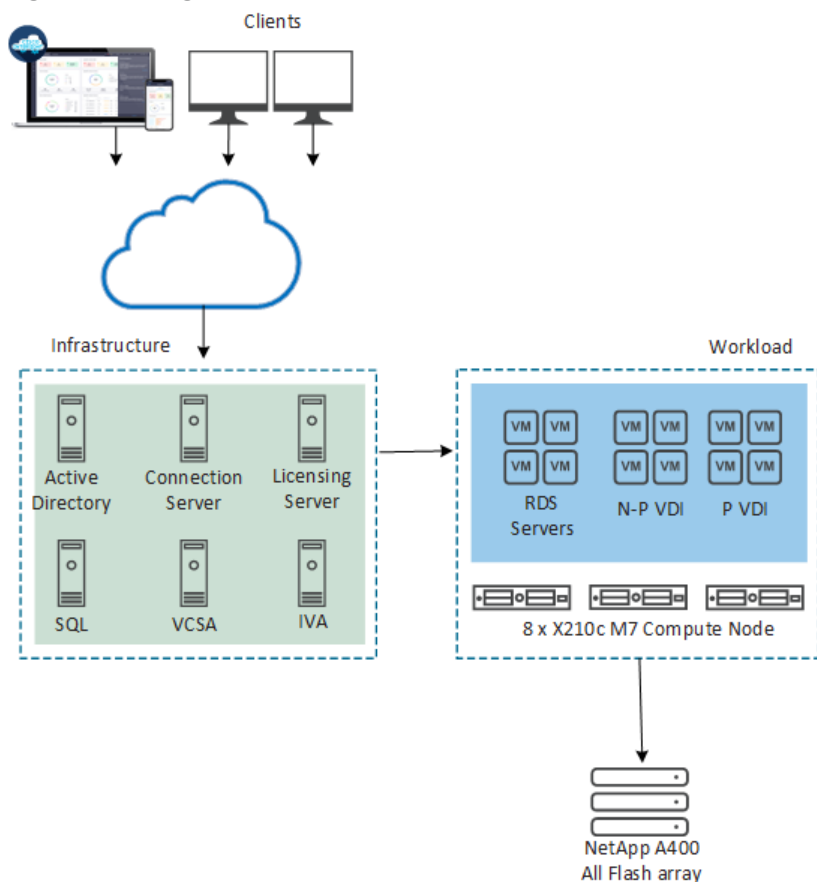
Table 6. Software and Firmware Versions

Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	5.2(0.230041) bundle release
Cisco	Intersight	5.2(0.230041) bundle release
Cisco	UCSX 210C M7 Compute Nodes	5.2(0.230041) bundle release
Cisco	VIC 15231	5.2(0.230041) bundle release
Cisco	Cisco Nexus N9K-C93180YC-FX3S	9.3(7a)
NetApp	AFF A400	ONTAP 9.13.1P3
NetApp	ONTAP Tools for VMWare vSphere	9.12
NetApp	NetApp NFS Plug-in for VMWare VAAI	2.0.1
VMware	vCenter Server Appliance	8.0
VMware	vSphere 8.0 U1	8.0.1, 21813344
VMware	Tools	12.2.0.21223074
VMware	VMware Horizon 8 2212 Connection server	8.8.0-21073894
VMware	VMware Horizon 8 2212 Agent	8.8.0-21067308
Microsoft	FSLogix 2210 hotfix 1	2.9.8440.42104

Logical Architecture

The logical architecture of the validated solution which is designed to support up to 3300 users on a single chassis containing three Cisco UCS X210c M7 Compute Node Servers, with physical redundancy for the Compute Node servers for each workload type is illustrated in [Figure 31](#).

Figure 31. Logical Architecture Overview



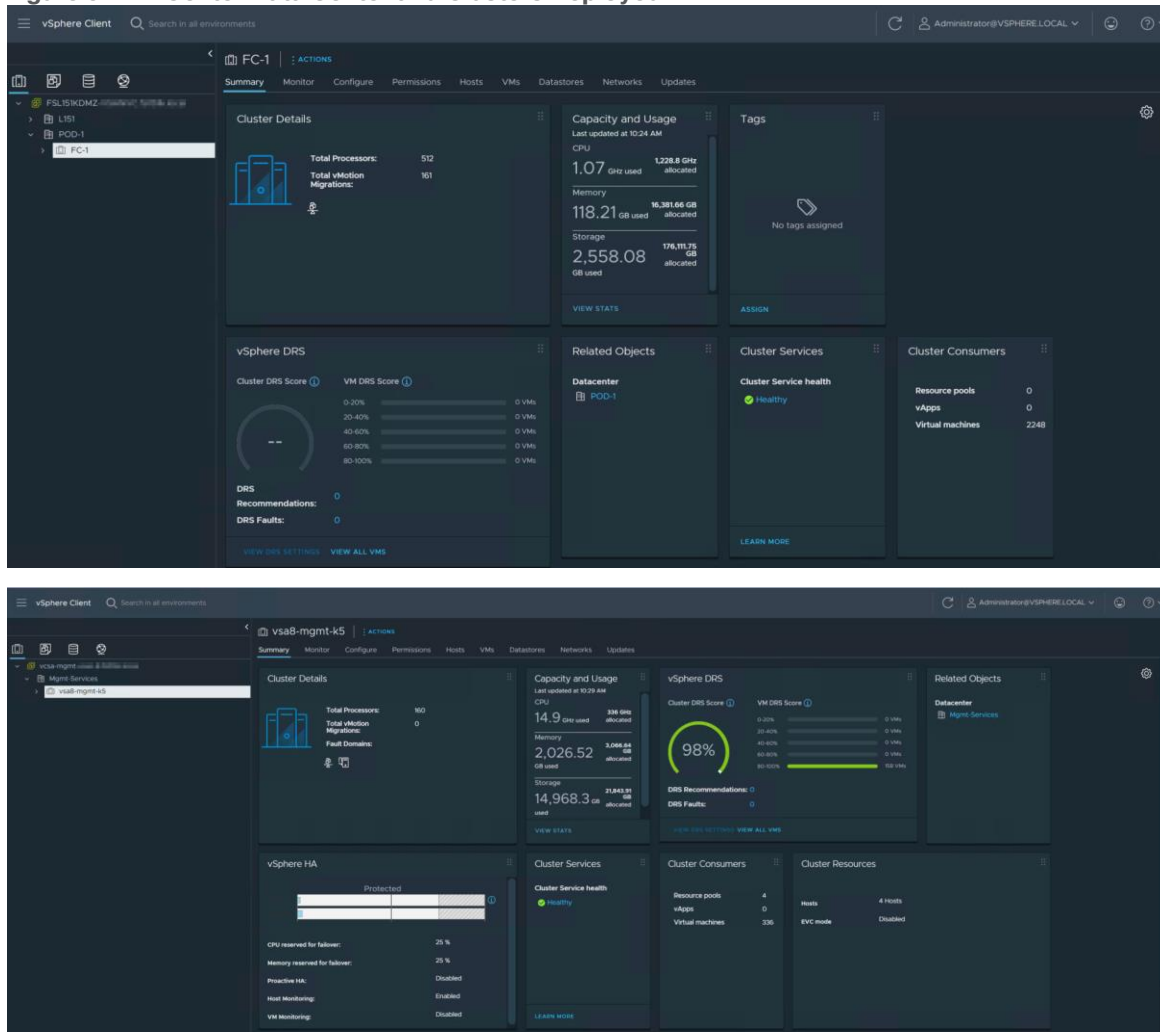
VMware Clusters

Two VMware Clusters in separate vCenter datacenters were utilized to support the solution and testing environment:

- VDI Cluster FlexPod Data Center with Cisco UCS
 - FC-1: VMware Horizon desktop VMs (Windows Server 2019 instant clones, VMware Horizon Windows 11 Instant Clones and Persistent desktops)
- Common Services and Launcher Cluster
 - Vsa8-mgmt: Infrastructure VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, VMware Connection and Replica Servers, and NetApp ONTAP Tools for VMware vSphere, ActiveIQ Unified Manager, VMs, and so on)

Note: The management components and LoginVSI Test infrastructure are hosted on a separate vSphere cluster and are not a part of the physical topology of this solution. The cluster was connected to the FlexPod using the same set of the Nexus switches.

Figure 32. vCenter Data Center and Clusters Deployed



Configuration Guidelines

The VMware Horizon solution described in this document provides details for configuring a fully redundant, highly-available configuration. Configuration guidelines are provided that refer to which redundant component is being configured with each step, whether that be A or B. For example, Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are configured similarly.

This document is intended to allow you to configure the environment for the VMware Horizon 8 2212 workload as a stand-alone solution.

VLANs

The VLAN configuration recommended for the environment includes a total of nine VLANs as listed in [Table 7](#).

Table 7. VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt	30	In-Band management interfaces

VLAN Name	VLAN ID	VLAN Purpose
CIFS-VLAN	32	CIFS Storage access
NFS- VLAN	33	VLAN for NFS storage traffic
VCC/VM-Network	34	RDSH, VDI Persistent and Non-Persistent
vMotion	35	vMotion for DRS and migration
InBand-Mgmt_70	70	Common Services In-Band management interfaces
Infra-Mgmt_71	71	Common Services/Infrastructure Virtual Machines
VCC/VM-Network_54	54	LVSI Launchers
OOB-Mgmt	132	Out-of-Band management interfaces

Solution Configuration

This chapter contains the following:

- [Solution Cabling](#)
- [Network Switch Configuration](#)
- [FlexPod Cisco Nexus Switch Configuration](#)

Solution Cabling

The following sections detail the physical connectivity configuration of the FlexPod VMware VDI environment.

The information provided in this section is a reference for cabling the physical equipment in this Cisco Validated Design environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain the details for the prescribed and supported configuration of the NetApp Storage AFF A400 storage array to the Cisco UCS 6536 Fabric Interconnects through Cisco Nexus Switches.

Note: This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

IMPORTANT! Be sure to follow the cabling directions in this section. Failure to do so will result in unnecessary changes to the deployment since specific port locations are mentioned.

[Figure 33](#) details the cable connections used in the validation lab for FlexPod topology based on the Cisco UCS 6536 fabric interconnect. Also, 40Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the NetApp AFF A400 controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each All Flash Array controller has a connection to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB), In-Band (IB) Management and VM-Network Subnets.

The architecture is divided into three distinct layers:

1. Cisco UCS Compute Platform
2. Network Access layer and LAN
3. Storage Access to the NetApp AFF A400

Figure 33. FlexPod Solution Cabling Diagram

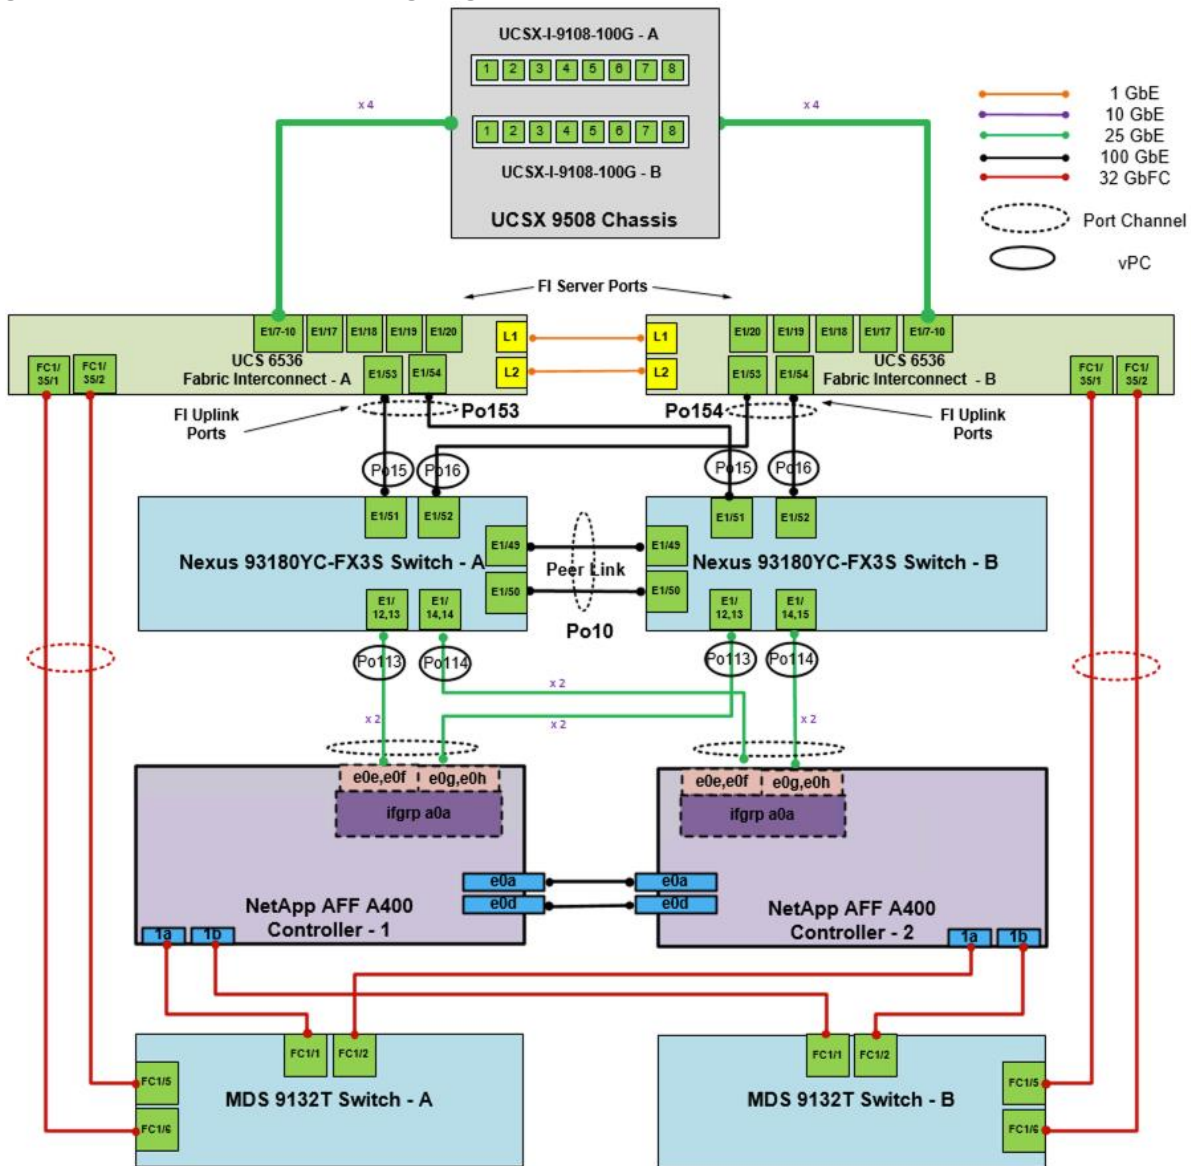


Table 8. Cisco Nexus N9K-C93180YC-FX3S-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus N9K-C93180YC-FX3S A	Eth1/19	25GbE	NetApp Controller 2	e0e Controller 1
	Eth1/20	25GbE	NetApp Controller 2	e0f Controller 1
	Eth1/17	25GbE	NetApp Controller 1	e0g Controller 2
	Eth1/18	25GbE	NetApp Controller 1	e0h Controller 2
	Eth1/49	100GbE	Cisco UCS fabric interconnect A	P0 1/21
	Eth1/50	100GbE	Cisco UCS fabric interconnect A	P0 1/23
	Eth1/51	100GbE	Cisco UCS fabric interconnect B	P0 1/21

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/52	100GbE	Cisco UCS fabric interconnect B	P0 1/23
	Eth1/53	100GbE	Cisco Nexus N9K-C93180YC-FX3S B	Eth1/53
	Eth1/54	100GbE	Cisco Nexus N9K-C93180YC-FX3S B	Eth1/54
	MGMT0	GbE	GbE management switch	Any

Note: For devices requiring GbE connectivity, use the GbE Copper SFP+s (GLC-T=).

Table 9. Cisco Nexus N9K-C93180YC-FX3S-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus N9K-C93180YC-FX3S B	Eth1/17	25GbE	NetApp Controller 1	e1a Controller 2
	Eth1/18	25GbE	NetApp Controller 1	e1b Controller 2
	Eth1/19	25GbE	NetApp Controller 2	e1c Controller 1
	Eth1/20	25GbE	NetApp Controller 2	e1d Controller 1
	Eth1/49	100GbE	Cisco UCS fabric interconnect A	P0 1/21
	Eth1/50	100GbE	Cisco UCS fabric interconnect A	P0 1/23
	Eth1/51	100GbE	Cisco UCS fabric interconnect B	P0 1/21
	Eth1/52	100GbE	Cisco UCS fabric interconnect B	P0 1/23
	Eth1/53	100GbE	Cisco Nexus N9K-C93180YC-FX3S A	Eth1/53
	Eth1/54	100GbE	Cisco Nexus N9K-C93180YC-FX3S A	Eth1/54
	MGMT0	GbE	GbE management switch	Any

Table 10. NetApp Controller-1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp AFF400 Node 1	e0M	1GbE	1GbE management switch	Any
	e0s	GbE	GbE management switch	Any
	e0a	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	e0a
	e0b	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	e0b
	e0c	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	e0c
	e0d	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	e0d
	e1a	25GbE	Data Ports	N9K-A Eth1/17

Local Device	Local Port	Connection	Remote Device	Remote Port
	e1c	25GbE	Data Ports	N9K-A Eth1/18
	e1d	25GbE	Data Ports	N9K-B Eth1/17
	e1b	25GbE	Data Ports	N9K-B Eth1/18

Note: When the term e0M is used, the physical Ethernet port to which the table is referring is the port indicated by a wrench icon on the rear of the chassis.

Table 11. NetApp Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp AFF400 Node 2	e0M	1GbE	1GbE management switch	Any
	e0s	GbE	GbE management switch	Any
	e0a	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	e0a
	e0b	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	e0b
	e0c	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	e0a
	e0d	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	e0b
	e1a	25GbE	Data Ports	N9K-A Eth1/19
	e1b	25GbE	Data Ports	N9K-A Eth1/20
	e1c	25GbE	Data Ports	N9K-B Eth1/19
	e1d	25GbE	Data Ports	N9K-B Eth1/20

Table 12. Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/41	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	Eth1/41
	Eth1/42	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	Eth1/42
	Eth1/43	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	Eth1/41
	Eth1/44	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	Eth1/42
	MGMT0	GbE	GbE Management switch	
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 13. Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth1/41	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	Eth1/43
	Eth1/42	25GbE	Cisco Nexus N9K-C93180YC-FX3S B	Eth1/44
	Eth1/43	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	Eth1/43
	Eth1/44	25GbE	Cisco Nexus N9K-C93180YC-FX3S A	Eth1/44
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect A	L1
	L2	GbE	Cisco UCS fabric interconnect A	L2

Network Switch Configuration

This section contains the following procedures:

- [Set Up Initial Configuration on Cisco Nexus A](#)
- [Set Up Initial Configuration on Cisco Nexus B](#)

This section provides a detailed procedure for configuring the Cisco Nexus N9K-C93180YC-FX3S switches for use in a FlexPod environment. The Cisco Nexus N9K-C93180YC-FX3S will be used LAN switching in this solution.

IMPORTANT! Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in section [Solution Cabling](#).

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 9.3(7a), the Cisco suggested Nexus switch release at the time of this validation.

The following procedure includes the setup of NTP distribution on both the mgmt0 port and the in-band management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes that the default VRF is used to route the in-band management VLAN.

Note: In this validation, port speed and duplex are hard set at both ends of every 100GE connection.

Procedure 1. Set Up Initial Configuration on Cisco Nexus A

Set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>

Step 1. On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass
password and basic configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
```

```

poap: Rolling back, please wait... (This may take 5-15 minutes)
      ---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: yes
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

Step 2. Review the configuration summary before enabling the configuration.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

Procedure 2. Set Up Initial Configuration on Cisco Nexus B

Set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>.

Step 1. On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```

Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass
password and basic configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)
      ---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes

```

```
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: yes
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

Step 2. Review the configuration summary before enabling the configuration.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

FlexPod Cisco Nexus Switch Configuration

This section contains the following procedures:

- [Enable Features on Cisco Nexus A and Cisco Nexus B](#)
- [Set Global Configurations on Cisco Nexus A and Cisco Nexus B](#)
- [Create VLANs on Cisco Nexus A and Cisco Nexus B](#)
- [Add NTP Distribution Interface on Cisco Nexus A](#)
- [Add NTP Distribution Interface on Cisco Nexus B](#)
- [Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus A](#)
- [Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus B](#)
- [Create Port Channels on Cisco Nexus A and Cisco Nexus B](#)
- [Configure Port Channel Parameters on Cisco Nexus A and Cisco Nexus B](#)
- [Configure Virtual Port Channels on Cisco Nexus A](#)
- [Configure Virtual Port Channels on Cisco Nexus B](#)

Procedure 1. Enable Features on Cisco Nexus A and Cisco Nexus B

SAN switching requires both the SAN_ENTERPRISE_PKG and FC_PORT_ACTIVATION_PKG licenses. Please ensure these licenses are installed on each Cisco Nexus N9K-C93180YC-FX3S switch.

Step 1. Log in as admin.

Step 2. Since basic FC configurations were entered in the setup script, feature-set fcoe has been automatically installed and enabled. Run the following commands:

```
config t
feature udd
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

Procedure 2. Set Global Configurations on Cisco Nexus A and Cisco Nexus B

Step 1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day> <end-month> <end-time> <offset-minutes>
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

Tech tip

It is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summer time, see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3\(x\)](#). Sample clock commands for the United States Eastern timezone are:

```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60.
```

Procedure 3. Create VLANs on Cisco Nexus A and Cisco Nexus B

Step 1. From the global configuration mode, run the following commands:

```
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
```

```
vlan <infra-nfs-vlan-id>
name Infra-NFS-VLAN
vlan <CIFS-VLAN>
name CIFS-VLAN
exit
```

Procedure 4. Add NTP Distribution Interface on Cisco Nexus A

Step 1. From the global configuration mode, run the following commands:

```
interface vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-b-ntp-ip> use-vrf default
```

Procedure 5. Add NTP Distribution Interface on Cisco Nexus B

Step 1. From the global configuration mode, run the following commands:

```
interface vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-a-ntp-ip> use-vrf default
```

Procedure 6. Add Port Profiles on Cisco Nexus A and Cisco Nexus B

This version of the FlexPod solution uses port profiles for virtual port channel (vPC) connections to NetApp Storage, Cisco UCS, and the vPC peer link.

Step 1. From the global configuration mode, run the following commands:

```
port-profile type port-channel FP-ONTAP-Storage
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled

port-profile type port-channel FP-UCS
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-
id>, <vm-traffic-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled

port-profile type port-channel vPC-Peer-Link
```

```
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-nfs-vlan-id>, <vmotion-vlan-id>,
<vm-traffic-vlan-id>
spanning-tree port type network
speed 50000
duplex full
state enabled
```

Procedure 7. Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus A

Note: In this procedure and in the following sections, configure the AFF nodename <st-node> and Cisco UCS 6536 fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

Step 1. From the global configuration mode, run the following commands:

```
interface Eth1/41
description <ucs-clustername>-a:1/43
udld enable
interface Eth1/42
description <ucs-clustername>-a:1/44
udld enable
interface Eth1/43
description <ucs-clustername>-b:1/43
udld enable
interface Eth1/44
description <ucs-clustername>-b:1/44
udld enable
```

Step 2. For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected. If you have fibre optic connections, do not enter the `udld enable` command.

```
interface Eth1/17
description <st-clustername>-01:e0e
interface Eth1/18
description <st-clustername>-01:e0f
interface Eth1/19
description <st-clustername>-02:e0e
interface Eth1/20
description <st-clustername>-02:e0f
interface Eth1/53
description <nexus-b-hostname>:1/53
interface Eth1/54
description <nexus-b-hostname>:1/54
exit
```

Procedure 8. Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces on Cisco Nexus B

Step 1. From the global configuration mode, run the following commands:

```
interface Eth1/41
description <ucs-clustername>-a:1/41
udld enable
interface Eth1/42
description <ucs-clustername>-a:1/42
udld enable
interface Eth1/43
description <ucs-clustername>-b:1/41
udld enable
interface Eth1/44
description <ucs-clustername>-b:1/42
udld enable
```

Step 2. For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected.

```
interface Eth1/17
description <st-clustername>-01:e0g
interface Eth1/18
description <st-clustername>-01:e0h
interface Eth1/19
description <st-clustername>-02:e0g
interface Eth1/20
description <st-clustername>-02:e0h
interface Eth1/53
description <nexus-a-hostname>:1/53
interface Eth1/54
description <nexus-a-hostname>:1/54
exit
```

Procedure 9. Create Port Channels on Cisco Nexus A and Cisco Nexus B

Step 1. From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
interface Eth1/53-54
channel-group 10 mode active
no shutdown
interface Po117
description <st-clustername>-01
interface Eth1/17-18
channel-group 117 mode active
no shutdown
interface Po119
```

```
description <st-clustername>-02
interface Eth1/19-20
channel-group 119 mode active
no shutdown
interface Po141
description <ucs-clustername>-a
interface Eth1/41-42
channel-group 121 mode active
no shutdown
interface Po142
description <ucs-clustername>-b
interface Eth1/43-44
channel-group 123 mode active
no shutdown
exit
copy run start
```

Procedure 10. Configure Port Channel Parameters on Cisco Nexus A and Cisco Nexus B

Step 1. From the global configuration mode, run the following commands:

```
interface Po10
inherit port-profile vPC-Peer-Link

interface Po117
inherit port-profile FP-ONTAP-Storage
interface Po119
inherit port-profile FP-ONTAP-Storage

interface Po141
inherit port-profile FP-UCS
interface Po142
inherit port-profile FP-UCS

exit
copy run start
```

Procedure 11. Configure Virtual Port Channels on Cisco Nexus A

Step 1. From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
```



```
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po141
vpc 121
interface Po142
vpc 123
exit
copy run start
```

Procedure 12. Configure Virtual Port Channels on Cisco Nexus B

Step 1. From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po117
vpc 117
interface Po119
vpc 119
interface Po141
vpc 121
interface Po142
vpc 123
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlexPod environment into the infrastructure. The previously described

procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Switch Testing Commands

The following commands can be used to check for correct switch configuration:

Note: Some of these commands need to run after further configuration of the FlexPod components are complete to see complete results.

```
show run
show vpc
show port-channel summary
show ntp peer-status
show cdp neighbors
show lldp neighbors
show run int
show int
show udd neighbors
show int status
```

Storage Configuration

This chapter contains the following:

- [NetApp Hardware Universe](#)
- [NetApp ONTAP 9.13.1P3](#)

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It also provides configuration information for all the NetApp storage appliances currently supported by ONTAP software and a table of component compatibilities.

To confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install, follow these steps at the [NetApp Support](#) site.

1. Access the [HWU application](#) to view the System Configuration guides. Click the Products tab to select the Platforms menu to view the compatibility between different versions of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Utilities and select Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers found here: <https://docs.netapp.com/us-en/ontap-systems/index.html>.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported by the AFF A400 and AFF A800 is available at the [NetApp Support](#) site.

When using SAS disk shelves with NetApp storage controllers, refer to: <https://docs.netapp.com/us-en/ontap-systems/sas3/install-new-system.html> for proper cabling guidelines.

When using NVMe drive shelves with NetApp storage controllers, refer to: <https://docs.netapp.com/us-en/ontap-systems/ns224/hot-add-shelf.html> for installation and servicing guidelines.

NetApp ONTAP 9.13.1P3

This section contains the following procedures:

- [Configure Node 01](#)
- [Configure Node 02](#)
- [Set Up Node](#)
- [Log into the Cluster](#)
- [Verify Storage Failover](#)
- [Set Auto-Revert on Cluster Management](#)
- [Zero All Spare Disks](#)
- [Set Up Service Processor Network Interface](#)

-
- [Create Manual Provisioned Aggregates \(Optional\)](#)
 - [Remove Default Broadcast Domains](#)
 - [Disable Flow Control on 25/100GbE Data Ports](#)
 - [Enable Cisco Discovery Protocol](#)
 - [Enable Link-layer Discovery Protocol on all Ethernet Ports](#)
 - [Enable FIPS Mode on the NetApp ONTAP Cluster \(Optional\)](#)
 - [Configure Timezone](#)
 - [Configure Simple Network Management Protocol \(SNMP\)](#)
 - [Configure SNMPv3 Access](#)
 - [Configure login banner for the NetApp ONTAP Cluster](#)
 - [Remove insecure ciphers from the NetApp ONTAP Cluster](#)
 - [Create Management Broadcast Domain](#)
 - [Create NFS Broadcast Domain](#)
 - [Create CIFS Broadcast Domain](#)
 - [Create Interface Groups](#)
 - [Change MTU on Interface Groups](#)
 - [Create VLANs](#)
 - [Create an infrastructure SVM](#)
 - [Configure CIFS Servers](#)
 - [Create Load-Sharing Mirrors of a SVM Root Volume](#)
 - [Configure HTTPS Access to the Storage Controller](#)
 - [Set password for SVM vsadmin user and unlock the user](#)
 - [Configure login banner for the SVM](#)
 - [Remove insecure ciphers from the SVM](#)
 - [Configure Export Policy Rule](#)
 - [Create CIFS Export Policy](#)
 - [Create a NetApp FlexVol Volume](#)
 - [Disable Volume Efficiency on swap volume](#)
 - [Create CIFS Shares](#)
 - [Create NFS LIFs](#)
 - [Create CIFS LIFs](#)
 - [Add Infrastructure SVM Administrator and SVM Administration LIF to In-band Management Network](#)
 - [Configure Auto-Support](#)

Complete Configuration Worksheet

Before running the setup script, complete the [Cluster setup worksheet](#) in the ONTAP 9 Documentation Center. You must have access to the [NetApp Support](#) site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [Software setup](#) section of the [ONTAP 9 Documentation Center](#) to learn about configuring ONTAP. [Table 14](#) lists the information needed to configure two ONTAP nodes. Customize the cluster-detail values with the information applicable to your deployment.

Table 14. ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
ONTAP 9.13.1P3 URL (http server hosting ONTAP software)	<url-boot-software>

Procedure 1. Configure Node 01

Step 1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press **Ctrl-C** to exit the autoboot loop when the following message displays:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

Step 2. Allow the system to boot up.

```
autoboot
```

Step 3. Press **Ctrl-C** when prompted.

Note: If NetApp ONTAP 9.13.1P3 is not the version of software being booted, continue with the following steps to install new software. If NetApp ONTAP 9.13.1P3 is the version being booted, select option 8 and y to reboot the node. Continue with section [Set Up Node](#).

Step 4. To install new software, select option **7** from the menu.

Step 5. Enter **y** to continue the installation.

Step 6. Select **e0M** for the network port for the download.

Step 7. Enter **n** to skip the reboot.

Step 8. Select option **7** from the menu: `Install new software first`

Step 9. Enter **y** to continue the installation.

Step 10. Enter the **IP address**, **netmask**, and **default gateway** for **e0M**.

```
Enter the IP address for port e0M: <node01-mgmt-ip>
Enter the netmask for port e0M: <node01-mgmt-mask>
```

```
Enter the IP address of the default gateway: <node01-mgmt-gateway>
```

Step 11. Enter the **URL** where the software can be found.

Step 12. The e0M interface should be connected to the management network and the web server must be reachable (using ping) from node 01.

```
<url-boot-software>
```

Step 13. Press **Enter** for the user name, indicating no user name.

Step 14. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

Step 15. Enter `yes` to reboot the node.

```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y ←

Please answer yes or no

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} yes
```

Note: When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

Note: During the ONTAP installation, a prompt to reboot the node requests a Y/N response. The prompt requires the entire Yes or No response to reboot the node and continue the installation.

Step 16. Press **Ctrl-C** when the following message displays:

```
Press Ctrl-C for Boot Menu
```

Step 17. Select option **4** for Clean Configuration and Initialize All Disks.

Step 18. Enter `y` to zero disks, reset config, and install a new file system.

Step 19. Enter `yes` to erase all the data on the disks.

Note: When the initialization and creation of the root aggregate is complete, the storage system reboots. You can continue with the configuration of node 02 while the initialization and creation of the root aggregate for node 01 is in progress. For more information about root aggregate and disk partitioning, refer to the ONTAP documentation on root-data partitioning: <https://docs.netapp.com/us-en/ontap/concepts/root-data-partitioning-concept.html>.

Procedure 2. Configure Node 02

Step 1. Connect to the storage system console port. You should see a Loader-B prompt. However, if the storage system is in a reboot loop, press **Ctrl-C** to exit the autoboot loop when the following message displays: Starting

```
AUTOBOOT press Ctrl-C to abort...
```

Step 2. Allow the system to boot up.

```
autoboot
```

Step 3. Press **Ctrl-C** when prompted.

Note: If NetApp ONTAP 9.13.1P3 is not the version of software being booted, continue with the following steps to install new software. If NetApp ONTAP 9.13.1P3 is the version being booted, select option 8 and y to reboot the node, then continue with section [Set Up Node](#).

Step 4. To install new software, select option 7.

Step 5. Enter y to continue the installation.

Step 6. Select e0M for the network port you want to use for the download.

Step 7. Enter n to skip the reboot.

Step 8. Select option **7: Install new software first**

Step 9. Enter y to continue the installation

Step 10. Enter the **IP address, netmask, and default gateway** for e0M.

```
Enter the IP address for port e0M: <node02-mgmt-ip>
Enter the netmask for port e0M: <node02-mgmt-mask>
Enter the IP address of the default gateway: <node02-mgmt-gateway>
```

Step 11. Enter the **URL** where the software can be found.

Step 12. The web server must be reachable (ping) from node 02.

```
<url-boot-software>
```

Step 13. Press **Enter** for the username, indicating no user name.

Step 14. Enter **y** to set the newly installed software as the default to be used for subsequent reboots.

Step 15. Enter **yes** to reboot the node.

```
Do you want to set the newly installed software as the default to be used for
subsequent reboots? {y|n} y

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y ←
Please answer yes or no

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} yes
```

Note: When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-B prompt. If these actions occur, the system might deviate from this procedure.

Note: During the ONTAP installation a prompt to reboot the node requests a Y/N response. The prompt requires the entire Yes or No response to reboot the node and continue the installation.

Step 16. Press **Ctrl-C** when you see this message:

```
Press Ctrl-C for Boot Menu
```

Step 17. Select option **4** for Clean Configuration and Initialize All Disks.

Step 18. Enter **y** to zero disks, reset config, and install a new file system.

Step 19. Enter **yes** to erase all the data on the disks.

Procedure 3. Set Up Node

Step 1. From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.13.1P3 boots on the node for the first time.

Step 2. Follow the prompts to set up node 01.

```

Welcome to node setup.
You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup."
To accept a default or omit a question, do not enter a value.
This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on
your system.
For further information on AutoSupport, see: http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created
Use your web browser to complete cluster setup by accessing https://<node01-mgmt-ip>. Otherwise press Enter to
complete cluster setup using the command line interface.

```

Step 3. To complete cluster setup, open a web browser and navigate to <https://<node01-mgmt-ip>>.

Table 15. Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
Cluster Admin SVM	<cluster-adm-svm>
Infrastructure Data SVM	<infra-data-svm>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-sp-ip>
Node 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>

Cluster Detail	Cluster Detail Value
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>
NTP server B IP address	<switch-b-ntp-ip>
SNMPv3 User	<snmp-v3-usr>
SNMPv3 Authentication Protocol	<snmp-v3-auth-PROTO>
SNMPv3 Privacy Protocol	<snmpv3-priv-PROTO>

Note: Cluster setup can also be performed using the CLI. This document describes the cluster setup using the NetApp ONTAP System Manager guided setup.

Step 4. Complete the required information on the Initialize Storage System screen:

- a. In the Cluster screen, enter the **cluster name** and **administrator password**.
- b. Complete the Networking information for the cluster and each node.

Tech tip

The nodes should be discovered automatically; if they are not, Refresh the browser page. By default, the cluster interfaces are created on all the new factory shipping storage controllers.

If all the nodes are not discovered, then configure the cluster using the command line.

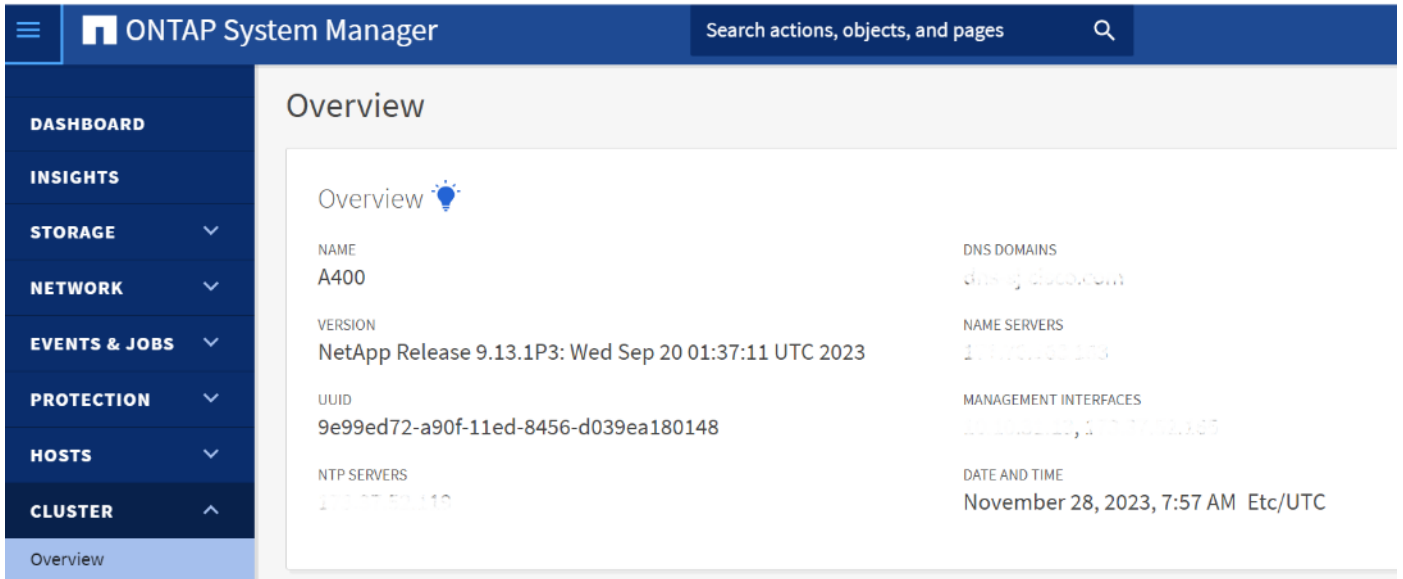
The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

- c. Click **Submit**.

Step 5. A few minutes will pass while the cluster is configured. When prompted, login to **ONTAP System Manager** to continue the cluster configuration.

Step 6. From the Dashboard click the **Cluster** menu and click **Overview**.

Step 7. Click the **More** button in the Overview pane and click **Edit**.

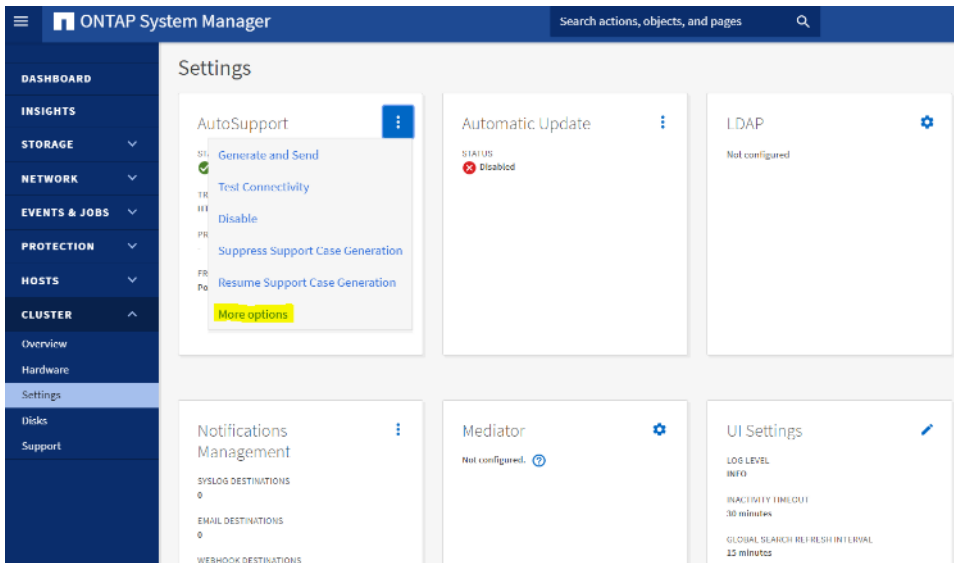


Step 8. Add additional cluster configuration details and click **Save** to make the changes persistent:

- Cluster location
- DNS domain name
- DNS server IP addresses
- DNS server IP addresses can be added individually or with a comma separated list on a single line.

Step 9. Click **Save** to make the changes persistent.

Step 10. Select the **Settings** menu under the Cluster menu.



Step 11. If AutoSupport was not configured during the initial setup, click the **ellipsis** in the AutoSupport tile and select **More options**.

Step 12. To enable AutoSupport click the **slider**.

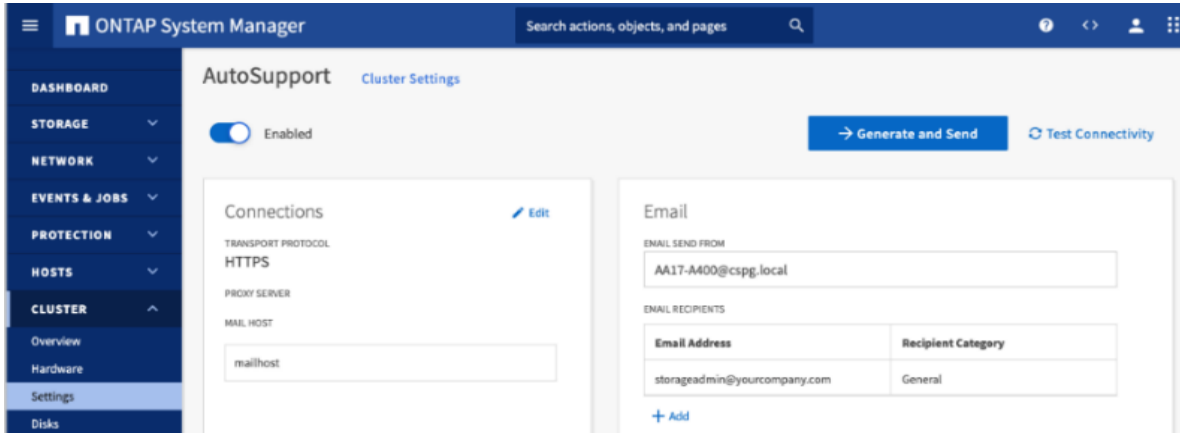
Step 13. Click **Edit** to change the transport protocol, add a proxy server address and a mail host as needed.

Step 14. Click **Save** to enable the changes.

Step 15. In the Email tile, click **Edit** and enter the desired email information:

- Email send from address
- Email recipient addresses
- Recipient Category

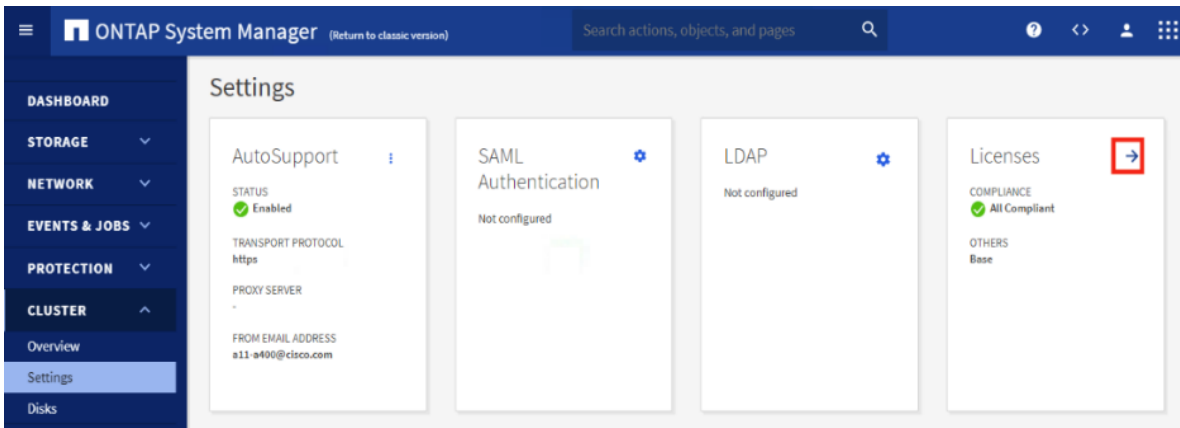
Step 16. Click **Save** when complete.




Step 17. Select **CLUSTER > Settings** to return to the cluster settings page.

Step 18. Locate the Licenses tile on the right and click the arrow.

Step 19. Add the desired licenses to the cluster by clicking **Add** and entering the license keys in a comma separated list.

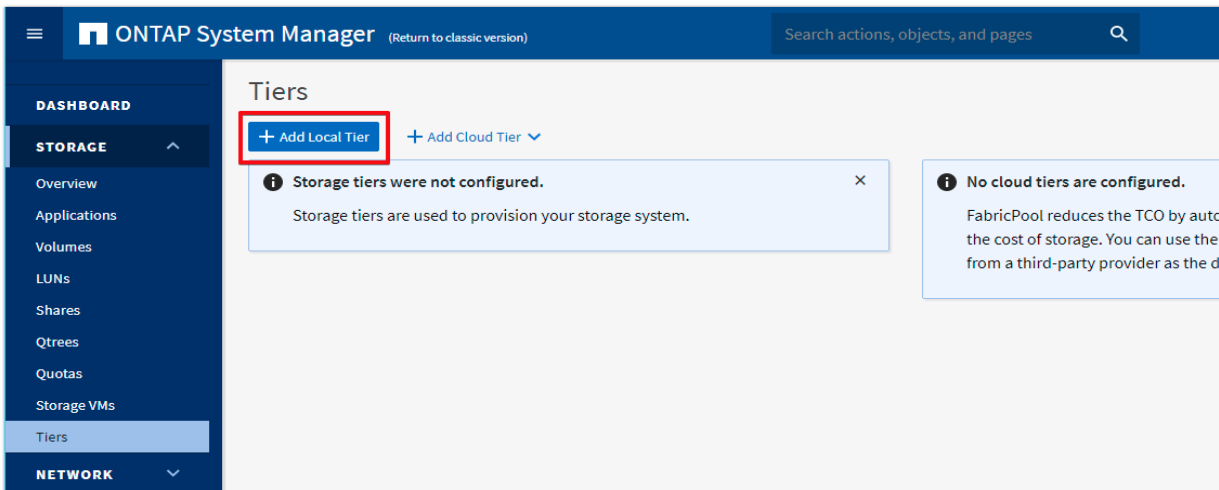




Note: NetApp ONTAP 9.10.1 and later for FAS/AFF storage systems uses a new file-based licensing solution to enable per-node NetApp ONTAP features. The new license key format is referred to as a NetApp License File, or NLF. For more information, go to: [NetApp ONTAP 9.10.1 and later Licensing Overview - NetApp Knowledge Base](#)

Step 20. Configure storage aggregates by selecting the **Storage** menu and selecting **Tiers**.

Step 21. Click **Add Local Tier** and allow ONTAP System Manager to recommend a storage aggregate configuration.




Step 22. ONTAP will use best practices to recommend an aggregate layout. Click the **Recommended** details link to view the aggregate information.

Step 23. Optionally, enable NetApp Aggregate Encryption (NAE) by checking the box for Configure Onboard Key Manager for encryption.

Step 24. Enter and confirm the **passphrase** and save it in a secure location for future use.

Step 25. Click **Save** to make the configuration persistent.



Note: Aggregate encryption may not be supported for all deployments. Please review the [NetApp Encryption Power Guide](#) and the [Security Hardening Guide for NetApp ONTAP 9 \(TR-4569\)](#) to help determine if aggregate encryption is right for your environment.

Procedure 4. Log into the Cluster

- Step 1.** Open an **SSH** connection to either the cluster IP or the host name.
- Step 2.** Log into the **admin user** with the **password** you provided earlier.

Procedure 5. Verify Storage Failover

Step 1. Verify the status of the storage failover:

```
A400-VDI-Cluster::> storage failover show
                        Takeover
Node      Partner      Possible State Description
-----
A400-VDI-Cluster-01
           A400-VDI-    true      Connected to A400-VDI-Cluster-02
           Cluster-02
A400-VDI-Cluster-02
           A400-VDI-    true      Connected to A400-VDI-Cluster-01
           Cluster-01
2 entries were displayed.
```

Note: Both <st-node01> and <st-node02> must be capable of performing a takeover. Continue with Step 2 if the nodes can perform a takeover.

Step 2. Enable failover on one of the two nodes if it was not completed during the installation:

```
storage failover modify -node <st-node01> -enabled true
```

Note: Enabling failover on one node enables it for both nodes.

Step 3. Verify the HA status for a two-node cluster:

Note: This step is not applicable for clusters with more than two nodes.

```
A400-VDI-Cluster ::> cluster ha show
```

```
High-Availability Configured: true
```

Note: If HA is not configured use the following commands. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

Step 4. Verify that hardware assist is correctly configured:

```
A400-VDI-Cluster::> storage failover hwassist show
Node
-----
A400-VDI-Cluster-01
           Partner: A400-VDI-Cluster-02
           Hwassist Enabled: true
           Hwassist IP: 192.X.X.84
           Hwassist Port: 162
           Monitor Status: active
           Inactive Reason: -
           Corrective Action: -
           Keep-Alive Status: healthy
A400-VDI-Cluster-02
           Partner: A400-VDI-Cluster-01
           Hwassist Enabled: true
           Hwassist IP: 192.X.X.85
           Hwassist Port: 162
           Monitor Status: active
           Inactive Reason: -
           Corrective Action: -
           Keep-Alive Status: healthy
2 entries were displayed.
```

Procedure 6. Set Auto-Revert on Cluster Management

Step 1. Set the `auto-revert` parameter on the cluster management interface:

Note: A storage virtual machine (SVM) is referred to as a Vserver or vserver in the GUI and CLI.

```
network interface modify -vserver <clustername> -lif cluster_mgmt_lif_1 -auto-revert true
```

Procedure 7. Zero All Spare Disks

Step 1. Zero all spare disks in the cluster by running the following command:

```
disk zerospares
```

Tech tip

Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an AFF configuration. Disk autoassign should have assigned one data partition to each node in an HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Procedure 8. Set Up Service Processor Network Interface

Step 1. Assign a static IPv4 address to the Service Processor on each node by running the following commands:

```
system service-processor network modify -node <st-node01> -address-family IPv4 -enable true -dhcp none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>
```

```
system service-processor network modify -node <st-node02> -address-family IPv4 -enable true -dhcp none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```

Note: The Service Processor IP addresses should be in the same subnet as the node management IP addresses.

Procedure 9. Create Manual Provisioned Aggregates - Optional

Note: An aggregate containing the root volume is created during the ONTAP setup process. To manually create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it should contain.

Step 1. Create new aggregates by running the following commands:

```
storage aggregate create -aggregate <aggr1_node01> -node <st-node01> -diskcount <num-disks> -diskclass solid-state
storage aggregate create -aggregate <aggr1_node02> -node <st-node02> -diskcount <num-disks> -diskclass solid-state
```

Note: You should have the minimum number of hot spare disks for the recommended hot spare disk partitions for their aggregate.

Note: For all-flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.

Tech tip

In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all, but one remaining disk (spare) assigned to the controller.

Note: The aggregate cannot be created until disk zeroing completes. Run the storage aggregate show command to display the aggregate creation status. Do not proceed until both aggr1_node01 and aggr1_node02 are online.

Procedure 10. Remove Default Broadcast Domains

Note: By default, all network ports are included in separate default broadcast domain. Network ports used for data services (for example, e0e, e0f, and so on) should be removed from their default broadcast domain and that broadcast domain should be deleted.

Step 1. Run the following commands:

```
network port broadcast-domain delete -broadcast-domain <Default-N> -ipspace Default
network port broadcast-domain show
```

Step 2. Delete the Default broadcast domains with Network ports (Default-1, Default-2, and so on). This does not include Cluster ports and management ports.

Procedure 11. Disable Flow Control on 25/100GbE Data Ports

Step 1. Disable the flow control on 25 and 100GbE data ports by running the following command to configure the ports on node 01:

```
network port modify -node <st-node01> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
```

Step 2. Run the following command to configure the ports on node 02:

```
network port modify -node <st-node02> -port e0e,e0f,e0g,e0h -flowcontrol-admin none
```

Note: Disable flow control only on ports that are used for data traffic.

Procedure 12. Enable Cisco Discovery Protocol

Step 1. Enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers by running the following command to enable CDP in ONTAP:

```
node run -node * options cdpd.enable on
```

Procedure 13. Enable Link-layer Discovery Protocol on all Ethernet Ports

Step 1. Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches, on all ports, of all nodes in the cluster, by running the following command:

```
node run * options lldp.enable on
```

Procedure 14. Enable FIPS Mode on the NetApp ONTAP Cluster - Optional

NetApp ONTAP is compliant in the Federal Information Processing Standards (FIPS) 140-2 for all SSL connections. When SSL FIPS mode is enabled, SSL communication from NetApp ONTAP to external client or server components outside of NetApp ONTAP will use FIPS compliant crypto for SSL.

Step 1. To enable FIPS on the NetApp ONTAP cluster, run the following commands:

```
set -privilege advanced  
security config modify -interface SSL -is-fips-enabled true
```

Note: If you are running NetApp ONTAP 9.8 or earlier manually reboot each node in the cluster one by one. Beginning in NetApp ONTAP 9.9.1, rebooting is not required.

Procedure 15. Configure Timezone

Step 1. To configure time synchronization on the cluster, follow this step:

```
timezone -timezone <timezone>
```

Note: For example, in the eastern United States, the time zone is America/New_York.

Procedure 16. Configure Simple Network Management Protocol - SNMP

Note: If you enabled FIPS then look at the following points while configuring SNMP.

Note: The SNMP users or SNMP traphosts that are non-compliant with FIPS will be deleted automatically. “Configure SNMP traphosts” configuration will be non-compliant with FIPS.

Note: The SNMPv1 user, SNMPv2c user (After configuring SNMP community) or SNMPv3 user (with none or MD5 as authentication protocol or none or DES as encryption protocol or both) is non-compliant with FIPS.

Step 1. Configure basic SNMP information, such as the location and contact. When polled, this information is visible as the sysLocation and sysContact variables in SNMP.

```
snmp contact <snmp-contact>  
snmp location "<snmp-location>"  
snmp init 1  
options snmp.enable on
```

Step 2. Configure SNMP traps to send to remote hosts, such as an Active IQ Unified Manager server or another fault management system.

Note: This step works when FIPS is disabled.

Note: An SNMPv1 traphost or SNMPv3 traphost (configured with an SNMPv3 user non-compliant to FIPS) is non-compliant to FIPS.

```
snmp traphost add <oncommand-um-server-fqdn>
```

Step 3. Configure SNMP community.

Note: This step works when FIPS is disabled.

Note: SNMPv1 and SNMPv2c are not supported when cluster FIPS mode is enabled.

```
system snmp community add -type ro -community-name -vserver
```

Note: In new installations of NetApp ONTAP, SNMPv1 and SNMPv2c are disabled by default. SNMPv1 and SNMPv2c are enabled after you create an SNMP community.

Note: NetApp ONTAP supports read-only communities.

Procedure 17. Configure SNMPv3 Access

Note: SNMPv3 offers advanced security by using encryption and passphrases. The SNMPv3 user can run SNMP utilities from the traphost using the authentication and privacy settings that you specify.

Note: Configure SNMP traps to send to remote hosts, such as an Active IQ Unified Manager server or another fault management system.

Note: When FIPS is enabled, below are the supported/compliant options for authentication and privacy protocol:

- Authentication Protocol: sha, sha2-256
- Privacy protocol: aes128

Step 1. Configure the SNMPv3 access by running the following command:

```
security login create -user-or-group-name <<snmp-v3-usr>> -application snmp -authentication-method usm
Enter the authoritative entity's EngineID [local EngineID]:
Which authentication protocol do you want to choose (none, md5, sha, sha2-256) [none]: <<snmp-v3-auth-proto>>
Enter the authentication protocol password (minimum 8 characters long):
Enter the authentication protocol password again:
Which privacy protocol do you want to choose (none, des, aes128) [none]: <<snmpv3-priv-proto>>
Enter privacy protocol password (minimum 8 characters long):
Enter privacy protocol password again:
```

Note: Refer to the NetApp [SNMP Configuration Express Guide](#) for additional information when configuring SNMPv3 security users.

Procedure 18. Configure login banner for the NetApp ONTAP Cluster

Step 1. To create login banner for the NetApp ONTAP cluster, run the following command:

```
security login banner modify -message "Access restricted to authorized users" -vserver <clustername>
```

Procedure 19. Remove insecure ciphers from the NetApp ONTAP Cluster

Step 1. Ciphers with the suffix CBC are considered insecure. To remove the CBC ciphers, run the following NetApp ONTAP command:

```
security ssh remove -vserver <clustername> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

Procedure 20. Create Management Broadcast Domain

Step 1. If the management interfaces are required to be on a separate VLAN, create a new broadcast domain for those interfaces by running the following command:

```
network port broadcast-domain create -broadcast-domain IB-MGMT -mtu 1500
```

Procedure 21. Create NFS Broadcast Domain

Step 1. To create an NFS, data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands to create a broadcast domain for NFS in ONTAP:

```
network port broadcast-domain create -broadcast-domain Infra-NFS -mtu 9000
```

Procedure 22. Create CIFS Broadcast Domain

Step 1. To create a CIFS data broadcast domain with a maximum transmission unit (MTU) of 9000, run the following commands to create a broadcast domain for CIFS in ONTAP:

```
network port broadcast-domain create -broadcast-domain CIFS-VLAN -mtu 9000
```

Procedure 23. Create Interface Groups

Step 1. To create the LACP interface groups for the 25GbE data interfaces, run the following commands:

```
network port ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0e
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0f
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node01> -ifgrp a0a -port e0h
network port ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0e
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0f
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0g
network port ifgrp add-port -node <st-node02> -ifgrp a0a -port e0h
```

Step 2. To verify, run the following:

```
A400-VDI-Cluster::> network port ifgrp show
      Port      Distribution      Active
Node   IfGrp      Function      MAC Address      Ports      Ports
-----
A400-VDI-Cluster-01
      a0a      port          d2:39:ea:40:53:32 full      e0e, e0f, e0g, e0h
A400-VDI-Cluster-02
      a0a      port          d2:39:ea:45:2d:36 full      e0e, e0f, e0g, e0h
2 entries were displayed.
```

Procedure 24. Change MTU on Interface Groups

Step 1. To change the MTU size on the base interface-group ports before creating the VLAN ports, run the following commands:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
```

Procedure 25. Create VLANs

Step 1. Create the management VLAN ports and add them to the management broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<ib-mgmt-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<ib-mgmt-vlan-id>
network port broadcast-domain add-ports -broadcast-domain IB-MGMT -ports <st-node01>:a0a-<ib-mgmt-vlan-id>,<st-node02>:a0a-<ib-mgmt-vlan-id>
```

Step 2. Create the NFS VLAN ports and add them to the Infra-NFS broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>
network port broadcast-domain add-ports -broadcast-domain Infra-NFS -ports <st-node01>:a0a-<infra-nfs-vlan-id>,<st-node02>:a0a-<infra-nfs-vlan-id>
```

Step 3. Create the CIFS VLAN ports and add them to the Infra-CIFS broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<cifs-vlan>
network port vlan create -node <st-node02> -vlan-name a0a-<cifs-vlan>
```

```
network port broadcast-domain add-ports -broadcast-domain CIFS-VLAN -ports <st-node01>:a0a-<infra-cifs-vlan-id>,<st-node02>:a0a-<infra-cifs-vlan-id>
```

Step 4. To verify, run the following command:

```
A400-VDI-Cluster::> network port vlan show
      Network Network
Node  VLAN Name Port   VLAN ID  MAC Address
-----
A400-VDI-Cluster-01
  a0a-31  a0a    31      d2:39:ea:40:53:32
  a0a-32  a0a    32      d2:39:ea:40:53:32
  a0a-33  a0a    33      d2:39:ea:40:53:32

A400-VDI-Cluster-02
  a0a-31  a0a    31      d2:39:ea:45:2d:36
  a0a-33  a0a    32      d2:39:ea:45:2d:36
  a0a-33  a0a    33      d2:39:ea:45:2d:36

8 entries were displayed.
```

Procedure 26. Create an Infrastructure SVM

Step 1. Run the `vserver create` command:

```
vserver create -vserver Infra-SVM
```

Step 2. Add the required data protocols to the SVM:

```
Vserver add-protocols -protocols nfs,cifs,fc -vserver Infra-SVM
```

Note: For FC-NVMe configuration, add “fc” and “nvme” protocols to the SVM.

Note: For NVMe/TCP configuration with iSCSI booting, add “nvme” and “iscsi” protocols to the SVM.

Step 3. Remove the unused data protocols from the SVM:

```
Vserver remove-protocols -vserver Infra-SVM -protocols iscsi
```

Note: It is recommended to remove iSCSI or FCP protocols if the protocol is not in use (in this case iSCSI was removed).

Step 4. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp ONTAP Tools:

```
vserver modify -vserver Infra-SVM -aggr-list <aggr1_node01>,<aggr1_node02>
```

Step 5. Enable and run the NFS protocol in the Infra-SVM:

```
vserver nfs create -vserver Infra-SVM -udp disabled -v3 enabled -v4.1 enabled -vstorage enabled
```

Note: If the NFS license was not installed during the cluster configuration, make sure to install the license before starting the NFS service.

Step 6. Verify the NFS `vstorage` parameter for the NetApp NFS VAAI plug-in was enabled:

```
A400-VDI-Cluster::> vserver nfs show -fields vstorage
vserver  vstorage
-----
Infra-SVM enabled
```

Procedure 27. Configure CIFS Servers

Note: You can enable and configure CIFS servers on storage virtual machines (SVMs) with NetApp FlexVol® volumes to let SMB clients access files on your cluster. Each data SVM in the cluster can be bound to exactly one Active Directory domain. However, the data SVMs do not need to be bound to the same domain. Each data SVM can be bound to a unique Active Directory domain.

Step 1. To configure DNS for the Infra-SVM, run the following command:

```
dns create -vserver <vserver-name> -domains <dns-domain> -nameserve <dns-servers>
```

Example:

```
dns create -vserver Infra-SVM -domains flexpodb4.cisco.com -nameservers 10.102.1.151,10.102.1.152
```

Note: The node management network interfaces should be able to route to the Active Directory domain controller to which you want to join the CIFS server. Alternatively, a data network interface must exist on the SVM that can route to the Active Directory domain controller.

Step 2. Create a network interface on the IB-MGMT VLAN:

```
network interface create -vserver Infra-SVM -lif svm-mgmt -service-policy default-management -home-node <st-node01> -home-port a0a-<ib-mgmt-vlan-id> -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true
```

Step 3. Create the CIFS service:

```
vserver cifs create -vserver Infra-SVM -cifs-server Infra-CIFS -domain <domain.com>
```

In order to create an Active Directory machine account for the CIFS server, you must supply the name and password of a Windows account with sufficient privileges to add computers to the "CN=Computers" container within the "DOMAIN.COM" domain.

Enter the user name: Administrator@active_directory.local

Enter the password:

Procedure 28. Create Load-Sharing Mirrors of a SVM Root Volume

Step 1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node:

```
volume create -vserver Infra-SVM -volume Infra-SVM_root_m01 -aggregate <aggr1_node01> -size 1GB -type DP
volume create -vserver Infra-SVM -volume Infra-SVM_root_m02 -aggregate <aggr1_node02> -size 1GB -type DP
```

Step 2. Create a job schedule to update the root volume mirror relationships every 15 minutes:

```
job schedule interval create -name 15min -minutes 15
```

Step 3. Create the mirroring relationships:

```
snapmirror create -source-path Infra-SVM:Infra_SVM_root -destination-path Infra-SVM:Infra_SVM_root_m01 -type LS -schedule 15min
```

```
snapmirror create -source-path Infra-SVM:Infra_SVM_root -destination-path Infra-SVM:Infra_SVM_root_m02 -type LS -schedule 15min
```

Step 4. Initialize the mirroring relationships:

```
snapmirror initialize-ls-set -source-path Infra-SVM:Infra_SVM_root
```

Step 5. To verify, run the following:

```
A400-VDI-Cluster::> snapmirror show -type ls
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Healthy	Progress Last Updated
A400-VDI-Cluster://Infra-SVM/Infra_SVM_root	LS	A400-VDI-Cluster://Infra-SVM/Infra_SVM_root_m01	Snapmirrored	Idle	-	true	-
A400-VDI-Cluster://Infra-SVM/Infra_SVM_root	LS	A400-VDI-Cluster://Infra-SVM/Infra_SVM_root_m02	Snapmirrored	Idle	-	true	-

2 entries were displayed.

Procedure 29. Configure HTTPS Access to the Storage Controller

Step 1. Increase the privilege level to access the certificate commands:

```
set -privilege diag
```

```
Do you want to continue? {y|n}: y
```

Step 2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, the <serial-number>) by running the following command:

```
security certificate show
```

Step 3. For each SVM shown, the certificate common name should match the DNS fully qualified domain name (FQDN) of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial <serial-number>
```

Note: Deleting expired certificates before creating new certificates is a best practice. Run the security certificate delete command to delete the expired certificates. In the following command, use TAB completion to select and delete each default certificate.

Step 4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Infra-SVM and the cluster SVM. Use TAB completion to aid in the completion of these commands:

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

Step 5. To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the security certificate show command.

Step 6. Enable each certificate that was just created by using the -server-enabled true and -client-enabled false parameters. Use TAB completion to aid in the completion of these commands:

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

Step 7. Disable HTTP cluster management access:

```
network interface service-policy remove-service -vserver <clustername> -policy default-management -service management-http
```

Note: It is normal for some of these commands to return an error message stating that the entry does not exist.

Step 8. Return to the normal admin privilege level and verify that the system logs are available in a web browser:

```
set -privilege admin
https://<node01-mgmt-ip>/spi
https://<node02-mgmt-ip>/spi
```

Procedure 30. Set password for SVM vsadmin user and unlock the user

Step 1. Set a password for the SVM vsadmin user and unlock the user using the following commands:

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver Infra-SVM
```

Procedure 31. Configure login banner for the SVM

Step 1. To create login banner for the SVM, run the following command:

```
Security login banner modify -vserver Infra-SVM -message "This Infra-SVM is reserved for authorized users only!"
```

Procedure 32. Remove insecure ciphers from the SVM

Step 1. Ciphers with the suffix CBC are considered insecure. To remove the CBC ciphers, run the following NetApp ONTAP command:

```
security ssh remove -vserver Infra-SVM -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
```

Procedure 33. Configure Export Policy Rule

Step 1. Create a new rule for the infrastructure NFS subnet in the default export policy:

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -clientmatch <infra-nfs-subnet-cidr> -rorule sys -rwrule sys -superuser sys -allow-suid true
```

Step 2. Assign the FlexPod export policy to the infrastructure SVM root volume:

```
volume modify -vserver Infra-SVM -volume Infra_SVM_root -policy default
```

Procedure 34. Create CIFS Export Policy

Note: Optionally, you can use export policies to restrict CIFS access to files and folders on CIFS volumes. You can use export policies in combination with share level and file level permissions to determine effective access rights.

Step 1. Run the following command to create an export policy that limits access to devices in the domain:

```
export-policy create -vserver Infra-SVM -policyname cifs_policy  
export-policy rule create -vserver Infra-SVM -policyname cifs_policy -clientmatch <domain_name> -rorule krb5i,krb5p -rwrule krb5i,krb5p
```

Procedure 35. Create a NetApp FlexVol Volume

The following information is required to create a NetApp FlexVol volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

Step 1. To create FlexVols for datastores, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore -aggregate <aggr1_node01> -size 1TB -state online -policy default -junction-path /infra_datastore -space-guarantee none -percent-snapshot-space 0
```

Note: If you are going to setup and use SnapCenter to backup the infra_datastore volume, add “- snapshot-policy none” to the end of the volume create command for the infra_datastore volume.

Step 2. Create vCLS datastores to be used by the vSphere environment to host vSphere Cluster Services (vCLS) VMs using the command below:

```
volume create -vserver Infra-SVM -volume vCLS -aggregate <aggr1_node02> -size 100GB -state online -policy default -junction-path /vCLS -space-guarantee none -percent-snapshot-space 0
```

Step 3. To create swap volumes, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_swap -aggregate <aggr1_node01> -size 200GB -state online -policy default -junction-path /infra_swap -space-guarantee none -percent-snapshot-space 0 -snapshot-policy none.
```

Step 4. To create a FlexVol for the boot LUNs of servers, run the following command:

```
volume create -vserver Infra-SVM -volume esxi_boot -aggregate <aggr1_node01> -size 320GB -state online -policy default -space-guarantee none -percent-snapshot-space 0
```

Step 5. Run the following command to create a FlexVol for storing SVM audit log configuration:

```
volume create -vserver Infra-SVM -volume audit_log -aggregate <aggr1_node01> -size 50GB -state online -policy default -junction-path /audit_log -space-guarantee none -percent-snapshot-space 0
```

Procedure 36. Create a NetApp FlexGroup Volume

Step 1. To create CIFS volumes, run the following commands:

```
volume create -vserver Infra-SVM -volume cifs_vol_01 -aggr-list aggr01_node01,aggr01_node02 -aggr-list-multiplier4-state online -policy cifs_policy -size 2TB -junction-path /cifs_vol_01 -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume cifs_vol_02 -aggr-list aggr01_node01,aggr01_node02 -aggr-list-multiplier4-state online -policy cifs_policy -size 800GB -junction-path /cifs_vol_02 -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume cifs_vol_03 -aggr-list aggr01_node01,aggr01_node02 -aggr-list-multiplier4-state online -policy cifs_policy -size 800GB -junction-path /cifs_vol_03 -space-guarantee none -percent-snapshot-space 0
```

Step 2. Update set of load-sharing mirrors using the command below:

```
snapmirror update-ls-set -source-path Infra-SVM:Infra_SVM_root
```

Note: If you are going to setup and use SnapCenter to backup the infra_datastore volume, add “-snapshot-policy none” to the end of the volume create command for the infra_datastore volume.

Procedure 37. Disable Volume Efficiency on swap volume

Step 1. On NetApp AFF systems, deduplication is enabled by default. To disable the efficiency policy on the infra_swap volume, run the following command:

```
volume efficiency off -vserver Infra-SVM -volume infra_swap
```

Procedure 38. Create CIFS Shares

Note: A CIFS share is a named access point in a volume that enables CIFS clients to view, browse, and manipulate files on a file server.

Step 1. Run the following commands to create CIFS shares:

```
cifs share create -vserver Infra-SVM -share-name <CIFS_share_1> -path /cifs_vol_01 -share properties oplocks,browsable,continuously-available,showsnapshot
cifs share create -vserver Infra-SVM -share-name <CIFS_share_2> -path /cifs_vol_02 -share properties oplocks,browsable,continuously-available,showsnapshot
cifs share create -vserver Infra-SVM -share-name <CIFS_share_3> -path /cifs_vol_03 -share properties oplocks,browsable,continuously-available,showsnapshot
```

Procedure 39. Create NFS LIFs

Step 1. Run the following commands to create NFS LIFs:

```
network interface create -vserver Infra-SVM -lif nfs-lif-01 -service-policy default-data-files -home-node <st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs-lif-01-ip> -netmask <node01-nfs-lif-01-mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true

network interface create -vserver Infra-SVM -lif nfs-lif-02 -service-policy default-data-files -home-node <st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs-lif-02-ip> -netmask <node02-nfs-lif-02-mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true
```

Step 2. Run the following commands to verify:

```
A400-VDI-Cluster::> network interface show -vserver Infra-SVM -service-policy default-data-files
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Infra-SVM
nfs-lif-01   up/up       10.10.33.113/24  A400-VDI-Cluster-01
                                     a0a-33      true
nfs-lif-02   up/up       10.10.33.114/24  A400-VDI-Cluster-02
                                     a0a-33      true
2 entries were displayed.
```

Procedure 40. Create CIFS LIFs

Step 1. Run the following commands to create CIFS LIFs:

```
network interface create -vserver Infra-SVM -lif cifs_lif01 -service-policy default-data-files -home-node
<st-node01> -home-port a0a-<infra-cifs-vlan-id> -address <node01-cifs_lif01-ip> -netmask <node01-cifs_lif01-
mask> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true

network interface create -vserver Infra-SVM -lif cifs_lif02 -service-policy default-data-files -home-node
<st-node02> -home-port a0a-<infra-cifs-vlan-id> -address <node02-cifs_lif02-ip> -netmask <node02-cifs_lif02-
mask>> -status-admin up -failover-policy broadcast-domain-wide -auto-revert true
```

Procedure 41. Add Infrastructure SVM Administrator and SVM Administration LIF to In-band Management Network

Step 1. Create a default route that enables the SVM management interface to reach the outside world:

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-gateway>
```

Step 2. To verify, run the following:

```
A400-VDI-Cluster::> network route show -vserver Infra-SVM
Vserver      Destination      Gateway          Metric
-----
Infra-SVM    0.0.0.0/0       10.10.31.1      20
```

Note: A cluster serves data through at least one and possibly several SVMs. By completing these steps, you have created a single data SVM. You can create additional SVMs depending on their requirement.

Procedure 42. Configure Auto-Support

Step 1. NetApp AutoSupport sends support summary information to NetApp through HTTPS. To configure AutoSupport using command-line interface, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts -transport https -support enable -from
<storage-admin-email> -to <storage-admin-email>
```


Cisco Intersight Managed Mode Configuration

This chapter contains the following:

- [Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects](#)
- [Configure a Cisco UCS Domain Profile](#)
- [Cisco UCS Domain Configuration](#)
- [Configure Server Profile Template](#)
- [Management Configuration](#)

The Cisco Intersight platform is a management solution delivered as a service with embedded analytics for Cisco and third-party IT infrastructures. The Cisco Intersight managed mode (also referred to as Cisco IMM or Intersight managed mode) is a new architecture that manages Cisco Unified Computing System (Cisco UCS) fabric interconnect-attached systems through a Redfish-based standard model. Cisco Intersight managed mode standardizes both policy and operation management Cisco UCS X210c M7 compute nodes used in this deployment guide.

Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

This section contains the following procedures:

- [Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects](#)
- [Set up a new Cisco Intersight account](#)
- [Set up Cisco Intersight account and associate it with Cisco Smart Licensing](#)
- [Set up Cisco Intersight Resource Group](#)
- [Set up Cisco Intersight Organization](#)
- [Claim Cisco UCS Fabric Interconnects in Cisco Intersight](#)
- [Verify the addition of Cisco UCS Fabric Interconnects to Cisco Intersight](#)

Procedure 1. Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

Note: The Cisco UCS fabric interconnects need to be set up to support Cisco Intersight managed mode. When converting an existing pair of Cisco UCS fabric interconnects from Cisco UCS Manager mode to Intersight Managed Mode (IMM), first erase the configuration and reboot your system.

WARNING! Converting fabric interconnects to Cisco Intersight Managed Mode is a disruptive process, and configuration information will be lost. You are encouraged to make a backup of their existing configuration.

Step 1. Configure Fabric Interconnect A (FI-A). On the Basic System Configuration Dialog screen, set the management mode to Intersight. All the remaining settings are like those for the Cisco UCS Manager managed mode (UCSM-Managed).

Cisco UCS Fabric Interconnect A
To configure the Cisco UCS for use in a FlexPod environment in intersight managed mode, follow these steps:

Step 2. Connect to the console port on the first Cisco UCS fabric interconnect.

```
Enter the configuration method. (console/gui) ? console
Enter the management mode. (ucsm/intersight)? intersight
You have chosen to setup a new Fabric interconnect in "intersight" managed mode. Continue? (y/n): y
```

```

Enforce strong password? (y/n) [y]: Enter

Enter the password for "admin": <password>
Confirm the password for "admin": <password>

Enter the switch fabric (A/B) []: A

Enter the system name: <ucs-cluster-name>

Physical Switch Mgmt0 IP address : <ucsa-mgmt-ip>

Physical Switch Mgmt0 IPv4 netmask : <ucsa-mgmt-mask>

IPv4 address of the default gateway : <ucsa-mgmt-gateway>

Configure the DNS Server IP address? (yes/no) [n]: y

    DNS IP address : <dns-server-1-ip>

Configure the default domain name? (yes/no) [n]: y

    Default domain name : <ad-dns-domain-name>
<SNIP>

Verify and save the configuration.

```

Step 3. After applying the settings, make sure you can ping the fabric interconnect management IP address. When Fabric Interconnect A is correctly set up and is available, Fabric Interconnect B will automatically discover Fabric Interconnect A during its setup process as shown in the next step.

Step 4. Configure Fabric Interconnect B (FI-B). For the configuration method, choose console. Fabric Interconnect B will detect the presence of Fabric Interconnect A and will prompt you to enter the admin password for Fabric Interconnect A. Provide the management IP address for Fabric Interconnect B and apply the configuration.

```

Cisco UCS Fabric Interconnect A
Enter the configuration method. (console/gui) ? console

    Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added
to the cluster. Continue (y/n) ? y

Enter the admin password of the peer Fabric interconnect: <password>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucsa-mgmt-mask>

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address

Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>

Local fabric interconnect model(UCS-FI-6536)
Peer fabric interconnect is compatible with the local fabric interconnect. Continuing with the installer...

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

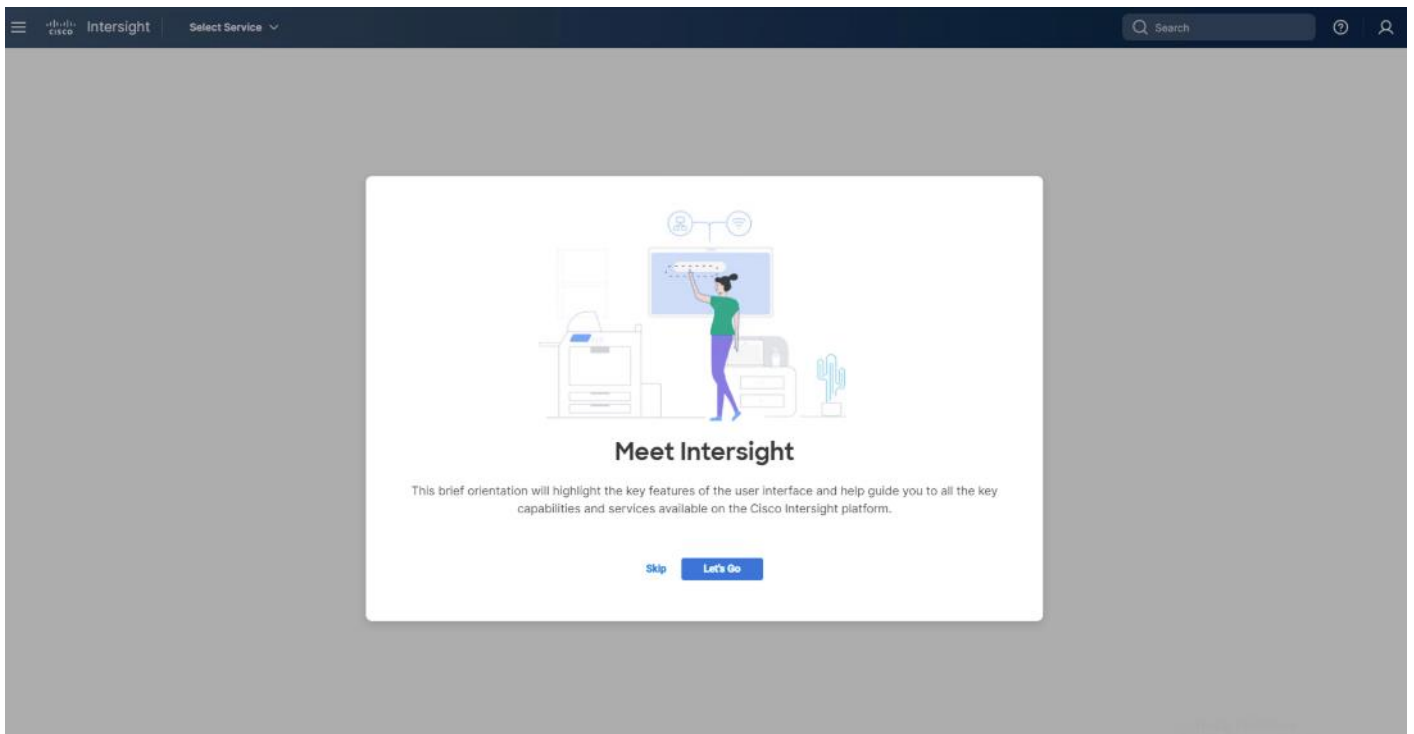
Procedure 2. Set up a new Cisco Intersight account

Step 1. Go to <https://intersight.com> and click **Create an account**.

Step 2. Read and accept the license agreement. Click **Next**.

Step 3. Provide an Account Name and click **Create**.

On successful creation of the Intersight account, following page will be displayed:



Note: You can also choose to add the Cisco UCS FIs to an existing Cisco Intersight account.

Procedure 3. Set up Cisco Intersight account and associate it with Cisco Smart Licensing

Note: When setting up a new Cisco Intersight account (as described in this document), the account needs to be enabled for Cisco Smart Software Licensing.

Step 1. Log into the **Cisco Smart Licensing portal**:

https://software.cisco.com/software/cs/ws/platform/home?locale=en_US#module/SmartLicensing.

Step 2. Verify that the correct virtual account is selected.

Step 3. Under **Inventory > General**, generate a new token for product registration.

Step 4. Copy this newly created token.


Create Registration Token

This will create a token that is used to register product instances, so that they can use licenses from this virtual account. Once it's created, go to the Smart Licensing configuration for your products and enter the token, to register them with this virtual account.


Virtual Account:	SJ-VDI-Lab
Description:	<input type="text" value="VDI Solutions Lab"/>
* Expire After:	<input type="text" value="30"/> Days
	<i>Between 1 - 365, 30 days recommended</i>
Max. Number of Uses:	<input type="text"/>
	<i>The token will be expired when either the expiration or the maximum uses is reached</i>

Allow export-controlled functionality on the products registered with this token ⓘ

Step 5. Log into the **Cisco Intersight portal** and click the **drop-down list**. Click **System**.



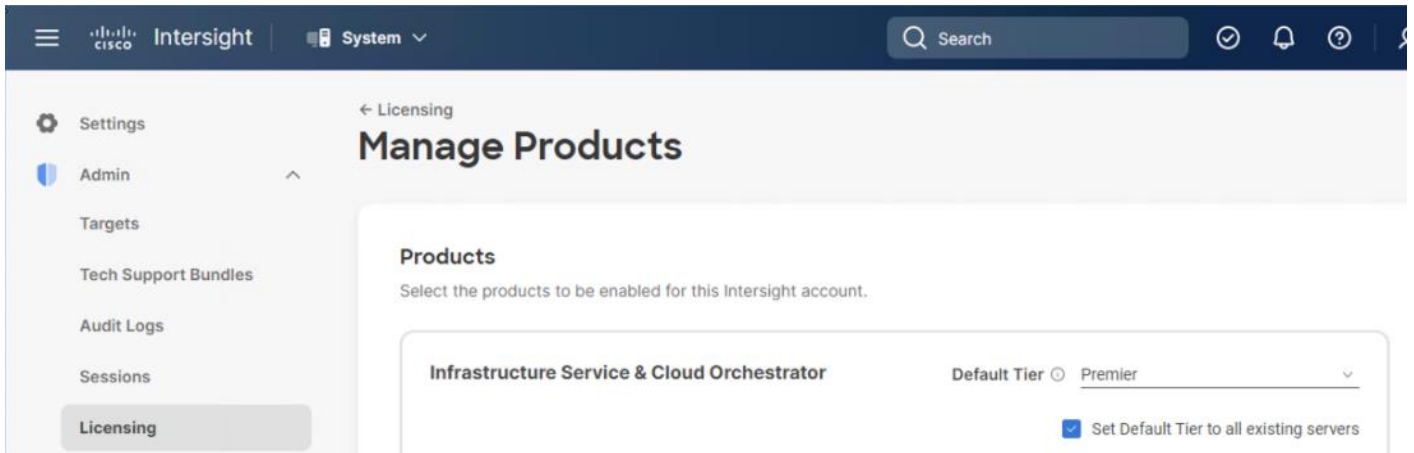
Step 6. Under **Cisco Intersight > Licensing**, click **Register**.



Step 7. Enter the copied token from the Cisco Smart Licensing portal. Click **Next**.

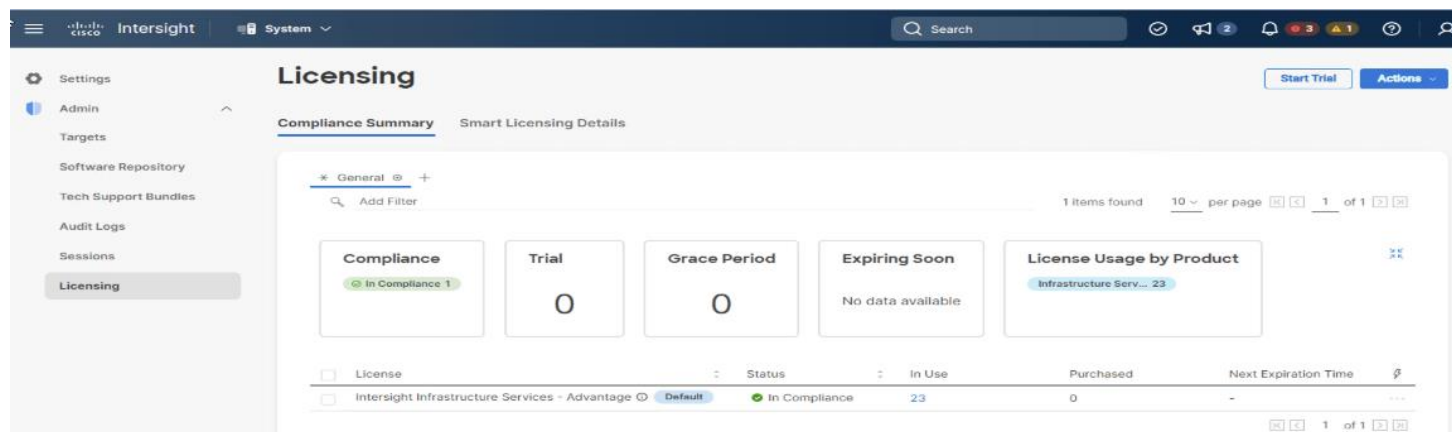
Step 8. From the drop-down list, select the pre-selected Default Tier * and select the license type (for example, Premier).

Step 9. Select **Move All Servers to Default Tier**.



Step 10. Click **Proceed**.

When the registration is successful (takes a few minutes), the information about the associated Cisco Smart account and default licensing tier selected in the last step is displayed.

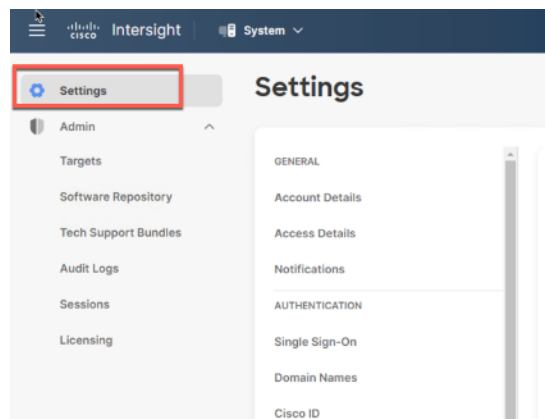


Procedure 4. Set up Cisco Intersight Resource Group

Note: In this procedure, a Cisco Intersight resource group is created where resources such as targets will be logically grouped. In this deployment, a single resource group is created to host all the resources, but you can choose to create multiple resource groups for granular control of the resources.


Step 1. Log into **Cisco Intersight**.

Step 2. Click **Settings** (the gear icon) and click **Settings**.




Step 3. Click **Resource Groups**.

Step 4. Click **+ Create Resource Group**.



Step 5. Provide a name for the Resource Group (for example, VDI-Lab).



Step 6. Under Memberships, click **Custom**.


Step 7. Click **Create**.

Procedure 5. Set up Cisco Intersight Organization

Note: In this procedure, a Cisco Intersight organization is created where all Cisco Intersight managed mode configurations including policies are defined.


Step 1. Log into the **Cisco Intersight portal**.

Step 2. Click **Settings** (the gear icon) and select **Settings**.



Step 3. Click **Organizations**.

Step 4. Click **+ Create Organization**.



Step 5. Provide a name for the organization (for example, VDI).

Step 6. Select the Resource Group created in the last step (for example, VDI Lab).

Step 7. Click **Create**.

← Organizations

Create Organization

Create Organization
Create an organization to manage and access the resources associated with Resource Groups.

General

Name *
VDI Description

Resource Groups

Select the Resource Groups to be associated with the Organization. Organization created will provide access to the resources in the selected Resource Groups.

2 items found 10 per page 1 of 1

Add Filter

<input type="checkbox"/>	Name	Used Organizations	Description
<input type="checkbox"/>	default	default	The Default Resource Group
<input type="checkbox"/>	X-Series	default, X-Series ... (2)	-

1 of 1

Cancel Create

Procedure 6. Claim Cisco UCS Fabric Interconnects in Cisco Intersight

Note: Make sure the initial configuration for the fabric interconnects has been completed. Log into Fabric Interconnect A using a web browser to capture the Cisco Intersight connectivity information.

Step 1. Use the management IP address of Fabric Interconnect A to access the device from a web browser and the previously configured admin password to Log into the device.


Step 2. Under **DEVICE CONNECTOR**, the current device status will show “Not claimed.” Note, or copy, the Device ID, and Claim Code information for claiming the device in Cisco Intersight.

Step 3. Log into **Cisco Intersight**.

Step 4. Click **Targets**.

Step 5. Click **Claim New Target**.

Step 6. Select **Cisco UCS Domain (Intersight Managed)** and click **Start**.



Step 7. Enter the Device ID and Claim Code captured from the Cisco UCS FI.

Step 8. Select the previously created Resource Group and click **Claim**.

System ▾ Search [icon] [icon] 2 [icon] 3 [icon] 1 [icon] [icon]

← Targets

Claim a New Target

Claim Cisco UCS Domain (Intersight Managed) Target

To claim your target, provide the Device ID, Claim Code and select the appropriate Resource Groups.

General

Device ID * Claim Code *

Resource Groups

• Select the Resource Groups if required. However, this selection is not mandatory as one or more Resource Group type is 'All'. The claimed target will be part of all Organizations with the Resource Group type 'All'.

1 items found 24 per page [icon] [icon] 1 of 1 [icon] [icon] [icon]

<input type="checkbox"/>	Name	Usage	Description
<input type="checkbox"/>	Series	default, :	...(2)

On successfully device claim, Cisco UCS FI appears as a target in Cisco Intersight.

Settings Admin **Targets** Claim a New T

<input type="checkbox"/>	Name	Status	Type	Claimed Time	Claimed By	[icon]
<input type="checkbox"/>	vdi-tme	Connected	Intersight Managed Dor	2021 10:53 AM		...
<input type="checkbox"/>	K6-FP-B	Connected	Intersight Managed Dor	2022 10:12 AM		...

Procedure 7. Verify the addition of Cisco UCS Fabric Interconnects to Cisco Intersight

Step 1. Log back into the web GUI of the Cisco UCS fabric interconnect and click **Refresh**.

The fabric interconnect status should now be set to **Claimed**.

Configure a Cisco UCS Domain Profile

This subject contains the following procedures:

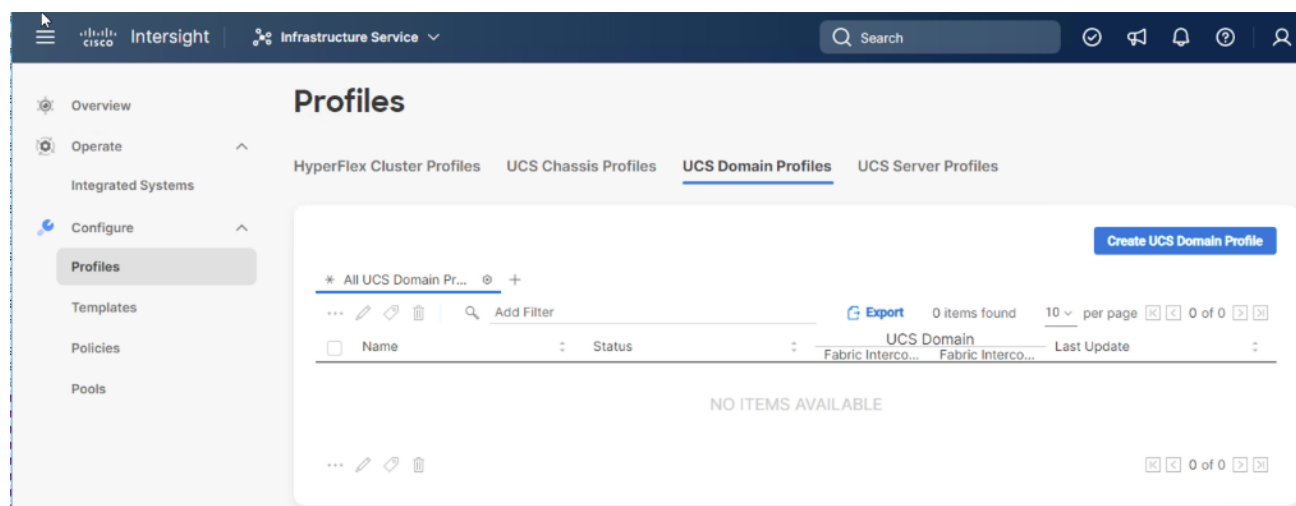
- [Create a Cisco UCS Domain Profile](#)
- [General Configuration](#)

- [Cisco UCS Domain Assignment](#)
- [Create and apply the VLAN Policy](#)
- [Configure the Ports on the Fabric Interconnects](#)

A Cisco UCS domain profile configures a fabric interconnect pair through reusable policies, allows configuration of the ports and port channels, and configures the VLANs in the network. It defines the characteristics of and configured ports on fabric interconnects. The domain-related policies can be attached to the profile either at the time of creation or later. One Cisco UCS domain profile can be assigned to one fabric interconnect domain.

Procedure 1. Create a Cisco UCS Domain Profile

- Step 1.** Log into the **Cisco Intersight** portal.
- Step 2.** From the drop-down list, click **Infrastructure Services**.
- Step 3.** Under **CONFIGURE** in the left pane and click **Profiles**.
- Step 4.** In the main window, click **UCS Domain Profiles** and click **Create UCS Domain Profile**.



Procedure 2. General Configuration

- Step 1.** From the drop-down list, select the organization (for example, VDI).
- Step 2.** Provide a name for the domain profile (for example, VDI-Domain-Profile).
- Step 3.** Provide an optional Description.



General

Add a name, description and tag for the UCS domain profile.

Organization *
VDI

Name *
VDI-Domain-Profile

Set Tags

Description
=< 1024

Step 4. Click Next.

Procedure 3. Cisco UCS Domain Assignment

Step 1. To assign the Cisco UCS domain to this new domain profile, click **Assign Now** and select the previously added Cisco UCS domain (for example, vdi-tme).

← Profiles

Create UCS Domain Profile

- General
- 2 UCS Domain Assignment**
- 3 VLAN & VSAN Configuration
- 4 Ports Configuration
- 5 UCS Domain Configuration
- 6 Summary

UCS Domain Assignment

Choose to assign a Fabric Interconnect pair to the profile now or later.

i Choose to assign a Fabric Interconnect pair now or later. If you choose Assign Now, select a pair that you want to assign and if you choose Assign Later, click Next to proceed to policy selection.

Show Assigned

Add Filter 3 Items found

Domain N...	Model	Fabric Interconnect A			
		Serial	Bundle Version	Model	
<input type="radio"/>					
<input type="radio"/>	FIA-K8-VDI	UCS-FI-6536	FDO27131330	4.2(3h)	UCS-FI-6536

Step 2. Click Next.

Procedure 4. Create and apply the VLAN Policy

Step 1. In this procedure, a single VLAN policy is created for both fabric interconnects

Step 2. Click **Select Policy** next to VLAN Configuration under Fabric Interconnect A.

VLAN & VSAN Configuration
Create or select a policy for the fabric interconnect pair.

^ **Fabric Interconnect A** 0 of 2 Policies Configured

VLAN Configuration [Select Policy](#)

VSAN Configuration [Select Policy](#)

^ **Fabric Interconnect B** 0 of 2 Policies Configured

VLAN Configuration [Select Policy](#)

VSAN Configuration [Select Policy](#)

Step 3. Click **Create New**.

Step 4. Verify correct organization is selected from the drop-down list (for example, default) and provide a name for the policy (for example, VDI-VLAN).

General

Add a name, description and tag for the policy.

Organization *

default

Name *

VDI-VLAN

Set Tags

Description

VDI VLAN Policy for both FIs

<= 1024

Step 5. Click **Next**.

Step 6. Click **Add VLANs**.

Step 7. Provide a name and VLAN ID for the native VLAN.

Add VLANs

Add VLANs to the policy

▲ VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix *

VLAN IDs *

Native-VLAN

2

Auto Allow On Uplinks

Enable VLAN Sharing

Multicast Policy *

[Select Policy](#)

Step 8. Make sure **Auto Allow On Uplinks** is enabled.

Step 9. To create the required Multicast policy, click **Select Policy** under Multicast*.

Step 10. In the window on the right, click **Create New** to create a new Multicast Policy.

Step 11. Provide a Name for the Multicast Policy (for example, VDI-MCAST-Pol).

Step 12. Provide optional Description and click **Next**.

Step 13. Leave the Snooping State selected and click **Create**.

↩

General

2 Policy Details

Policy Details

Add policy details

Multicast Policy

Snooping State

Querier State

Step 14. Click **Add** to add the VLAN.

Step 15. Select **Set Native VLAN ID** and enter the VLAN number (for example, 2) under VLAN ID.

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANS

Add VLANs

Show VLAN Ranges

2 items found 50 per page 1 of 1

<input type="checkbox"/>	VLAN ID	Name	Sharing Ty...	Primary VL...	Multicast Policy	Auto Allow On U...	
<input type="checkbox"/>	1	default	None			Yes	...
<input type="checkbox"/>	2	Native-VLAN_2	None		VDI-MCAST-Pol	Yes	...

Selected 1 of 2 [Show Selected](#) [Unselect All](#) 1 of 1

Set Native VLAN ID

VLAN ID

2

Step 16. Add the remaining VLANs for FlexPod by clicking Add VLANs and entering the VLANs one by one. Reuse the previously created multicast policy for all the VLANs.

The VLANs created during this validation are shown below:

6 items found 10 per page 1 of 1

<input type="checkbox"/>	VLAN ID	Name	Sharing Type	Primary VLAN ID	Multicast Policy	Auto Allow On Uplinks	
<input checked="" type="checkbox"/>	1	default	None			Yes	...
<input type="checkbox"/>	30	Mgmt_30	None		FP-Multicast	Yes	...
<input type="checkbox"/>	31	Inband Mgmt_31	None		FP-Multicast	Yes	...
<input type="checkbox"/>	32	CIFs_32	None		FP-Multicast	Yes	...
<input type="checkbox"/>	33	NFS_33	None		FP-Multicast	Yes	...
<input type="checkbox"/>	34	VM Public_34	None		FP-Multicast	Yes	...

Selected 1 of 6 [Show Selected](#) [Unselect All](#) 1 of 1

Step 17. Click **Create** to finish creating the VLAN policy and associated VLANs.

Step 18. Click **Select Policy** next to VLAN Configuration for Fabric Interconnect B and select the same VLAN policy.

Procedure 5. Configure the Ports on the Fabric Interconnects

Step 1. Click **Select Policy** for Fabric Interconnect A.

Step 2. Click **Create New** to define a new port configuration policy.

Note: Use two separate port policies for the fabric interconnects. Using separate policies provide flexibility when port configuration (port numbers or speed) differs between the two FIs.

Step 3. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-PortPol-A).

Step 4. Click **Next** to pass the Unified Port page.

Step 5. Click **Next** to pass the Breakout Options page.

Step 6. Select the ports that need to be configured as server ports by clicking the ports in the graphics (or select from the list below the graphic). When all ports are selected, click **Configure**.

Step 7. From the drop-down list, select **Server** as the role.

Configure (10 Ports)

Step 8. Click **Save**.

Step 9. Configure the Ethernet uplink port channel by selecting the Port Channel in the main pane and then clicking **Create Port Channel**.

Step 10. Select **Ethernet Uplink Port Channel** as the role, provide a port-channel ID (for example, 11), and select a value for Admin Speed from drop-down list (for example, Auto).

Note: You can create the Ethernet Network Group, Flow Control, Ling Aggregation or Link control policy for defining disjoint Layer-2 domain or fine tune port-channel parameters. These policies were not used in this deployment and system default values were utilized.

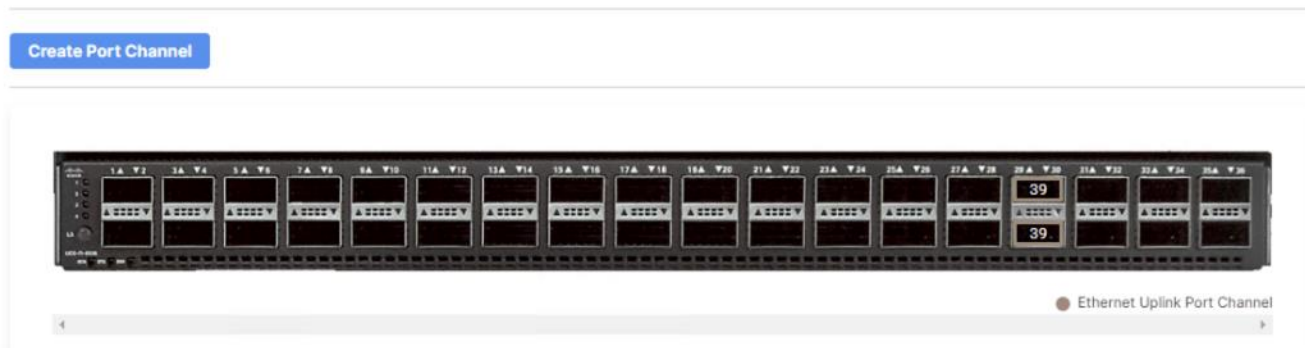
Step 11. Scroll down and select uplink ports from the list of available ports (for example, port 39)

Step 12. Click **Save**.

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

Port Roles Port Channels Pin Groups



Step 13. Click **Save** to create the port policy for Fabric Interconnect A.

Note: Use the summary screen to verify that the ports were selected and configured correctly.

Cisco UCS Domain Configuration

This subject contains the following procedures:

- [Configure NTP Policy for the Cisco UCS Domain](#)
- [Configure Network Connectivity Policy](#)
- [Configure System QoS Policy](#)
- [Verify Settings](#)
- [Deploy the Cisco UCS Domain Profile](#)
- [Verify Cisco UCS Domain Profile Deployment](#)

Note: Under UCS domain configuration, additional policies can be configured to setup NTP, Syslog, DNS settings, SNMP, QoS and UCS operating mode (end host or switch mode). For this deployment, three policies (NTP, Network Connectivity and System QoS) will be configured.

Procedure 1. Configure NTP Policy for the Cisco UCS Domain

Step 1. Click **Select Policy** next to NTP and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-NTPPol).

Step 3. Click **Next**.

Step 4. **Enable NTP**, provide the NTP server IP addresses, and select the **Timezone** from the drop-down list.

Step 5. If required, add a second NTP server by clicking **+** next to the first NTP server IP address.

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Domain

Enable NTP

NTP Servers *	
72.163.32.44	<input type="text"/>
+	

Timezone

America/Los_Angeles

Step 6. Click **Create**.

Procedure 2. Configure Network Connectivity Policy

Note: To define the Domain Name Service (DNS) servers for Cisco UCS, configure network connectivity policy.

Step 1. Click **Select Policy** next to Network Connectivity and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-NetConn-Pol).

Step 3. Provide DNS server IP addresses for Cisco UCS.

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Domain

Common Properties

IPv4 Properties

Preferred IPv4 DNS Server

10.81.72.40

Alternate IPv4 DNS Server

10.81.72.41

Enable IPv6

Step 4. Click **Create**.

Procedure 3. Configure System QoS Policy

To define the QoS settings for Cisco UCS, configure System QoS policy.

Step 1. Click **Select Policy** next to System QoS* and click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-QoSPol).

Step 3. Click **Next**.

Step 4. Change the MTU for Best Effort class to **9216**.

Step 5. Keep the default selections or change the parameters if necessary.

↑

Policy Details

Add policy details

- This policy is applicable only for UCS Domains

Configure Priorities

Platinum

Gold

Silver

Bronze

Best Effort

CoS	Weight	Allow Packet Drops	MTU
Any	5	<input checked="" type="checkbox"/>	9216
	0 - 10		1500 - 9216

Fibre Channel


CoS	Weight	Allow Packet Drops	MTU
3	5	<input type="checkbox"/>	2240
0 - 6	0 - 10		1500 - 9216

Step 6. Click **Create**.

Step 7. Click **Next**.

Procedure 4. Verify Settings

Step 1. Verify all the settings including the fabric interconnect settings, by expanding the settings and make sure that the configuration is correct.



Procedure 5. Deploy the Cisco UCS Domain Profile

Note: After verifying the domain profile configuration, deploy the Cisco UCS profile.

Step 1. From the UCS domain profile Summary view, click **Deploy**.

Step 2. Acknowledge any warnings and click **Deploy** again.

Step 3. The system will take some time to validate and configure the settings on the fabric interconnects. Log into the console servers to see when the Cisco UCS fabric interconnects have finished configuration and are successfully rebooted.

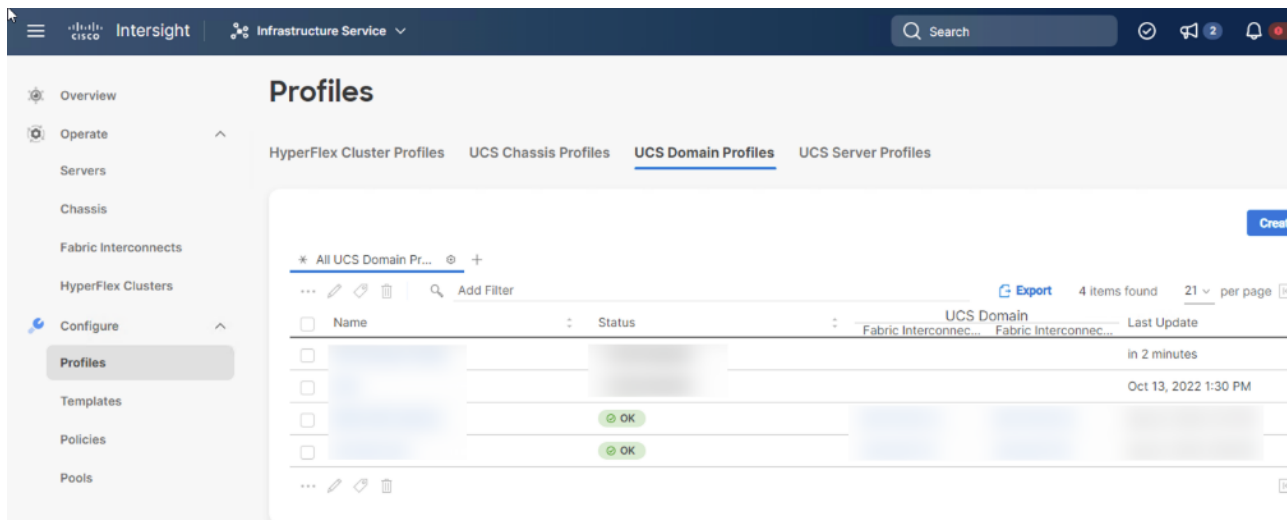
Procedure 6. Verify Cisco UCS Domain Profile Deployment

When the Cisco UCS domain profile has been successfully deployed, the Cisco UCS chassis and the Compute Nodes should be successfully discovered.

Note: It takes a while to discover the Compute Nodes for the first time. Watch the number of outstanding tasks in Cisco Intersight:



Step 1. Log into **Cisco Intersight**. Under **CONFIGURE > Profiles > UCS Domain Profiles**, verify that the domain profile has been successfully deployed.



Step 2. Verify that the servers have been successfully discovered and are visible under **OPERATE > Servers**.

Configure Server Profile Template

This subject contains the following procedures:

- [Configure a Server Profile Template](#)
- [Configure UUID Pool](#)
- [Configure BIOS Policy](#)
- [Local Boot Policy and vMedia policy for OS Install](#)

In the Cisco Intersight platform, a server profile enables resource management by simplifying policy alignment and server configuration. The server profiles are derived from a server profile template. The server profile template and its associated policies can be created using the server profile template wizard. After creating server profile template, you can derive multiple consistent server profiles from the template.

Procedure 1. Configure a Server Profile Template

Step 1. Log into **Cisco Intersight**.

Step 2. Go to **CONFIGURE > Templates** and click **Create UCS Server Profile Template**.

Step 3. Select the organization from the drop-down list (for example, VDI).

Step 4. Provide a name for the server profile template. The name used in this deployment is Local-Boot-Template.

Step 5. Click **UCS Server (FI-Attached)**.

Step 6. Provide an optional description.

General

Enter a name, description, tag and select a platform for the server profile template.

Organization *
default

Name *

Target Platform UCS Server (Standalone) UCS Server (FI-Attached)

Set Tags

Description
≤ 1024

Step 7. Click **Next**.

Procedure 2. Configure UUID Pool

Step 1. Click **Select Pool** under UUID Pool and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the UUID Pool (for example, VDI-UUID-Pool).

Step 3. Provide an optional Description and click **Next**.

Step 4. Provide a UUID Prefix (for example, a random prefix of 33FB3F9C-BF35-4BDE was used).

Step 5. Add a UUID block.

Configuration

Prefix *
33FB3F9C-BF35-4BDE

UUID Blocks

From	Size
0000-000A00000001	64

1 - 1024

Step 6. Click **Create**.

Procedure 3. Configure BIOS Policy

Step 1. Click **Select Policy** next to BIOS and click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-BIOSPol).

Step 3. Click **Next**.

Step 4. From the Policy Details screen, select appropriate values for the BIOS settings. In this deployment, the BIOS values were selected based on recommendations in the performance tuning guide for Cisco UCS M7 BIOS: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-Compute-Node-servers/performance-tuning-guide-ucs-M7-servers.html>.

Policy Details

Add policy details

▼ All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

▲ The BIOS settings will be applied only on next host reboot.

+ Boot Options

+ Intel Directed IO

+ LOM And PCIe Slots

+ Main

+ Memory

+ PCI

+ Power And Performance

+ Processor

+ QPI

Cancel

Back

Create

- LOM and PCIe Slot > CDN Support for LOM: Enabled
- Processor > Enhanced CPU performance: Auto
- Memory > NVM Performance Setting: Balanced Profile

Step 5. Click **Create**.

Step 6. Click **Next** to move to Management Configuration.

Procedure 4. Local Boot Policy and vMedia policy for OS Install

Step 1. Click **Select Policy** next to Boot Order and click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-LocalBoot or M7-Local).

Step 3. Select **UEFI** for Boot Mode.

Step 4. Enable the **Enable Secure Boot** toggle button.

Step 5. Use the drop down to add two Boot Devices (Local Disk and Virtual Media).

Step 6. Under Local Disk label it **M2** and type **MSTOR-RAID** by Slot.

Step 7. Under Virtual Media name it **dvd** and select **KVM MAPPED DVD** for the Sub-Type.

Step 8. Click **Next**.

Step 9. Click **Next** to go to Management Configuration.

Management Configuration

This section contains the following procedures:

- [Configure Cisco IMC Access Policy](#)
- [Configure IPMI Over LAN Policy](#)
- [Configure Local User Policy](#)
- [Storage Configuration](#)
- [Create MAC Address Pool for Fabric A and B](#)
- [Create Ethernet Network Group Policy for a vNIC](#)
- [Create Ethernet Network Control Policy](#)
- [Create Ethernet QoS Policy](#)
- [Create Ethernet Adapter Policy](#)
- [Verify Summary](#)
- [Derive Server Profile](#)

Three policies will be added to the management configuration:

- IMC Access to define the pool of IP addresses for compute node KVM access
- IPMI Over LAN to allow Intersight to manage IPMI messages
- Local User to provide local administrator to access KVM

Procedure 1. Configure Cisco IMC Access Policy

Step 1. Click **Select Policy** and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-IMC-Access).

Step 3. Click **Next**.

Note: You can select in-band management access to the compute node using an in-band management VLAN (for example, VLAN 30) or out-of-band management access via the Mgmt0 interfaces of the FIs. KVM Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Since these policies were not configured in this deployment, out-of-band management access was configured so that KVM access to compute nodes is not impacted by any potential switching issues in the fabric.

Step 4. Click **UCS Server (FI-Attached)**.

Step 5. **Enable** Out-Of-Band Configuration.

Step 6. Under IP Pool, click **Select IP Pool** and then click **Create New**.

Policy Details

Add policy details

All Platforms **UCS Server (FI-Attached)** UCS Chassis

• A minimum of one configuration must be enabled. Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Check here for more info, [Help Centre](#)

In-Band Configuration

Enabled

Out-Of-Band Configuration

Enabled

IP Pool *

[Select IP Pool](#)

Step 7. Verify the correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-IMC-OOB-Pool).

Step 8. Select **Configure IPv4 Pool** and provide the information to define a pool for KVM IP address assignment including an IP Block.

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Configuration

Netmask *	Gateway
255.255.255.0	19.81.72.254
Primary DNS	Secondary DNS
10.81.72.40	10.81.72.41

IP Blocks

From	Size	
10.81.72.150	16	1 - 1024

Note: The management IP pool subnet should be accessible from the host that is trying to open the KVM connection. In the example shown here, the hosts trying to open a KVM connection would need to be able to route to 10.81.72.0/24 subnet.

Step 9. Click **Next**.

Step 10. Deselect **Configure IPv6 Pool**.

Step 11. Click **Create** to finish configuring the IP address pool.

Step 12. Click **Create** to finish configuring the IMC access policy.

Procedure 2. Configure IPMI Over LAN Policy

Step 1. Click **Select Policy** and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, Enable-IPMIoLAN).

Step 3. Turn on **Enable IPMI Over LAN**.

Step 4. Click **Create**.

↕

Policy Details

Add policy details



All Platforms

UCS Server (Standalone)

UCS Server (FI-Attached)



Enable IPMI Over LAN

Procedure 3. Configure Local User Policy

Step 1. Click **Select Policy** and click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-LocalUser-Pol).

Step 3. Verify that **UCS Server (FI-Attached)** is selected.

Step 4. Verify that **Enforce Strong Password** is selected.

Policy Details

Add policy details

 All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

Password Properties

Enforce Strong Password 

Enable Password Expiry 

Password History

5   
0 - 5

Always Send User Password 

Local Users



- This policy will remove existing user accounts other than the ones configured with this policy. However, the default admin user account is not deleted from the endpoint device. You can only enable/disable or change account password for the admin account by creating a user with the user name and role as 'admin'. If there are no users in the policy, only the admin user account will be available on the endpoint device. By default, IPMI support is enabled for all users








[Add New User](#)

Step 5. Click **Add New User** and then click **+** next to the New User

Step 6. Provide the username (for example, fpadmin), choose a role for example, admin), and provide a password.

[Add New User](#)

— fpadmin (admin)  Enable 

Username *	Role
<input type="text" value="fpadmin"/> 	<input type="text" value="admin"/>  
Password *	Password Confirmation *
<input type="password" value="*****"/>  	<input type="password" value="*****"/>  

Note: The username and password combination defined here will be used to log into KVMs. The typical Cisco UCS admin username and password combination cannot be used for KVM access.

Step 7. Click **Create** to finish configuring the user.

Step 8. Click **Create** to finish configuring local user policy.

Step 9. Click **Next** to move to Storage Configuration.

Procedure 4. Storage Configuration

Note: Because we configured M2 local storage in the “Boot Order” section, we did not need to configure a SAN Connectivity policy. Local datastores will be file based using NFS and the OS will be local to M2 cards.

Procedure 5. Create MAC Address Pool for Fabric A and B

When creating the first vNIC, the MAC address pool has not been defined yet, therefore a new MAC address pool will need to be created. Two separate MAC address pools are configured for each Fabric. FP-MAC-A will be reused for all Fabric-A vNICs, and FP-MAC-B will be reused for all Fabric-B vNICs.

Table 16. MAC Address Pools

Pool Name	Starting MAC Address	Size	vNICs
FP-MAC-Pool-A	00:25:B5:17:0A:00	64*	eth0-eth2
FP-MAC-Pool-B	00:25:B5:17:0B:00	64*	eth1, eth3

Note: Each server requires 2 MAC addresses from the pool. Adjust the size of the pool according to your requirements.

Step 1. Click **Select Pool** and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the pool from [Table 16](#) depending on the vNIC being created (for example, FP-MAC-A for Fabric A).

Step 3. Click **Next**.

Step 4. Provide the starting MAC address from [Table 16](#) (for example, 00:25:B5:17:0A:00).

Note: For troubleshooting FlexPod, some additional information is always coded into the MAC address pool. For example, in the starting address 00:25:B5:17:0A:00, 17 is the rack ID and 0A indicates Fabric A.

Step 5. Provide the size of the MAC address pool from [Table 16](#) (for example, 64).

Pool Details

Collection of MAC Blocks.



Step 6. Click **Create** to finish creating the MAC address pool.

Step 7. From the Add vNIC window, provide vNIC Name, Slot ID, Switch ID, and PCI Order information from [Table 15](#).

Name *
 01-vSwitch0-A ⊙ Pin Group Name ⌵ ⊙

MAC

Pool Static

MAC Pool * ⊙

Selected Pool VDI-Mac | × | |

Placement

Simple Advanced

Slot ID *
 MLOM ⊙ PCI Link
 0 ⌵ ⊙
 0 - 1

Switch ID *
 A ⌵ ⊙

PCI Order
 2 ⌵ ⊙

Step 8. For Consistent Device Naming (CDN), from the drop-down list select **vNIC Name**.

Step 9. Verify that Failover is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and VDS.

Consistent Device Naming (CDN)

Source
 vNIC Name ⌵ ⊙

Failover

Enabled ⊙

Procedure 6. Create Ethernet Network Group Policy for a vNIC

The Ethernet Network Group policies will be created and reused on applicable vNICs as explained below. The Ethernet Network Group policy defines the VLANs allowed for a particular vNIC, therefore multiple network group policies will be defined for this deployment as follows:

Table 17. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs	VLANs
VDI-VDS0-NetGrp	Native-VLAN (2)	03-VDS0-A, 04-VDS0-B	VM Traffic, vMotion

Step 1. Click **Select Policy** under Ethernet Network Group Policy and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy from [Table 17](#) (for example, VDI-vSwitch0-NetGrp).

Step 3. Click **Next**.

Step 4. Enter the allowed VLANs from [Table 6](#) (for example, 30-36) and the native VLAN ID from [Table 6](#) (for example, 2).

Policy Details
Add policy details

VLAN Settings

Allowed VLANs	Native VLAN
30-36	2

1 - 4093

Step 5. Click **Create** to finish configuring the Ethernet network group policy.

Note: When ethernet group policies are shared between two vNICs, the ethernet group policy only needs to be defined for the first vNIC. For subsequent vNIC policy mapping, click Select Policy and select the previously defined ethernet group policy from the list.

Procedure 7. Create Ethernet Network Control Policy

The Ethernet Network Control Policy is used to enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) for the vNICs. A single policy will be created and reused for all the vNICs.

Step 1. Click **Select Policy** under Ethernet Network Control Policy and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-Enable-CDP-LLDP).

Step 3. Click **Next**.

Step 4. **Enable Cisco Discovery Protocol** and both **Enable Transmit** and **Enable Receive** under LLDP.

This policy is applicable only for UCS Servers (FI-Attached)


Enable CDP 

Mac Register Mode 

Only Native VLAN All Host VLANs

Action on Uplink Fail 

Link Down Warning

 Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost.

MAC Security

Forge 

Allow Deny

LLDP

Enable Transmit 

Enable Receive 

Step 5. Click **Create** to finish creating Ethernet network control policy.

Procedure 8. Create Ethernet QoS Policy

Ethernet QoS policy is used to enable jumbo maximum transmission units (MTUs) for all the vNICs. A single policy will be created and reused for all the vNICs.

Step 1. Click **Select Policy** under Ethernet QoS and click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-EthQos-Pol).

Step 3. Click **Next**.

Step 4. Change the MTU, Bytes value to **9000**.

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

QoS Settings

MTU, Bytes

9000

1500 - 9000

Rate Limit, Mbps

0

0 - 100000

Burst

10240

1 - 1000000

Priority

Best-effort

Enable Trust Host CoS

Step 5. Click **Create** to finish setting up the Ethernet QoS policy.

Procedure 9. Create Ethernet Adapter Policy

Ethernet adapter policy is used to set the interrupts and the send and receive queues. The values are set according to the best-practices guidance for the operating system in use. Cisco Intersight provides default VMware Ethernet Adapter policy for typical VMware deployments.

Optionally, you can configure a tweaked ethernet adapter policy for additional hardware receive queues handled by multiple CPUs in scenarios where there is a lot of vMotion traffic and multiple flows. In this deployment, a modified ethernet adapter policy, VDI-VMware-High-Traffic, is created and attached to the 03-VDS0-A and 04-VDS0-B interfaces which handle vMotion.

Table 18. Ethernet Adapter Policy association to vNICs

Policy Name	vNICs
VDI-EthAdapter-VMware	VDI-VDS0-NetGrp
VDI-VMware-High-Traffic	NFS, vMotion

Step 1. Click **Select Policy** under Ethernet Adapter and then click **Create New**.

Step 2. Verify correct organization is selected from the drop-down list (for example, VDI) and provide a name for the policy (for example, VDI-EthAdapter-VMware).

Step 3. Click **Select Default Configuration** under Ethernet Adapter Default Configuration.

General

Add a name, description and tag for the policy.

Organization *

default ▼

Name *

EA

Set Tags

Description ⌘

<= 1024

Ethernet Adapter Default Configuration

Select Default Configuration 

Step 4. From the list, select **VMware**.

Step 5. Click **Next**.

Step 6. For the VDI-EthAdapter-VMware policy, click **Create** and skip the rest of the steps in this section.

Step 7. For the optional VDI-VMware-High-Traffic policy (for VDS interfaces), make the following modifications to the policy:

- Increase Interrupts to 11
- Increase Receive Queue Count to 8
- Increase Completion Queue Count to 9
- Enable Receive Side Scaling

Interrupt Settings

Interrupts: 11 (1 - 1024)

Interrupt Mode: MSix

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type: Min

Receive

Receive Queue Count: 8 (1 - 1000)

Receive Ring Size: 4096 (64 - 16384)

Transmit

Transmit Queue Count: 1 (1 - 1000)

Transmit Ring Size: 256 (64 - 16384)

Completion

Completion Queue Count: 9 (1 - 2000)

Completion Ring Size: 1 (1 - 256)

Uplink Failback Timeout (seconds): 5 (0 - 600)

TCP Offload

- Enable Tx Checksum Offload
- Enable Rx Checksum Offload
- Enable Large Send Offload
- Enable Large Receive Offload

Receive Side Scaling

- Enable Receive Side Scaling
- Enable IPv4 Hash
- Enable IPv6 Extensions Hash
- Enable IPv6 Hash
- Enable TCP and IPv4 Hash
- Enable TCP and IPv6 Extensions Hash
- Enable TCP and IPv6 Hash
- Enable UDP and IPv4 Hash
- Enable UDP and IPv6 Hash

- Step 8.** Click **Create**.
- Step 9.** Click **Create** to finish creating the vNIC.
- Step 10.** Go back to [Step 1](#) and repeat vNIC creation for all four vNICs.
- Step 11.** Verify all four vNICs were successfully created.
- Step 12.** Click **Create** to finish creating the LAN Connectivity policy for hosts.

Procedure 10. Verify Summary

Step 1. When the LAN connectivity policy is created, click **Next** to move to the Summary screen.

Step 2. On the summary screen, verify policies mapped to various settings. The screenshots below provide summary view for a Local M.2 boot server profile template.

Summary

Verify details of the template and the policies, resolve errors and deploy.

^ General

Template Name
m7Tpl

Organization
default

Target Platform
UCS Server (FI-Attached)

Compute Configuration Management Configuration Storage Configuration Network Configuration Errors/Warnings (0)

BIOS	Cisco-UCS-BIOS-M7
Boot Order	m7-local
Firmware	M7-FW
Power	m7-PowerPolicy
UUID	FP-UUID

Compute Configuration Management Configuration Storage Configuration Network Configuration Errors/Warnings (0)

IMC Access	vdi-5108-IMC
IPMI Over LAN	IPMIEnable
Local User	FP-LocalUser

Compute Configuration Management Configuration Storage Configuration Network Configuration Errors/Warnings (0)

LAN Connectivity	FP-LAN-Connect
SAN Connectivity	FP-SAN-Connect

Procedure 11. Derive Server Profile

Step 1. From the Server profile template Summary screen, click **Derive Profiles**.

Note: This action can also be performed later by navigating to Templates, clicking “...” next to the template name and selecting **Derive Profiles**.

Step 2. Under the Server Assignment, select **Assign Now** and click **Cisco UCS X210c M7**. You can select one or more servers depending on the number of profiles to be deployed.

Step 3. Click **Next**.

Note: Cisco Intersight will fill the default information for the number of servers selected.

Step 4. Adjust the Prefix and number if needed.

Step 5. Click **Next**.

Step 6. Verify the information and click **Derive** to create the Server Profiles.

Step 7. Cisco Intersight will start configuring the server profiles and will take some time to apply all the policies. Use the Requests tab to see the progress.



Step 8. When the Server Profiles are deployed successfully, they will appear under the Server Profiles with the status of OK.

Profiles

HyperFlex Cluster Profiles UCS Chassis Profiles UCS Domain Profiles UCS Server Profiles

* All UCS Server Prof... +

Search k8 x Add Filter

Status	Inconsistency Reason	Target Platform	UCS Server Template
<input type="checkbox"/>			
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl
<input type="checkbox"/>		UCS Server (FI-Attached)	m7Tmpl

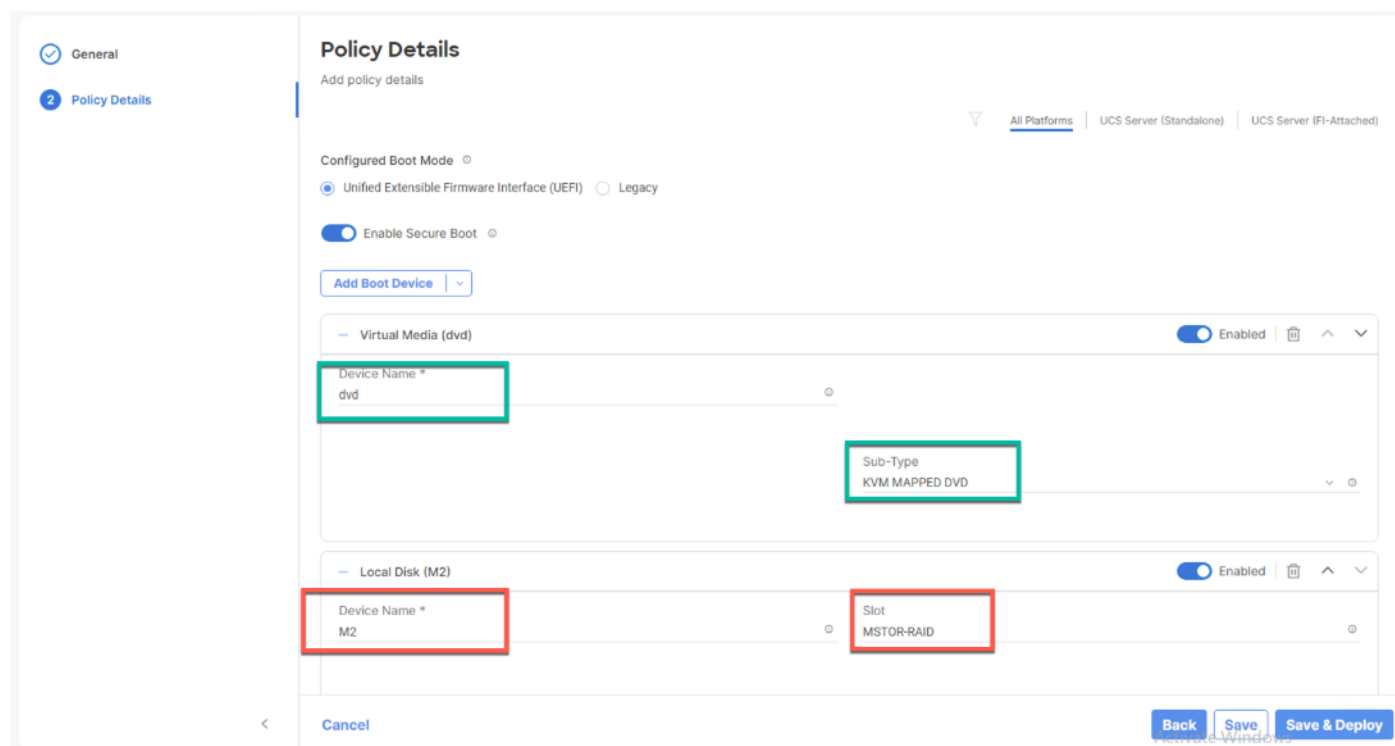
Install VMware ESXi 8.0 U1

This chapter contains the following:

- [Configure Local Boot to M.2](#)
- [Install VMware ESXi 8.0 U1](#)
- [VMware vCenter 8.0 U1](#)

Configure Local Boot to M.2

In this CVD we installed the ESXi operating system onto locally installed M2 cards with a vMedia policy.



Install VMware ESXi 8.0 U1

This section contains the following procedures:

- [Download ESXi 8.0 U1 from VMware](#)
- [Log into the Cisco UCS Environment using Cisco Intersight](#)
- [Prepare the Server for the OS Installation](#)
- [Install VMware ESXi to the Bootable M.2 card of the Hosts](#)
- [Set Up Management Networking for ESXi Hosts](#)
- [Install VMware and Cisco VIC Drivers for the ESXi Host](#)
- [Install VMware VIC Drivers and the NetApp NFS Plug-in for VMware VAAI](#)

This section provides detailed instructions for installing VMware ESXi 8.0 U1 in a FlexPod environment. After the procedures are completed, three booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco Intersight to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Procedure 1. Download ESXi 8.0 U1 from VMware

Step 1. Click the following link: [Cisco Custom ISO for UCS 4.3.1a](https://customerconnect.vmware.com/downloads/details?downloadGroup=OEM-ESXI80U1-CISCO&productId=1345). You will need a user id and password on vmware.com to download this software, <https://customerconnect.vmware.com/downloads/details?downloadGroup=OEM-ESXI80U1-CISCO&productId=1345>.

Note: The Cisco Custom ISO for UCS 4.3.1a should also be used for Cisco UCS software release 4.2.(3g) and VMware vSphere 8.0 U1.

Step 2. Download the .iso file.

Procedure 2. Log into the Cisco UCS Environment using Cisco Intersight

The Cisco UCS IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log into the Cisco UCS environment to run the IP KVM.

Step 1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco Intersight application.

Step 2. Click the **Launch Intersight** link to launch the HTML 5 Intersight GUI.

Step 3. If prompted to accept security certificates, accept, as necessary.

Step 4. When prompted, enter admin for the user name and enter the administrative password.

Step 5. To log into Cisco Intersight, click **Login**.

Step 6. From the main menu, click **Servers**.

Step 7. Click **Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-01**.

Step 8. In the Actions pane, click **KVM Console**.

Step 9. Follow the prompts to launch the HTML5 KVM console.

Step 10. Click **Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-02**.

Step 11. In the Actions pane, click **KVM Console**.

Step 12. Follow the prompts to launch the HTML5 KVM console.

Step 13. Go to **Servers > Service Profiles > root > Sub-Organizations > FlexPod Organization > VM-Host-Infra-03**.

Step 14. In the Actions pane, click **KVM Console**.

Step 15. Follow the prompts to launch the HTML5 KVM console.

Procedure 3. Prepare the Server for the OS Installation

Note: Skip this section if you're using vMedia policies; the ISO file will already be connected to KVM.

Step 1. In the KVM window, click **Virtual Media**.

Step 2. Select **Activate Virtual Devices**.

Step 3. If prompted to accept an Unencrypted KVM session, accept, as necessary.

Step 4. Click **Virtual Media** and select **Map CD/DVD**.

Step 5. Browse to the ESXi installer ISO image file and click **Open**.

Step 6. Click **Map Device**.

Step 7. Click the **KVM Console** tab to monitor the server boot.

Procedure 4. Install VMware ESXi to the Bootable M.2 card of the Hosts

Step 1. Boot the server by selecting **Boot Server** in the KVM and click **OK**, then click **OK** again.

On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.

Step 2. If the ESXi installer fails to load because the software certificates cannot be validated, reset the server, and when prompted, press **F2** to go into BIOS and set the system time and date to current. Now the ESXi installer should load properly.

Step 3. After the installer is finished loading, press **Enter** to continue with the installation.

Step 4. Read and accept the end-user license agreement (EULA). Press **F11** to accept and continue.

Note: It may be necessary to map function keys as User Defined Macros under the Macros menu in the Cisco UCS KVM console.

Step 5. Select the M.2 card that was previously set up for the installation disk for ESXi and press **Enter** to continue with the installation.

Step 6. Select the appropriate keyboard layout and press **Enter**.

Step 7. Enter and confirm the root password and press **Enter**.

Step 8. The installer issues a warning that the selected disk will be repartitioned. Press **F11** to continue with the installation.

Step 9. After the installation is complete, press **Enter** to reboot the server.

Note: The ESXi installation image will be automatically unmapped in the KVM when Enter is pressed.

Step 10. In Cisco Intersight, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

Procedure 5. Set Up Management Networking for ESXi Hosts

Note: Adding a management network for each VMware host is necessary for managing the host.

Step 1. After the server has finished rebooting, in the UCS KVM console, press **F2** to customize VMware ESXi.

Step 2. Log in as **root**, enter the corresponding password, and press **Enter** to log in.

Step 3. Use the down arrow key to select **Troubleshooting Options** and press **Enter**.

Step 4. Select **Enable ESXi Shell** and press **Enter**.

Step 5. Select **Enable SSH** and press **Enter**.


Step 6. Press **Esc** to exit the Troubleshooting Options menu.

Step 7. Select the **Configure Management Network** option and press **Enter**.

Step 8. Select **Network Adapters** and press **Enter**.

Step 9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field. If the numbers do not match, note the mapping of vmnic ports to vNIC ports for later use.

Step 10. Using the spacebar, select **vmnic1**.



Note: In lab testing, examples were seen where the vmnic and device ordering do not match. In this case, use the Consistent Device Naming (CDN) to note which vmnics are mapped to which vNICs and adjust the upcoming procedure accordingly.

Step 11. Press **Enter**.

Step 12. Select the **VLAN (Optional)** option and press **Enter**.

Step 13. Enter the <ib-mgmt-vlan-id> and press **Enter**.

Step 14. Select **IPv4 Configuration** and press **Enter**.

Step 15. Select the “**Set static IPv4 address and network configuration**” option by using the arrow keys and space bar.

Step 16. Move to the IPv4 Address field and enter the IP address for managing the ESXi host.

Step 17. Move to the Subnet Mask field and enter the subnet mask for the ESXi host.

Step 18. Move to the Default Gateway field and enter the default gateway for the ESXi host.

Step 19. Press **Enter** to accept the changes to the IP configuration.

Step 20. Select the **IPv6 Configuration** option and press **Enter**.

Step 21. Using the spacebar, select **Disable IPv6 (restart required)** and press **Enter**.

Step 22. Select the **DNS Configuration** option and press **Enter**.

Note: Since the IP address is assigned manually, the DNS information must also be entered manually.

Step 23. Using the spacebar, select “**Use the following DNS server addresses and hostname.**”

Step 24. Move to the Primary DNS Server field and enter the IP address of the primary DNS server.

Step 25. Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.

Step 26. Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.

Step 27. Press **Enter** to accept the changes to the DNS configuration.

Step 28. Press **Esc** to exit the Configure Management Network submenu.

Step 29. Press **Y** to confirm the changes and reboot the ESXi host.

Procedure 6. Install VMware and Cisco VIC Drivers for the ESXi Host

Step 1. Download the offline bundle for the Cisco UCS Tools Component and the NetApp NFS Plug-in for VMware VAAI to the Management workstation: [Cisco UCS Tools Component for ESXi 8.0 U1 1.3.1](#) (ucs-tool-esxi_1.3.1-1OEM.zip) (NetAppNasPluginV2.0.1.zip)

Note: This document describes using the driver versions shown above along with Cisco VIC nenic version 2.0.10.0 along with VMware vSphere version 8.0 U1.U3, Cisco UCS version 5.2(0.230041), and the latest patch NetApp ONTAP 9.13.1P3. These were the versions validated and supported at the time this document was published. This document can be used as a guide for configuring future versions of software. Consult the [Cisco UCS Hardware Compatibility List](#) and the [NetApp Interoperability Matrix Tool](#) to determine supported combinations of firmware and software.

Procedure 7. Install VMware VIC Drivers and the NetApp NFS Plug-in for VMware VAAI on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02

Step 1. Using an SCP program such as WinSCP, copy the two offline bundles referenced above to the /tmp directory on each ESXi host.

Step 2. Using a ssh tool such as PuTTY, ssh to each VMware ESXi host. Log in as root with the root password.

Step 3. Type `cd /tmp`.

Step 4. Run the following commands on each host:

```
esxcli software component apply -d /tmp/ucs-tool-esxi_1.3.1-1OEM.zip

esxcli software vib install -d /tmp/NetAppNasPlugin2.0.1.zip

reboot
```

Step 5. After reboot, log back into each host and run the following commands and ensure the correct version is installed:

```
esxcli software component list | grep ucs

esxcli software vib list | grep NetApp
```

VMware vCenter 8.0 U1

This subject contains the following:

- [Build the VMware vCenter Server Appliance](#)
- [Adjust vCenter CPU Settings](#)
- [Set up VMware vCenter Server](#)

Procedure 1. Build the VMware vCenter Server Appliance

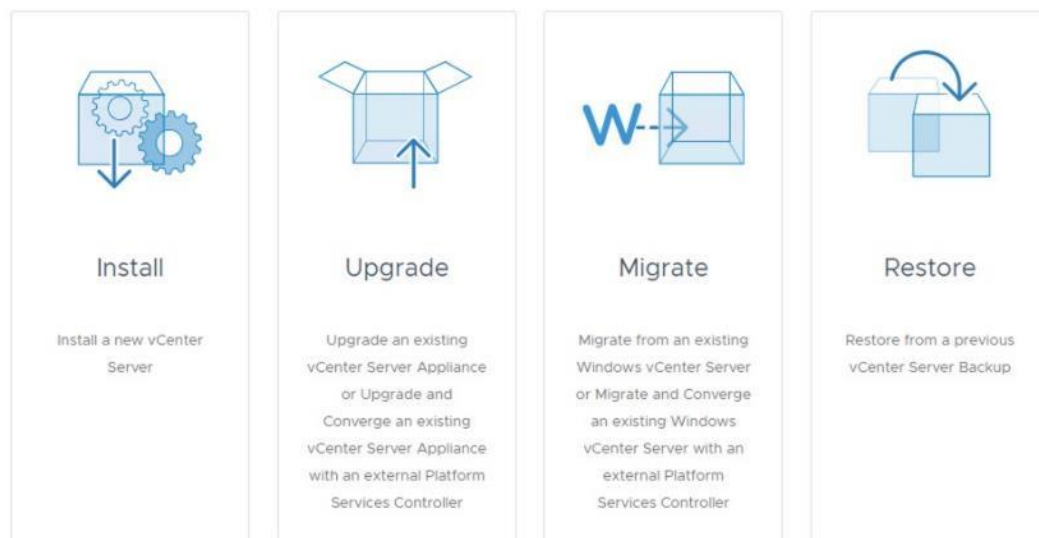
Note: The VCSA deployment consists of 2 stages: install and configuration.

Step 1. Locate and copy the **VMware-VCSA-all-8.0.1.00300-22088981.iso** file to the desktop of the management workstation. This ISO is for the VMware vSphere 8.0 U1 vCenter Server Appliance.

Note: It is important to use at a minimum VMware vCenter release 8.0 U1 to ensure access to all needed features.

Step 2. Using ISO mounting software, mount the ISO image as a disk on the management workstation.

Step 3. In the mounted disk directory, navigate to the **vcsa-ui-installer > win32 directory** and double-click `installer.exe`. The vCenter Server Appliance Installer wizard appears.




Step 4. Click **Install** to start the vCenter Server Appliance deployment wizard.

Step 5. Click **NEXT** in the Introduction section.

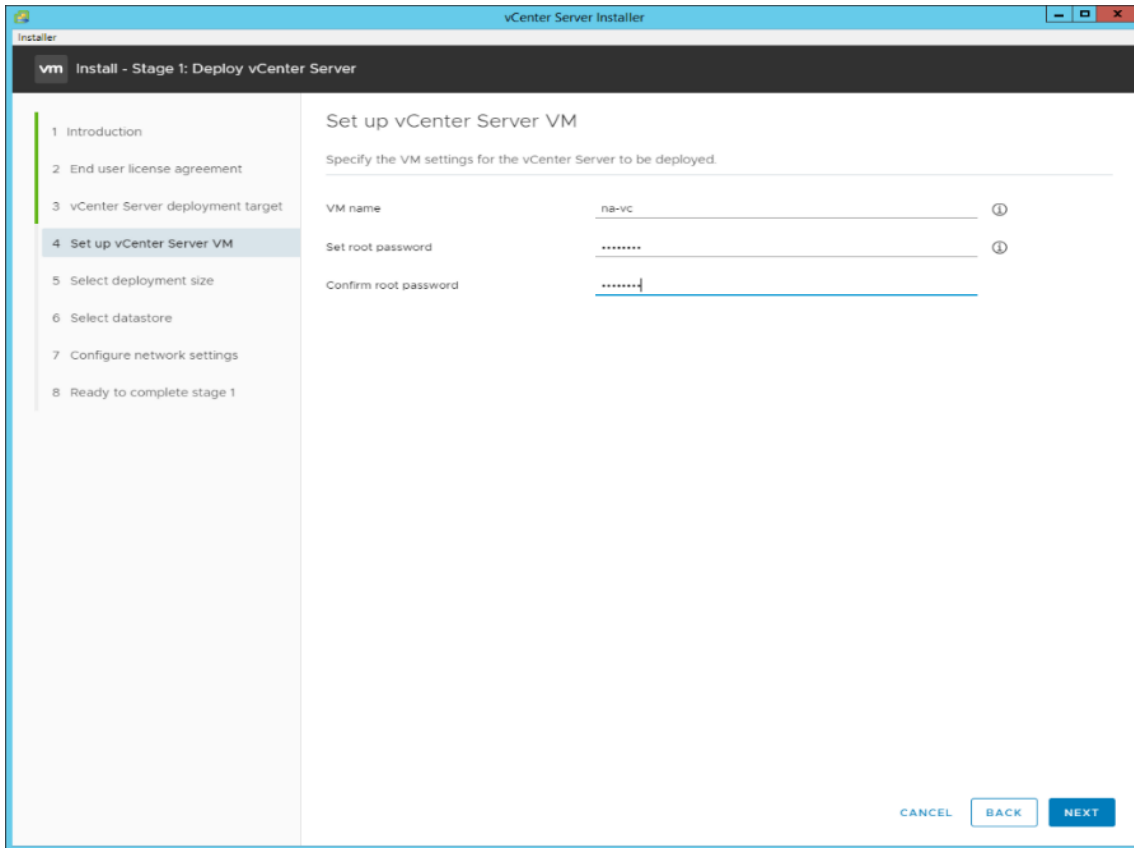
Step 6. Read and accept the license agreement and click **NEXT**.

Step 7. In the “vCenter Server deployment target” window, enter the host name or IP address of the first ESXi host, Username (root) and Password. Click **NEXT**.

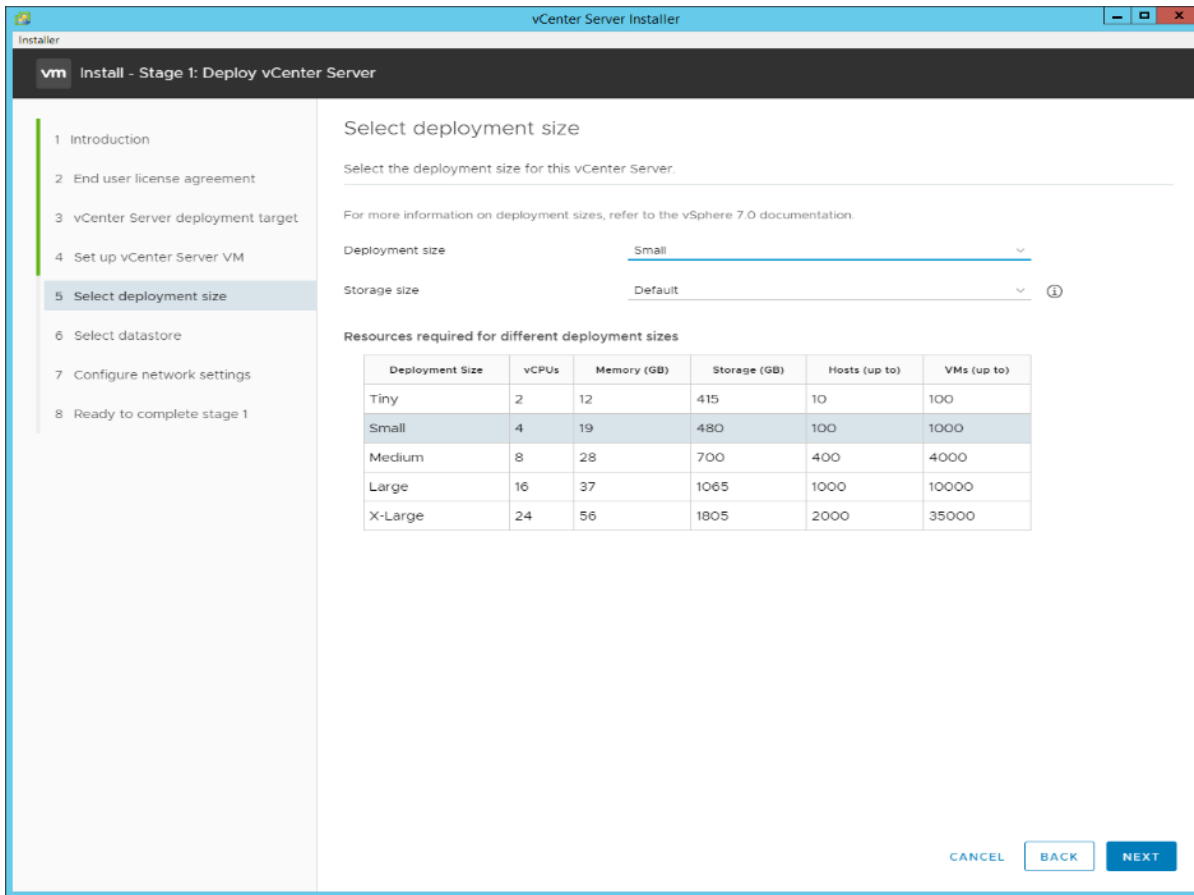


Step 8. Click **YES** to accept the certificate.


Step 9. Enter the Appliance VM name and password details in the Set up vCenter Server VM section. Click **NEXT**.



Step 10. In the Select deployment size section, click the **Deployment size** and **Storage size**. For example, click **Small** and **Default**. Click **NEXT**.



Step 11. Click **infra_datastore** for storage. Click **NEXT**.




Step 12. In the Network Settings section, configure the following settings:

- a. Click a Network: **IB-MGMT Network**.

Note: It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it not get moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, and it is attempted to bring up vCenter on a different host than the one it was running on before the shutdown, vCenter will not have a functional network connection. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. Moving vCenter to a different host on vSwitch0 to be brought up always occurs correctly without requiring vCenter to already be up and running.


- b. IP version: **IPV4**
- c. IP assignment: **static**
- d. FQDN: **<vcenter-fqdn>**
- e. IP address: **<vcenter-ip>**
- f. Subnet mask or prefix length: **<vcenter-subnet-mask>**
- g. Default gateway: **<vcenter-gateway>**
- h. DNS Servers: **<dns-server1>,<dns-server2>**



Step 13. Click **NEXT**.

Step 14. Review all values and click **FINISH** to complete the installation.

Note: The vCenter Server appliance installation will take a few minutes to complete.




Step 15. Click **CONTINUE** to proceed with stage 2 configuration.

Step 16. Click **NEXT**.


Step 17. In the vCenter Server configuration window, configure these settings:

- a. Time Synchronization Mode: Synchronize time with NTP servers.
- b. NTP Servers: <nexus-a-ntp-ip>,<nexus-b-ntp-ip>
- c. SSH access: **Enabled**.



Step 18. Click **NEXT**.

Step 19. Complete the SSO configuration as shown below or according to your organization's security policies:



Step 20. Click **NEXT**.

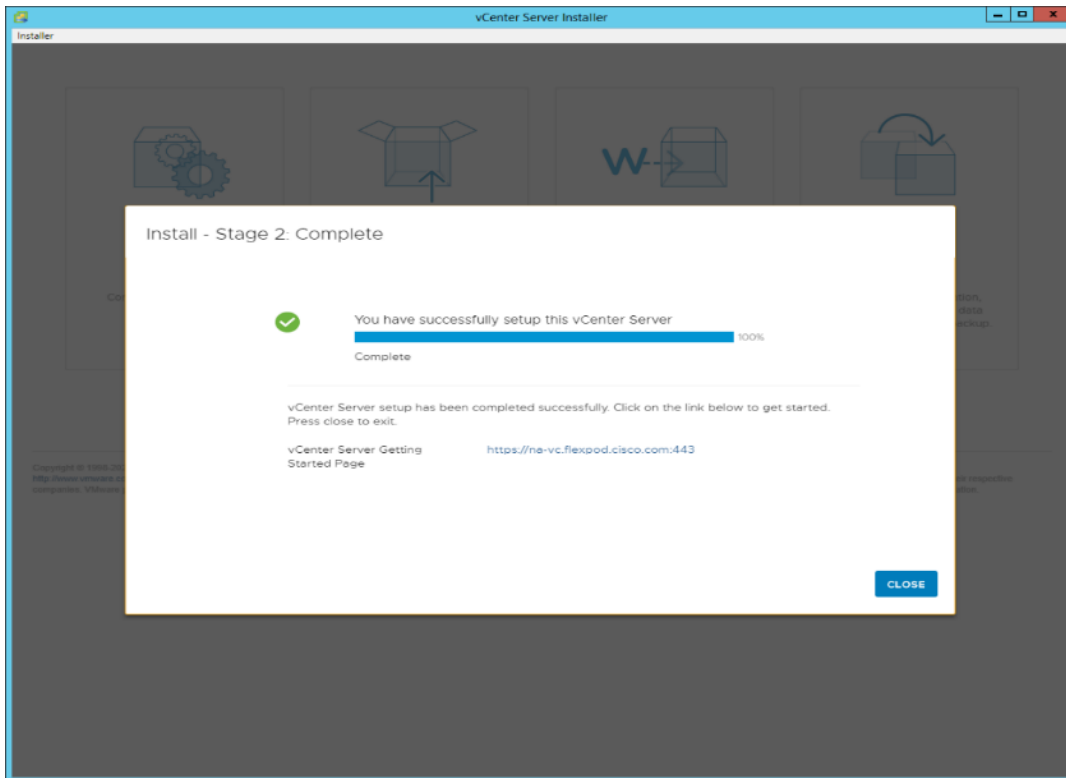
Step 21. Decide whether to join VMware's Customer Experience Improvement Program (CEIP).

Step 22. Click **NEXT**.

Step 23. Review the configuration and click **FINISH**.

Step 24. Click **OK**.

Note: The Server setup will take a few minutes to complete.



Step 25. Click **CLOSE**. Eject or unmount the VCSA installer ISO.

Procedure 2. Adjust vCenter CPU settings and resolve Admission Control issues

Note: If a vCenter deployment size Small or Larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration. Cisco UCS B and C-Series servers are normally 2-socket servers. In this validation, the Small deployment size was selected and vCenter was setup for a 4-socket server. This setup will cause issues in the VMware ESXi cluster Admission Control.

Step 1. Open a web browser on the management workstation and navigate to the **VM-Host-Infra-01 management IP address**.

Step 2. Enter **root** for the user name.

Step 3. Enter the **root** password.

Step 4. Click **Login** to connect.

Step 5. Click **Virtual Machines**.

Step 6. Right-click the **vCenter VM** and click **Edit settings**.

Step 7. In the Edit settings window, expand CPU and check the value of Sockets.

▼ CPU	8 ▼
Cores per Socket	1 ▼ Sockets: 8
CPU Hot Plug	<input checked="" type="checkbox"/> Enable CPU Hot Add

Step 8. If the number of Sockets does not match your server configuration, it will need to be adjusted. Click **Cancel**.

Step 9. If the number of Sockets needs to be adjusted:

- a. Right-click the **vCenter VM** and click **Guest OS > Shut down**. Click **Yes** on the confirmation.
- b. Once vCenter is shut down, right-click the **vCenter VM** and click **Edit settings**.
- c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to your server configuration (normally 2).
- d. Click **Save**.
- e. Right-click the **vCenter VM** and click **Power > Power on**. Wait approximately 10 minutes for vCenter to come up.

Procedure 3. Set up VMware vCenter Server

Step 1. Using a web browser, navigate to **https://<vcenter-ip-address>:5480**.

Step 2. Log into the VMware vCenter Server Management interface as root with the root password set in the vCenter installation.

Step 3. Click **Time**.

Step 4. Click **EDIT**.

Step 5. Select the appropriate Time zone and click **SAVE**.

Step 6. Click **Administration**.

Step 7. According to your Security Policy, adjust the settings for the root user and password.

Step 8. Click **Update**.

Step 9. Follow the prompts to STAGE AND INSTALL any available vCenter updates. In this validation, vCenter version 8.0 U1.2.00500 was installed.

Step 10. Go to **root > Logout** to logout of the Appliance Management interface.


Step 11. Using a web browser, navigate to **https://<vcenter-fqdn>**. You will need to navigate security screens.

Note: With VMware vCenter 8.0 U1.U3 the use of the vCenter FQDN is required.

Step 12. Click **LAUNCH VSPHERE CLIENT (HTML5)**.


Note: Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is the only option in vSphere 7 and will be used going forward.

Step 13. Log in using the Single Sign-On username (administrator@vsphere.local) and password created during the vCenter installation. Dismiss the Licensing warning at this time.



Step 14. Click **ACTIONS > New Datacenter.**

Step 15. Type “FlexPod-DC” in the Datacenter name field.



Step 16. Click **OK.**

Create vSphere Host Cluster

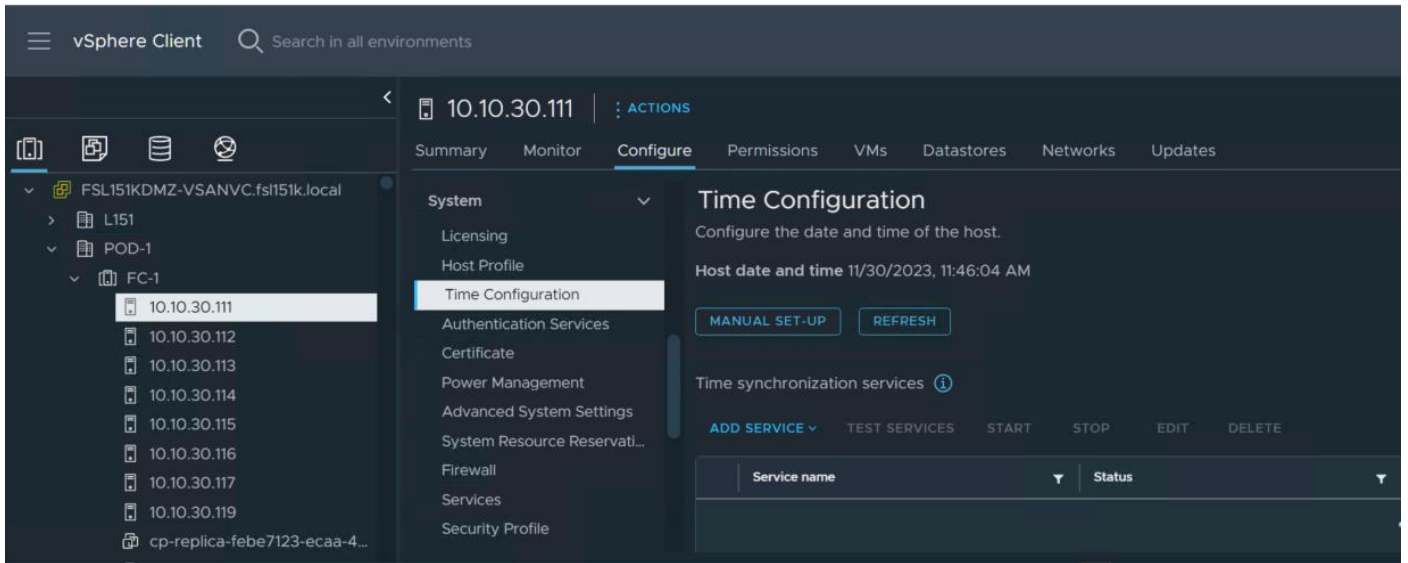
Procedure 1. Create vSphere Cluster

Step 1. Create a vSphere host cluster.

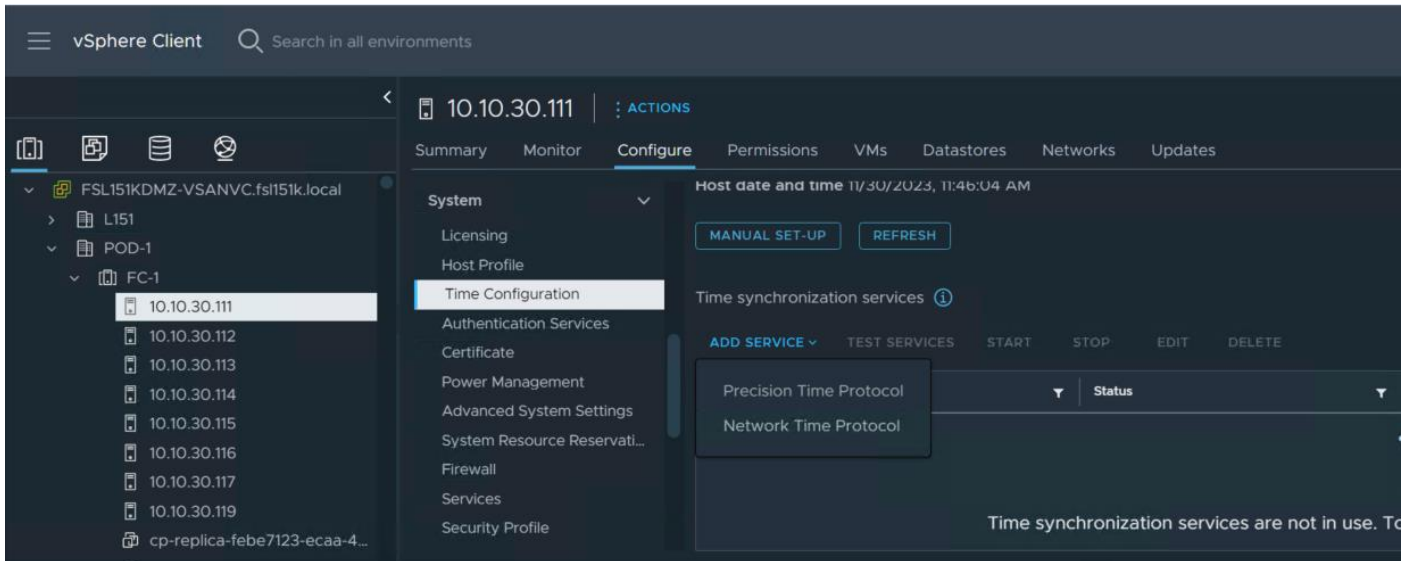
Step 2. Configure NTP for all hosts in your vSphere cluster.

Procedure 2. Configure NTP

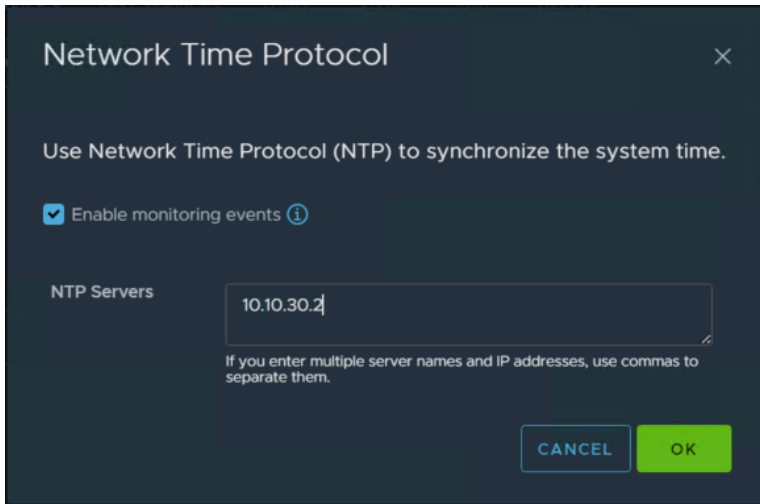
Step 1. On the vSphere Web Client home screen, select the Host object from the list on the left. From the Configure tab System area click **Time Configuration.**



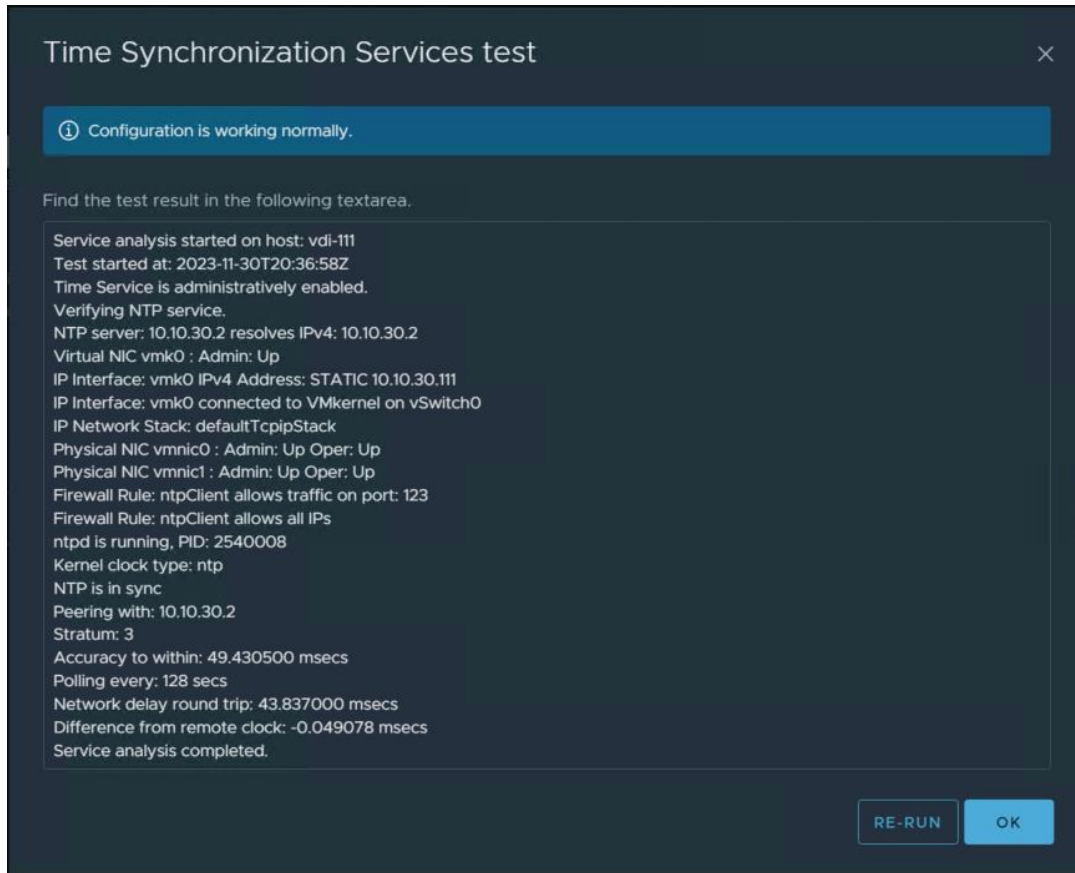
Step 2. From Add Service drop-down list select **Network Time Protocol**.



Step 3. Provide the IP addresses for the NTP servers in your environment. Click **OK**.



Step 4. Test the service configuration.

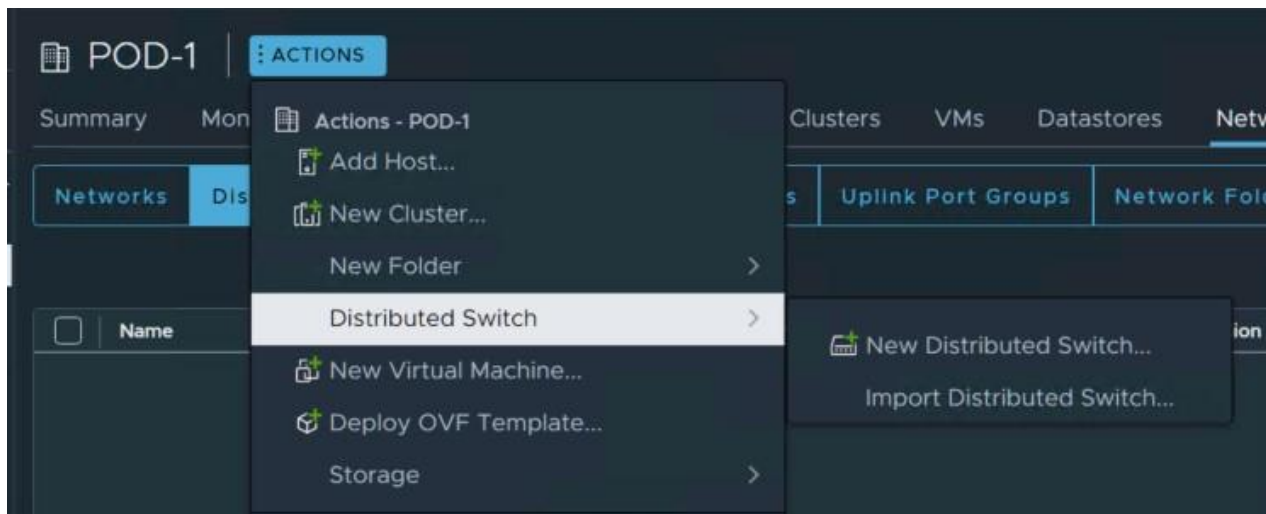


Note: SSH was also enabled to support data collection.

Step 5. Create Distributed Switch.

Procedure 3. Create a new vDS

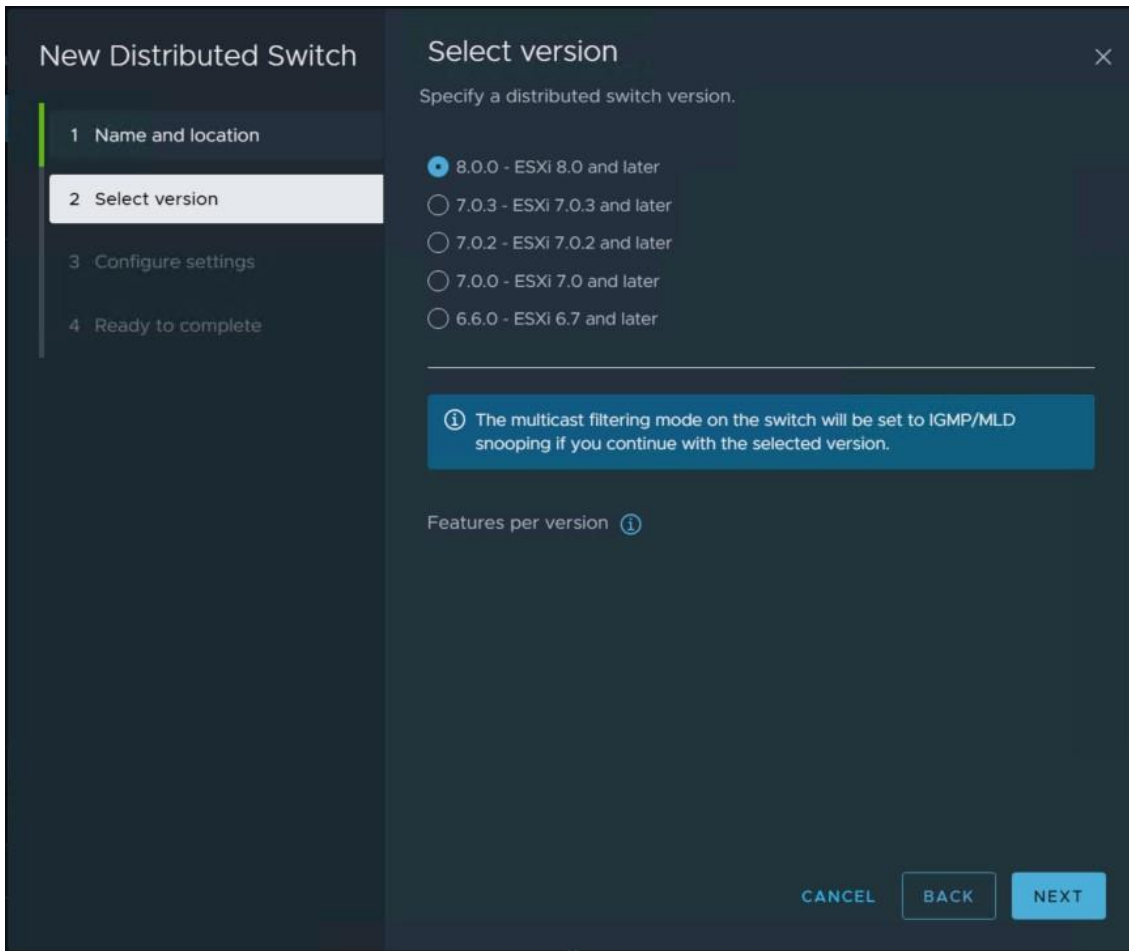
Step 1. On the vSphere Web Client home screen, select the vCenter object from the list on the left. From the Inventory Lists area right-click on **datacenter**, then select **Distributed Switch** and click **New Distributed Switch**.



Step 2. Provide a name for the new distributed switch and select the location within the vCenter inventory where you would like to store the new vDS (a data center object or a folder). Click **NEXT**.

The screenshot shows a dark-themed wizard window titled "New Distributed Switch". On the left, a vertical sidebar contains three steps: "1 Name and location" (highlighted), "2 Select version", and "3 Ready to complete". The main area is titled "Name and location" and contains the instruction "Specify distributed switch name and location." Below this, there are two fields: "Name" with the value "vDS-1" and "Location" with a folder icon and the value "POD-1". At the bottom right, there are two buttons: "CANCEL" and "NEXT".

Step 3. Select the version of the vDS to create. Click **NEXT**.



Step 4. Specify the Network Offloads compatibility as **None**, and number of uplink ports as **2**. Uncheck the Create a default port group box. Click **NEXT**.

New Distributed Switch

- 1 Name and location
- 2 Select version
- 3 Configure settings
- 4 Ready to complete

Configure settings

Specify network offloads compatibility, number of uplink ports, resource allocation and default port group.

Network Offloads compatibility: **None** ⓘ

Number of uplinks: **2**

Network I/O Control: **Enabled**

Default port group: Create a default port group

Port group name: **DPortGroup**

[CANCEL](#) [BACK](#) [NEXT](#)

Step 5. Click **Finish**.

The screenshot shows a dark-themed wizard window titled "New Distributed Switch" with a close button (X) in the top right corner. On the left, a vertical sidebar lists four steps: "1 Name and location", "2 Select version", "3 Configure settings", and "4 Ready to complete", with the fourth step highlighted. The main area is titled "Ready to complete" and contains the text "Review your settings selections before finishing the wizard." Below this, a table lists the following settings:

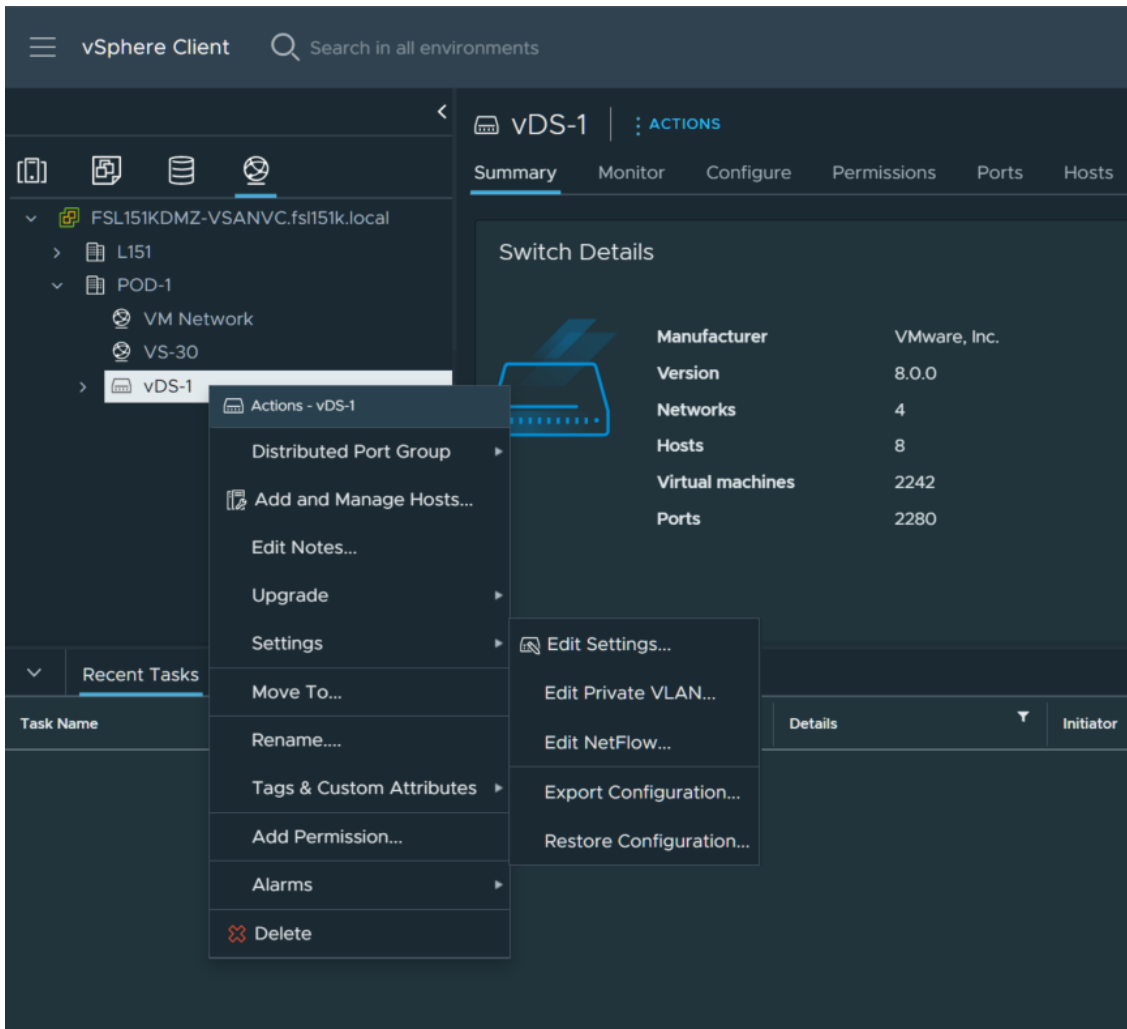
Name	vDS-1
Version	8.0.0
Network Offloads compatibility	None
Number of uplinks	2
Network I/O Control	Enabled

Below the table, there is a section titled "Suggested next actions" with a downward arrow icon. It contains two items:

- New Distributed Port Group
- Add and Manage Hosts

A note below the actions states: "These actions will be available in the Actions menu of the new distributed switch." At the bottom right of the wizard, there are three buttons: "CANCEL", "BACK", and "FINISH".

Step 6. Right-click the new distributed switch in the list of objects and select **Settings > Edit Settings....**



Step 7. In the Distributed Switch-Edit Settings dialog box Advanced tab, set the MTU to **9000**, Discovery protocol to **Link Layer Discovery Protocol** and Operation to **Both**. Click **OK**.

Distributed Switch - Edit Settings | vDS-1

General **Advanced** Uplinks

MTU (Bytes) 9000

Multicast filtering mode IGMP/MLD snooping

Discovery protocol

Type Cisco Discovery Protocol

Operation Both

Administrator contact

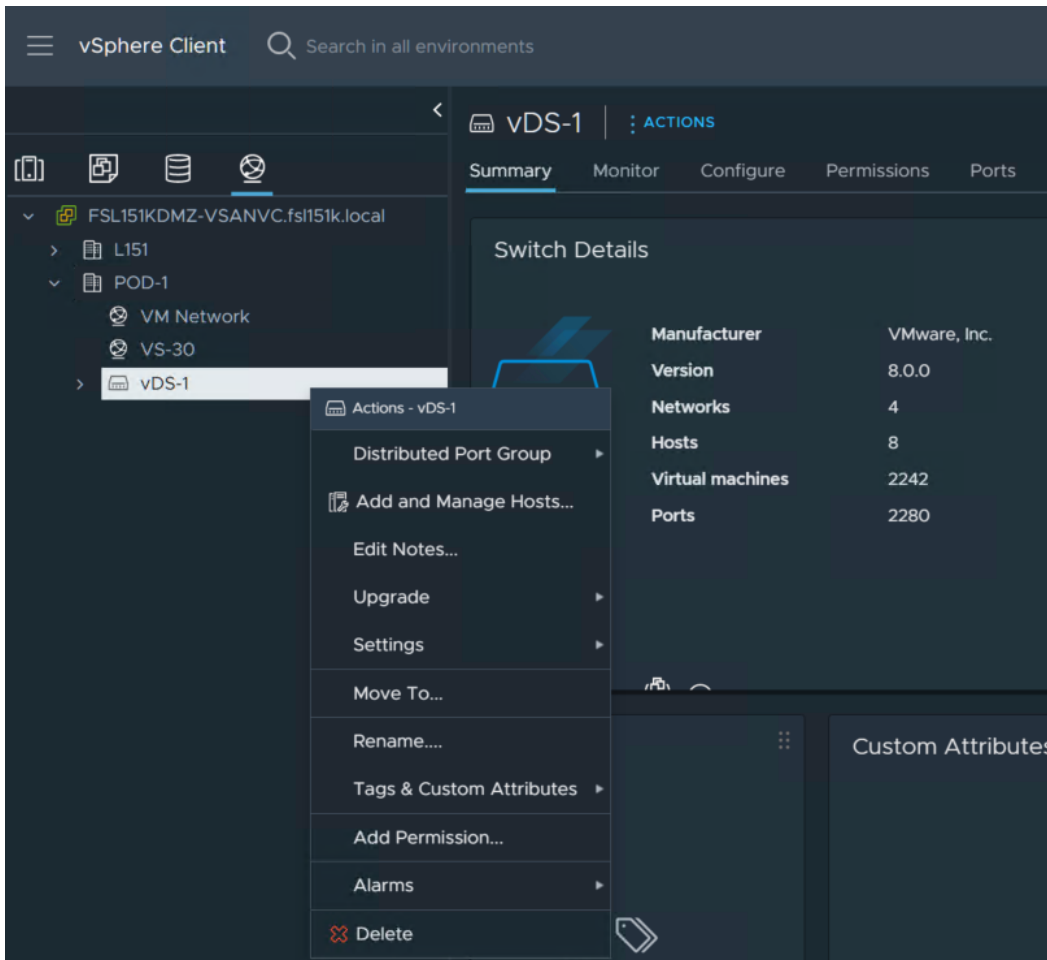
Name

Other details

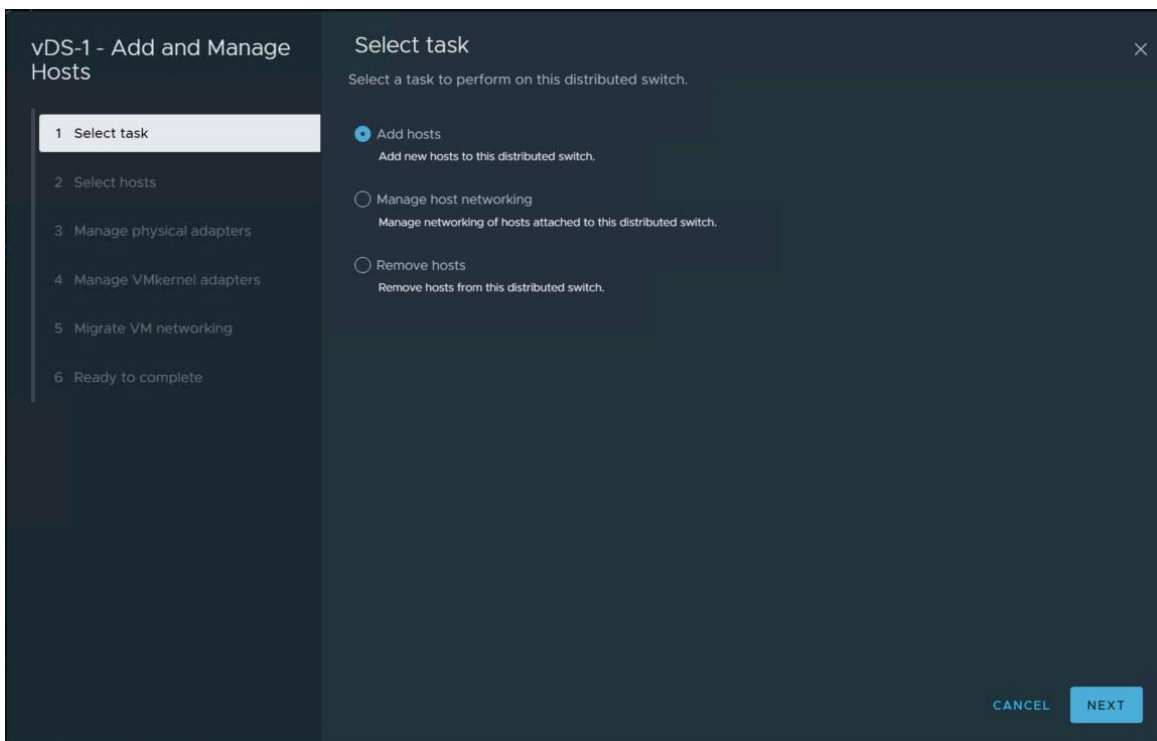
CANCEL OK

Note: In server profiles vmnic2 and vmnic3 are created for the use as vDS uplinks.

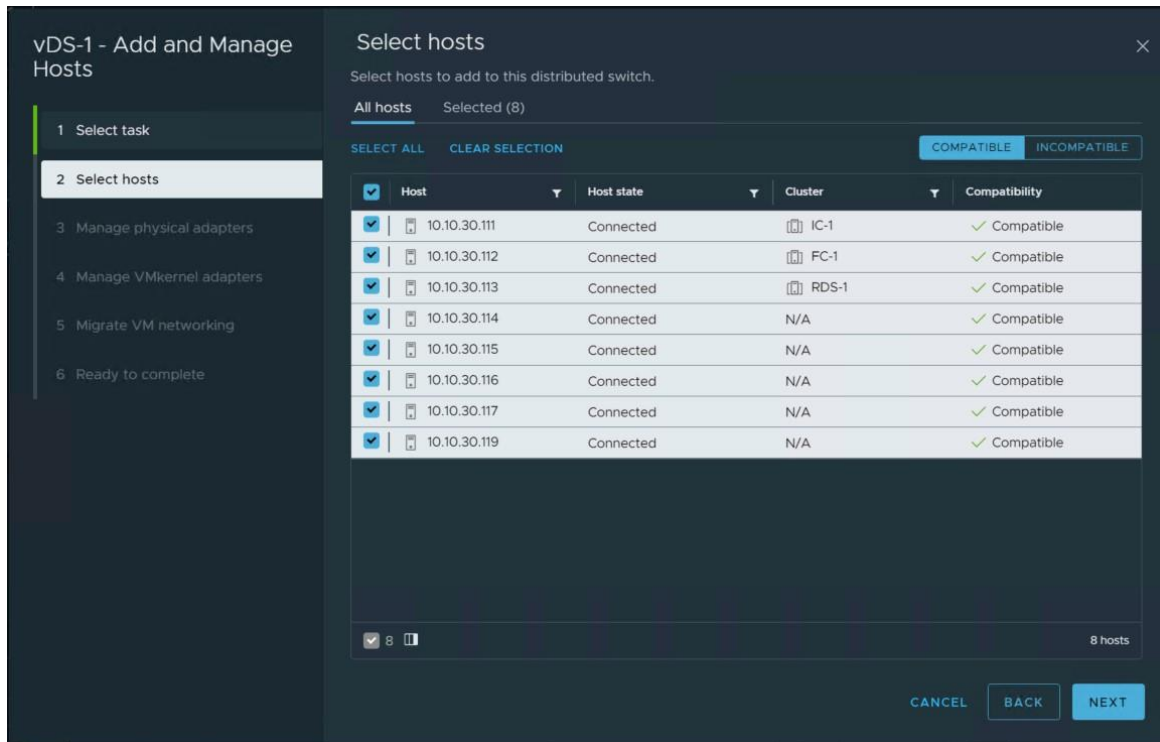
Step 8. Right-click ONTAP Boot storage setup is not required for this solution as it uses local boot using M.2 drives. Click the new distributed switch in the list of objects and select **Add and Manage Hosts** from the Actions menu.



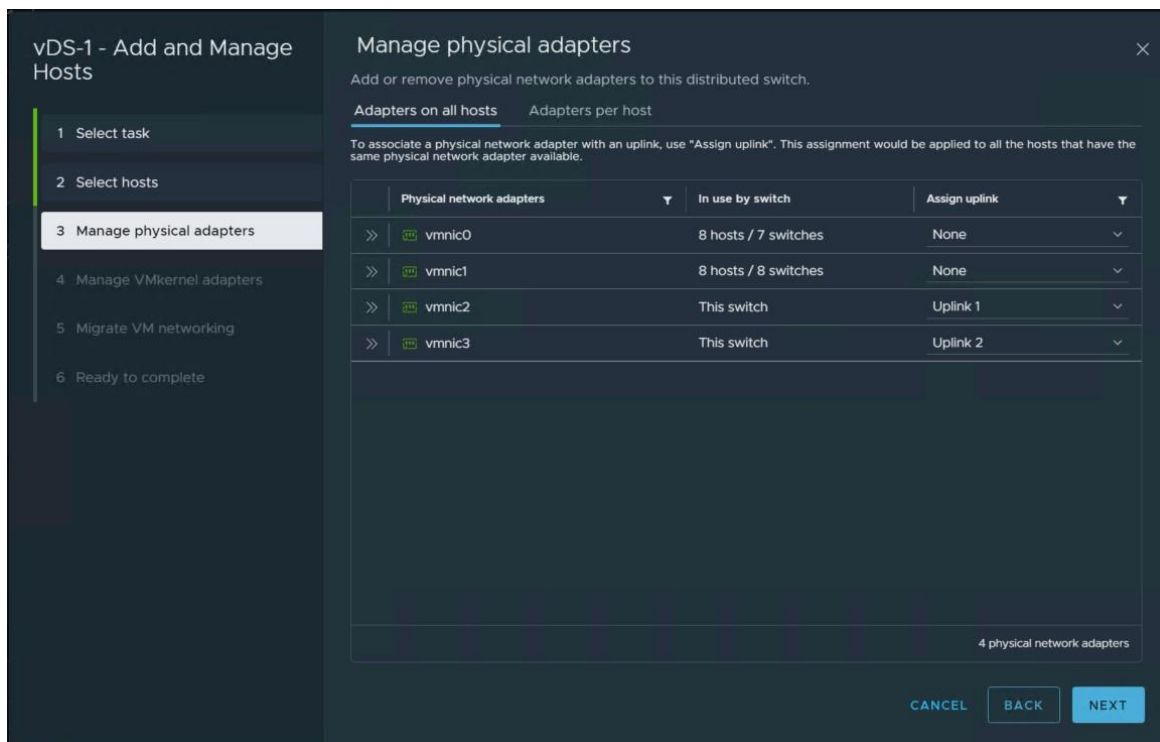
Step 9. Select the **Add hosts** and click **NEXT**.



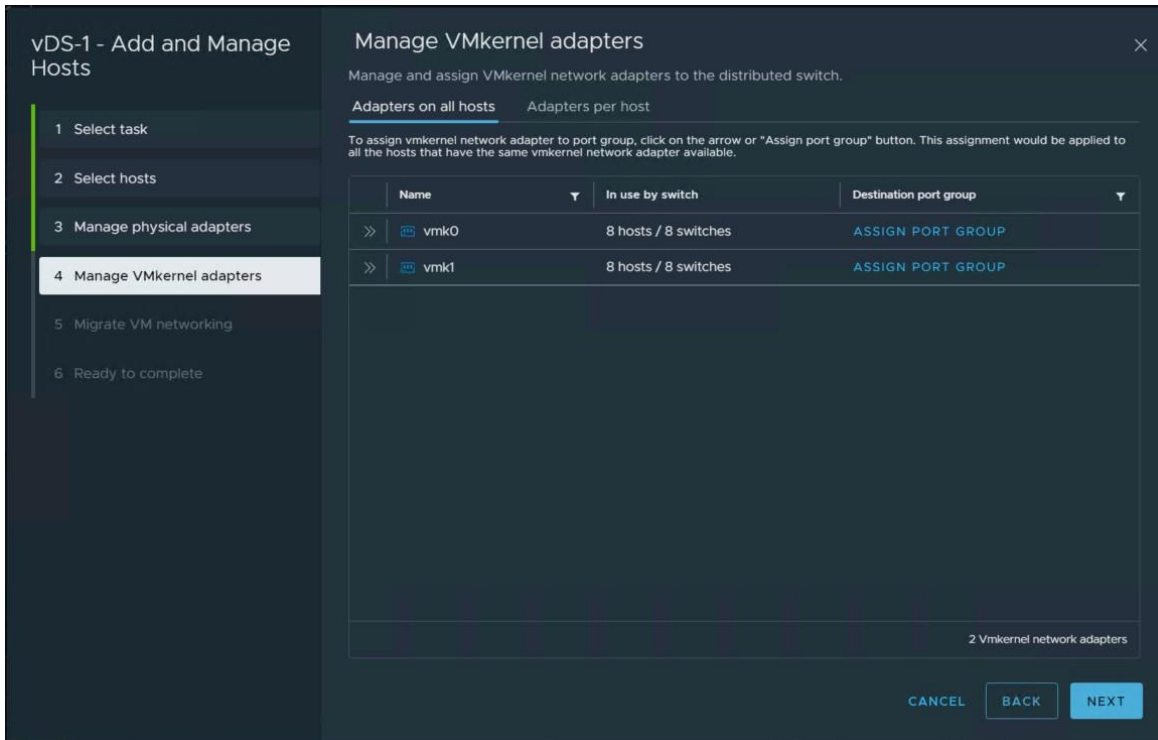
Step 10. From the list of the new hosts, check the boxes with the names of each ESXi host you would like to add to the VDS. Click **NEXT**.



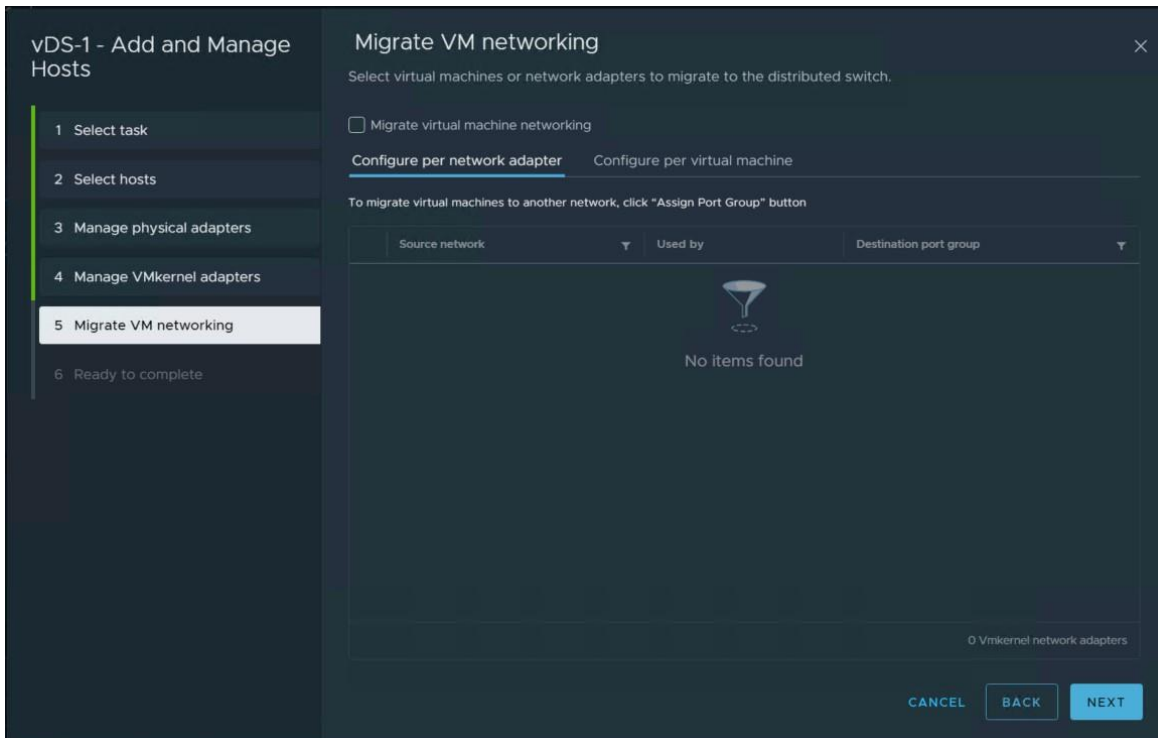
Step 11. In the next Manage physical adapters menu, click **Adapters on all hosts** and configure the adapters (in this case - vmnic2 and vmnic3) in an ESXi host as Uplink 1 and Uplink 2 for the vDS.



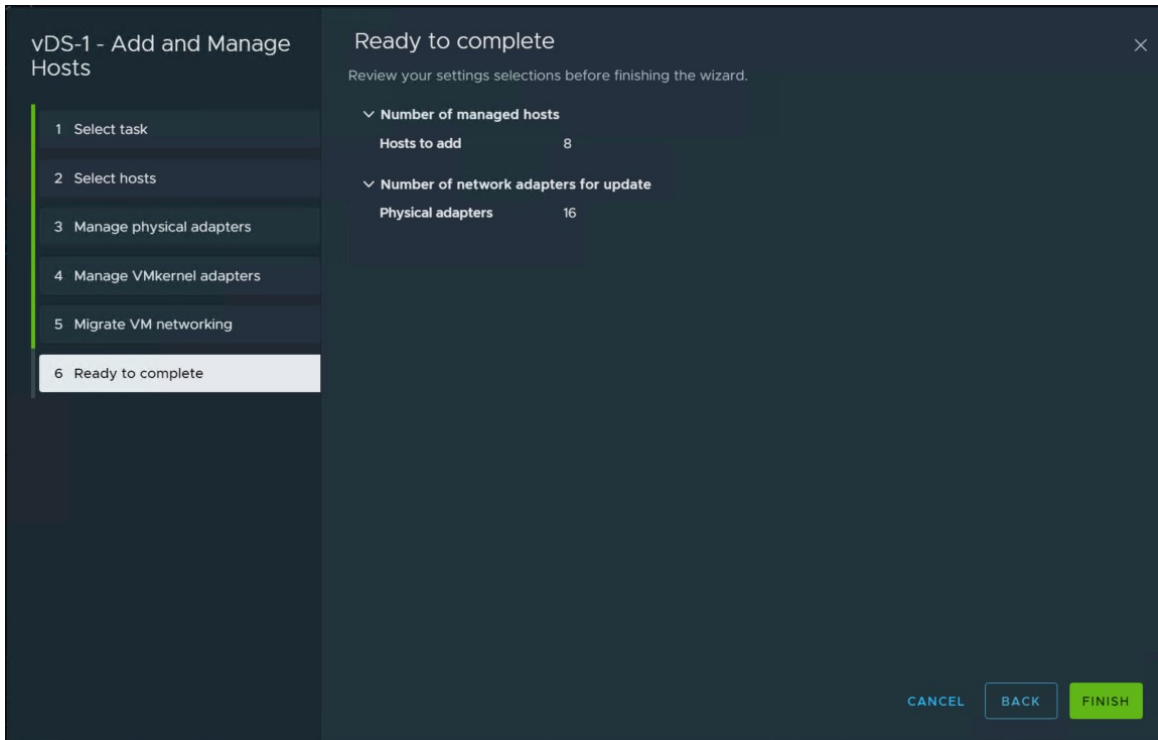
Step 12. In the Manage VMkernel adapters and Migrate VM networking menus, click **NEXT** to continue.



Step 13. In the Manage VMkernel adapters and Migrate VM networking menus, click **NEXT** to continue.

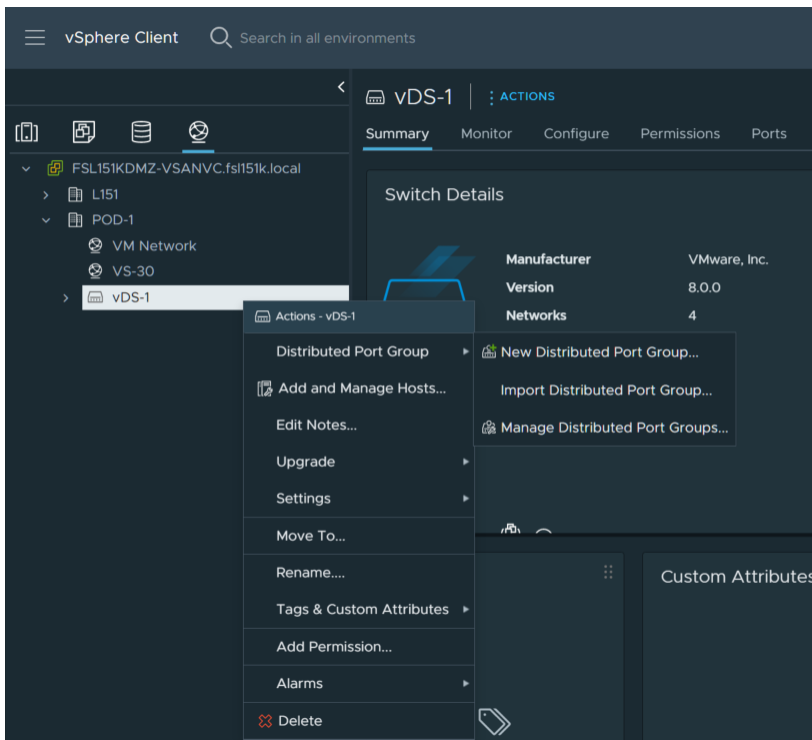


Step 14. Click **FINISH**.



Procedure 4. Create a Distributed Port Group for vMotion traffic

Step 1. Right-click **Distributed switch** and select **Distributed Port Group** then click **New Distributed Port Group**.



Step 2. On the New Distributed Port Group dialog box, enter a Name (for example vMotion), and click **NEXT**.

New Distributed Port Group

Name and location ×

Specify distributed port group name and location.

1 Name and location
2 Configure settings
3 Ready to complete

Name vMotion-32

Location vDS-1

CANCEL NEXT

Step 3. In the VLAN type field, select **VLAN**, and set the VLAN ID to your VLAN (–for example 73). Check the **Customize default policies configuration** checkbox and click **NEXT**.

New Distributed Port Group

Configure settings ×

Set general properties of the new port group.

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Port binding Static binding ▾

Port allocation Elastic ▾ ⓘ

Number of ports 8

Network resource pool (default) ▾

VLAN

VLAN type VLAN ▾

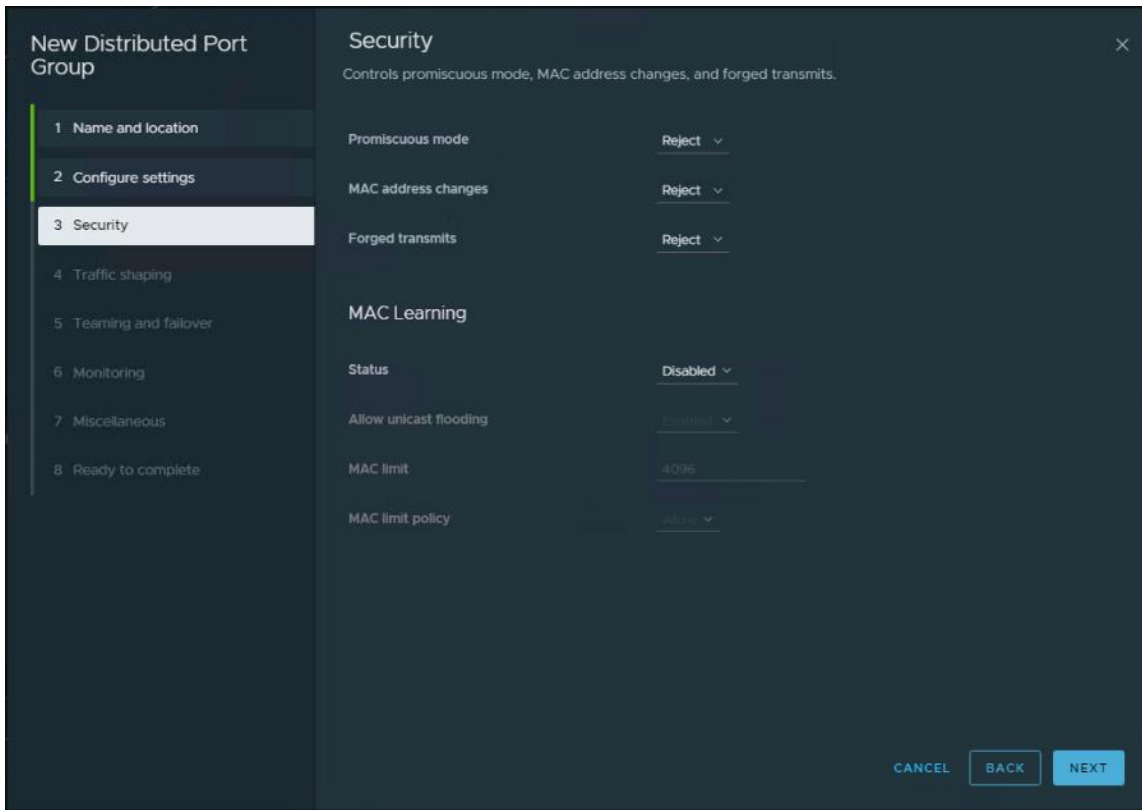
VLAN ID 32

Advanced

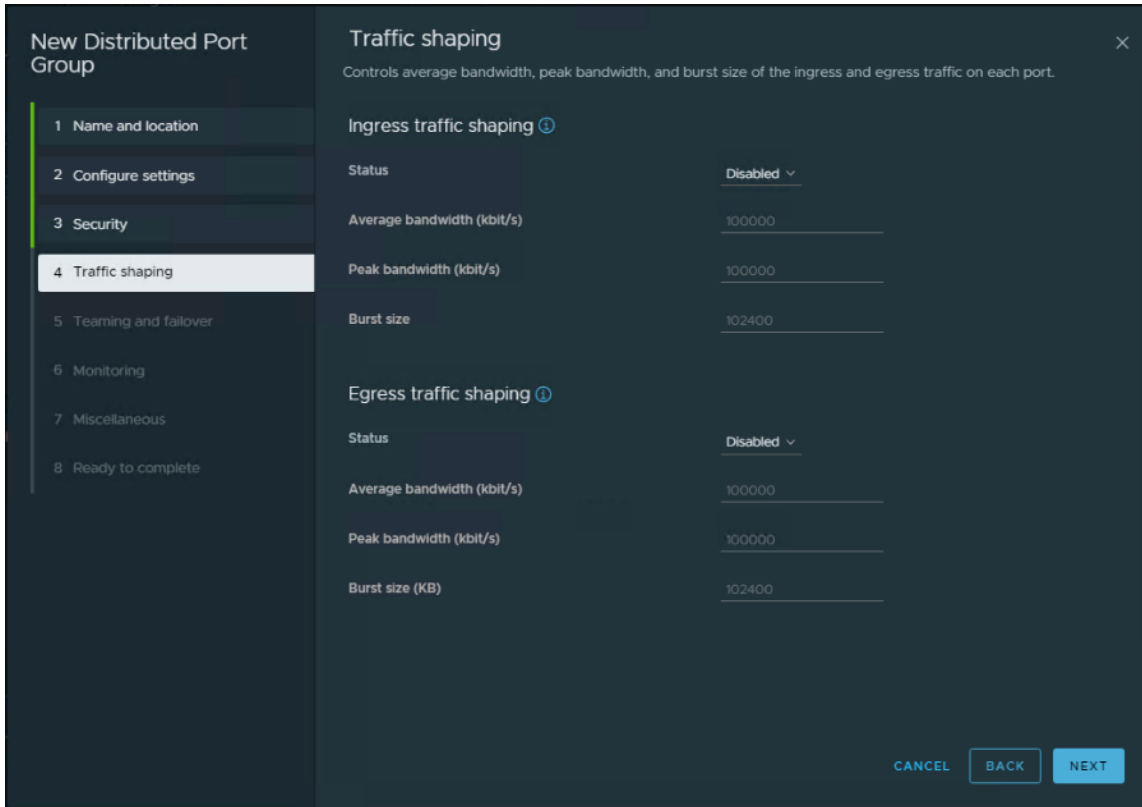
Customize default policies configuration

CANCEL BACK NEXT

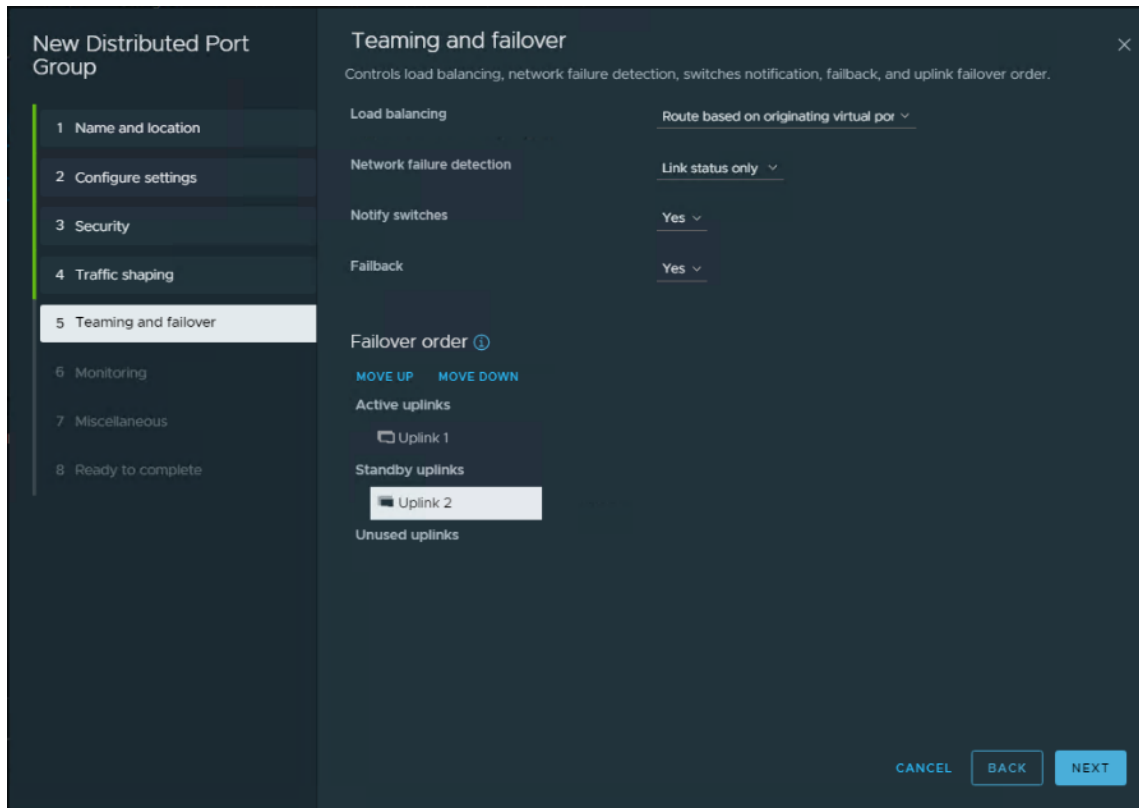
Step 4. On the Security dialog box, click **NEXT**.



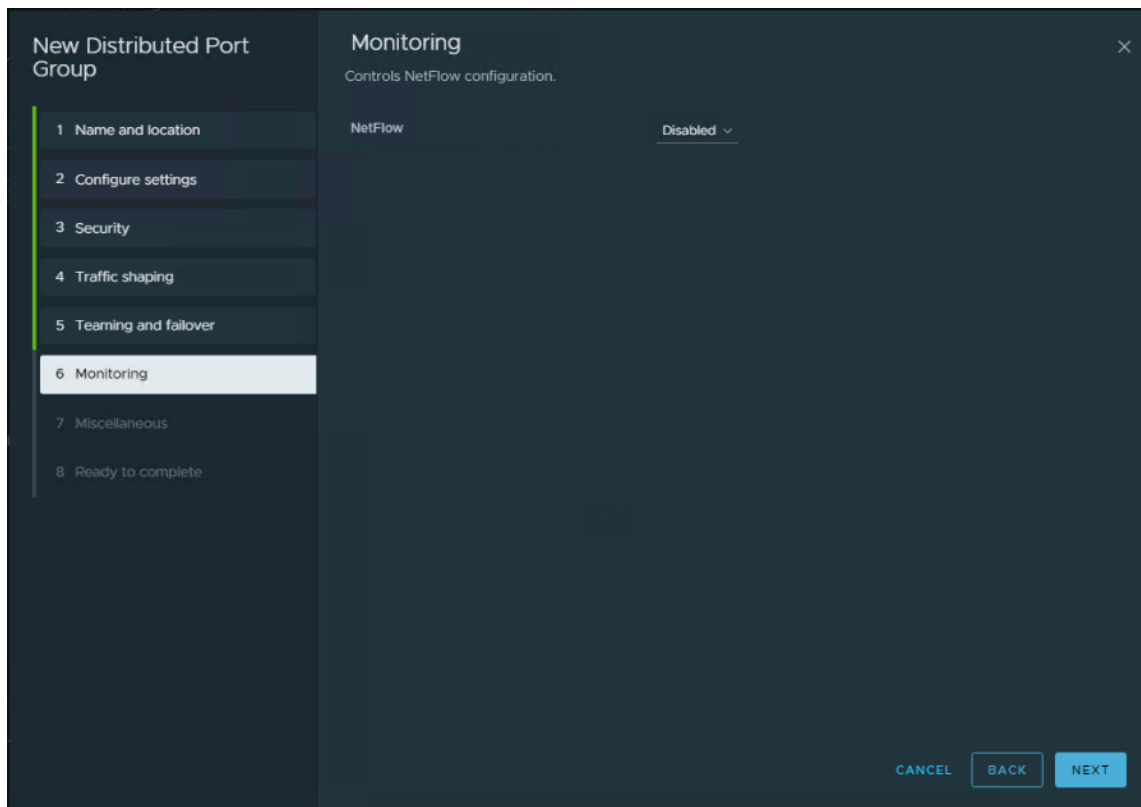
Step 5. On the Traffic shaping dialog box, click **NEXT**.



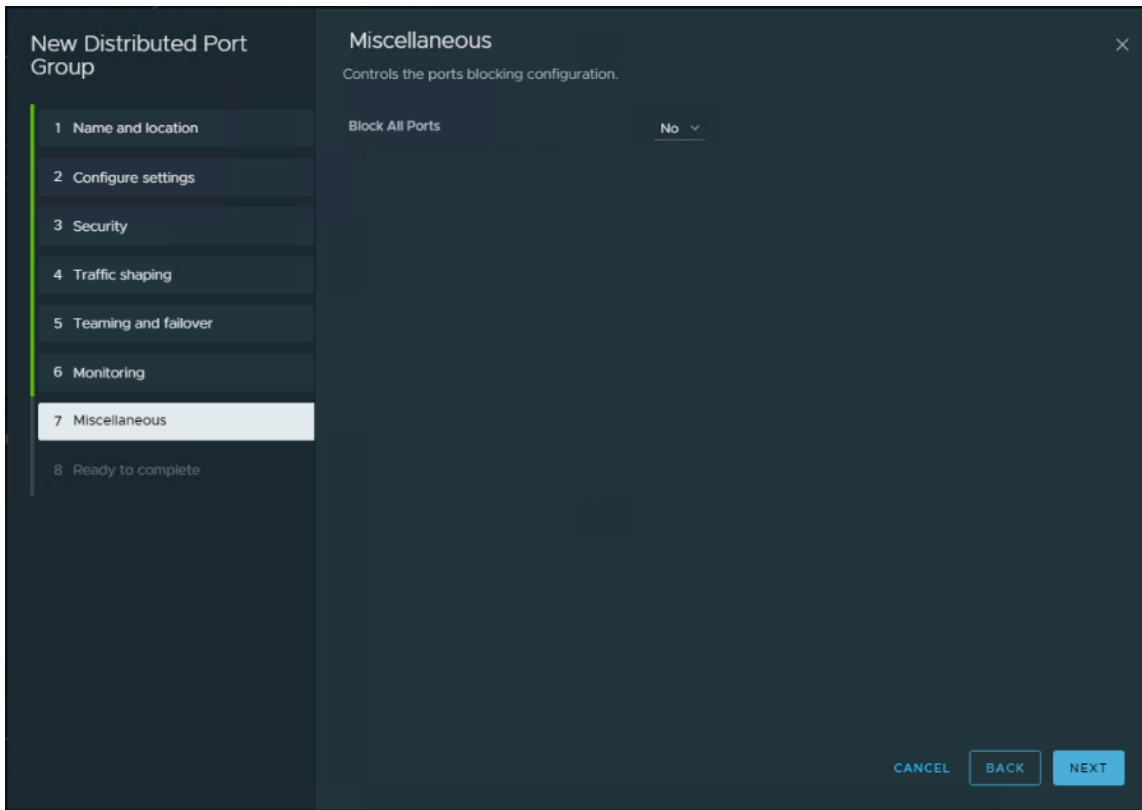
Step 6. In the Teaming and failover dialog box, select Uplink 1 as active uplink, and set Uplink 2 to be the standby uplink. Click **NEXT**.



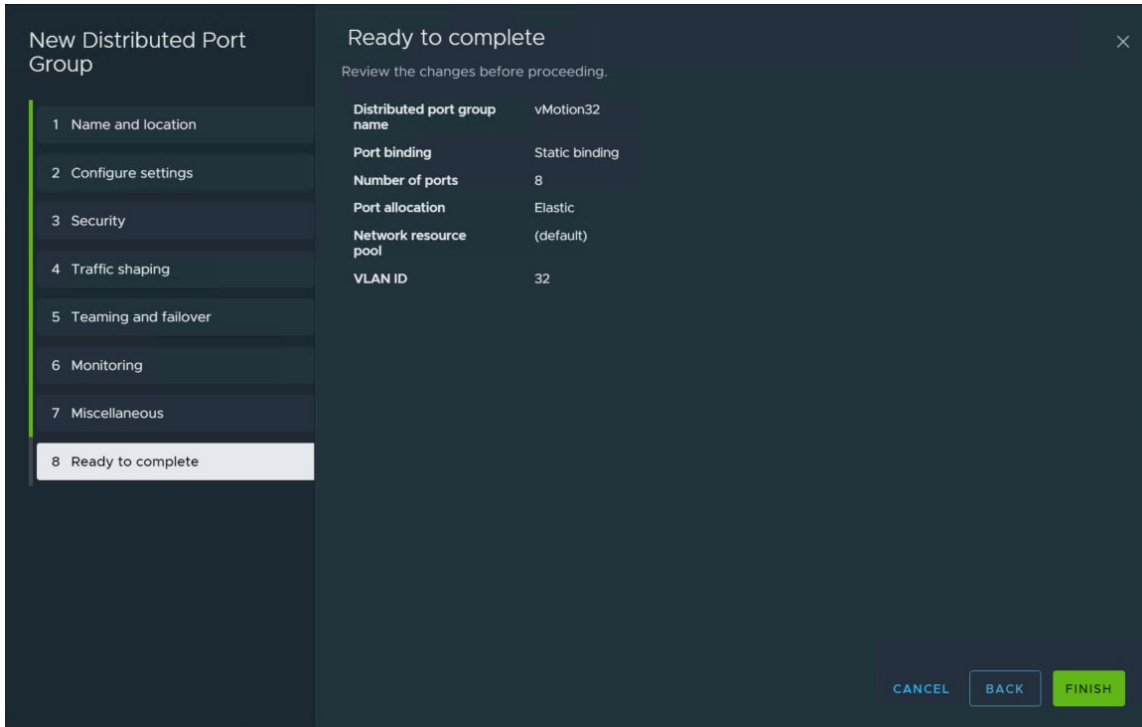
Step 7. In the Monitoring dialog box, set NetFlow to **Disabled**, and click **NEXT**.



Step 8. In the Miscellaneous dialog box, set Block All Ports to **No**, and click **NEXT**.

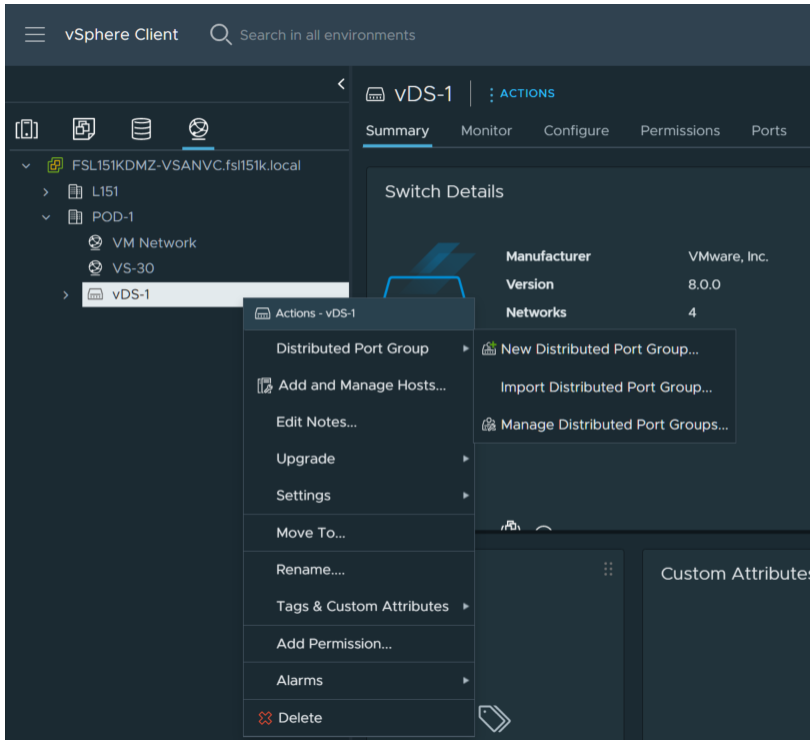


Step 9. In the Ready to complete dialog box, review all the changes, and click **FINISH**.

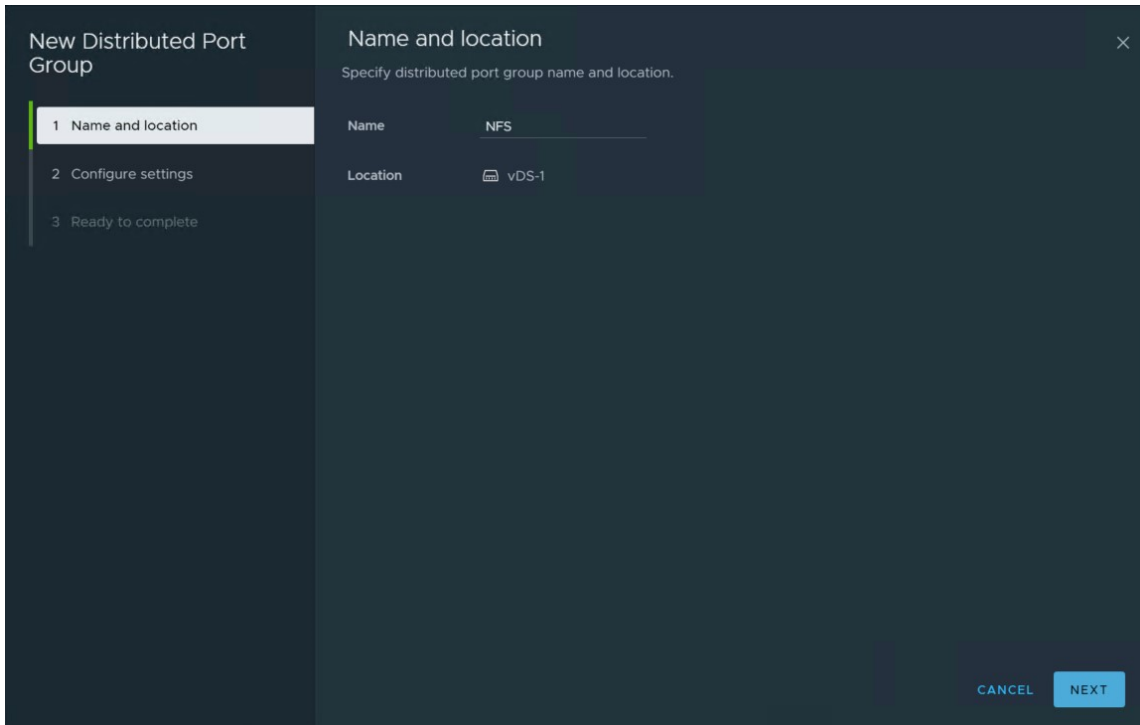


Procedure 5. Creating a Distributed Port Group for Storage traffic

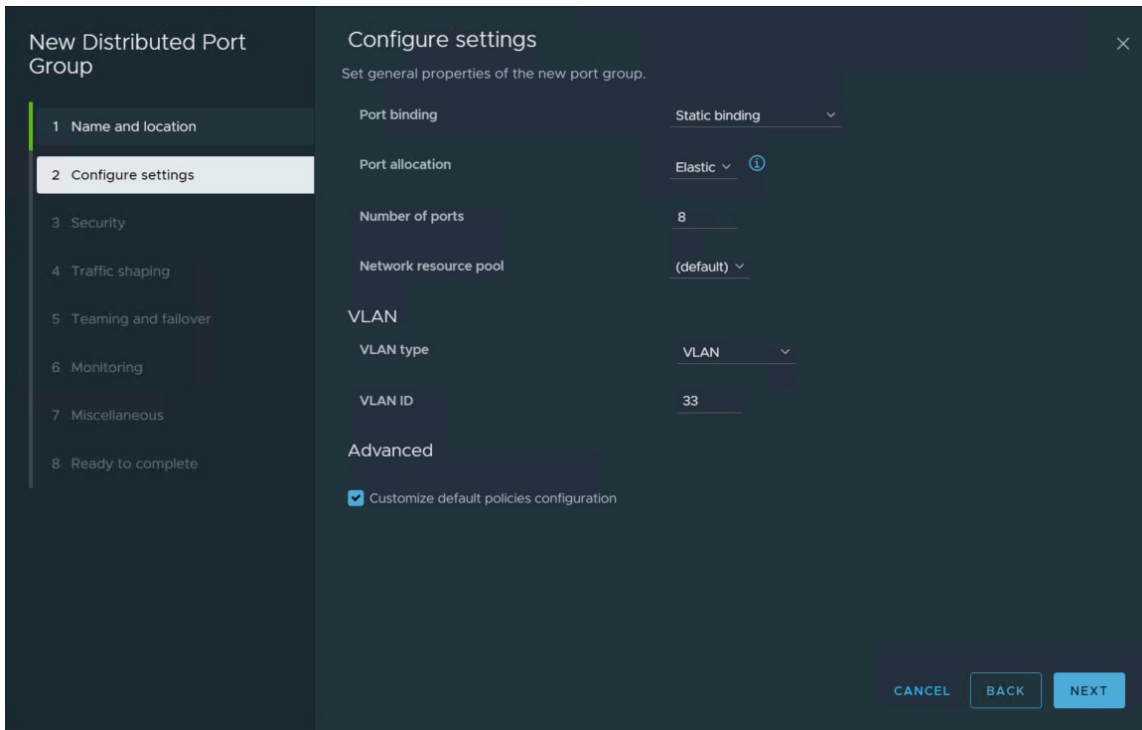
Step 1. Right-click **Distributed switch** and select **Distributed Port Group** then click **New Distributed Port Group**.



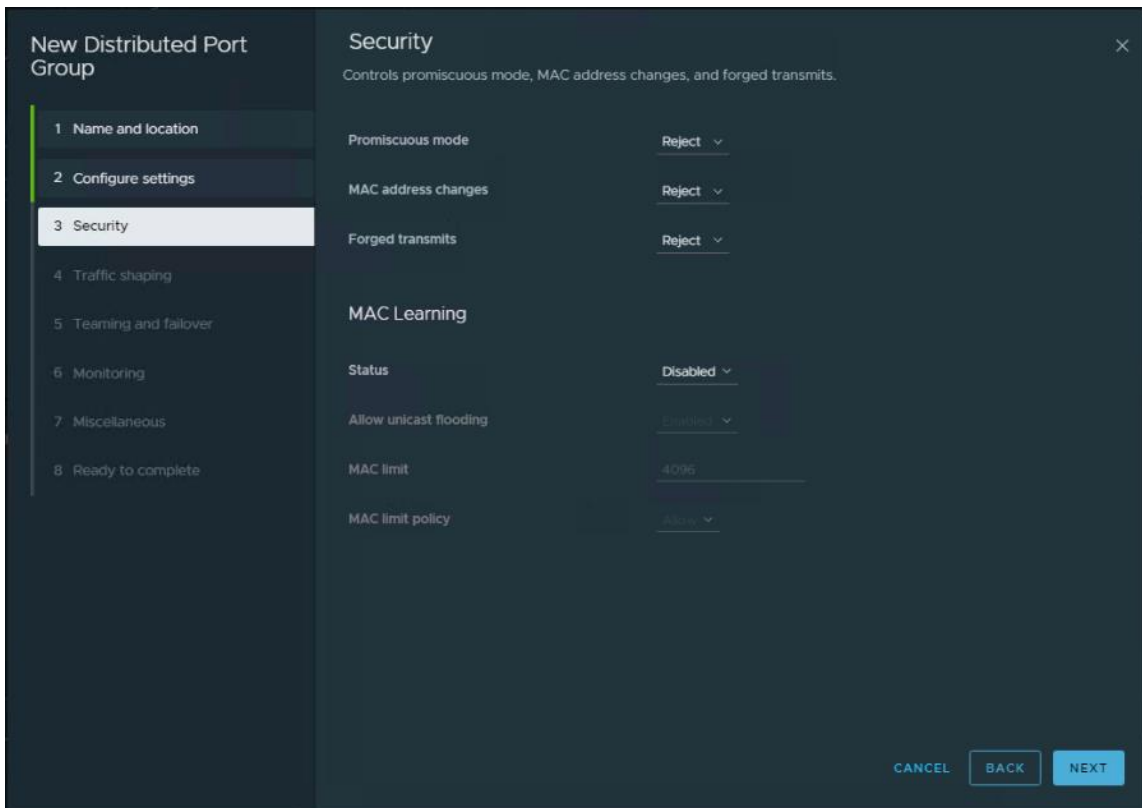
Step 2. On the New Distributed Port Group dialog box, enter a Name (for example NFS), and click **NEXT**.



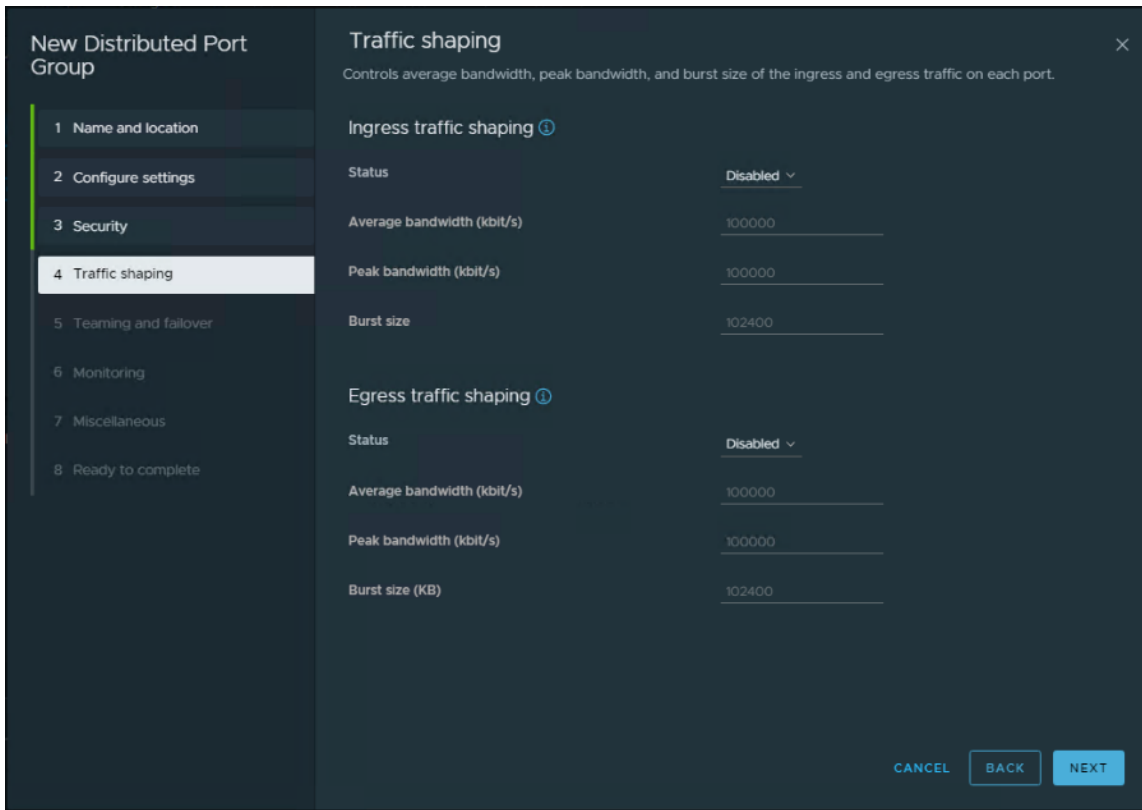
Step 3. In the VLAN type field, select **VLAN**, and set the VLAN ID to your VLAN (–for example 33). Check the **Customize default policies configuration** checkbox and click **NEXT**.



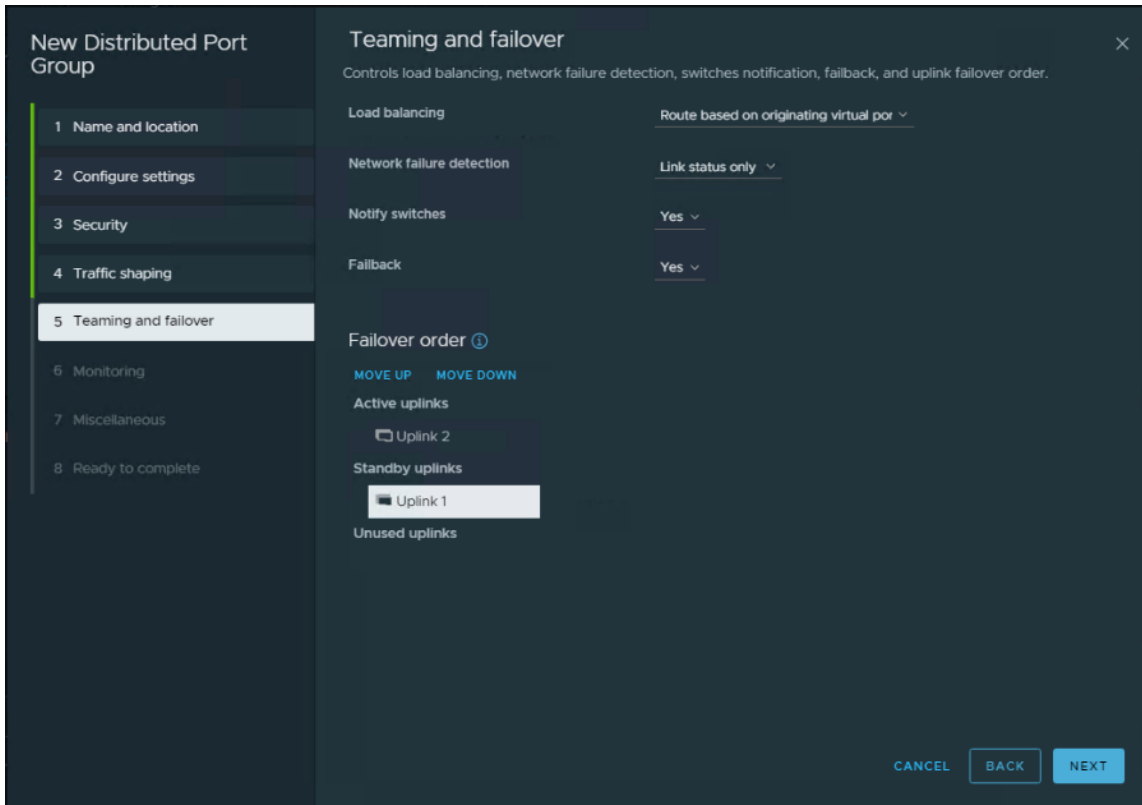
Step 4. On the Security dialog box, click **NEXT**.




Step 5. On the Traffic shaping dialog box, click **NEXT**.



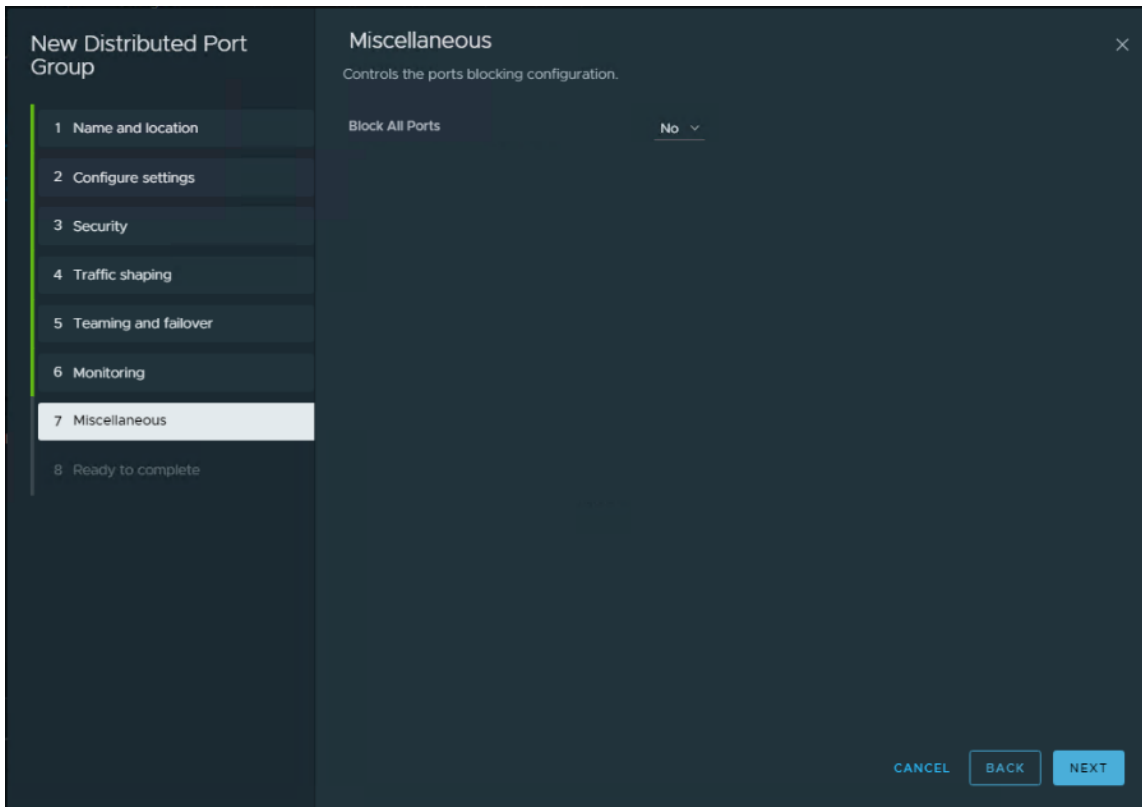
Step 6. In the Teaming and failover dialog box, select Uplink 2 as active uplink, and set Uplink 1 to be the standby uplink. Click **NEXT**.



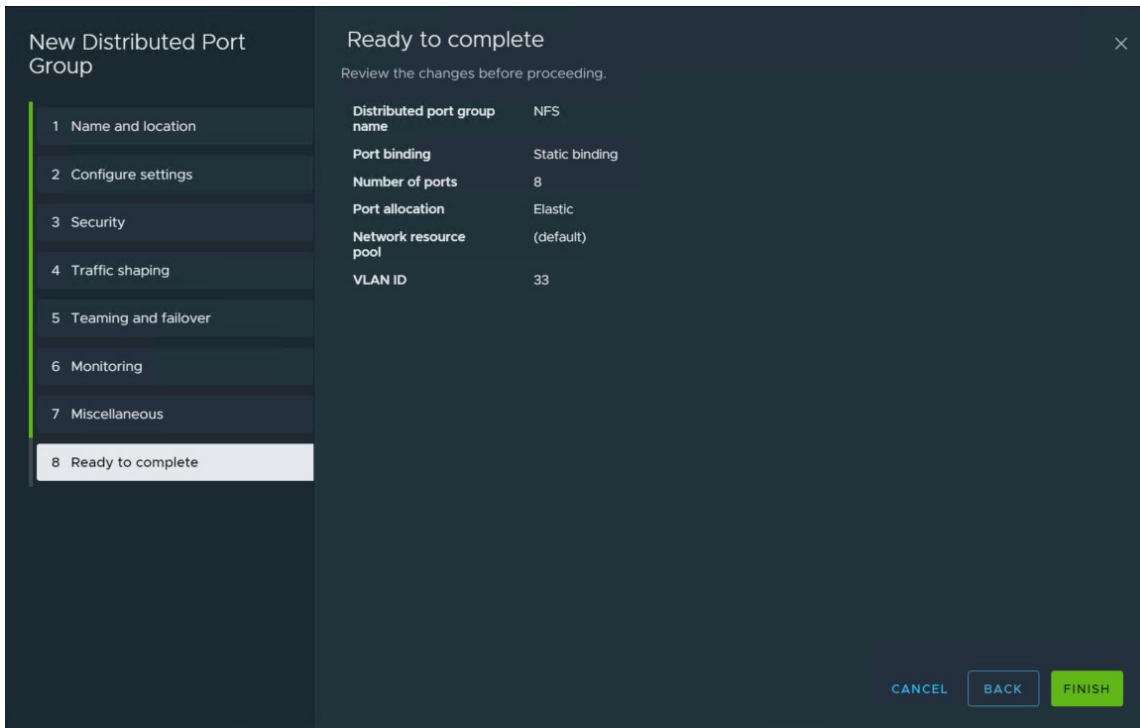
Step 7. In the Monitoring dialog box, set NetFlow to **Disabled**, and click **NEXT**.



Step 8. In the Miscellaneous dialog box, set Block All Ports to **No**, and click **NEXT**.

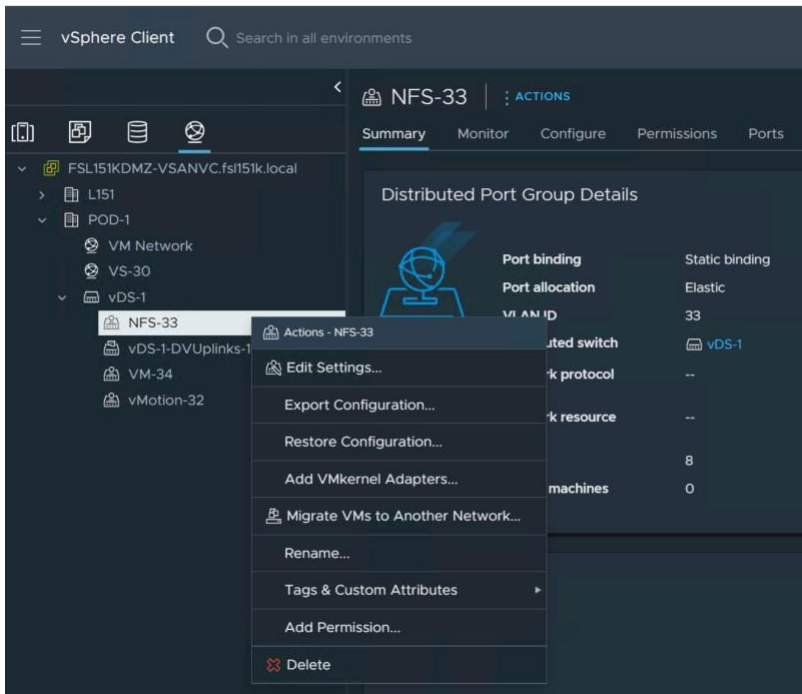


Step 9. In the Ready to complete dialog box, review all the changes, and click **FINISH**.

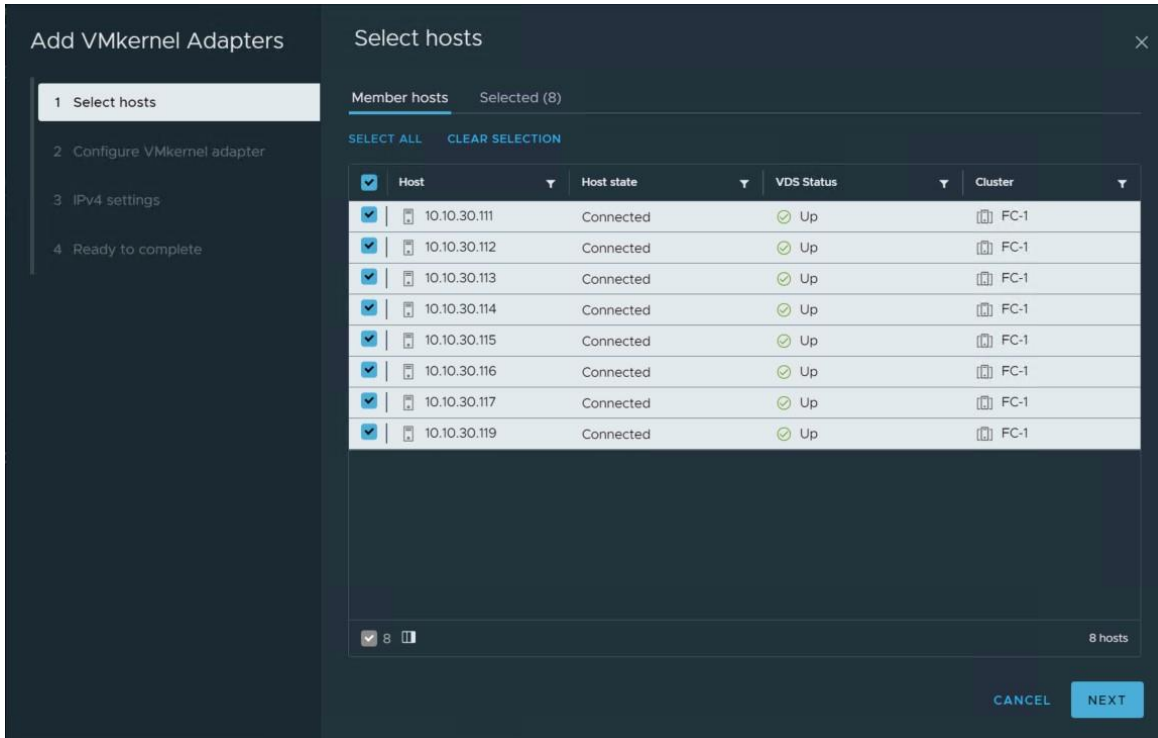


Procedure 6. Adding a VMkernel Adapters to distributed port groups

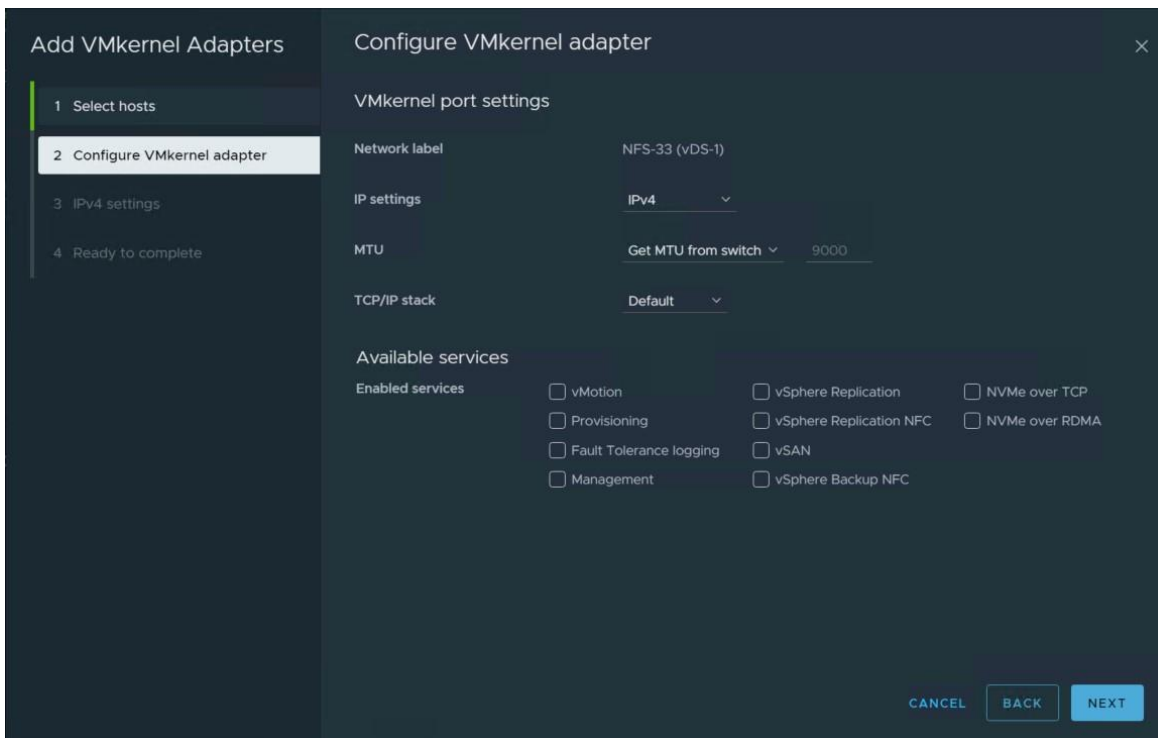
Step 1. Right-click the distributed port group and select **Add VMkernel Adapters....**



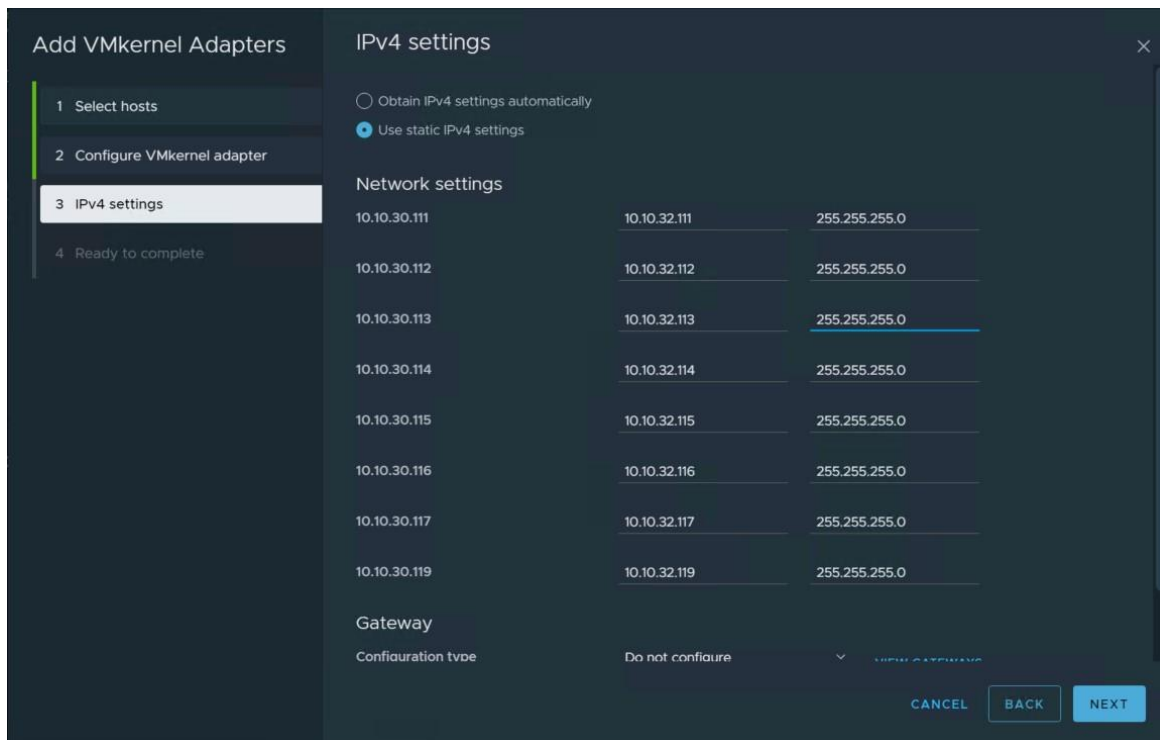
Step 2. Select **Attached Hosts** and click **NEXT**.



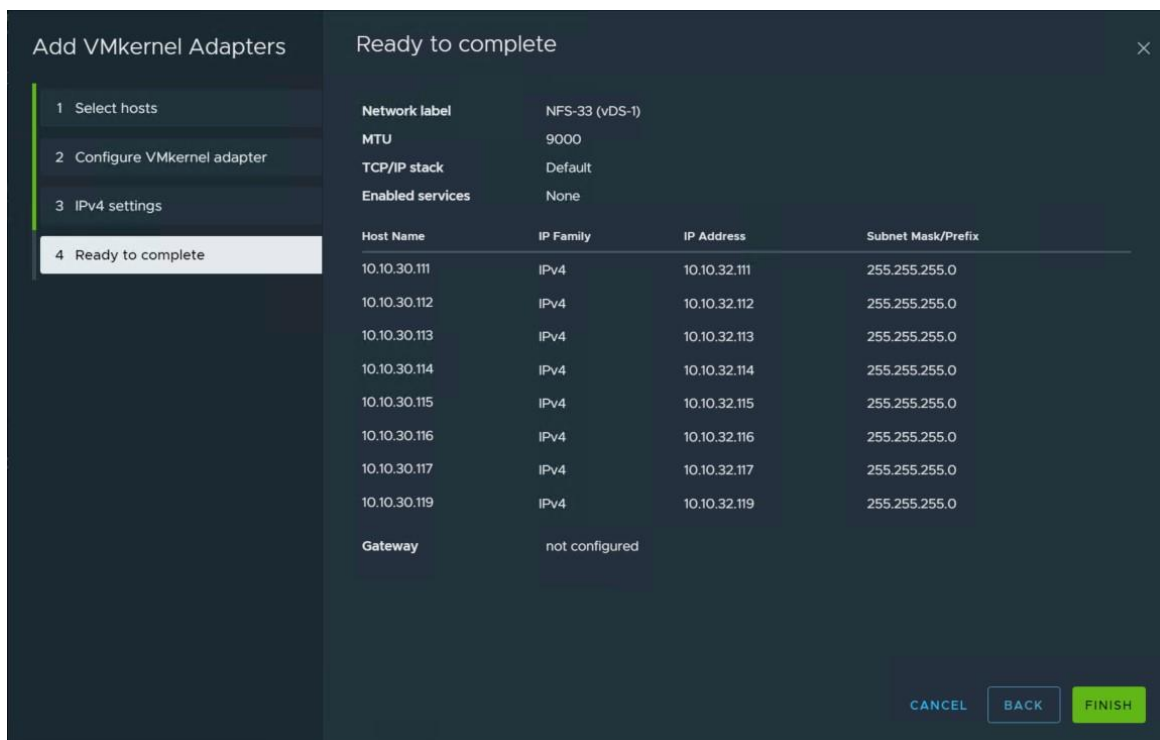
Step 3. Select service (for example vMotion) in Available services and click **NEXT**.



Step 4. Enter the Network Settings and Gateway details and click **NEXT**.



Step 5. Click **FINISH**.

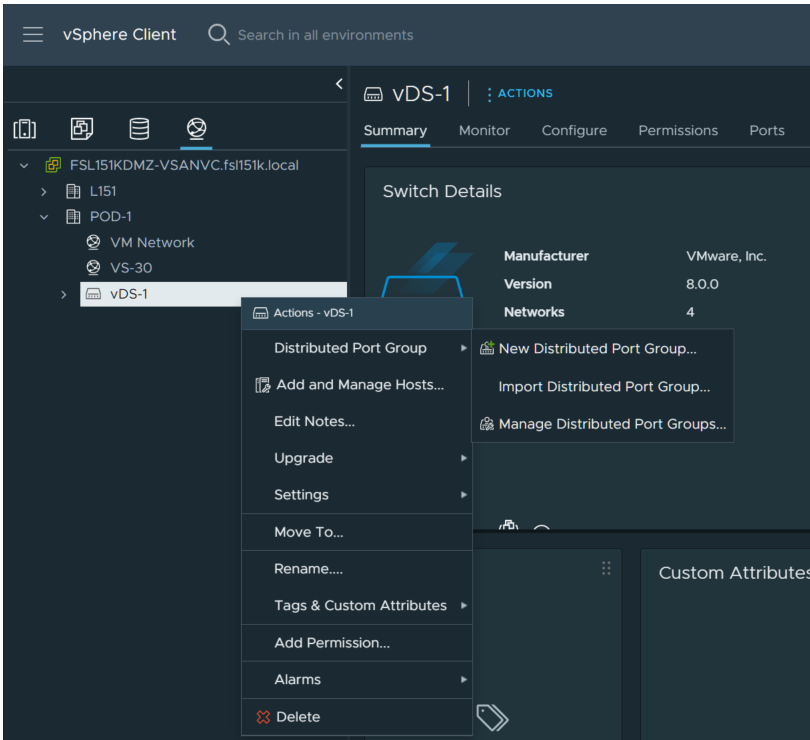


Step 6. Repeat the Add VMkernel Adapters... steps for the vMotion traffic, and any other required port groups.

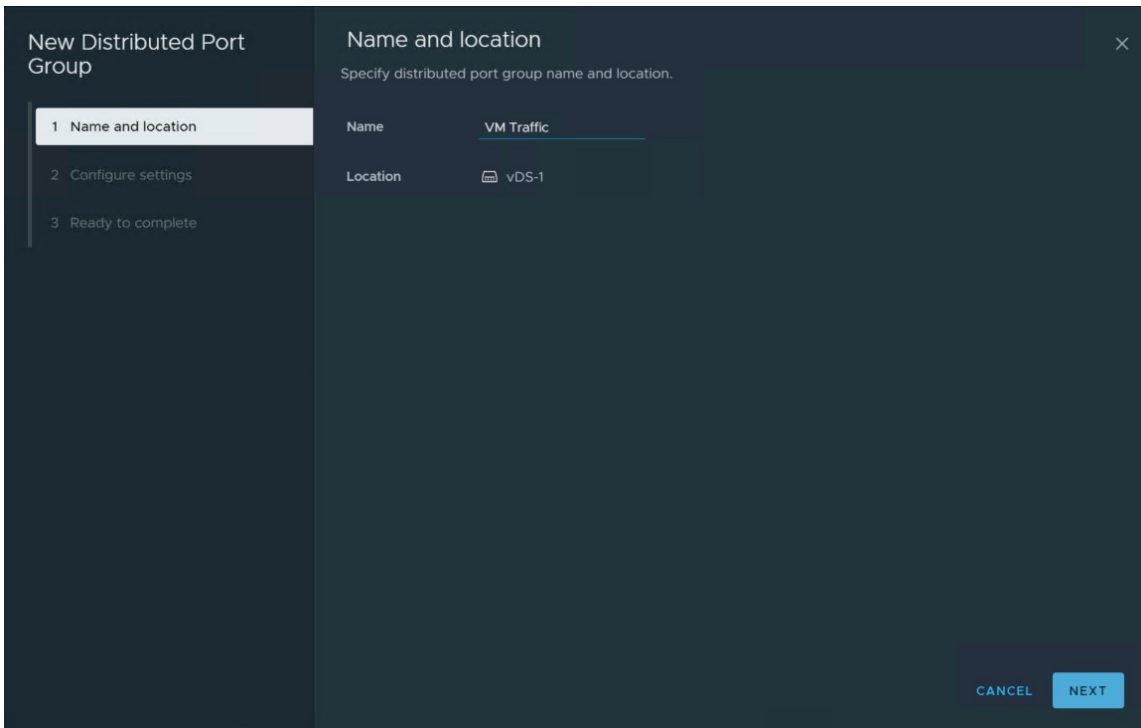
Note: The completed configuration can be verified under the Distributed Switch Topology tab.

Procedure 7. Creating a Distributed Port Group for desktop traffic

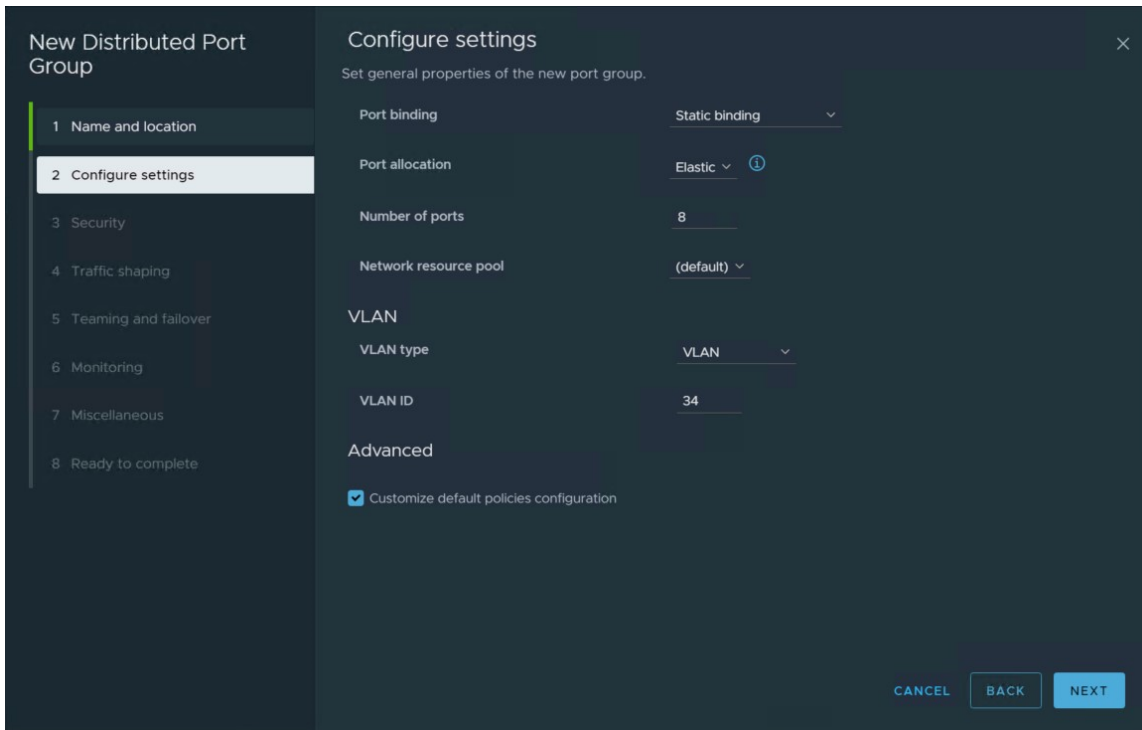
Step 1. Right-click **Distributed switch** and select **Distributed Port Group** then click **New Distributed Port Group**.



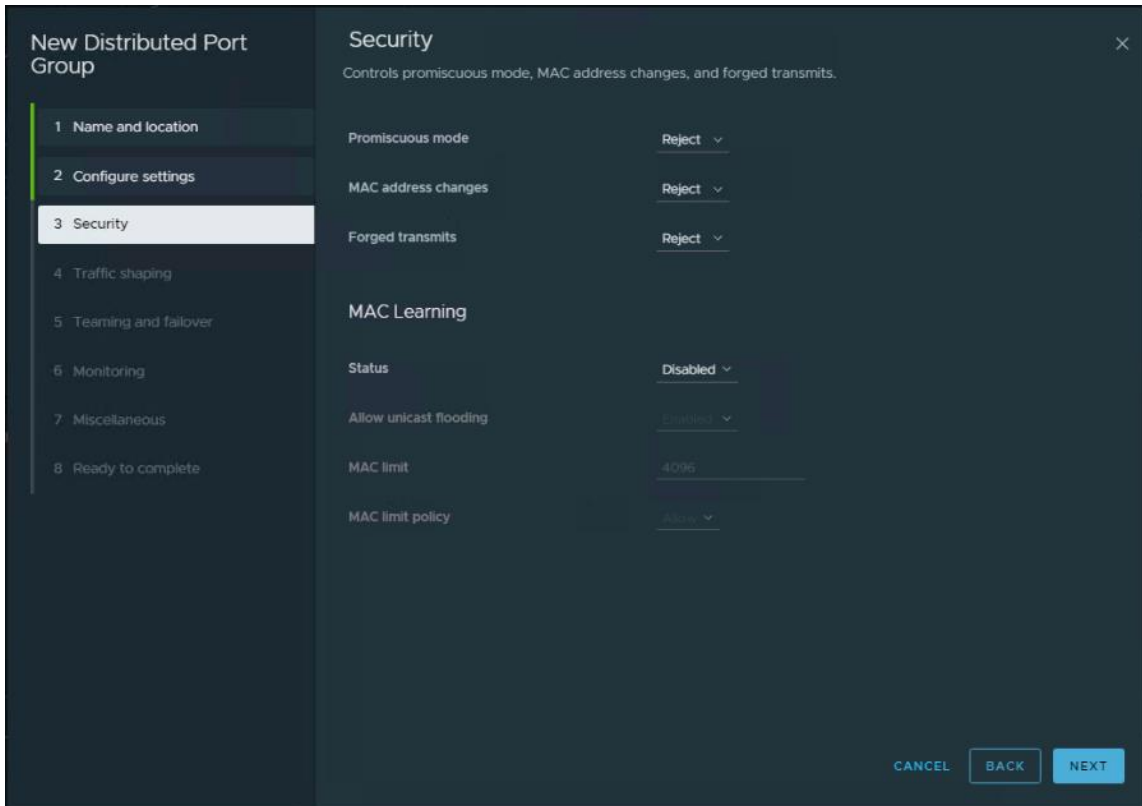
Step 2. On the New Distributed Port Group dialog box, enter a Name (for example VM Traffic), and click **NEXT**.




Step 3. In the VLAN type field, select **VLAN**, and set the VLAN ID to your VLAN (-for example 74). Check the **Customize default policies configuration** checkbox and click **NEXT**.



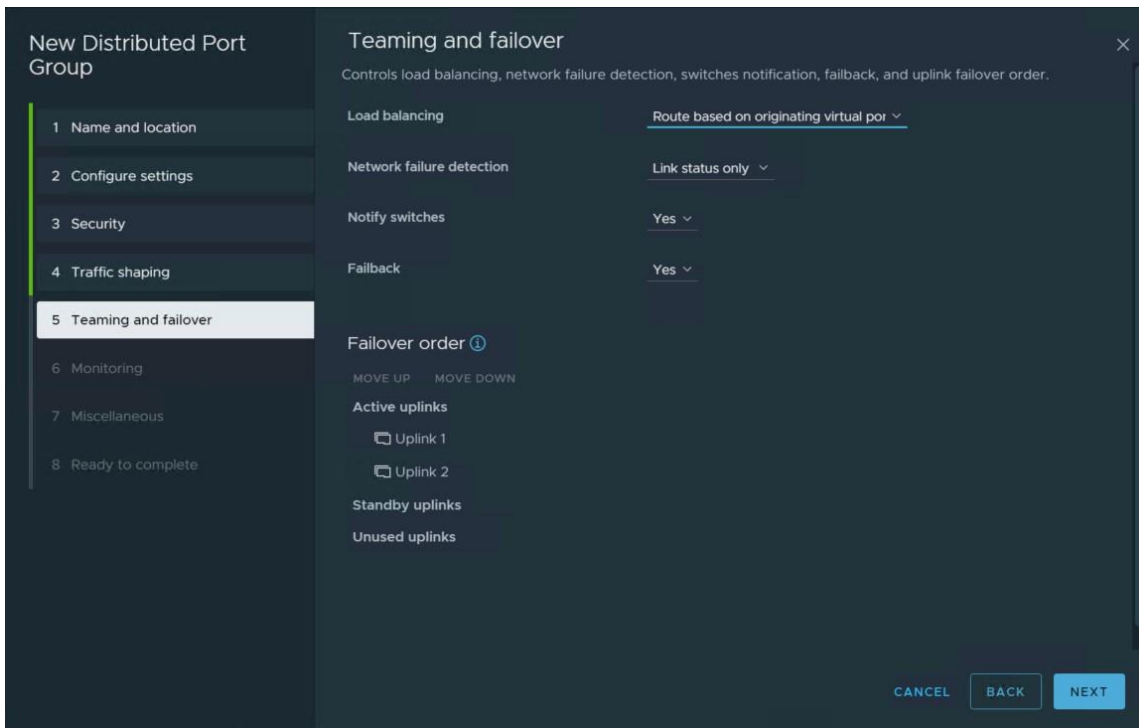
Step 4. On the Security dialog box, click **NEXT**.



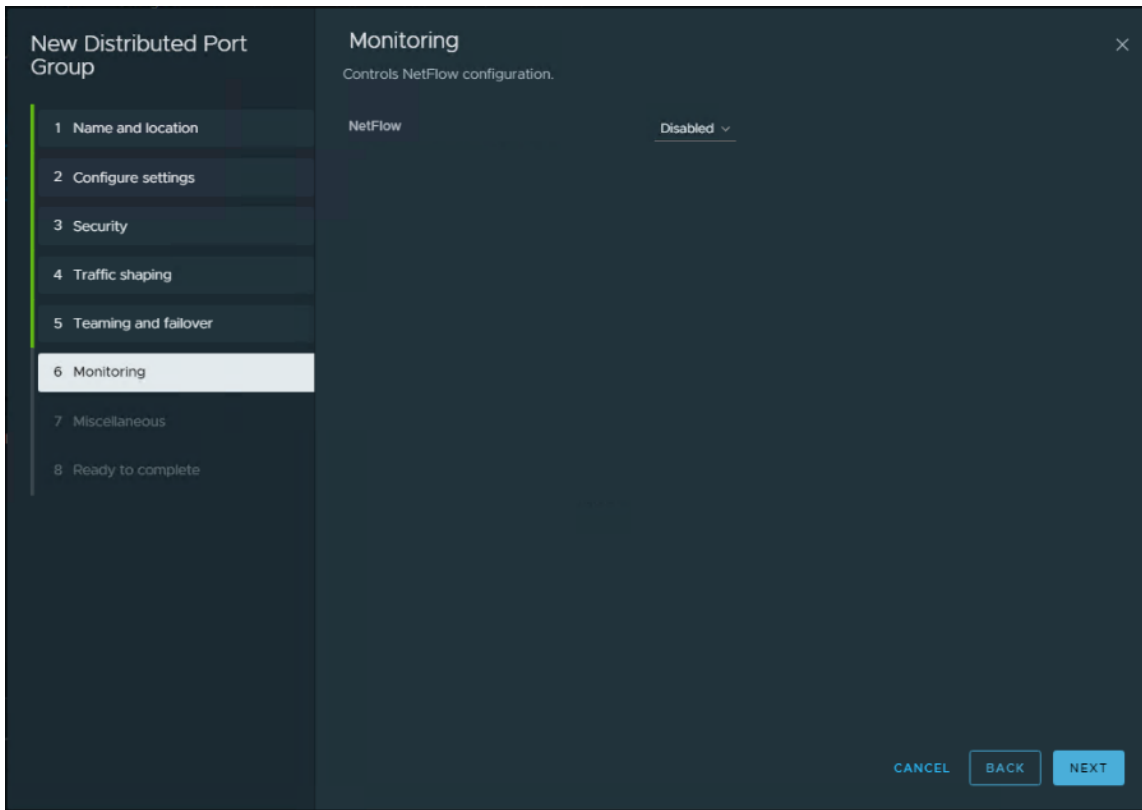
Step 5. On the Traffic shaping dialog box, click **NEXT**.



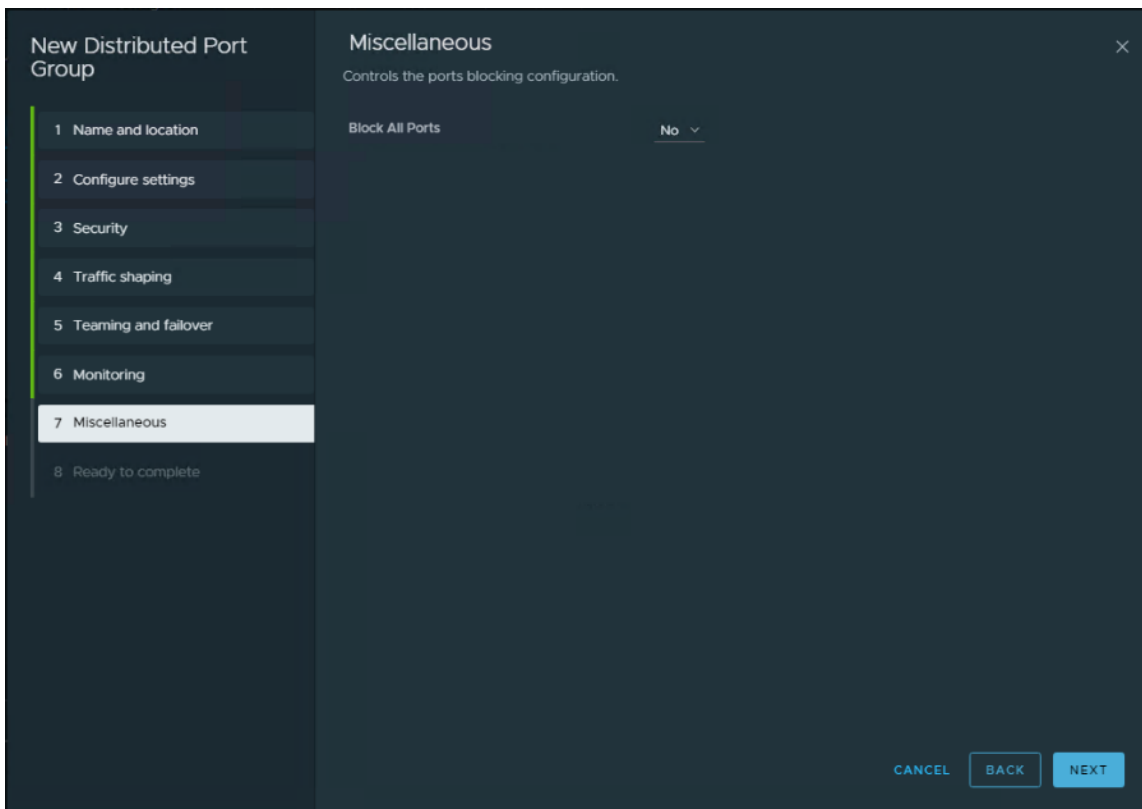
Step 6. In the Teaming and failover dialog box, select Uplink 2 as active uplink, and set Uplink 1 to be the standby uplink. Click **NEXT**.



Step 7. In the Monitoring dialog box, set NetFlow to **Disabled**, and click **NEXT**.



Step 8. In the Miscellaneous dialog box, set Block All Ports to **No**, and click **NEXT**.



Step 9. In the Ready to complete dialog box, review all the changes, and click **FINISH**.

New Distributed Port Group

- 1 Name and location
- 2 Configure settings
- 3 Security
- 4 Traffic shaping
- 5 Teaming and failover
- 6 Monitoring
- 7 Miscellaneous
- 8 Ready to complete**

Ready to complete

Review the changes before proceeding.

Distributed port group name	VM Traffic
Port binding	Static binding
Number of ports	8
Port allocation	Elastic
Network resource pool	(default)
VLAN ID	34

CANCEL BACK FINISH

Finalize the NetApp ONTAP Storage Configuration

This chapter contains the following:

- [Configure DNS for infrastructure SVM](#)
- [Create and enable auditing configuration for the SVM](#)
- [Delete the residual default broadcast domains with ifgroups \(Applicable for 2-node cluster only\)](#)
- [Configure and Test AutoSupport](#)

This chapter contains the procedures to finalize the NetApp controller configuration.

Note: ONTAP Boot storage setup is not required for this solution as it uses local boot using M.2 drives

Procedure 1. Configure DNS for infrastructure SVM

Step 1. To configure DNS for the Infra-SVM, run the following command:

```
dns create -vserver <vserver-name> -domains <dns-domain> -nameserve <dns-servers>
```

Example:

```
dns create -vserver Infra-SVM -domains flexpodb4.cisco.com -nameservers 10.102.1.151,10.102.1.152
```

Procedure 2. Create and enable auditing configuration for the SVM

Step 1. To create auditing configuration for the SVM, run the following command:

```
vserver audit create -vserver Infra-SVM -destination /audit_log
```

Step 2. Run the following command to enable audit logging for the SVM:

```
vserver audit enable -vserver Infra-SVM
```

Note: It is recommended that you enable audit logging so you can capture and manage important support and availability information. Before you can enable auditing on the SVM, the SVM's auditing configuration must already exist.

Note: If you do not perform the configuration steps for the SVM, you will see a warning in AIQUM stating "Audit Log is disabled."

Procedure 3. Delete the residual default broadcast domains with ifgroups (Applicable for 2-node cluster only)

Step 1. To delete the residual default broadcast domains that are not in use, run the following commands:

```
broadcast-domain delete -broadcast-domain <broadcast-domain-name>
```

```
broadcast-domain delete -broadcast-domain Default-1  
broadcast-domain delete -broadcast-domain Default-2
```

Procedure 4. Configure and Test AutoSupport

Note: NetApp AutoSupport sends support summary information to NetApp through HTTPS.

Step 1. Test the AutoSupport configuration by sending a message from all nodes of the cluster:

```
autosupport invoke -node * -type all -message "FlexPod storage configuration completed"
```

Configuration and Installation

This chapter contains the following:

- [FlexPod Automated Deployment with Ansible](#)


FlexPod Automated Deployment with Ansible

If using the published Ansible playbooks to configure the FlexPod infrastructure, follow the procedures detailed in this section.

Ansible Automation Workflow and Solution Deployment

The Ansible automated FlexPod solution uses a management workstation (control machine) to run Ansible playbooks to configure Cisco Nexus, NetApp ONTAP Storage, Cisco UCS, Cisco MDS, and VMware ESXi.

Figure 34. High-level FlexPod Automation



To use the Ansible playbooks demonstrated in this document [Getting Started with Red Hat Ansible](#), the management workstation must also have a working installation of Git and access to the Cisco DevNet public GitHub repository. The Ansible playbooks used in this document are cloned from the public repositories, located at the following links:

- Cisco DevNet: <https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/FlexPod-IMM-VMware/>
- GitHub repository for FlexPod infrastructure setup: <https://github.com/ucs-compute-solutions/FlexPod-IMM-VMware>
- For more information on FlexPod setup using Ansible, go to https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_iac_e2d_deploy.html.

Procedure 1. Update Cisco VIC Drivers for ESXi

Note: When ESXi is installed from Cisco Custom ISO, you might have to update the Cisco VIC drivers for VMware ESXi Hypervisor to match the current [Cisco Hardware and Software Interoperability Matrix](#).

Note: The Cisco-nenic- 2.0.10.0 drivers were used in this Cisco Validated Design.

Step 1. Log into your **VMware** Account to download required drivers for FNIC and NENIC as per the recommendation.

Step 2. Enable **SSH** on ESXi to run following commands:

```
esxcli software vib update -d /path/offline-bundle.zip
```

Build the Virtual Machines and Environment for Workload Testing

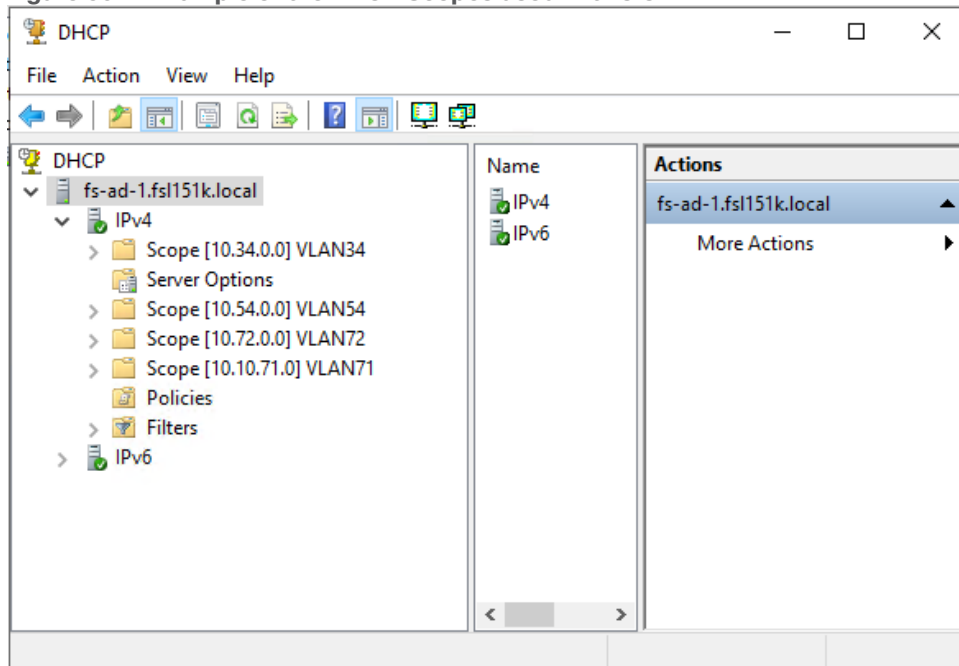
This chapter contains the following:

- [Prerequisites](#)
- [Software Infrastructure Configuration](#)
- [Prepare the Master Targets](#)
- [Install and Configure VMware Horizon](#)

Prerequisites

Create all necessary DHCP scopes for the environment and set the Scope Options.

Figure 35. Example of the DHCP Scopes used in this CVD



Software Infrastructure Configuration

This section explains how to configure the software infrastructure components that comprise this solution.

Install and configure the infrastructure virtual machines by following the process listed in [Table 19](#).

Table 19. Test Infrastructure Virtual Machine Configuration

Configuration	Microsoft Active Directory DCs Virtual Machines	vCenter Server Appliance Virtual Machine
Operating system	Microsoft Windows Server 2019	VCSA – SUSE Linux
Virtual CPU amount	4	8
Memory amount	8 GB	32 GB

Configuration	Microsoft Active Directory DCs Virtual Machines	vCenter Server Appliance Virtual Machine
Network	VMXNET3 Infra-Mgmt_71	VMXNET3 Infra-Mgmt_71
Disk size	40 GB	698.84 GB (across 13 VMDKs)

Configuration	Microsoft SQL Server Virtual Machine	VMware Horizon Connection Server Virtual Machine
Operating system	Microsoft Windows Server 2019 Microsoft SQL Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	8	8
Memory amount	24GB	16 GB
Network	VMXNET3 Infra-Mgmt_71	VMXNET3 Infra-Mgmt_71
Disk-1 (OS) size	40 GB	40 GB
Disk-2 size	100 GB SQL Databases\Logs	-

Prepare the Master Targets

This section provides guidance regarding creating the golden (or master) images for the environment. Virtual machines for the master targets must first be installed with the software components needed to build the golden images. Additionally, all available security patches as of October 2019 for the Microsoft operating systems and Microsoft Office 2021 were installed.

The single-session OS and multi-session OS master target virtual machines were configured as detailed in [Table 20](#).

Table 20. Single-session OS and Multi-session OS Virtual Machines Configurations

Configuration	Single-session OS Virtual Machine	Mutli-session OS Virtual Machine
Operating system	Microsoft Windows 11 64-bit 21H2 (19044.2006)	Microsoft Windows Server 2019 Standard 1809 (17763.3469)
Virtual CPU amount	2	8
Memory amount	3 GB reserve for all guest memory	24 GB reserve for all guest memory


Configuration	Single-session OS Virtual Machine	Mutli-session OS Virtual Machine
Network	VMXNET3 VCC/VM-Network	VMXNET3 VCC/VM-Network
vDisk size	48 GB	60 GB
Additional software used for testing	Microsoft Office 2021 Office Update applied Login VSI 4.1.40.6 Target Software (Knowledge Worker Workload)	Microsoft Office 2021 Office Update applied Login VSI 4.1.40.6 Target Software (Knowledge Worker Workload)
Additional Configuration	Configure DHCP Add to domain Install VMWare tools Install .Net 3.5 Activate Office Install Horizon Agent Install FSLogix 2210 hotfix 1	Configure DHCP Add to domain Install VMWare tools Install .Net 3.5 Activate Office Install Horizon Agent Install FSLogix 2210 hotfix 1

Procedure 1. Prepare the Master Virtual Machines

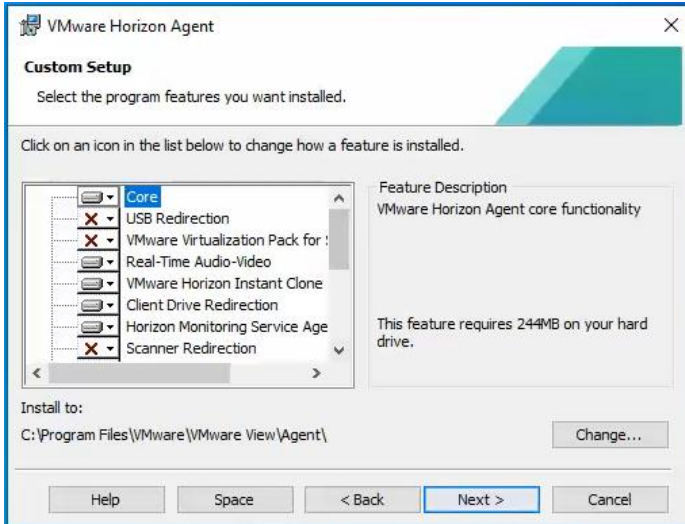
To prepare the master virtual machines, there are three major steps: installing the operating system and VMware tools, installing the application software, and installing the VMware Horizon Agent.

Note: For this CVD, the images contain the basics needed to run the Login VSI workload.

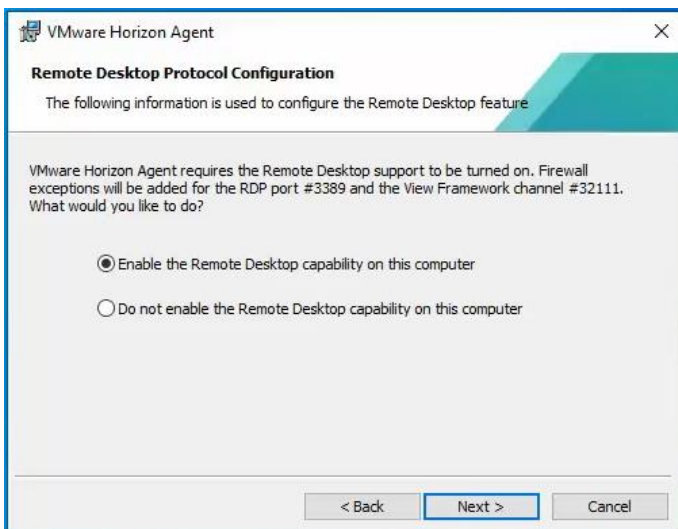
Step 1. During the VMware Horizon Agent installation, select **IPv4** for the network protocol.



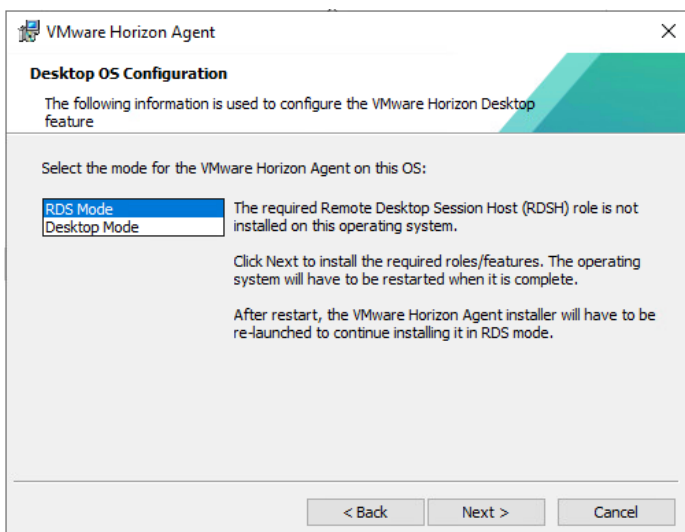
Step 2. On the Custom Setup screen, leave the defaults preparing the Instant clone master image. Deselect the VMware Horizon Instant Clone option for the Full clone master image.



Step 3. Enable the remote Desktop Protocol.



Step 4. During the VMware Horizon Agent installation on the Windows server, select RDS Mode otherwise select Desktop Mode.



The final step is to optimize the Windows OS. VMware OS Optimization Tool for Horizon includes customizable templates to enable or disable Windows system services and features using VMware recommendations and best practices across multiple systems. Because most Windows system services are enabled by default, the optimization tool can be used to easily disable unnecessary services and features to improve performance.

Note: In this CVD, the Windows OS Optimization Tool for VMware Horizon. Version 1.2.0 was used.

Home / Windows OS Optimization Tool for VMware Horizon

Download Product

Select Version	1.2.0 ▼
Documentation	Release Notes
Release Date	2023-03-30
Type	Product Binaries

Base images were optimized with Default template for Windows 11 (1809-21H2) and Windows 11 (21H2) or Server 2019 and Server 2019.

To successfully run the Login VSI knowledge worker workload the 'Disable animation in web pages - Machine Policy' option in Default Template under Programs -> Internet Explorer was disabled.

▼	Programs (75 items)			
>	.Net (5 items)			
▼	Internet Explorer (20 items)			
✓	⚠ Turn off suggestions for all user-installed providers - Machine Policy	0	NA	Turn off suggestions for all user-installed providers
✓	⚠ Disable Internet Explorer Enhanced Security - HKLM Registry	0	1	Disable Internet Explorer Enhanced Security.
✓	⚠ Disable IE Customization Wizard - Machine Policy	1	NA	Removes the customization wizard upon first launch of Internet Explor
✓	⚠ Disable Background Synchronization - Machine Policy	0	NA	Turn off background synchronization for feeds and Web Slices
✓	⚠ Turn off the next pre-loaded page of a website - Machine Policy	0	NA	Turn off the flip ahead with page prediction feature
<input type="checkbox"/>	⚠ Disable animations in web pages - Machine Policy	no	NA	Only animated GIF files are affected by this setting

Install and Configure FSLogix

FSLogix, a Microsoft tool, was used to manage user profiles in this Cisco Validated Design.

A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Profile management in VDI environments is an integral part of the user experience.

FSLogix allows you to:

- Roam user data between remote computing session hosts
- Minimize sign in times for virtual desktop environments
- Optimize file IO between host/client and remote profile store
- Provide a local profile experience, eliminating the need for roaming profiles.
- Simplify the management of applications and 'Gold Images'

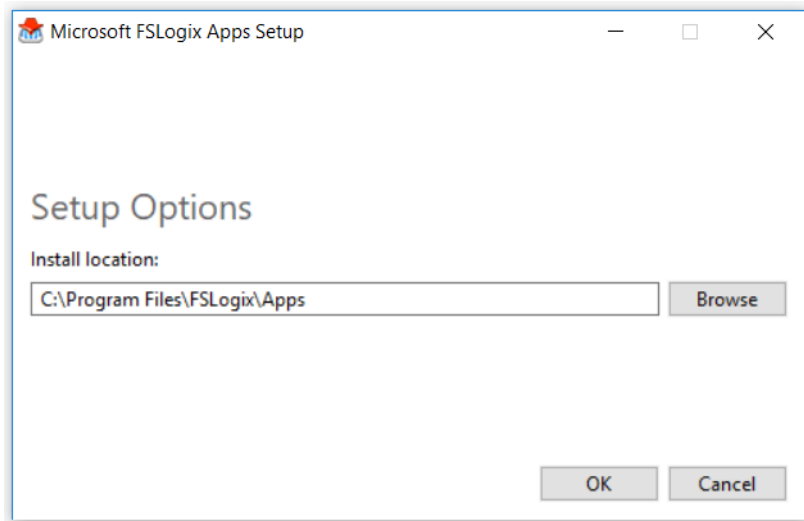
Additional documentation about the tool can be found [here](#).

Procedure 1. FSLogix Apps Installation

Step 1. Download the FSLogix file [here](#).

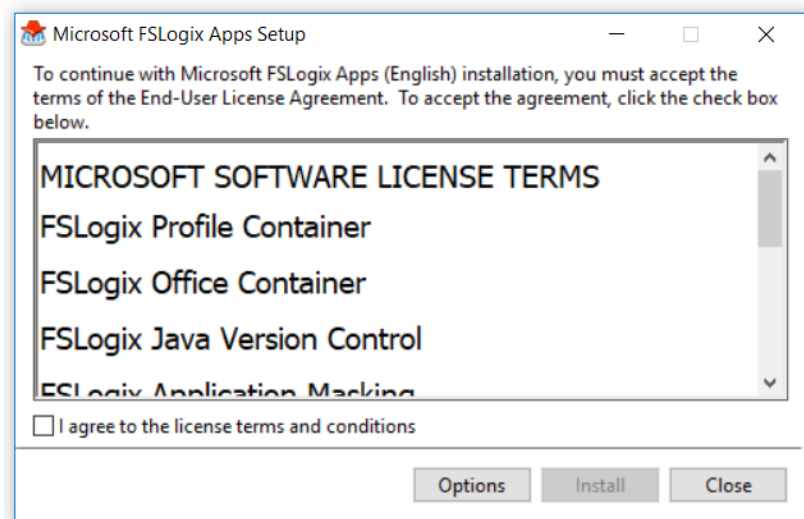
Step 2. Run FSLogixAppSetup.exe on VDI master image (32 bit or 64 bit depending on your environment).

Step 3. Click **OK** to proceed with default installation folder.



Step 4. Review and accept the license agreement.

Step 5. Click **Install**.



Step 6. Reboot.

Procedure 2. Configure Profile Container Group Policy

Step 1. Copy "fslogix.admx" to C:\Windows\PolicyDefinitions, and "fslogix.adml" to C:\Windows\PolicyDefinitions\en-US on Active Directory Domain Controllers.

Step 2. Create FSLogix GPO and apply to the desktops OU. Navigate to **Computer Configuration > Administrative Templates > FSLogix > Profile Containers**.

Step 3. Configure the following settings:

- Enabled - **Enabled**
- VHD location - **Enabled**, with the path set to \\<FileServer>\<Profiles Directory>

Note: Consider enabling and configuring FSLogix logging as well as limiting the size of the profiles and excluding additional directories.

Figure 36. Example of FSLogix Policy

FSLogix		
Policy	Setting	Comment
Days to keep log files	Enabled	
Days to keep log files	3	
Enable logging	Enabled	
Enable logging: FSLogix agent service	Enabled	Only specifically enabled logs
Enable logging: FSLogix agent service	Enabled	Enabled
Enable logging: Profiles	Enabled	
Enable logging: Profiles	Enabled	Enabled
FSLogix/Profile Containers		
Delete local profile when FSLogix Profile should apply	Enabled	
Delete local profile when FSLogix Profile should apply	Enabled	Enabled
Dynamic VHD(X) allocation	Enabled	
Dynamic VHD(X) allocation	Enabled	Enabled
Enabled	Enabled	
Enabled	Enabled	Enabled
Profile type	Enabled	
Profile type	Enabled	Try for read-write profile and fallback to read-only
Size in MBs	Enabled	
Size in MBs	Enabled	2048
VHD location	Enabled	
VHD location	Enabled	\\10.10.32.120\cfs_vol_04\VDI
FSLogix/Profile Containers/Advanced		
Locked VHD retry count	Enabled	
Locked VHD retry count	Enabled	3
Provide RedirXML file to customize redirections	Enabled	
Provide RedirXML file to customize redirections	Enabled	\\fa151k.local\NETLOGON\fslogirect.xml

Figure 37. FSLogix policy exclusions list

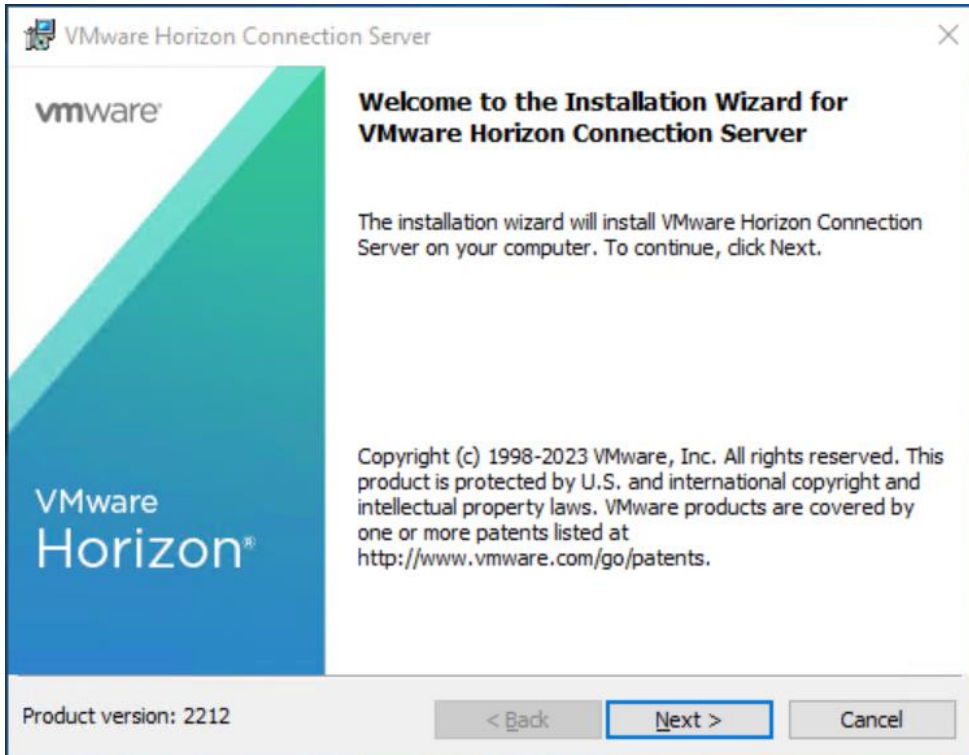
```
<?xml version="1.0" encoding="UTF-8"?>
- <FrxFolderRedirection ExcludeCommonFolders="###VALUE###">
  - <Excludes>
    <Exclude Copy="0">Videos</Exclude>
    <Exclude Copy="0">Saved Games</Exclude>
    <Exclude Copy="0">Contacts</Exclude>
    <Exclude Copy="0">Searches</Exclude>
    <Exclude Copy="0">Citrix</Exclude>
    <Exclude Copy="0">Tracing</Exclude>
    <Exclude Copy="0">Music</Exclude>
    <Exclude Copy="0">$Recycle.Bin</Exclude>
    <Exclude Copy="0">AppData\LocalLow\Adobe</Exclude>
    <Exclude Copy="0">AppData\LocalLow\Microsoft</Exclude>
    <Exclude Copy="0">AppData\Local\Apps</Exclude>
    <Exclude Copy="0">AppData\Local\Downloaded Installations</Exclude>
    <Exclude Copy="0">AppData\Local\assembly</Exclude>
    <Exclude Copy="0">AppData\Local\CEF</Exclude>
    <Exclude Copy="0">AppData\Local\Comms</Exclude>
    <Exclude Copy="0">AppData\Local\Deployment</Exclude>
    <Exclude Copy="0">AppData\Local\FSLogix</Exclude>
    <Exclude Copy="0">AppData\Local\Packages</Exclude>
    <Exclude Copy="0">AppData\Local\VirtualStore</Exclude>
    <Exclude Copy="0">AppData\Local\CrashDumps</Exclude>
    <Exclude Copy="0">AppData\Local\Package Cache</Exclude>
    <Exclude Copy="0">AppData\Local\D3DSCache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\TokenBroker\Cache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Notifications</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Internet Explorer\DOMStore</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Internet Explorer\Recovery</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\MSOIdentityCRL\Tracing</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Messenger</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Terminal Server Client</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\UEV</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\Application Shortcuts</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\Mail</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\WebCache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\WebCache.old</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\AppCache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\Explorer</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\GameExplorer</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\DNTException</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\IECompatCache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\iecompatuaCache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\Notifications</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\PRICache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\PrivacIE</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\RoamingTiles</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\SchCache</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\Temporary Internet Files</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\0030</Exclude>
    <Exclude Copy="0">AppData\Local\Microsoft\Windows\1031</Exclude>
    <Exclude Copy="0">AppData\Roaming\com.adobe.formscentral.FormsCentralForAcrobat</Exclude>
    <Exclude Copy="0">AppData\Roaming\Adobe\Acrobat\DC</Exclude>
    <Exclude Copy="0">AppData\Roaming\Adobe\SLData</Exclude>
    <Exclude Copy="0">AppData\Roaming\Microsoft\Document Building Blocks</Exclude>
    <Exclude Copy="0">AppData\Roaming\Microsoft\Windows\Network Shortcuts</Exclude>
    <Exclude Copy="0">AppData\Roaming\Microsoft\Windows\Printer Shortcuts</Exclude>
    <Exclude Copy="0">AppData\Roaming\ICAClient\Cache</Exclude>
    <Exclude Copy="0">AppData\Roaming\Macromedia\Flash Player\macromedia.com\support\flashplayer</Exclude>
  </Excludes>
</FrxFolderRedirection>
```

Install and Configure VMware Horizon

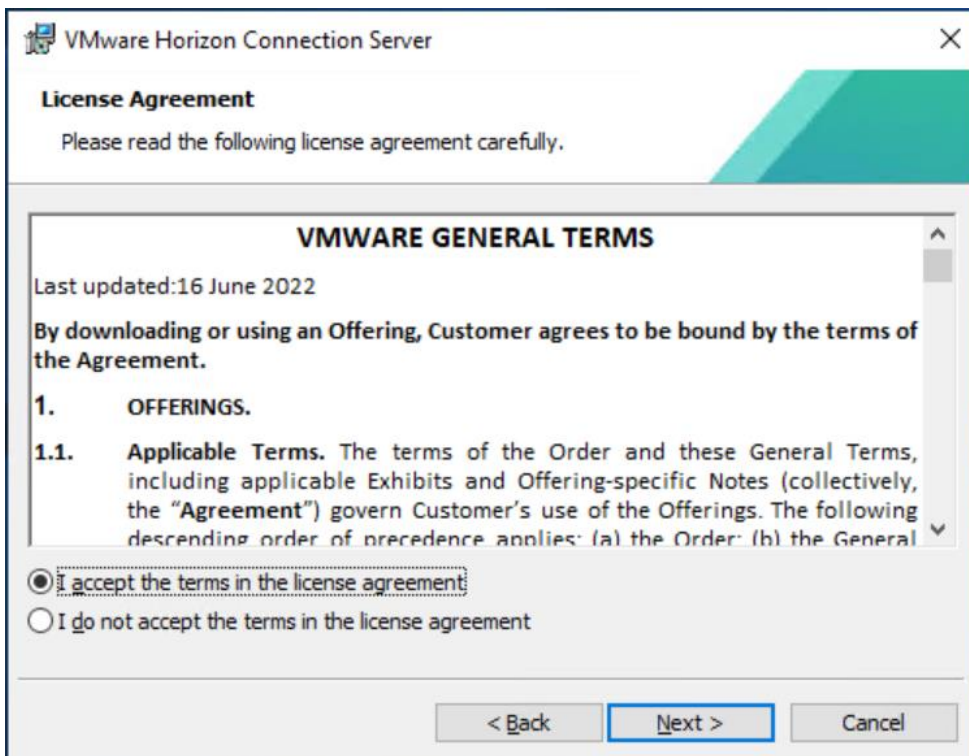
Procedure 1. Configure VMware Horizon Connection Server

Step 1. Download the Horizon Connection server installer from VMware and click **Install on the Connection Server Windows Server Image**. In this study, we used version Connection Server Horizon 8 2212 build 8.8.0. 21073894.


Step 2. Click **Next**.



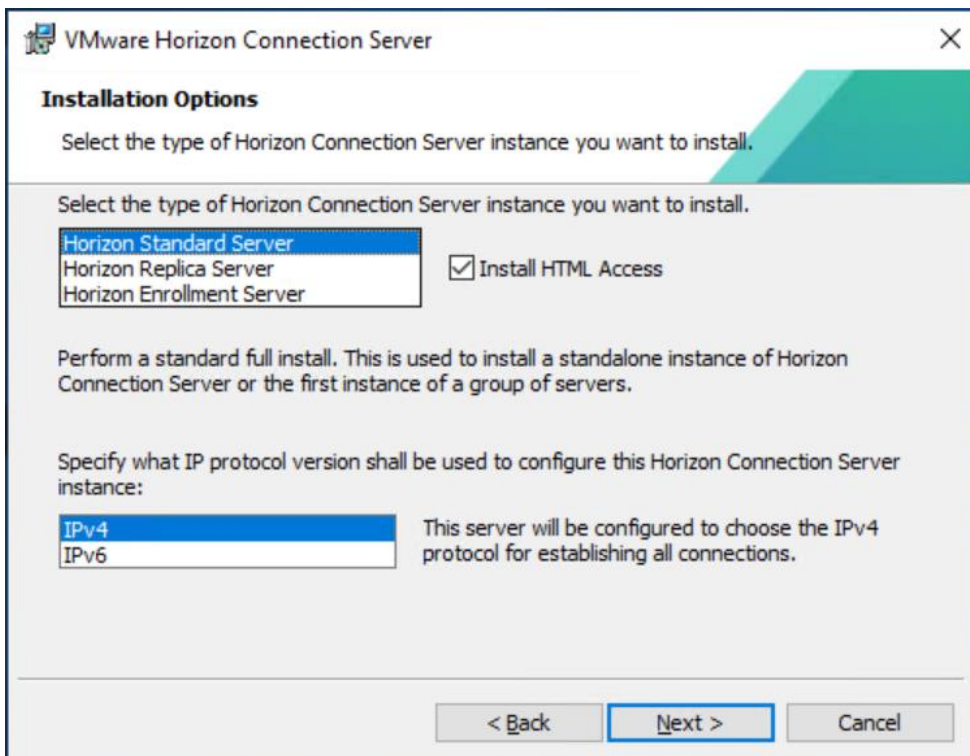
Step 3. Read and accept the End User License Agreement and click **Next**.



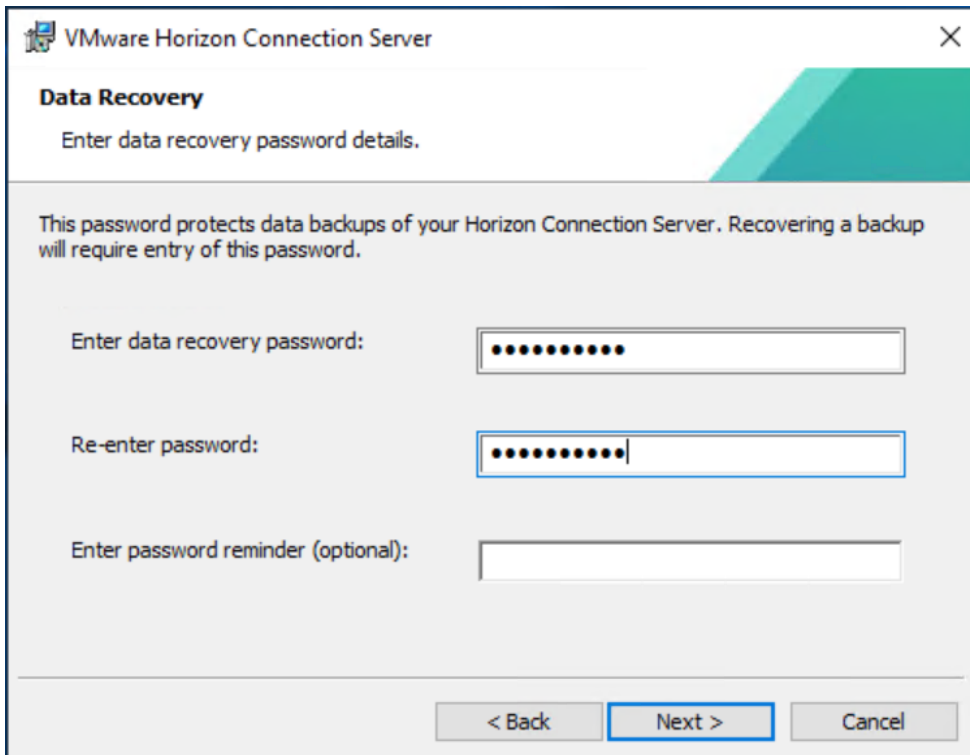
Step 4. Select the destination folder where you want to install the application and click **Next**.



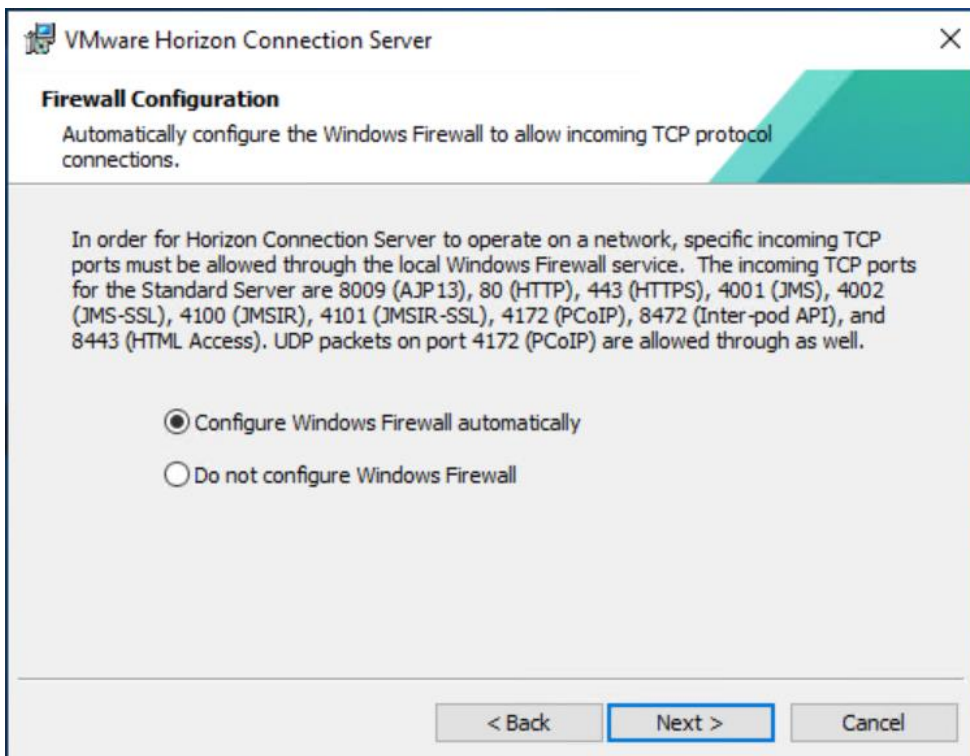
Step 5. Select the **Standard Server** and **IPv4** for the IP protocol version.



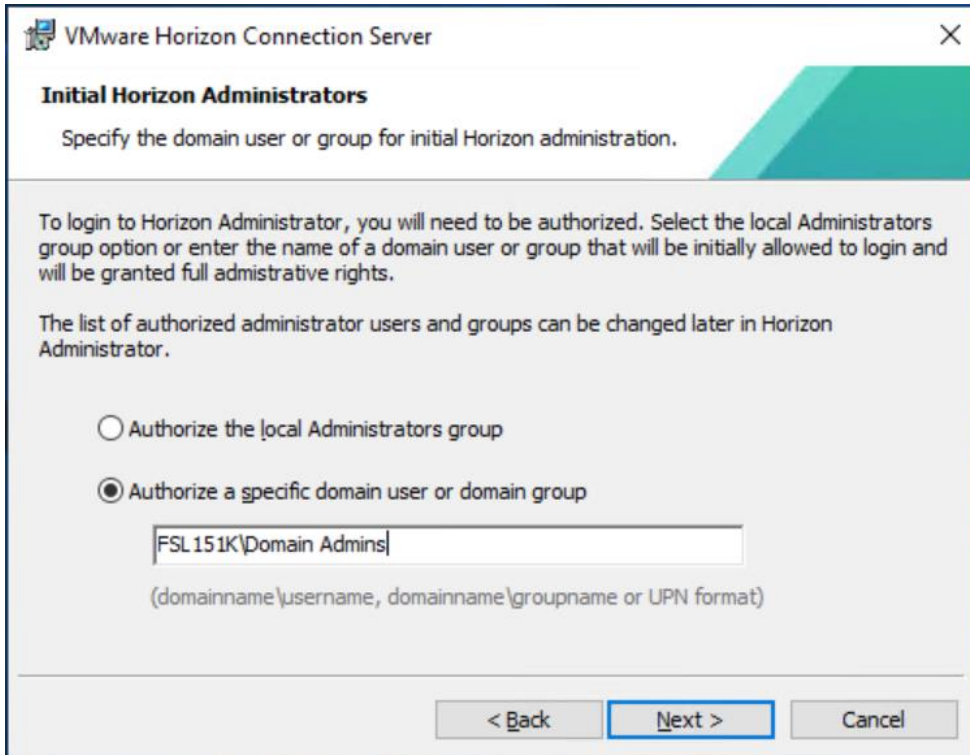
Step 6. Provide data recovery details.



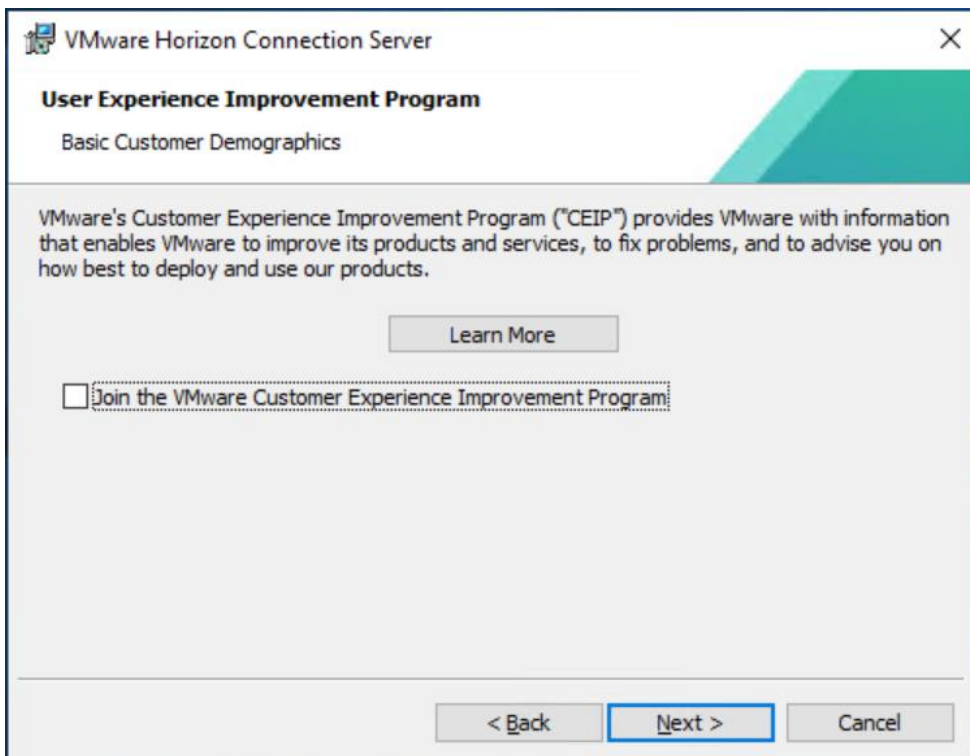
Step 7. Select **Configure Windows Firewall automatically**. Click **Next**.



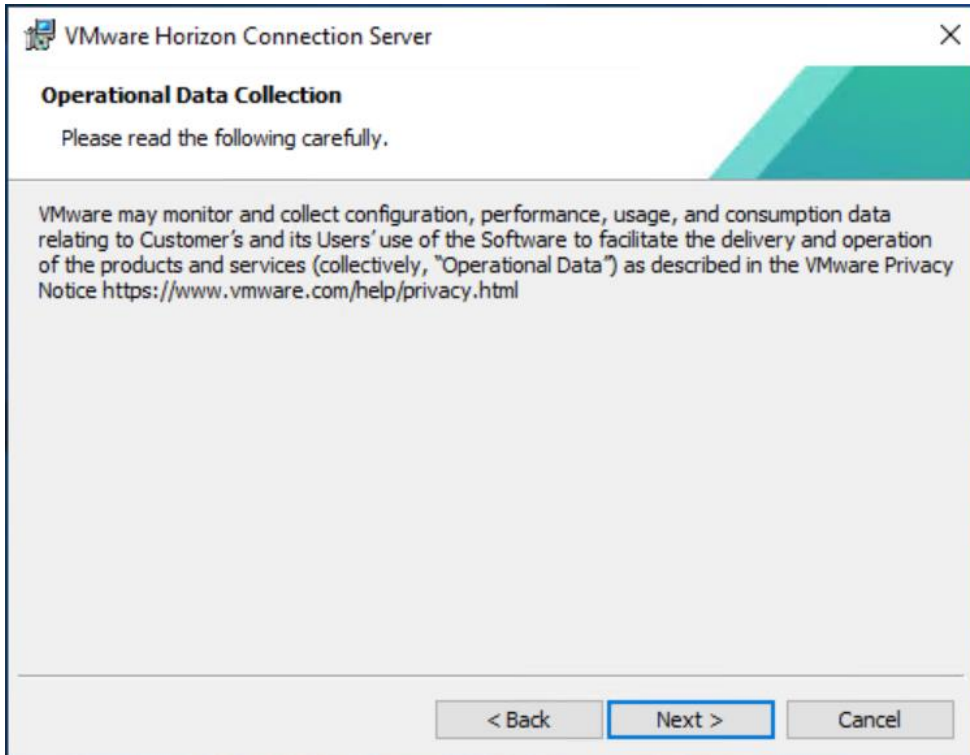
Step 8. Authorize Domain Admins to be VMware Horizon administrators.



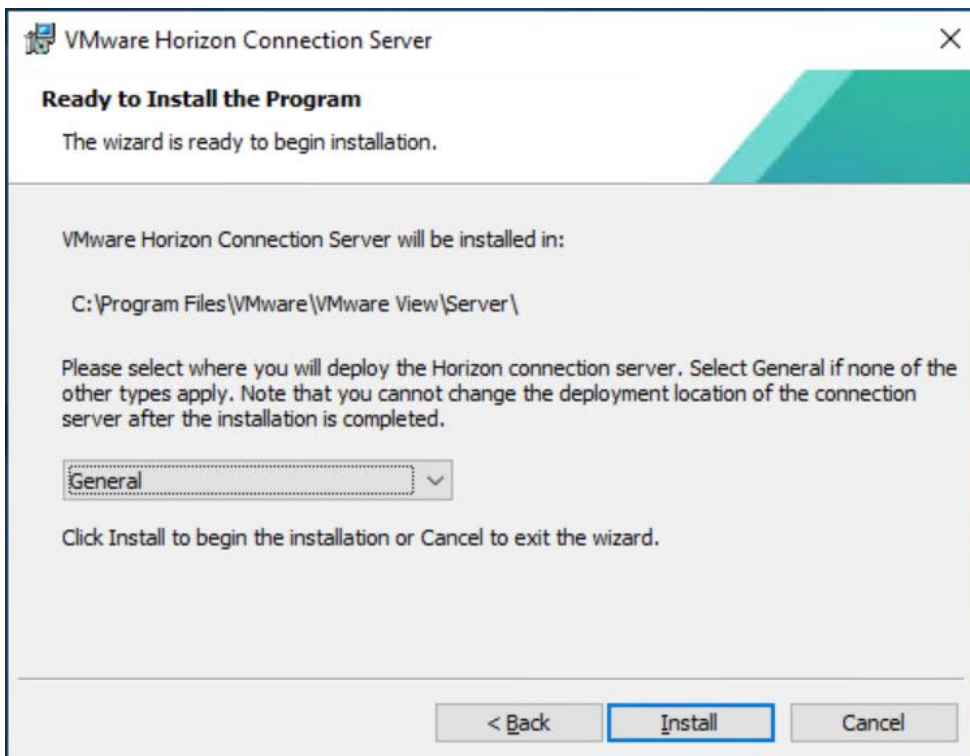
Step 9. (Optional) Join Customer Experience Program.




Step 10. Click Next.



Step 11. Select **General** for the type of the type of installation. Click **Install**.



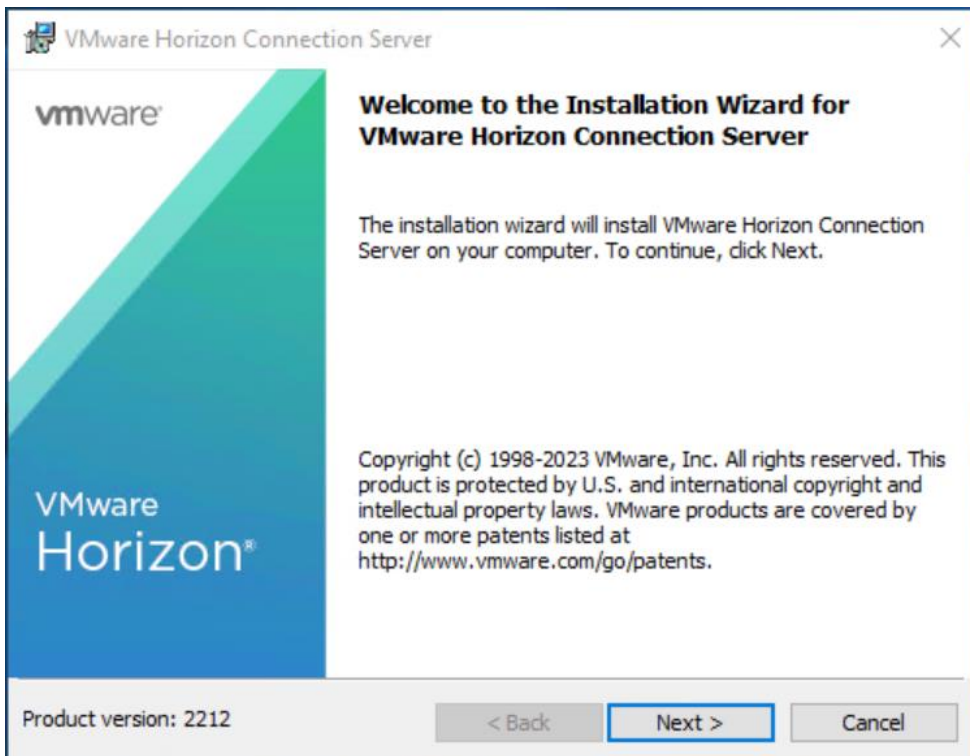
Step 12. After Horizon Connection Server installation is complete, click **Finish**.



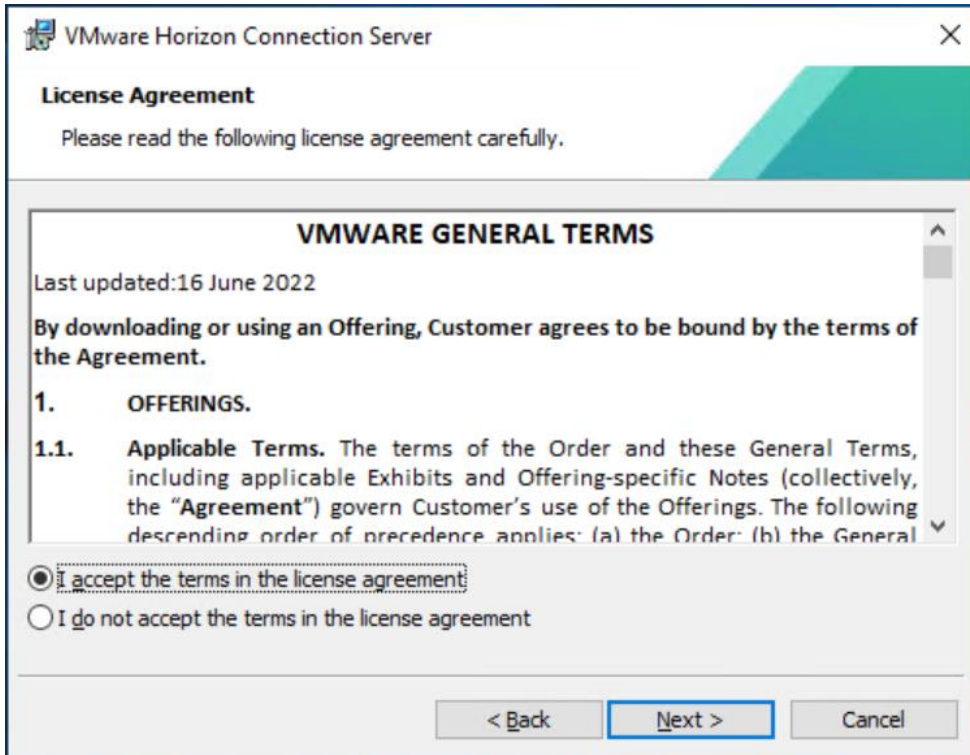
Procedure 2. Install VMware Horizon Replica Server

Step 1. Click the Connection Server installer based on your Operating System.


Step 2. Click **Next**.



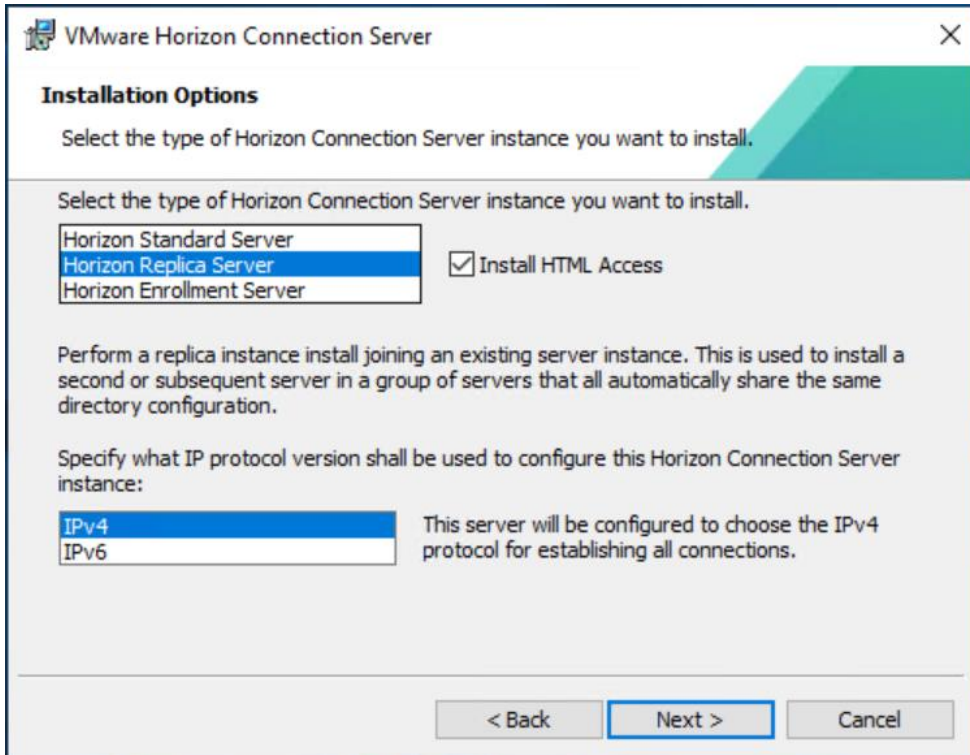
Step 3. Read and accept the End User License Agreement and click **Next**.



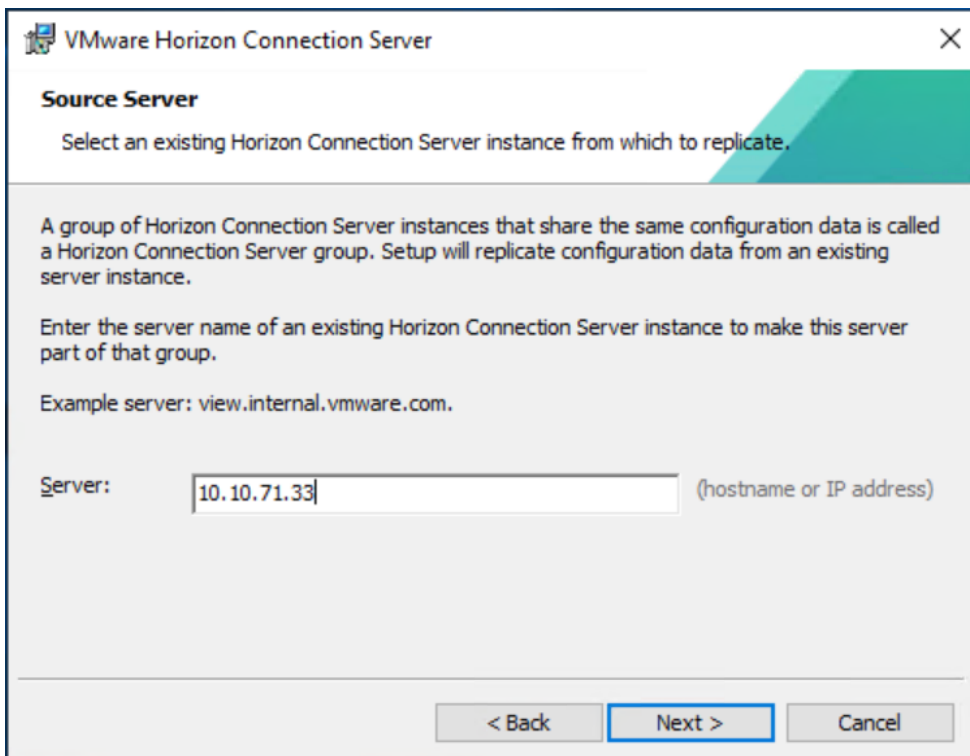
Step 4. Select the destination folder where you want to install the application and click **Next**.



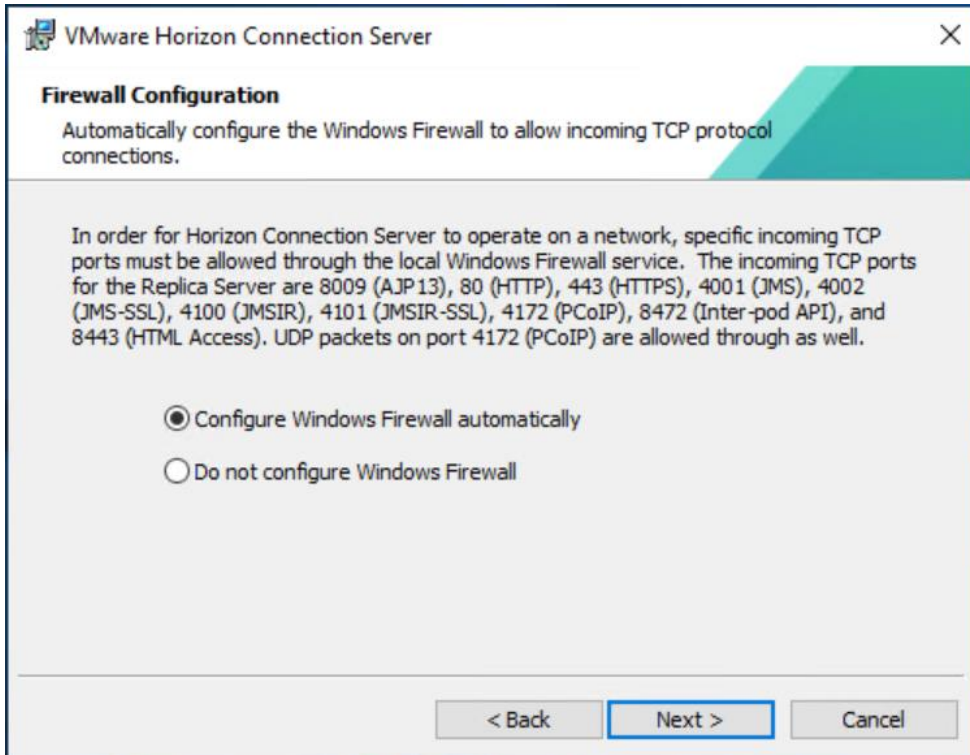
Step 5. Select the **Replica Server** and **IPv4** for the IP protocol version.




Step 6. Provide the existing Standard View Connection Server's FQDN or IP address and click **Next**.




Step 7. Select **Configure the Windows Firewall automatically**.



Step 8. Click **Install** to begin the installation process.



Step 9. After installation is complete, click **Finish**.



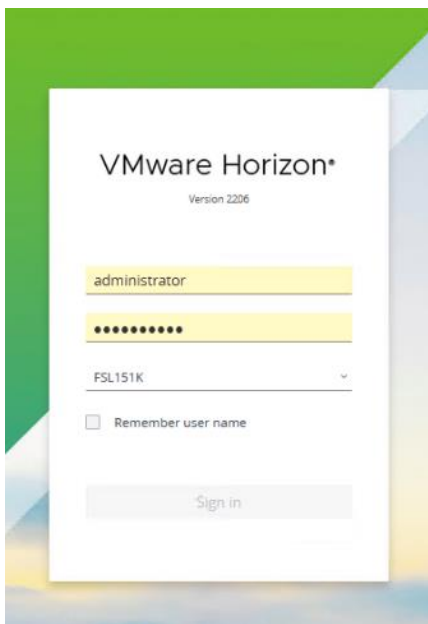
VMware Horizon Desktop Configuration

Management of the desktops, application pools and farms is accomplished in VMware Horizon Console (HTML5) or Horizon Administrator (Flex). We used Horizon Console to administer VMware Horizon environment in this validated design.


Note: VMware recommends using Horizon Console, an HTML5 based interface with enhanced security, capabilities, and performance.

Procedure 1. Configure VMware Horizon Desktop

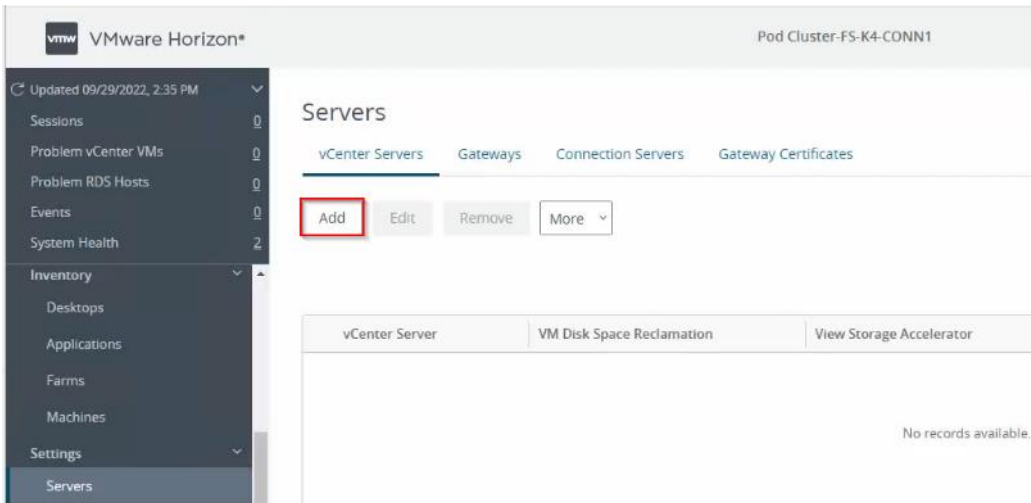
Step 1. Log into **Horizon Console 2212** via a web browser using Address or FQDN>/admin/#/login.



Step 2. In Horizon Console, expand Settings and click **Servers**.



Step 3. Select the **vCenter Settings** tab and click **Add**.



Step 4. Provide Server Address (IP or FQDN) and credentials that Horizon will use to login to vCenter, then click **Next**.

Note: In the environment used to deploy full clone virtual desktops, it is recommended to set up Max Provision value to 5 and limit the number of virtual machines provisioned at a time to a hundred to allow deduplication and compression engines to perform optimally.

Add vCenter Server

1 vCenter Information

2 Storage

3 Ready to Complete

Asterisk (*) denotes required field

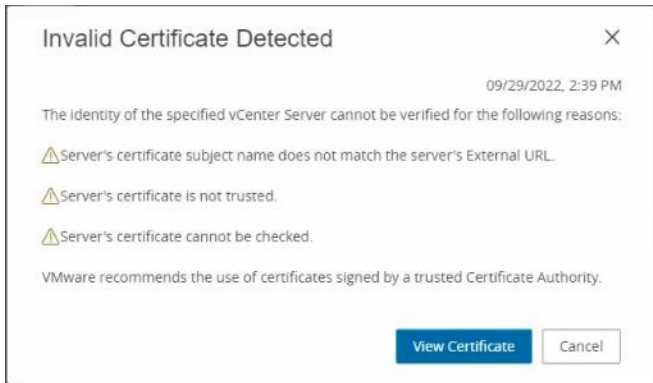
* Server address ⓘ

* User Name

* Password

Description

Step 5. If you receive a message stating an invalid certificate, click **View Certificate**.



Step 6. Click **Accept**.



Step 7. Keep the defaults, select **Reclaim VM disk space** and **Enable Horizon Storage Accelerator** with cache size of 1024MB. Click **Next**.

Add vCenter Server

1 vCenter Information

2 Storage

3 Ready to Complete

Storage Settings ⓘ

- Reclaim VM disk space
- Enable View Storage Accelerator

Default Host Cache Size MB

Cache must be between 100 MB and 32,768 MB.

Step 8. Review the information you provided and click **Submit**.

Add vCenter Server

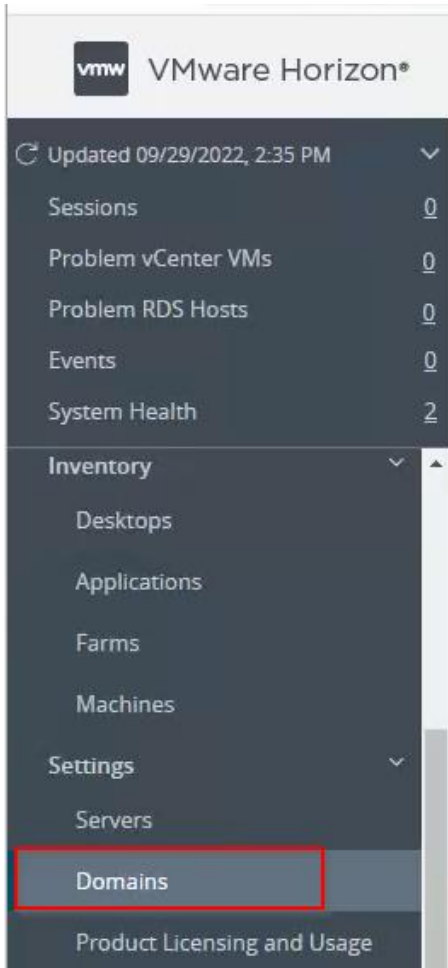
1 vCenter Information

2 Storage

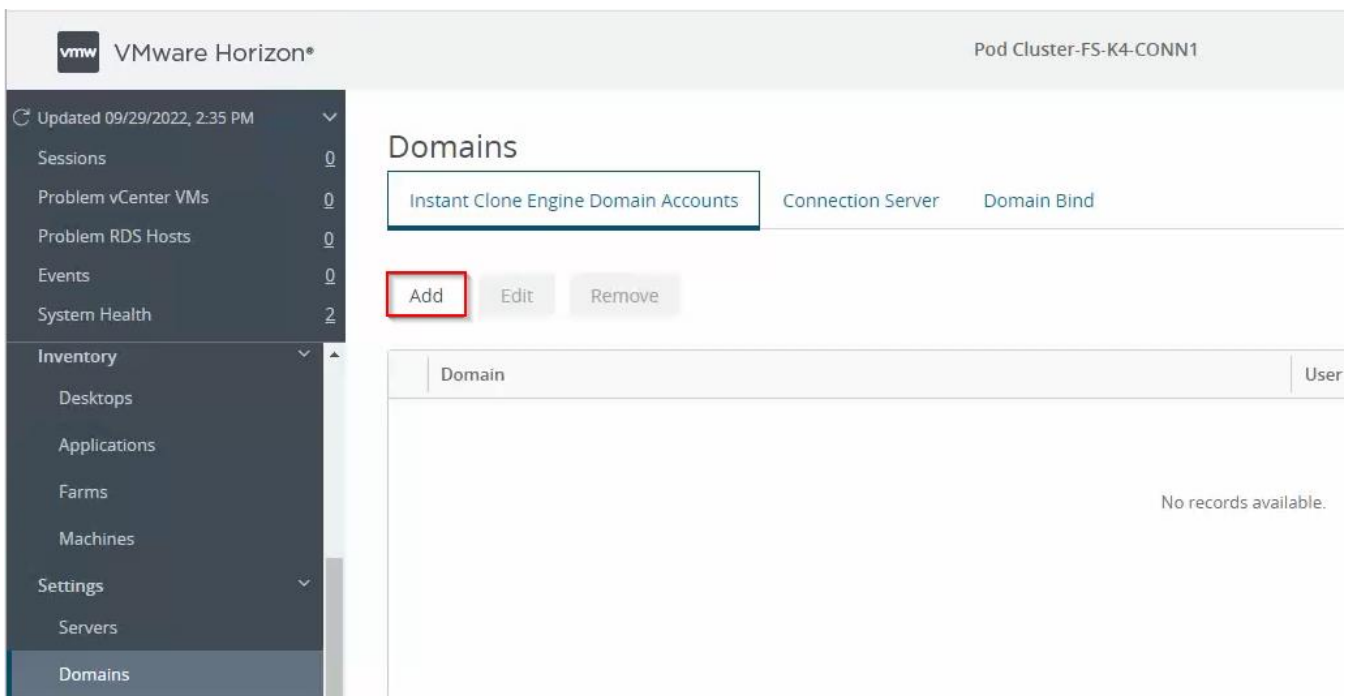
3 Ready to Complete

vCenter Server	10.10.70.32
User Name	administrator@vsphere.local
Password	*****
Description	.
Server Port	443
Max Provision	20
Max Power	50
Max concurrent maintenance operations	12
Max Instant Clone Engine Provision	20
Enable View Storage Accelerator	Yes
Default host cache size (MB)	1,024
VM Disk Space Reclamation	Yes
Deployment Type	General

Step 9. In Horizon Console, expand **Settings** and click **Domains**.



Step 10. Select the **Instant Clone Engine Domain Accounts** tab and click **Add**.



Step 11. Provide a domain name and credentials that Horizon will use to login to AD during Instant Clone management tasks, then click **OK**.

Add Domain Admin ✕

Asterisk (*) denotes required field

Full Domain Name

* Username

* Password

Procedure 2. Create VDI Instant Clone Desktop Pool

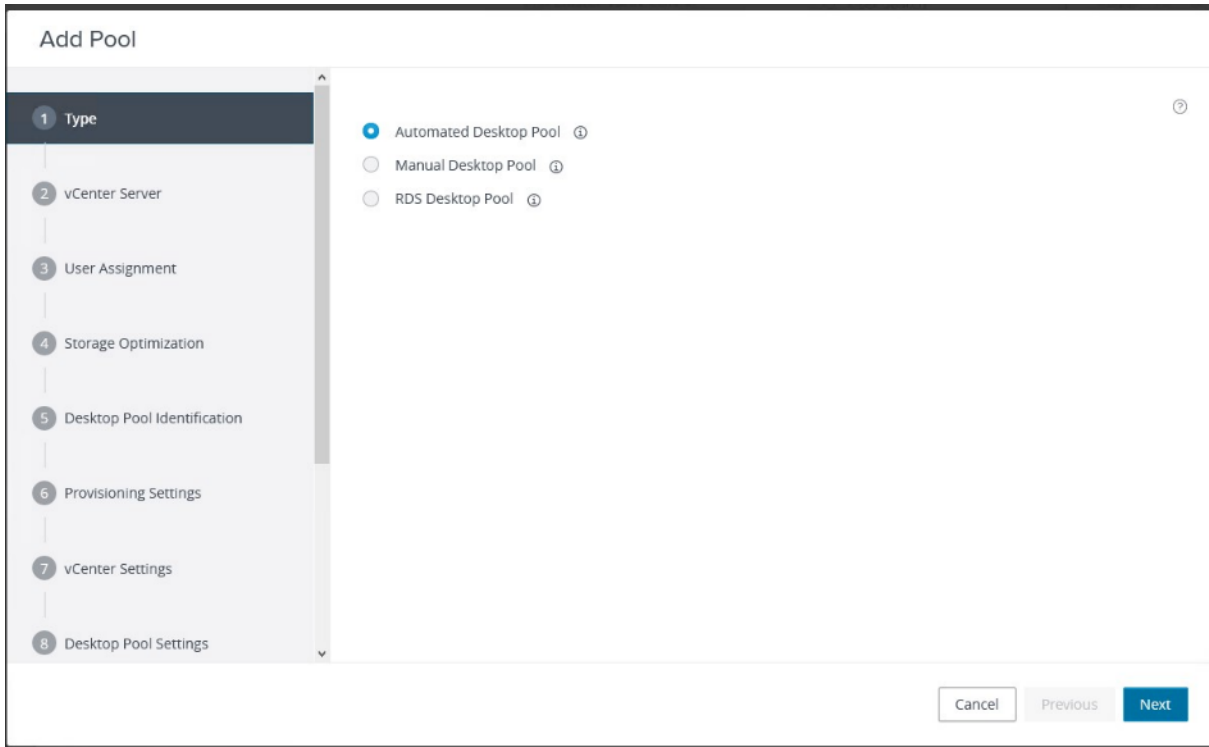
Step 1. In Horizon Console on the left pane, expand **Inventory**, select **Desktops**. Click **Add**.

Desktop Pools

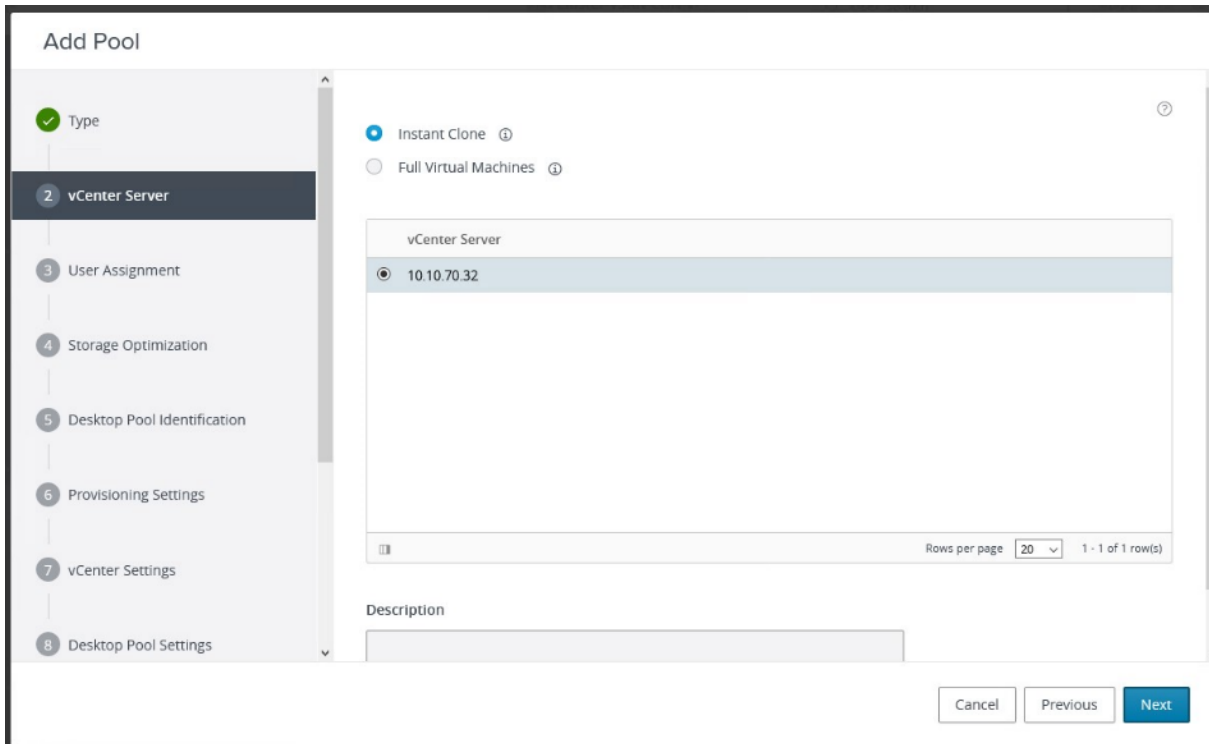
Access group

<input checked="" type="checkbox"/>	ID	Display name	Type	Source	User Assign...	vCenter Ser...	Entitled	Application...	Enabled	App Shortcuts	Sessions
No records available											

Step 2. Select Type of Desktop pool to be created. Click **Next**.

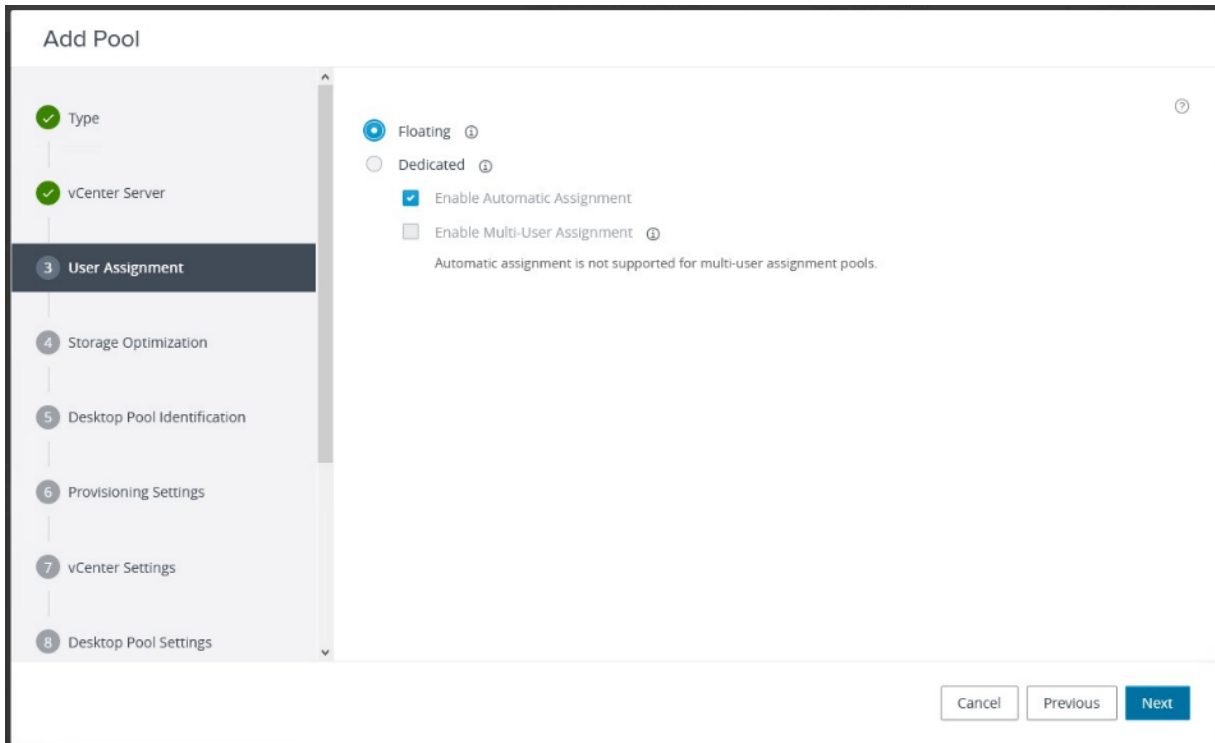


Step 3. Select the provisioning type as **Instant Clones** for the desktops in the pool. Click **Next**.

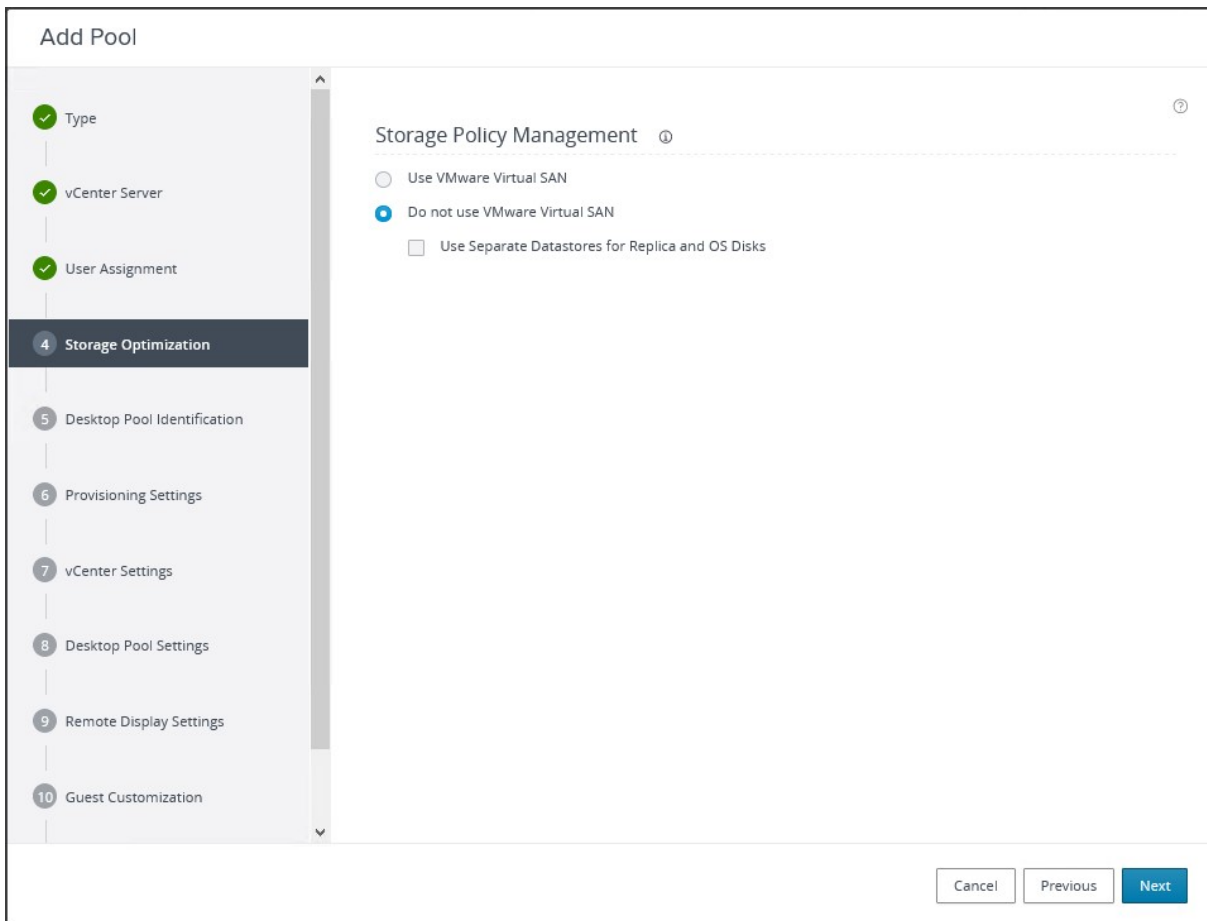


Step 4. Select the User assignment to be used by the desktop pool. Click **Next**.

Note: We used the Floating assignment for the Instant Clone pool.



Step 5. On Storage Optimization screen, select **Do not Use VMware Virtual SAN** as Storage Policy Management, and click **Next**.



Step 6. Provide Desktop Pool ID and virtual display name. Click **Next**.

The screenshot shows the 'Add Pool - Flex-IC' configuration wizard. The left sidebar lists the steps: 1. Type, 2. vCenter Server, 3. User Assignment, 4. Storage Optimization, 5. Desktop Pool Identification (highlighted), 6. Provisioning Settings, 7. vCenter Settings, 8. Desktop Pool Settings, 9. Remote Display Settings, and 10. Guest Customization. The main configuration area includes a legend 'Asterisk (*) denotes required field' and a help icon. The fields are: ID (required, value: Flex-IC), Display Name (value: Flex-IC), Access Group (value: /), and Description (empty). At the bottom right are 'Cancel', 'Previous', and 'Next' buttons.

Step 7. Provide the naming pattern and the number of desktops to be provisioned. Click **Next**.

Note: In this Cisco Validate Design, we used:

Single Server test:320 desktop pool

Cluster/Full scale test: 2 pools of 1120 desktops

Add Pool - Flex-IC

- Type
- vCenter Server
- User Assignment
- Storage Optimization
- Desktop Pool Identification
- 6 Provisioning Settings**
- vCenter Settings
- Desktop Pool Settings
- Remote Display Settings
- Guest Customization

Asterisk (*) denotes required field

Basic

Enable Provisioning ⓘ

Stop Provisioning on Error

Virtual Machine Naming

Specify Names Manually

0 names entered

Use a Naming Pattern ⓘ

+ Naming Pattern

W11-F-IC-

Provision Machines

Machines on Demand

Min Number of Machines

All Machines Up-Front

Desktop Pool Sizing

* Maximum Machines

200

* Spare (Powered On) Machines

1

Step 8. Provide the parent VM, snapshot and host/cluster info, and data store information for the virtual machines to create. Click **Next**.

Add Pool - Flex-IC

- ✓ Type
- ✓ vCenter Server
- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- 7 vCenter Settings**
- 8 Desktop Pool Settings
- 9 Remote Display Settings
- 10 Guest Customization

Default Image

Asterisk (*) denotes required field

* Golden Image in vCenter

* Snapshot

Virtual Machine Location

* VM Folder Location

Resource Settings

* Cluster

* Resource Pool

* Datastores
 4 selected

Network
 Golden Image's network configuration selected

Note: Four datastores were selected to load balance the virtual desktops between them.

Select Datastores

Select the Datastore Type

Select the datastores to use for this desktop pool. Only datastores that can be used by the selected host or cluster can be selected.

<input type="checkbox"/>	<input checked="" type="checkbox"/>	Datastore	Datastore Cluster	Capacity (GB)	Free Space (GB)	FS Type	Drive Type
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS04		20,480	20,426.24	NFS	Non-VMFS
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS03		20,480	20,426.18	NFS	Non-VMFS
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS02		20,480	20,421.58	NFS	Non-VMFS
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS01		20,480	20,363.72	NFS	Non-VMFS
<input type="checkbox"/>	<input checked="" type="checkbox"/>	M7_FC		87,039.75	86,627.91	VMFS6	SSD
<input type="checkbox"/>	<input checked="" type="checkbox"/>	datastore-112		766	742.26	VMFS6	SSD

4 Select all Pages Rows per page 1 - 6 of 6 rows

Free Space Selected 81,637.72 (A minimum of 32,000 GB is recommended for new virtual machines)

Step 9. Configure the State and Session Type for Desktop Pool Settings. Click **Next**.

Add Pool - Flex-IC

- ✔ Type
- ✔ vCenter Server
- ✔ User Assignment
- ✔ Storage Optimization
- ✔ Desktop Pool Identification
- ✔ Provisioning Settings
- ✔ vCenter Settings
- 8 Desktop Pool Settings
- 9 Remote Display Settings
- 10 Guest Customization

^ ?

State

Enabled ▼

Connection Server Restrictions

None Browse

Category Folder

None Browse

Client Restrictions Enabled

Session Types

Desktop ▼ ⓘ

Log Off After Disconnect

Never ▼

Allow Users to Restart Machines

No ▼

Allow Separate Desktop Sessions from Different Client Devices

No ▼ ⓘ

Cancel
Previous
Next

Step 10. Configure the Remote Display Protocol. Click **Next**.

Note: We used the VMware Blast for the Instant Clone pool.

Add Pool - Flex-IC

- ✓ Type
- ✓ vCenter Server
- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- 9 Remote Display Settings**
- 10 Guest Customization

Remote Display Protocol

Default Display Protocol
 VMware Blast

Allow Users to Choose Protocol
 No

3D Renderer
 Manage using vSphere Client

Allow Session Collaboration Enabled ⓘ
 Requires VMware Blast Protocol.

Cancel Previous **Next**

Step 11. Select the AD Container for desktops to place in a Domain Controller computer location.

Add Pool - Flex-IC

- ✓ Type
- ✓ vCenter Server
- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- ✓ Remote Display Settings
- 10 Guest Customization**

Asterisk (*) denotes required field

Domain

* AD Container

Allow Reuse of Existing Computer Accounts ⓘ

Image Publish Computer Account
 ⓘ

Use ClonePrep

Power-Off Script Name
 ⓘ

Power-Off Script Parameters

 Example: p1 p2 p3

Post-Synchronization Script Name
 ⓘ

Post-Synchronization Script Parameters

 Example: p1 p2 p3

Use a customization specification (SysPrep)

Step 12. Review the deployment specifications and click **Submit** to complete the deployment.

Add Pool - Flex-IC

- vCenter Server
- User Assignment
- Storage Optimization
- Desktop Pool Identification
- Provisioning Settings
- vCenter Settings
- Desktop Pool Settings
- Remote Display Settings
- Guest Customization

11 Ready to Complete

<input type="checkbox"/> Entitle Users After Adding Pool	
Type	Automated Desktop Pool
User Assignment	Floating Assignment
vCenter Server	10.10.70.32
Unique ID	Flex-IC
Description	-
Display Name	Flex-IC
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Client Restrictions	Disabled
Log Off After Disconnect	Never
Connection Server Restrictions	None
Category Folder	None

Cancel Previous **Submit**

Step 13. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

Add Entitlements

Add new users and groups who can use the selected pool(s).

Add Remove

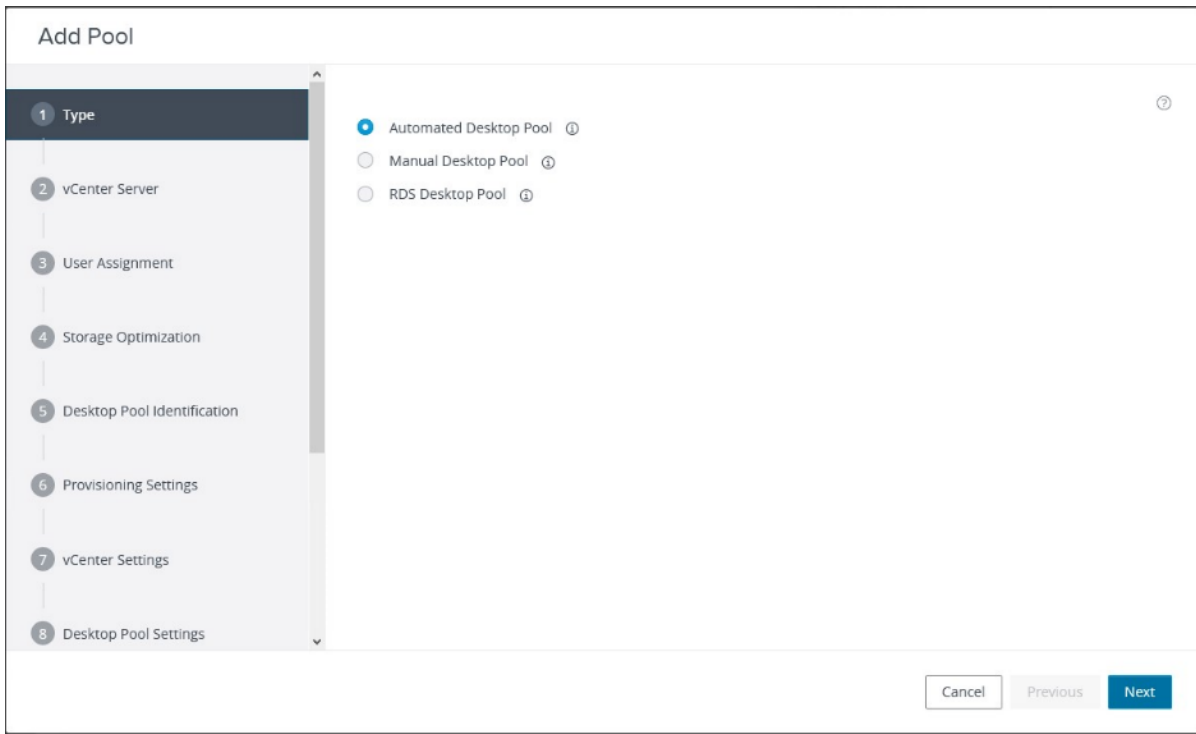
<input type="checkbox"/> Name	Domain	Email
<input type="checkbox"/> Domain Users	FSL151K.LOCAL	

Select all Pages Rows per page: 20 1 - 1 of 1 row(s)

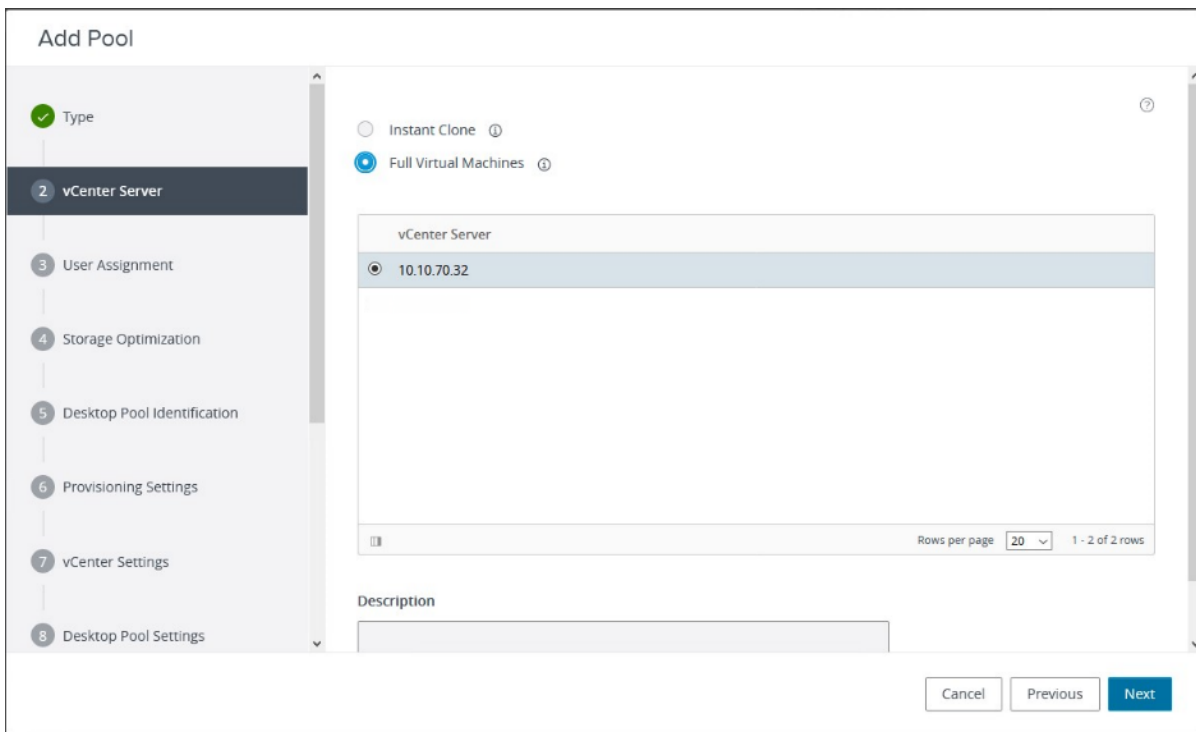
Cancel **OK**

Procedure 3. Create VDI Full Clone Desktop Pool

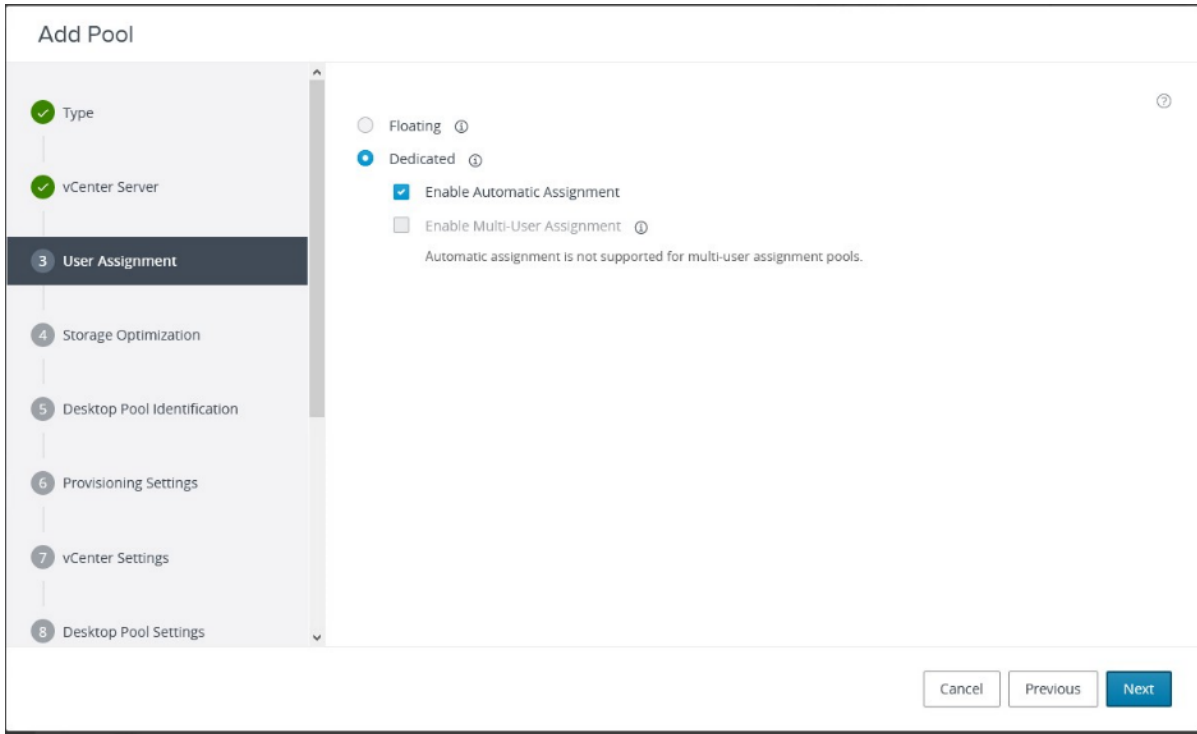
Step 1. Select Type of Desktop pool to be created. Click **Next**.



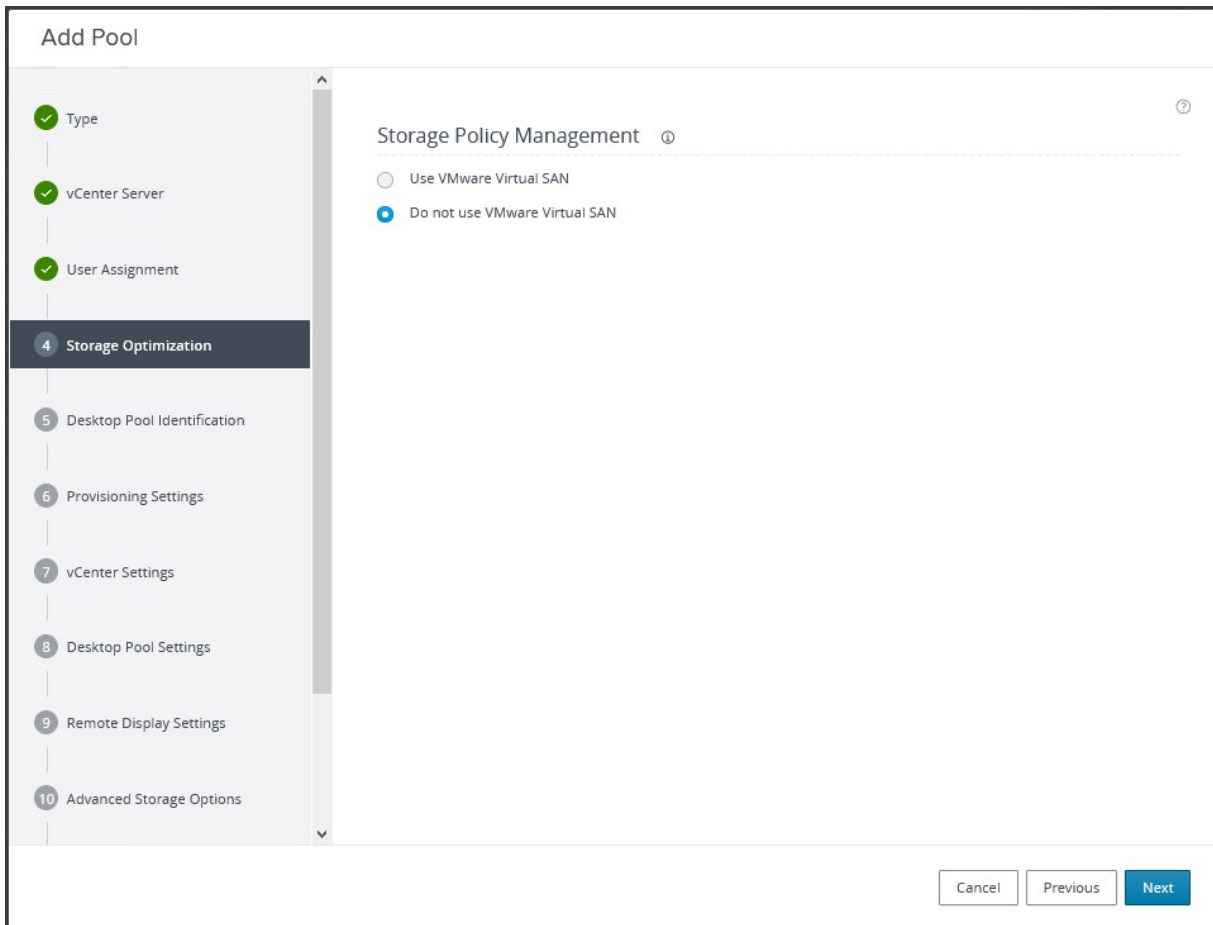
Step 2. Select the provisioning type as Full Virtual Machines for the desktops in the pool. Click **Next**.



Step 3. Select the User assignment to be used by the desktop pool, (we used Dedicated assignment for Full Cone pool). Click **Next**.



Step 4. On Storage Optimization screen, select **Do not Use VMware Virtual SAN** for the Storage Policy Management and click **Next**.



Step 5. Provide the Desktop Pool ID and Display Name. Click **Next**.

The screenshot shows the 'Add Pool - Flex-FC' configuration wizard. On the left is a vertical navigation pane with steps 1 through 10. Step 5, 'Desktop Pool Identification', is highlighted. The main area contains the following fields:

- Type:** Flex-FC
- Display Name:** Flex-FC
- Access Group:** /
- Description:** (Empty text area)

At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

Step 6. Provide the naming pattern and the number of desktops to be provisioned. Click **Next**.

Note: We used the following in this validated design for the VDI full clone pool:

Single Server test:320 desktop pool

Cluster/Full scale test: 2 pools of 1120 desktops

Add Pool - Flex-FC

- Type
- vCenter Server
- User Assignment
- Storage Optimization
- Desktop Pool Identification
- 6 Provisioning Settings**
- vCenter Settings
- Desktop Pool Settings
- Remote Display Settings
- Advanced Storage Options

Asterisk (*) denotes required field

Basic

Enable Provisioning ⓘ

Stop Provisioning on Error

Virtual Machine Naming

Specify Names Manually

0 names entered

Enter Names

Start machines in maintenance mode

Unassigned Machines Kept Powered On

1

Use a Naming Pattern ⓘ

* Naming Pattern

W11-F-FC-

Provision Machines

Machines on Demand

Min Number of Machines

All Machines Up-Front

Desktop Pool Sizing

* Maximum Machines

Cancel
Previous
Next

Step 7. Provide the parent VM, snapshot and host/cluster info, data store information for the virtual machines to create.

Add Pool - Flex-FC

- ✓ Type
- ✓ vCenter Server
- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- 7 vCenter Settings**
- 8 Desktop Pool Settings
- 9 Remote Display Settings
- 10 Advanced Storage Options

Virtual Machine Template

Asterisk (*) denotes required field

* Template

Virtual Machine Location

* VM Folder Location

Resource Settings

* Host or Cluster

* Resource Pool

* Datastores
 4 selected

Note: Four datastores were selected to load balance the virtual desktops between them.

Select Datastores

Select the Datastore Type:

Select the datastores to use for this desktop pool. Only datastores that can be used by the selected host or cluster can be selected.

<input type="checkbox"/>	<input checked="" type="checkbox"/>	Datastore	Datastore Cluster	Capacity (GB)	Free Space (GB)	FS Type	Drive Type
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS04		20,480	20,426.24	NFS	Non-VMFS
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS03		20,480	20,426.18	NFS	Non-VMFS
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS02		20,480	20,421.58	NFS	Non-VMFS
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	NFS01		20,480	20,363.72	NFS	Non-VMFS
<input type="checkbox"/>	<input checked="" type="checkbox"/>	M7_FC		87,039.75	86,627.91	VMFS6	SSD
<input type="checkbox"/>	<input checked="" type="checkbox"/>	datastore-112		766	742.26	VMFS6	SSD

4 Select all Pages
 Rows per page: 1 - 6 of 6 rows

Free Space Selected 81,637.72 (A minimum of 32,000 GB is recommended for new virtual machines)

Step 8. Configure Desktop Pool settings.

Add Pool - Flex-FC

- ✔ Type
- ✔ vCenter Server
- ✔ User Assignment
- ✔ Storage Optimization
- ✔ Desktop Pool Identification
- ✔ Provisioning Settings
- ✔ vCenter Settings
- 8 Desktop Pool Settings
- 9 Remote Display Settings
- 10 Advanced Storage Options

?

State

Enabled ▼

Connection Server Restrictions

None Browse

Category Folder

None Browse

Client Restrictions Enabled

Session Types

Desktop ▼ ⓘ

Remote Machine Power Policy

Take no power action ▼ ⓘ

Log Off After Disconnect

Never ▼

Allow Users to Restart Machines

No ▼

Show Assigned Machine Name ⓘ

Show Machine Alias Name ⓘ

Cancel Previous Next

Step 9. Provide the customizations to remote display protocol to be used by the desktops in the pool.

Note: We used the VMware Blast and the defaults for the rest options for Full Clone pool.

Add Pool - Flex-FC

- Type
- vCenter Server
- User Assignment
- Storage Optimization
- Desktop Pool Identification
- Provisioning Settings
- vCenter Settings
- Desktop Pool Settings
- 9 Remote Display Settings**
- 10 Advanced Storage Options

Remote Display Protocol

Default Display Protocol: VMware Blast

Allow Users to Choose Protocol: No

3D Renderer: Disabled

VRAM Size: 96 MB
More VRAM can improve 3D performance.

Maximum number of monitors: 2
Might require power cycle of related virtual machines.

Maximum Resolution of Any One Monitor: 1920x1200
Might require power cycle of related virtual machines.

Allow Session Collaboration: Enabled
Requires VMware Blast Protocol.

Cancel Previous **Next**

Step 10. Click **Next** on Advanced Storage Options tab.

Note: For Advanced Storage Options, we used defaults in this deployment.

Add Pool - Flex-FC

- ✓ Type
- ✓ vCenter Server
- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- ✓ Remote Display Settings
- 10 Advanced Storage Options**

Advanced Storage Options ⓘ

The following features are recommended based on your resource selection. Options that are not supported by the selected hardware are disabled.

Asterisk (*) denotes required field

Use View Storage Accelerator

 * Regenerate Storage Accelerator After

Days

Blackout Times

Storage accelerator regeneration and VM disk space reclamation do not occur during blackout times. The same blackout policy applies to both operations.

Day	Time
No records available.	

Rows per page 0 rows

Transparent Page Sharing Scope

Step 11. Select the AD Container for desktops to place in a Domain Controller computer location and the VM Customization Specification to be used during deployment. Click **Next**.

Add Pool - Flex-FC

- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- ✓ Remote Display Settings
- ✓ Advanced Storage Options
- 11 Guest Customization**
- 12 Ready to Complete

Asterisk (*) denotes required field

Select AD Container from Domain Accounts

Domain: FSL151K.LOCAL(administrator)

* AD Container: OU=VDI,OU=Target,OU=Computers,OU=LoginVSI [Find] [Browse]

None - Customization will be done manually

Do not Power on Virtual Machines After Creation

Use this customization specification

Allow Reuse of Existing Computer Accounts ⓘ

Name	Guest OS	Description
<input checked="" type="radio"/> W11 spec	Windows	
<input type="radio"/> Win10 Spec	Windows	
<input type="radio"/> Win2019 Spec	Windows	
<input type="radio"/> Windows 11	Windows	

Cancel Previous **Next**

Note: The following VM Customization Specifications were used:

Name	W11 spec
Description	
OS type	Windows
OS options	Generate new security ID
Registration info	Owner name: labuser Organization: cisco
Computer name	Use Virtual Machine name
Windows license	
Product key	No product key specified
Server license information	Not included
Server license mode	
Log in	
Administrator log in	Do not log in automatically as Administrator
Time zone	(UTC-08:00) Pacific Time (US & Canada)
Network type	Standard
Workgroup/Domain	Windows Server domain: FSL151K.local

Step 12. Review all the deployment specifications and click Submit to complete the deployment.

Add Pool - Flex-FC

- ✓ User Assignment
- ✓ Storage Optimization
- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- ✓ Remote Display Settings
- ✓ Advanced Storage Options
- ✓ Guest Customization

12 Ready to Complete

<input type="checkbox"/> Entitle Users After Adding Pool	
Type	Automated Desktop Pool
User Assignment	Dedicated Assignment
Assign on First Login	Yes
Enable Multi-User Assignment	No
vCenter Server	10.10.70.32
Unique ID	Flex-FC
Description	-
Display Name	Flex-FC
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Show Assigned Machine Name	Disabled
Show Machine Alias Name	Disabled

Cancel Previous **Submit**

Step 13. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

Add Entitlements

Add new users and groups who can use the selected pool(s).

Add Remove

<input type="checkbox"/>	Name	Domain	Email
<input type="checkbox"/>	Domain Users	FSL151K.LOCAL	

Select all Pages Rows per page: 20 1 - 1 of 1 row(s)

Cancel **OK**

Note: Alternatively, you can use the esxi to create full clones and manual desktop pool to serialize deployment.

Deployment script example

```
$TargetCluster = Get-Cluster -Name "FC-1"
$SourceVMTemplate = Get-Template -Name "W11-Flex-FC-Tmpl"
$SourceCustomSpec = Get-OSCustomizationSpec -Name "W11 spec"
$basename = "W11-F2-FC"
```

```

$vmList = @()

foreach ($i in 1..1120) {

foreach-object {

    $vm = New-Object System.Object

    $name = "${basename}-${i}"

    $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
}

$vmList = $vmList + $vm
}

$counter = 0
$createJobList = @()
#create 4 VMs asynchronously and place them on differen datastore
$vmList | foreach-object {
    $counter++
    $vmname = $_.Name
    $datastore = Get-Datastore -Name NFS0$counter
    write-host "Creating VM $vmname from template $sourceVMtemplate on datastore
$datastore" -ForegroundColor Green
    new-vm -name $vmname -resourcepool $TargetCluster -template $SourceVMtemplate -
OSCustomizationSpec $SourceCustomSpec -datastore $datastore -DiskStorageFormat Thin -
RunAsync -OutVariable createJob

    $createJob | add-member -MemberType NoteProperty -name "VM" -value "$vmname"

    $createJobList = $createJobList + $createJob

    if ($counter -eq 4) {
        # Pause execution
        Write-Host "Pausing for 5 seconds..." -ForegroundColor Yellow
        Start-Sleep -Seconds 5
        $createJobList | foreach-object {
            $vmname = $_.VM

            write-host " Waiting for VM $vmname clone process to finish " -
ForegroundColor Green

```

```
do { start-sleep -seconds 5 }
until ( (get-task -id $_.ID).state -eq "Success" )

write-host " Starting VM $vmname " -ForegroundColor Green
start-vm -vm $vmname -runAsync

}

# Reset the counter and job list
$counter = 0
    $createJobList = @()
}
}
```

Procedure 4. Create RDSH Farm and Pool

Step 1. Select **FARM** when creating the RDS Pool.

Note: You can entitle the user access at the RDS Pool level for RDS users/groups who can access the RDS VMs.

Step 2. Select Type of Farm (we used Automated Farm the RDS desktops in this design). Click **Next**.

Add Farm

1 Type

2 vCenter Server

3 Storage Optimization

4 Identification and Settings

5 Load Balancing Settings

6 Provisioning Settings

7 vCenter Settings

8 Guest Customization

9 Ready to Complete

Automated Farm ⓘ

Manual Farm ⓘ

Cancel Previous Next

Step 3. Select the provisioning type and vCenter Server for the desktops in the pool. Click **Next**.

Add Farm

- 1 Type
- 2 vCenter Server**
- 3 Storage Optimization
- 4 Identification and Settings
- 5 Load Balancing Settings
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Guest Customization
- 9 Ready to Complete

Instant Clone ⓘ

vCenter Server
<input checked="" type="radio"/> 10.10.70.32

Rows per page 20 1 - 1 of 1 row(s)

Description

Cancel Previous **Next**

Step 4. Select **Use VMware Virtual SAN** on the Storage Optimization screen click **Next**.

Add Farm

- ✓ Type
- ✓ vCenter Server
- 3 Storage Optimization**
- 4 Identification and Settings
- 5 Load Balancing Settings
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Guest Customization
- 9 Ready to Complete

Storage Policy Management ⓘ

- Use VMware Virtual SAN
- Do not use VMware Virtual SAN
- Use Separate Datastores for Replica and OS Disks

Cancel Previous **Next**

Step 5. Provide the ID and Description for RDS FARM. Select the Display Protocol which is required for users to connect to the RDS Sessions. Click **Next**.

Note: We used Microsoft RDP in this CVD environment.

Add Farm - W2019-Farm

- ✔ Type
- ✔ vCenter Server
- ✔ Storage Optimization
- 4 Identification and Settings**
- 5 Load Balancing Settings
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Guest Customization
- 9 Ready to Complete

Asterisk (*) denotes required field

*** ID**

Description

Access Group

Farm Settings

Default Display Protocol ⓘ

Allow Users to Choose Protocol

3D Renderer ⓘ

vSphere doesn't support 3D option other than NVIDIA Grid vGPU for Windows Server OS

Pre-launch Session Timeout (Applications Only) ⓘ

minutes

Empty Session Timeout (Applications Only) ⓘ

minutes

When Timeout Occurs

Logoff Disconnected Sessions

Bypass Session Timeout ⓘ

Allow Session Collaboration Enabled ⓘ

Step 6. Select Load Balancing Settings. Click **Next**.

Add Farm - W2019-Farm

- ✔ Type
- ✔ vCenter Server
- ✔ Storage Optimization
- ✔ Identification and Settings
- 5 Load Balancing Settings
- 6 Provisioning Settings
- 7 vCenter Settings
- 8 Guest Customization
- 9 Ready to Complete

Use Custom Script Enabled ⓘ

Include Session Count Enabled ⓘ

Asterisk (*) denotes required field

* CPU Usage Threshold ⓘ

* Memory Usage Threshold ⓘ

* Disk Queue Length Threshold ⓘ

* Disk Read Latency Threshold ⓘ

* Disk Write Latency Threshold ⓘ

* Connecting Session Threshold ⓘ

* Load Index Threshold ⓘ

Step 7. Provide naming pattern and a number of virtual machines to create. Click **Next**.

Note: In this validated design for RDS farms we used:
 Single Server - 32
 Cluster - 224

Add Farm - W2019-Farm

Asterisk (*) denotes required field

Basic

Enable Provisioning ⓘ

Stop Provisioning on Error

Virtual Machine Naming

* Naming Pattern

W2019-F-

Farm Sizing

* Maximum Machines

32

* Minimum Number of Ready (Provisioned) Machines during Instant Clone Maintenance Operations

0

6 Provisioning Settings

7 vCenter Settings

8 Guest Customization

9 Ready to Complete

Cancel Previous **Next**

Step 8. Select the previously created golden image to be used as RDS host. Select the datastore where the RDS hosts will be deployed. Click **Next**.



Add Farm - W2019-Farm

- Type
- vCenter Server
- Storage Optimization
- Identification and Settings
- Load Balancing Settings
- Provisioning Settings
- 7 vCenter Settings**
- 8 Guest Customization
- 9 Ready to Complete

Default Image

Asterisk (*) denotes required field

- * Golden Image in vCenter
- * Snapshot

Virtual Machine Location

- * VM Folder Location

Resource Settings

- * Cluster
- * Resource Pool
- * Datastores
4 selected

Network

Golden Image's network configuration selected

VM Compute Profile Settings

Review the default VM Compute Profile settings and modify if needed.

- * CPU
 ⓘ
- * RAM
 GB
- Cores per Socket
 ⓘ

Four datastores were selected to load balance the virtual desktops between them:

Select Instant Clone Datastores ✕

Select the instant clone datastores to use for this Automated Farm. Only datastores that can be used by the selected host or cluster can be selected.

Show all datastores (including local datastores) ⓘ ↻

	Datastore	Capacity (GB)	Free Space (GB)	FS Type	Drive Type	Storage Overcommit
<input checked="" type="checkbox"/>	NFS01	20,480	20,018.78	NFS	Non-VMFS	Unbounded
<input checked="" type="checkbox"/>	NFS02	20,480	19,902.03	NFS	Non-VMFS	Unbounded
<input checked="" type="checkbox"/>	NFS03	20,480	19,895.14	NFS	Non-VMFS	Unbounded
<input checked="" type="checkbox"/>	NFS04	20,480	19,901.16	NFS	Non-VMFS	Unbounded

4 Deselect all Pages Rows per page: 20 1 - 4 of 4 rows

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% Utilization (GB)	Max Recommended (GB)
Instant clones	79,717.11	948	1,620	3,060

Rows per page: 20 1 - 1 of 1 row(s)

Step 9. Select the AD Container for desktops to place in a Domain Controller computer location.

Add Farm - W2019-Farm

- Type
- vCenter Server
- Storage Optimization
- Identification and Settings
- Load Balancing Settings
- Provisioning Settings
- vCenter Settings
- 8 Guest Customization**
- Ready to Complete

Asterisk (*) denotes required field

Domain

*** AD Container**

Allow Reuse of Existing Computer Accounts ⓘ

Image Publish Computer Account

 ⓘ

Use ClonePrep

Power-Off Script Name

 ⓘ

Power-Off Script Parameters

Example: p1 p2 p3

Post-Synchronization Script Name

 ⓘ

Post-Synchronization Script Parameters

Example: p1 p2 p3

Use a customization specification (SysPrep)

	Name	Guest OS	Description
<input checked="" type="radio"/>	W11 spec	Windows	
<input type="radio"/>	Win10 Spec	Windows	
<input type="radio"/>	Win2019 Spec	Windows	
<input type="radio"/>	Windows 11	Windows	

Step 10. Review the Farm information and click **Submit** to complete the RDS Farm creation.

Add Farm - W2019-Farm

- Type
- vCenter Server
- Storage Optimization
- Identification and Settings
- Load Balancing Settings
- Provisioning Settings
- vCenter Settings
- Guest Customization
- 9 Ready to Complete

ID	W2019-Farm
Description	-
Access Group	/
<hr/>	
Farm Settings	
Default Display Protocol	Microsoft RDP
Allow Users to Choose Protocol	Yes
3D Renderer	Manage using vSphere Client
Pre-launch Session Timeout (Applications Only)	10 minutes
Empty Session Timeout (Applications Only)	1 minute
When Timeout Occurs	Disconnect
Logoff Disconnected Sessions	Never
Bypass Session Timeout	Disabled
Allow Session Collaboration	Disabled
CPU	4
RAM	24 GB
Cores per Socket	1
Load Balancing Settings	
Use Custom Script	Disabled
Include Session Count	Enabled
CPU Usage Threshold	0

Procedure 5. Create RDS Pool

When the RDS FARM is created, you need to create an RDS Pool to absorb the RDS VMs FARM into the Pool for further managing the RDS pool.

Step 1. Select type as **RDS Desktop Pool**.

Add Pool

1 Type

2 Desktop Pool Identification

3 Desktop Pool Settings

4 Select RDS Farms

5 Ready to Complete

Automated Desktop Pool ⓘ

Manual Desktop Pool ⓘ

RDS Desktop Pool ⓘ

Cancel Previous Next

Step 2. Provide an ID and Display Name for the Pool. Click **Next**.

Add Pool - Flex-RDS-1

✓ Type

2 Desktop Pool Identification

3 Desktop Pool Settings

4 Select RDS Farms

5 Ready to Complete

Asterisk (*) denotes required field

* ID ⓘ

Flex-RDS-1

Display Name ⓘ

Flex-RDS-1

Description

Cancel Previous **Next**

Step 3. Leave the default settings for the Desktop Pool Settings. Click **Next**.

Add Pool - Flex-RDS-1

- ✓ Type
- ✓ Desktop Pool Identification
- 3 Desktop Pool Settings**
- 4 Select RDS Farms
- 5 Ready to Complete

State ⓘ
Enabled

Connection Server Restrictions
None

Category Folder
None

Client Restrictions Enabled

Allow Separate Desktop Sessions from Different Client Devices ⓘ
No

Step 4. Select the RDS Farm. Select the farm which was already created for this desktop pool. Click **Next**.

Add Pool - Flex-RDS-1

1 Type

2 Desktop Pool Identification

3 Desktop Pool Settings

4 **Select RDS Farms**

5 Ready to Complete

Create a new RDS farm

Select an RDS farm for this desktop pool

Filter

Farm ID	Description	RDS Hosts	Max Number of Connections	Status
<input checked="" type="checkbox"/> W2019-Farm1		32	Unlimited	Farm disabled

Rows per page 20 1 - 2 of 2 rows

Cancel Previous **Next**

Step 5. Review the RDS Pool deployment specifications and click **Next** to complete the RDS pool deployment.

Add Pool - Flex-RDS-1

- ✔ Type
- ✔ Desktop Pool Identification
- ✔ Desktop Pool Settings
- ✔ Select RDS Farms
- 5 Ready to Complete

Entitle Users After Adding Pool ⓘ

Type	RDS Desktop Pool
Unique ID	Flex-RDS-1
Description	.
Display Name	Flex-RDS-1
Desktop Pool State	Enabled
Client Restrictions	Disabled
Connection Server Restrictions	None
Category Folder	None
Allow Separate Desktop Sessions from Different Client Devices	No
RDS Farm	W2019-Farm1
Number of RDS Hosts in the Farm	32

Step 6. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.



Add Entitlements ✕

Add new users and groups who can use the selected pool(s).

Add Remove

<input type="checkbox"/>	Name	Domain	Email
<input type="checkbox"/>	Domain Users	FSL151K.LOCAL	

Select all Pages Rows per page: 20 1 - 1 of 1 row(s)

Cancel **OK**

Hybrid Cloud Deployment for Disaster Recovery

This chapter contains the following:

- [Access NetApp BlueXP](#)
- [Deploy Connector](#)
- [Deploy Cloud Volumes ONTAP in Microsoft Azure Cloud](#)
- [Set up SnapMirror relationship](#)

Disasters can come in many forms for a business and protecting data with disaster recovery (DR) is a critical goal for businesses continuity. DR allows organizations to failover their business operations to a secondary location and later recover and failback to the primary site efficiently and reliably.

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. As its name implies, SnapMirror creates a replica, or mirror, of your working data in secondary storage from which you can continue to serve data in the event of a catastrophe at the primary site.

Procedure 1. Access NetApp BlueXP

Step 1. To access NetApp BlueXP and other cloud services, you need to sign up here: [NetApp Cloud Central](#).

Procedure 2. Deploy Connector

Step 1. To deploy Connector in Microsoft Azure Cloud, go to: <https://docs.netapp.com/us-en/cloud-manager-setup-admin/task-creating-connectors-azure.html#overview>.

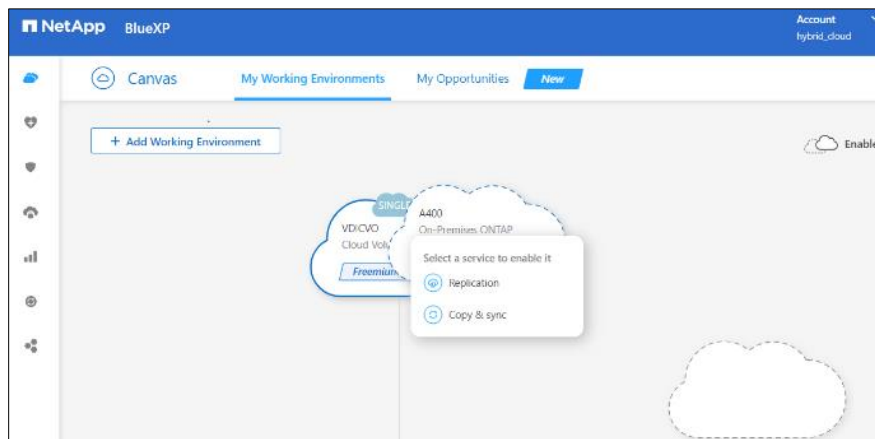
Procedure 3. Deploy Cloud Volumes ONTAP in Microsoft Azure Cloud

Step 1. To deploy CVO in Microsoft Azure Cloud, go to: <https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/task-getting-started-azure.html>.

Procedure 4. Set up SnapMirror relationship

Step 1. In the BlueXP window, select your **workspace** and **connector**.

Step 2. Drag and drop the source **on-prem ONTAP** to the **Azure Cloud Volumes ONTAP** instance which will be your destination.



Step 3. In the Source Peering Setup select the **intercluster LIF**.

Source Peering Setup

Select the source LIFs you would like to use for cluster peering setup.
Replication requires an initial connection between the two working environments which is called a cluster peer relationship.
For more information about LIF selections, see Cloud Manager documentation.

cluster1_inter_lif1
 A400-01 : a0a-31
 10.10.31.15/24 | up

cluster1_inter_lif2
 A400-02 : a0a-31
 10.10.31.16/24 | up

Step 4. In Source Volume Selection select the **volume** to replicate.

Step 5. In Destination Disk Type and Tiering select the default **Destination Disk Type**.

Destination Disk Type and Tiering

Destination Disk Type

Premium SSD

Standard SSD

Standard HDD

Blob Tiering [What are storage tiers?](#)

Enabled Disabled

Note: If you enable Blob tiering, thin provisioning must be enabled on volumes created in this aggregate.

Continue

Step 6. In Destination Volume Name, accept all the defaults.

Destination Volume Name

Destination Volume Name

nfs_test_vol_copy

Destination Aggregate

Automatically select the best aggregate ▼

Step 7. Click **Continue**.

Step 8. Accept the default for the Max Transfer Rate.

Max Transfer Rate

You should limit the transfer rate. An unlimited rate might negatively impact the performance of other applications and it might impact your Internet performance.

Limited to: MB/s

Unlimited (recommended for DR only machines)

Step 9. Click **Continue**.

Step 10. Click **Mirror** to select a Replication Policy.

Replication Policy

Default Policies Additional Policies

Mirror

Typically used for disaster recovery

[More info](#)

Mirror and Backup (1 month retention)

Configures disaster recovery and long-term retention of backups on the same destination volume

[More info](#)

Note: To control the frequency of the replication updates, you can select a replication schedule. If you scroll down the window, you will see there are many from which to choose.

Replication Setup Schedule

↑ Previous Step Select a replication schedule

One-time copy

No schedule

hourly

Every hour
Minutes: 5th minute

5min

Every hour
Minutes: 0th, 5th, 10th, 15th...

10min

Every hour
Minutes: 0th, 10th, 20th, 30th...

8hour

Every day
Hours: 2 AM, 10 AM and 6 PM
Minutes: 15th minute

daily

Every day
Hours: 12 AM
Minutes: 10th minute

6-hourly

Every day
Hours: 12 AM, 6 AM, 12 PM...
Minutes: 15th minute

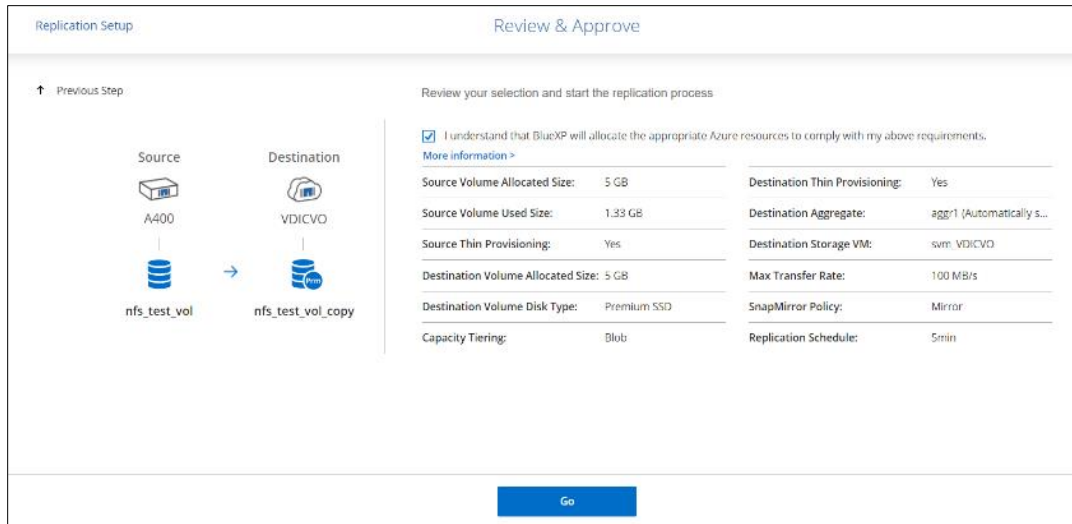
12-hourly

Every day
Hours: 12 AM and 12 PM
Minutes: 15th minute

weekly

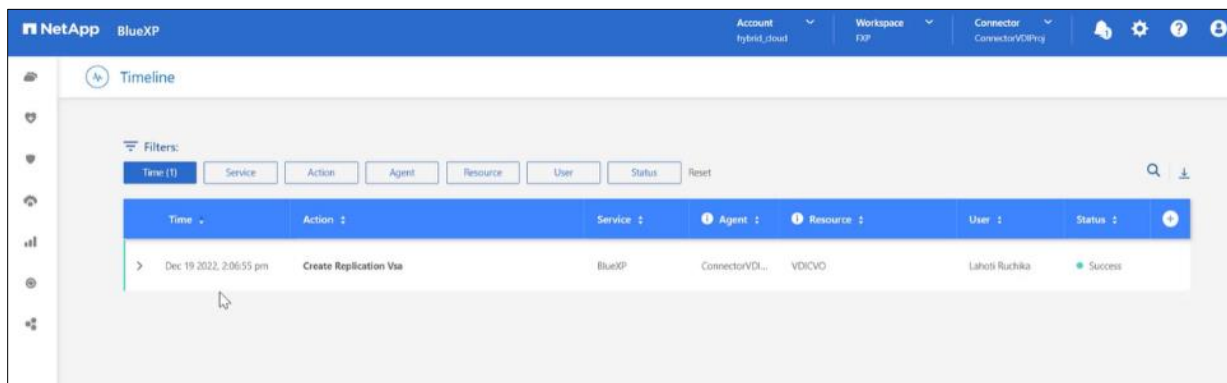
Every week
Days: Sun
Hours: 12 AM
Minutes: 15th minute

Step 11. Review and accept the choices by clicking the checkbox next to **I understand**




Step 12. Click **Go**.

Step 13. Click **Timeline** to monitor when the creation of the protection relationship is complete.



Step 14. Click the **CVO instance** and go to replication tab. Wait for the Mirror State to change to snapmirrored.



Test Setup and Configuration

This chapter contains the following:

- [Cisco UCS Test Configuration for Single Compute Node Scalability](#)
- [Testing Methodology and Success Criteria](#)
- [Login Enterprise Testing](#)
- [Single-Server Recommended Maximum Workload](#)

Note: In this solution, we tested a single Cisco UCS X210c M7 Compute Node Server to validate against the performance of one Compute Node and eight Cisco UCS X210c M7 Compute Node Servers in a single chassis to illustrate linear scalability for each workload use case studied.

Cisco UCS Test Configuration for Single Compute Node Scalability

This test case validates the Recommended Maximum Workload per host server using VMware Horizon 8 2212 with 480 Multi-session OS sessions and 320 Single-session OS sessions.

Figure 38. Test Configuration for Single Server Scalability VMware Horizon 8 2212 Non-persistent (NP) Single-session OS machine VDAs

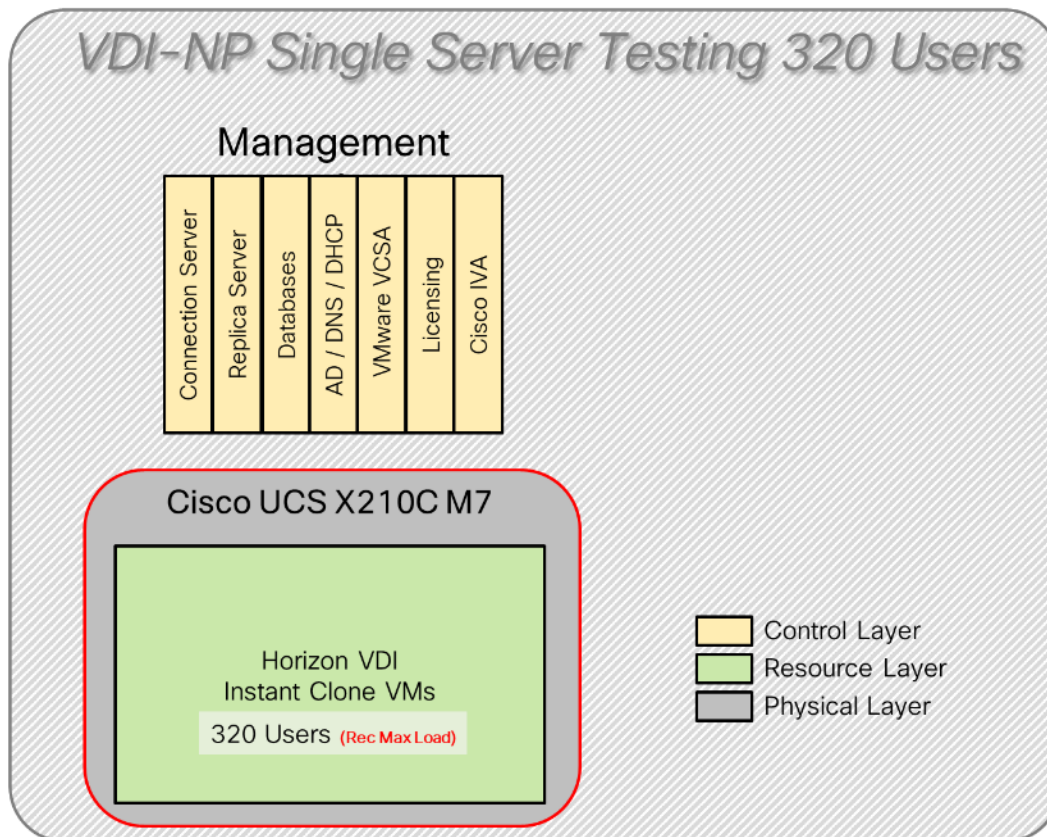


Figure 39. Test configuration for Single Server Scalability VMware Horizon 8 2212 Persistent (P) Single-session OS machine VDAs

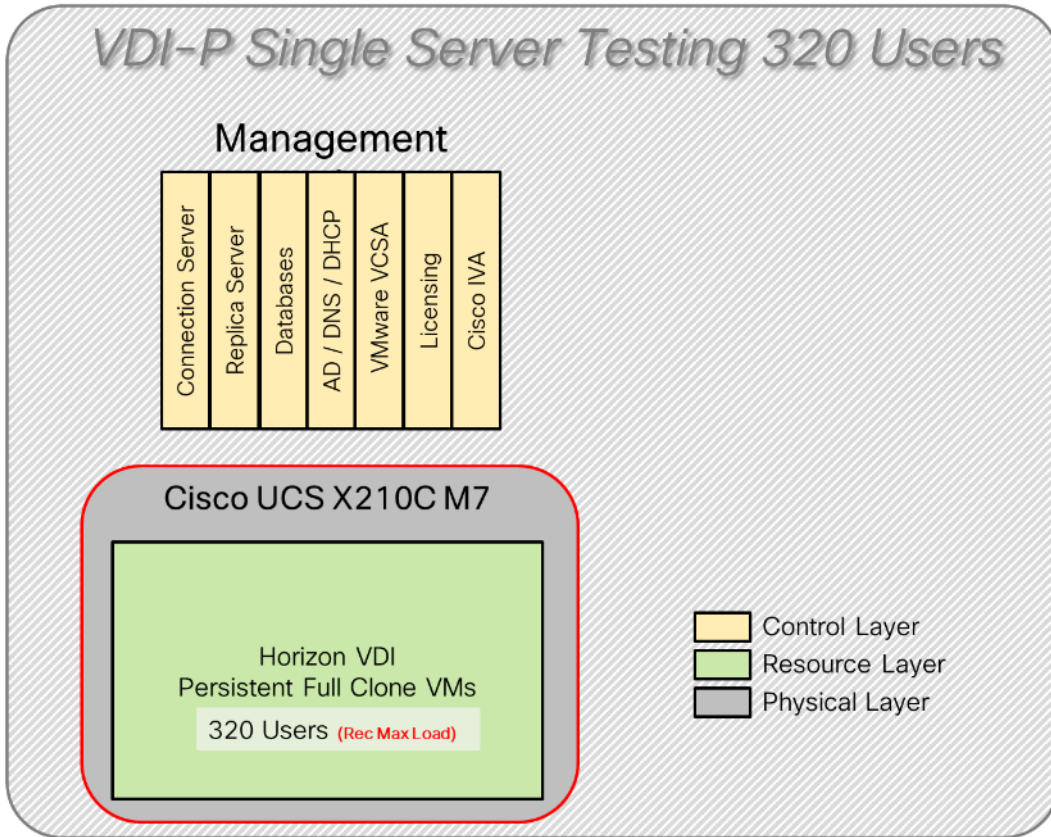
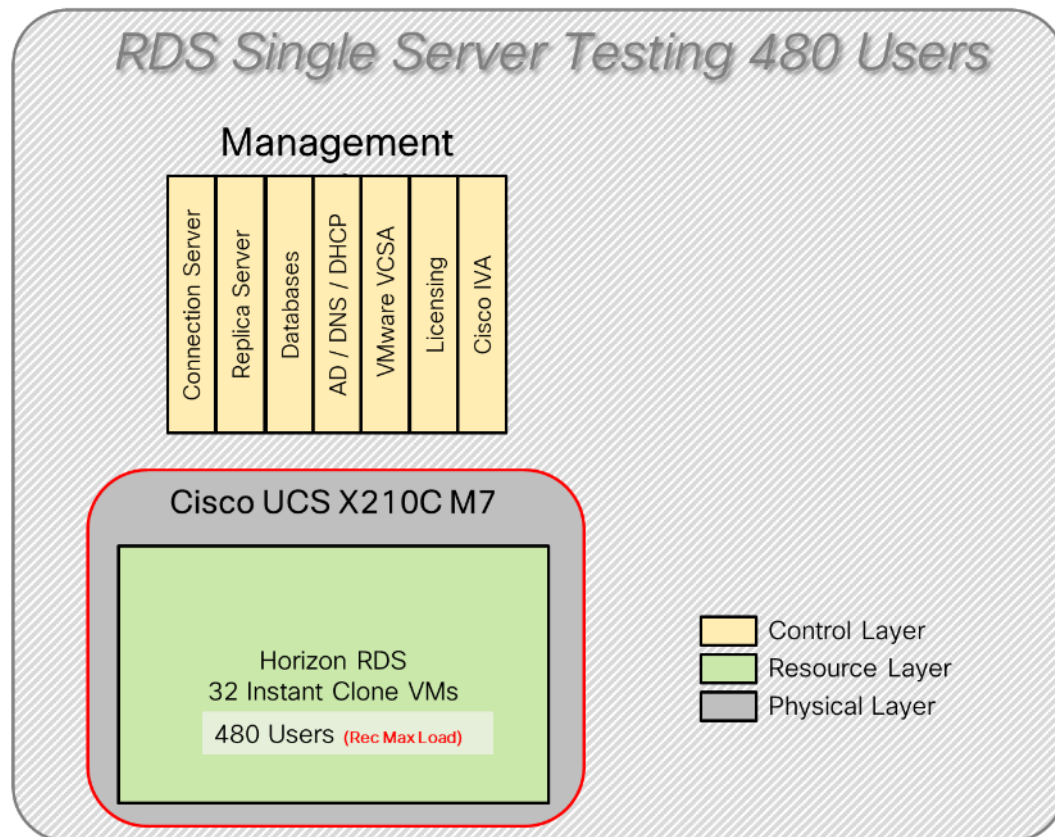


Figure 40. Test configuration for Single Server Scalability VMware Horizon 8 2212 Instant-clone Multi-session OS machine VDAs



Hardware components:

- 1 Cisco UCS X9508 Chassis
- Cisco UCS 6536 5th Generation Fabric Interconnects
- 1 Cisco UCS X210c M7 Compute Node Servers with Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 2TB 4800 MHz RAM for all host blades
- Cisco UCS VIC 15231 CNA (1 per blade)
- 2 Cisco Nexus N9K-93180YC-FX3S Access Switches
- NetApp AFF A400

Software components:

- Cisco UCS firmware 5.2(0.230041)
- NetApp ONTAP 9.13. 1P3VMware ESXi 8.0 U1 for host blades
- VMware Horizon 8 2212
- Microsoft SQL Server 2019
- Microsoft Windows 11 64 bit (21H2), 2vCPU, 4 GB RAM, 96 GB HDD (master)
- Microsoft Windows Server 2019 (1809), 4vCPU, 24GB RAM, 90 GB vDisk (master)
- Microsoft Office LTSC Standard 2021
- FSLogix 2210 hotfix 1

- Login VSI 4.1.40 Knowledge Worker Workload
- NAbbox3.3

Cisco UCS Test Configuration for Full Scale Testing

These test cases validate eight blades in a cluster hosting three distinct workloads using VMware Horizon 8 2212 with:

- 2240 VDI-NP (instant clones) Single-session OS sessions
- 2240 VDI-P (full clones) Single-session OS sessions
- 3360 RDS sessions on Instant clone virtual machines

Note: Server N+1 fault tolerance is factored into this solution for each cluster/workload.

Figure 41. Test Configuration for Full Scale VMware Horizon 8 2212 non-persistent Single-session OS machine VDAs

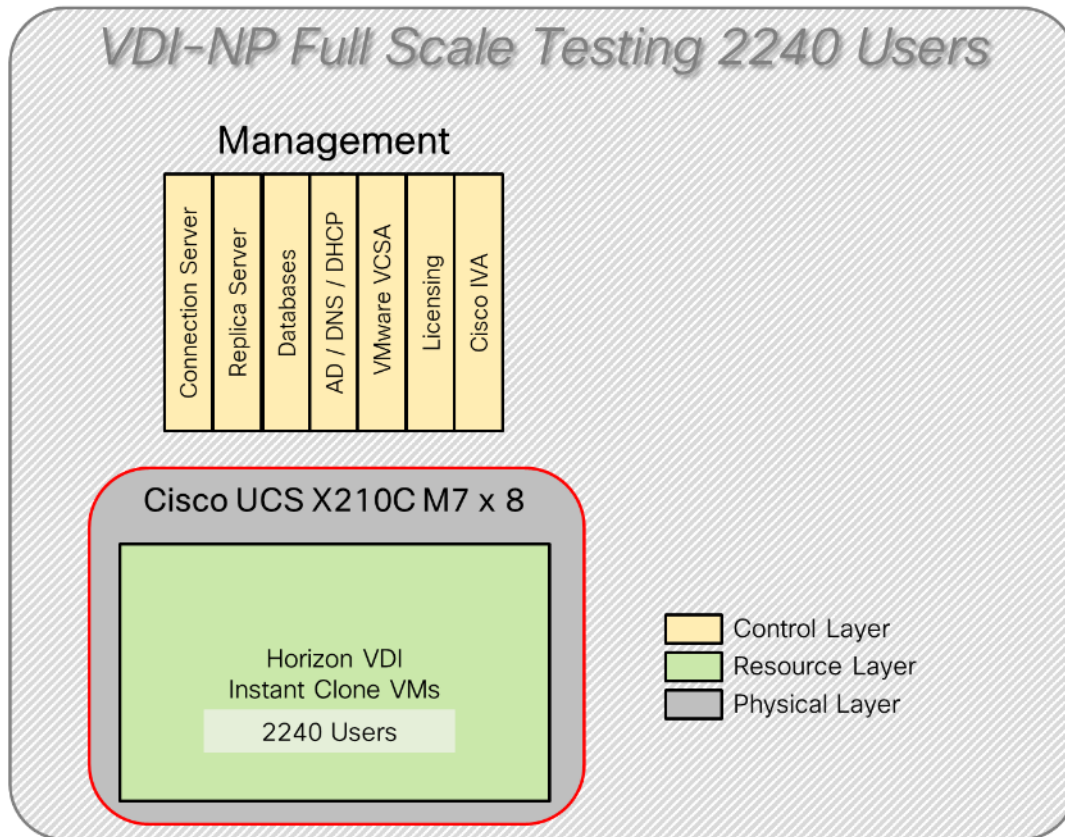


Figure 42. Test Configuration for Full Scale VMware Horizon 8 2212 persistent Single-session OS machine VDAs

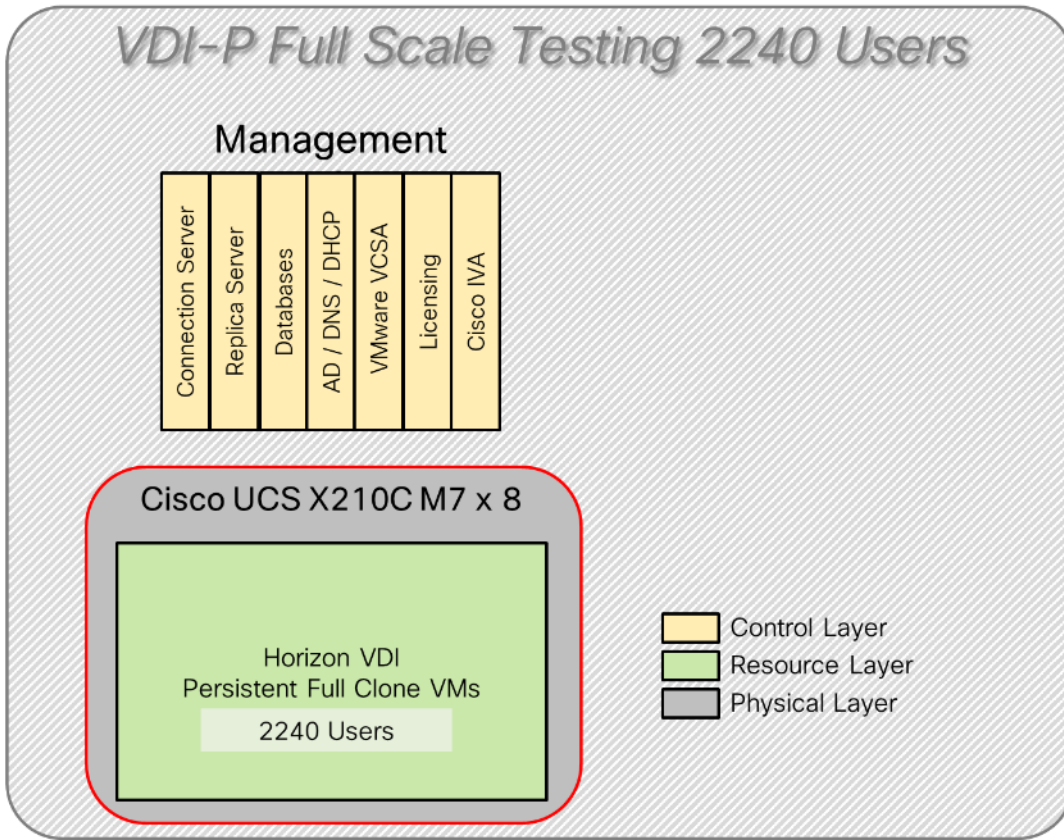
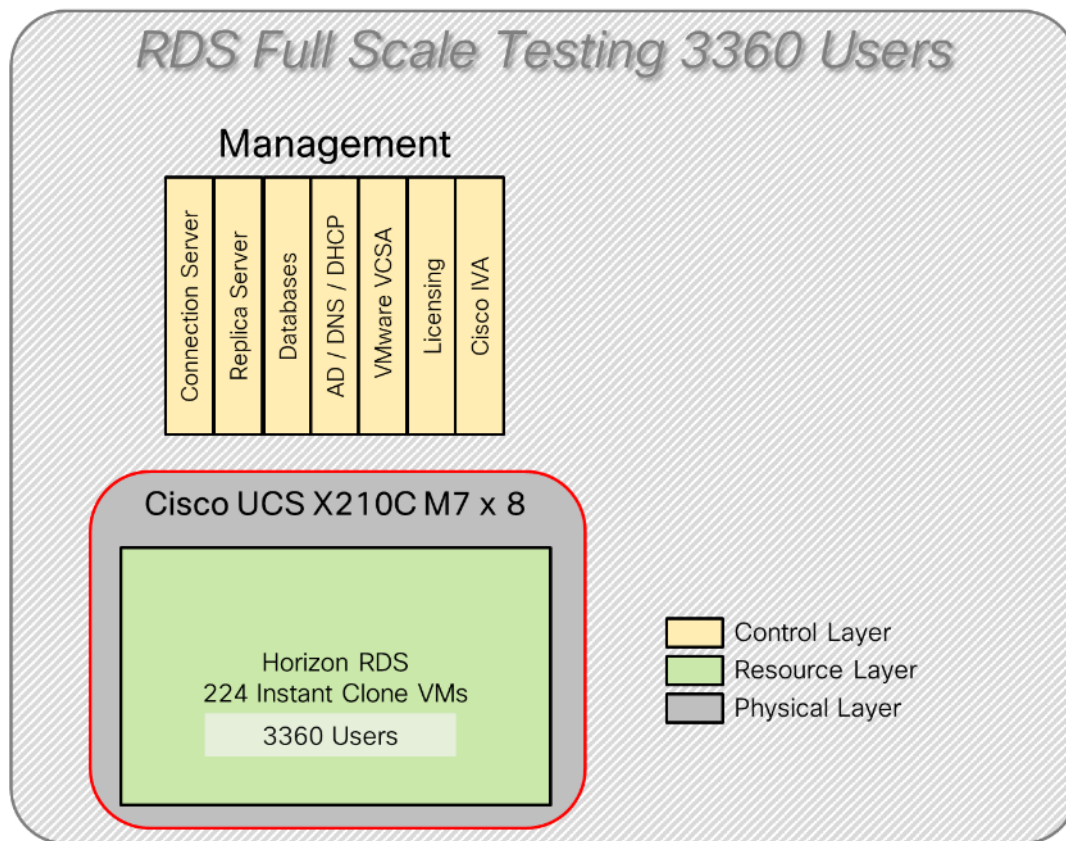


Figure 43. Test Configuration for Full Scale VMware Horizon 8 2212 instant-clones Multi-session OS machine VDAs



Hardware components:

- 1 Cisco UCS X9508 Chassis
- Cisco UCS 6536 5th Generation Fabric Interconnects
- 1 Cisco UCS X210c M7 Compute Node Servers with Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 2TB 4800 MHz RAM for all host blades
- Cisco UCS VIC 15231 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- NetApp A400

Software components:

- Cisco UCS firmware 5.2(0.230041)
- Netapp ONTAP 9.12. 1P3VMware ESXi 8.0 U1 for host blades
- VMware Horizon 8 2212
- Microsoft SQL Server 2019
- Microsoft Windows 11 64 bit (21H2), 2vCPU, 4 GB RAM, 96 GB HDD (master)
- Microsoft Windows Server 2019 (1809), 4vCPU, 24GB RAM, 90 GB vDisk (master)
- Microsoft Office LTSC Standard 2021
- FSLogix 2210 hotfix 1

- Login VSI 4.1.40 Knowledge Worker Workload
- NAbbox 3.3

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the VMware Horizon Virtual Desktop and RDS sessions under test.

Test metrics were gathered from the hypervisor, virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from <http://www.loginvsi.com>.

Testing Procedure

This section contains the following procedure:

- [Test Run Protocol](#)

The following protocol was used for each test cycle in this study to ensure consistent results.

Pre-Test Setup for Single and Multi-Compute Node Testing

All virtual machines were shut down utilizing the VMware Horizon console and vCenter.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

All VMware ESXi VDI host Compute Nodes to be tested were restarted prior to each test cycle.

Procedure 1. Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, it's required to start all sessions, whether single server users or full scale test users to become active within two minutes after the last session is launched.

In addition, Cisco requires that the Login VSI Benchmark method be used for all single server and scale testing. This assures that our tests represent real-world scenarios. For each of the three consecutive runs on single server tests, the same process was followed.

Time 0:00:00 Start Esxtop Logging on the following systems:

- a. Infrastructure and VDI Host Compute Nodes used in the test run
- b. vCenter used in the test run
- c. All Infrastructure VMs used in test run (AD, SQL, brokers, image mgmt., and so on)

Time 0:00:10 Start Storage Partner Performance Logging on Storage System

Time 0:05: Boot Virtual Desktops/RDS Virtual Machines using VMware Horizon Console

Note: The boot rate should be around 10-12 VMs per minute per server.

Time 0:06 First machines boot

Time 0:30 Single Server or Scale target number of desktop VMs booted on 1 or more Compute Nodes

Note: No more than 30 minutes for boot up of all virtual desktops is allowed.

Time 0:35 Single Server or Scale target number of desktop VMs desktops registered on VMware Horizon Console

Virtual machine settling time.

Note: No more than 60 Minutes of rest time is allowed after the last desktop is registered on the VMware Horizon Console . Typically, a 30-40 minute rest period is sufficient.

Time 1:35 Start Login VSI 4.1.40 Office Worker Benchmark Mode Test, setting auto-logoff time at 900 seconds, with Single Server or Scale target number of desktop VMs utilizing sufficient number of Launchers (at 20-25 sessions/Launcher)

Time 2:23 Single Server or Scale target number of desktop VMs desktops launched (48 minute benchmark launch rate)

Time 2:25 All launched sessions must become active

Note: All sessions launched must become active for a valid test run within this window.

Time 2:40 Login VSI Test Ends (based on Auto Logoff 900 Second period designated above.)

Time 2:55 All active sessions logged off

Time 2:57 All logging terminated; Test complete

Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows machines

Time 3:30 Reboot all hypervisor hosts.

Time 3:45 Ready for the new test sequence.

Success Criteria

Our pass criteria for this testing follows:

- Cisco will run tests at a session count level that effectively utilizes the Compute Node capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 4.1 Office Worker workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.
- The VMware Horizon Console should be monitored throughout the steady state to make sure of the following:
 - All running sessions report In Use throughout the steady state
 - No sessions move to unregistered, unavailable, or available state at any time during steady stateWithin 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Stuck sessions define a test failure condition.

- Cisco requires three consecutive runs with results within +/-1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with the proposed white paper.)
- We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing. FlexPod Data Center with Cisco UCS and VMware Horizon 8 desktops and RDS session on VMware ESXi 8.0 U12 Update 1 Test Results

The purpose of this testing is to provide the data needed to validate VMware Horizon RDS Hosted Shared Desktop (RDS) and VMware Horizon Desktop (VDI) randomly assigned, non-persistent VMware Horizon instant clones and VMware Horizon Desktop (VDI) statically assigned, persistent full-clones models using ESXi and vCenter to virtualize Microsoft Windows 11 desktops and Microsoft Windows Server 2019 sessions on X-Series Compute Nodes M7 Compute Node Servers using a NetApp AFF400 storage system.

The information contained in this section provides data points that you may reference in designing your own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of VMware vSphere and Horizon.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single Compute Node performance and multi-Compute Node, linear scalability.

VSImax 4.1.x Description

The philosophy behind Login VSI is different from conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put: the faster the execution time or the bigger the throughput, the faster the system is according to the benchmark.

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for SBC or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer, and Adobe PDF reader. By gradually increasing the amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what is its true maximum user capacity.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSImax. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time, it will be clear the response times escalate at saturation point.

This VSImax is the “Virtual Session Index (VSI).” With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system and are initiated at logon within the simulated user’s desktop session context.

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solution, for a benchmark like Login VSI.

Calculate VSImax v4.1.x

The simulated desktop workload is scripted in a 48 minute loop when a simulated Login VSI user is logged on, performing generic Office worker activities. After the loop is finished it will restart automatically. Within each loop, the response times of sixteen specific operations are measured in a regular interval: sixteen times in within each loop. The response times of these five operations are used to determine VSImax.

The five operations from which the response times are measured are:

- Notepad File Open (NFO)

Loading and initiating VSINotepad.exe and opening the openfile dialog. This operation is handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Notepad Start Load (NSLD)

Loading and initiating VSINotepad.exe and opening a file. This operation is also handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Zip High Compression (ZHC)

This action copy's a random file and compresses it (with 7zip) with high compression enabled. The compression will very briefly spike CPU and disk IO.

- Zip Low Compression (ZLC)

This action copy's a random file and compresses it (with 7zip) with low compression enabled. The compression will very briefly disk IO and creates some load on the CPU.

- CPU

Calculates a large array of random data and spikes the CPU for a short period of time.

These measured operations within Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, and so on. These operations are specifically short by nature. When such operations become consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. If such operations consistently consume multiple seconds, the user will regard the system as slow and unresponsive.

Figure 44. Sample of a VSI max response time graph, representing a normal test





Figure 45. Sample of a VSI test response time graph where there was a clear performance issue



When the test is finished, VSI_{max} can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSI_{max} is not reached, and the amount of sessions ran successfully.

The response times are very different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. These response time of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSI_{max} models, this weighting much better represents system performance. All actions have very similar weight in the VSI_{max} total. The following weighting of the response times is applied.

The following actions are part of the VSI_{max} v4.1 calculation and are weighted as follows (US notation):

- Notepad File Open (NFO): 0.75
- Notepad Start Load (NSLD): 0.2
- Zip High Compression (ZHC): 0.125
- Zip Low Compression (ZLC): 0.2
- CPU: 0.75

This weighting is applied on the baseline and normal Login VSI response times.

With the introduction of Login VSI 4.1, we also created a new method to calculate the basephase of an environment. With the new workloads (Taskworker, Powerworker, and so on) enabling 'basephase' for a more

reliable baseline has become obsolete. The calculation is explained below. In total 15 lowest VSI response time samples are taken from the entire test; the lowest 2 samples are removed, and the 13 remaining samples are averaged. The result is the Baseline.

Calculate the Basephase

Take the lowest 15 samples of the complete test. From those 15 samples remove the lowest 2. Average the 13 results that are left is the baseline.

The VSImax average response time in Login VSI 4.1.x is calculated on the amount of active users that are logged on the system.

Always a 5 Login VSI response time samples are averaged + 40 percent of the number of “active” sessions. For example, if the active sessions is 60, then latest 5 + 24 (=40 percent of 60) = 31 response time measurement are used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5 percent, and bottom 5 percent of the VSI response time samples are removed from the average calculation, with a minimum of 1 top and 1 bottom sample. As a result, with 60 active users, the last 31 VSI response time sample are taken. From those 31 samples the top 2 samples are removed and lowest 2 results are removed (5 percent of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

VSImax v4.1.x is reached when the VSImax + a 500 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSImax response time can grow 2 - 3x the baseline average. In end-user computing, a 3x increase in response time in comparison to the baseline is typically regarded for the maximum performance degradation to be considered acceptable.

In VSImax v4.1.x this latency threshold is fixed to 500ms, this allows better and fairer comparisons between two different systems, especially when they have different baseline results. Ultimately, in VSImax v4.1.x, the performance of the system is not decided by the total average response time, but by the latency is has under load. For all systems, this is now 500ms (weighted).

The threshold for the total response time is: average weighted baseline response time + 500ms.

When the system has a weighted baseline response time average of 1500ms, the maximum average response time may not be greater than 500ms (1500+500). If the average baseline is 4000 the maximum average response time may not be greater than 4000ms (4000+500).

When the threshold is not exceeded by the average VSI response time during the test, VSImax is not hit, and the amount of sessions ran successfully. This approach is fundamentally different in comparison to previous VSImax methods, as it was always required to saturate the system beyond VSImax threshold.

Lastly, VSImax v4.1.x is now always reported with the average baseline VSI response time result. For example: “The VSImax v4.1 was 125 with a baseline of 1526ms”. This helps considerably in the comparison of systems and gives a more complete understanding of the system. The baseline performance helps to understand the best performance the system can give to an individual user. VSImax indicates what the total user capacity is for the system. These two are not automatically connected and related:

When a server with a very fast dual core CPU, running at 3.6 GHz, is compared to a 10 core CPU, running at 2,26 GHz, the dual core machine will give an individual user better performance than the 10 core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

However, the server with the slower 10 core CPU will easily have a larger capacity than the faster dual core system. This is indicated by VSI_{max} v4.1.x, and the higher VSI_{max} is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSI_{max} method is introduced: VSI_{max} v4.1. This methodology gives much better insight into system performance and scales to extremely large systems.

Login Enterprise Testing

In this deployment guide, End User Experience testing was introduced using the Login Enterprise testing platform from Login VSI. This platform, along with the traditional LoginVSI test workloads gives multiple views into the outstanding performance provided by the AMD FlexPod solution.

EUX Score

As of the 4.11 Login Enterprise release, the 2023 EUX Score (End User Experience Score) represents the performance of any windows machine (virtual, physical, cloud or on premises) and ranks it between 0 and 10 as experienced by the virtual user. The EUX score can be determined with any number of users (1 and up).

Using VDI systems with shared hardware, the EUX Score will decrease as more users are added to the platform as performance will go down as systems get more crowded.

Determining the EUX Score

The EUX score is built up from several actions, or timers, which reflect typical user operations that we found correlate well with a human users experience while working on the system. These actions allow us to measure application responsiveness, keyboard input processing, CPU heavy actions and storage I/O latency.

Additionally, the EUX Score for load testing also accounts for application and session failures and will reduce the score when these events happen.

[Table 21](#) lists the 2023 EUX timers.

Table 21. 2023 EUX Timers

EUX Timers	Description
My Documents I/O score	Perform disk read and write operations (mostly sequential) in the 'My Documents' folder with caching disabled. We measure IOPS and latency.
Local AppData I/O score	Perform disk read and write operations (mostly random) in the 'Local AppData' folder with caching disabled. We measure IOPS and latency.
CPU score	Perform a series of mixed CPU operations and see how many can be done in a fixed period.
Mixed CPU I/O score	Perform a mix of cached and not cached, compression and decompression operations
Generic Application score and User Input score	Start a proprietary purpose built text editor application. This application performs a series of actions on startup that are similar to Microsoft Office, but shorter. Measure the time from start to ready for user input. Then type a sequence of characters and measure number of characters per second.

VSIMax

The VSIMax defines the EUX tipping point of your system. If you go beyond the VSIMax your EUX Score will start to degrade significantly.

The EUX performance input metrics were chosen to reflect the user experience. Additionally, we carefully chose and tuned the timers to also reflect the load that a virtualization system experiences when increasing the number of sessions that run on that system.

VSIMax is triggered when the user experience dips below a certain value. We determined two criteria for this threshold:

- An absolute value, which indicates at which point the experience starts to degrade
- A relative value, based on the performance of the system when it is running optimally

The reason for the absolute value is obvious. The second criterion needs more explanation.

We ran capacity planning experiments on several systems with different configurations. A recurring pattern was found in nearly all test results. While adding more sessions, at a certain point the performance starts to degrade very rapidly. Ideally you want VSIMax to be triggered before the point of that collapse, or just when it starts to happen. To determine this point, the second criterion proved to be more reliable. When correlating our results with the hypervisor's internal metrics, we found that about 85 percent of the initial EUX score, we are at the point where we are generating too much load. Even though the exact percentage did vary in between systems, we found that 85 percent was a good number to base VSIMax.

Calculation of VSIMax

The VSIMax score is calculated in two steps. The first step determines the threshold from which to compare the EUX to as it changes over time. The second step determines the number of sessions that were running when we crossed that threshold.

Step 1: Determining the Threshold

The baseline represents the EUX score when the system under test is running without any load. The closest we get to that state, the better. To that end, we use the EUX score of the median response times.

To find the baseline, we scan the complete test run to determine the five consecutive minutes with the highest average. That average is the score we use as the baseline. We record where that took place as the starting point for the next step.

Step 2: Determining the VSIMax

In this step we are looking at which point the EUX score falls below the threshold found in the previous step for three consecutive minutes. The threshold will be set to 85% of the average that was found in the previous step or 5.5, whichever is highest.

We use the EUX score as our input. We return the session count that was found at the start of the stretch of three minutes where we dipped below the threshold.

If we encounter any missing data after we find our baseline, the result cannot be interpreted, and we will not get a VSIMax score.

Single-Server Recommended Maximum Workload

For both the VMware Horizon 8 2212 Virtual Desktop and VMware Horizon 8 2212 Remote Desktop Service Hosts (RDSH) use cases, a recommended maximum workload was determined by the Login VSI Knowledge

Worker Workload in VSI Benchmark Mode end user experience measurements and blade server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the Compute Node can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to insure that end user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.

Note: Memory should never be oversubscribed for Desktop Virtualization workloads.

Table 22. Phases of test runs

Test Phase	Description
Boot	Start all RDS and VDI virtual machines at the same time
Idle	The rest time after the last desktop is available in the VMware Horizon Console. (typically, a 30-40 minute, <60 min)
Logon	The Login VSI phase of the test is where sessions are launched and start executing the workload over a 48 minutes duration
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files (typically for 15-minute duration)
Logoff	Sessions finish executing the Login VSI workload and logoff

Test Results

This chapter contains the following:

- [Single-Server Recommended Maximum Workload Testing](#)
- [Single-Server Recommended Maximum Workload for non-persistent Single-session OS Random Sessions with 320 Users](#)
- [Single-Server Recommended Maximum Workload for RDS with 420 Users](#)
- [Single-Server Recommended Maximum Workload for VDI Non-Persistent with 325 Users](#)
- [Single-Server Recommended Maximum Workload for VDI Persistent with 325 Users](#)
- [Full-Scale RDS Workload Testing with 2500 Users](#)
- [Full-Scale Non-Persistent Workload Testing with 2000 Users](#)
- [Full-Scale Persistent Workload Testing with 2000 Users](#)
- [NetApp AFF A400 Storage Detailed Test Results for Cluster Scalability Test](#)
- [3360 Users RDS Windows 2019 Sessions](#)
- [2240 Users Persistent Desktops Cluster Test](#)
- [2240 Users Non-Persistent Desktops Cluster Test](#)
- [Scalability Considerations and Guidelines](#)
- [Scalability of VMware Horizon 8 Configuration](#)

Single-Server Recommended Maximum Workload Testing

This section shows the key performance metrics that were captured on the Cisco UCS Host Compute Nodes during the single server testing to determine the Recommended Maximum Workload per host server. The single server testing comprised of three tests: 480 RDS sessions, 320 VDI Non-Persistent sessions, and 320 VDI Persistent sessions.

This section shows the key performance metrics that were captured on the Cisco UCS host blades during the single server testing to determine the Recommended Maximum Workload per host server. The single server testing comprised of following three tests:


- 320 VDI Non-Persistent sessions (Instant Clone machines with floating assignment)
- 320 VDI Persistent sessions (Full Virtual machines with dedicated assignments)
- 480 instant clones Multi-session OS RDS sessions (floating assignment)

Single-Server Recommended Maximum Workload for Non-persistent Single-session OS Random Sessions with 320 Users

The recommended maximum workload for a Cisco UCS X210c M7 Compute Node server with dual Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 2TB 4800 MHz RAM is 320 Windows 11 64-bit non-persistent instant clones virtual machines with 2 vCPU and 4 GB RAM.

Login VSI performance data is shown below:

Figure 46. Single Server | VMware Horizon 8 2212 non-persistent Single-session OS machine VDAs | VSI Score



Performance data for the server running the workload is shown below:

Figure 47. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host CPU Utilization




Figure 48. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Memory Utilization





Figure 49. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Network Utilization




Single-Server Recommended Maximum Workload for Persistent Single-session OS dedicated sessions with 320 Users

The recommended maximum workload for a Cisco UCS X210c M7 Compute Node server with dual Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 2TB 4800 MHz RAM is 320 Windows 11 64-bit VDI Persistent virtual machines with 2 vCPU and 4GB RAM.

Login VSI performance data is as shown below:

Figure 50. Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Score



Performance data for the server running the workload is shown below:

Figure 51. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host CPU Utilization




Figure 52. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Memory Utilization





Figure 53. Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Network Utilization




Single-Server Recommended Maximum Workload for Non-persistent Multiple-session OS Random Sessions with 480 Users

The recommended maximum workload for a Cisco UCS X210c M7 Compute Node server with dual Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 2TB 4800 MHz RAM is 480 Windows Server 2019 sessions. The blade server ran 32 Windows Server 2019 Virtual Machines. Each virtual server was configured with 4 vCPUs and 24GB RAM.

LoginVSI data is shown below:

Figure 54. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Score



Performance data for the server running the workload is shown below:

Figure 55. Single Server Recommended Maximum Workload VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host CPU Utilization




Figure 56. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Memory Utilization





Figure 57. Single Server | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Network Utilization



Performance data for the RDS Virtual Machine running the workload is shown below:

Figure 58. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Virtual Machine CPU Utilization




Figure 59. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Virtual Machine Memory Utilization





Figure 60. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Network Utilization



Full Scale Workload Testing

This section describes the key performance metrics that were captured on the Cisco UCS, during the full-scale testing. Full Scale testing was done with the following Workloads using Cisco UCS X210c M7 Servers, configured in a single 8 ESXi Host cluster, and designed to support single Host failure (N+1 Fault tolerance).

- 2240 Non-persistent Single-session OS sessions
- 2240 Persistent Single-session OS sessions
- 3360 Non-persistent Multi-session OS sessions

To achieve the target, sessions were launched against each workload set at a time. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

Full Scale Recommended Maximum Workload Testing for Non-persistent Single-session OS Machine VDAs with 2240 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware cluster during the full-scale testing with 2240 Non-persistent Single-session OS machines using eight blades in a single cluster.

The workload for the test is 2240 Non-Persistent VDI users. To achieve the target, sessions were launched against all workload hosts concurrently. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

The configured system efficiently and effectively delivered the following results:

Figure 61. Full Scale | 2240 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | VSI Score




Figure 62. Full Scale | 2240 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | VSI Repeatability




Figure 63. Full Scale | 2240 Users | VMware Horizon 8 2212 Non-persistent single-session OS machine VDAs | Host CPU Utilization




Figure 64. Full Scale | 2240 Users | VMware Horizon 8 2212 Non-persistent single-session OS machine VDAs | Host Memory Utilization





Figure 65. Full Scale | 2240 Users | VMware Horizon 8 2212 Non-persistent single-session OS machine VDAs | Host Network Utilization



Full Scale Recommended Maximum Workload Testing for Persistent Single-session OS Machine VDAs with 2240 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware cluster during the persistent desktop full-scale testing with 2240 Persistent Single-session OS machines using eight blades in a single cluster.

The workload for the test is 2240 Persistent VDI users. To achieve the target, sessions were launched against all workload clusters concurrently. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

The configured system efficiently and effectively delivered the following results:

Figure 66. Full Scale | 2240 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Score




Figure 67. Full Scale | 2240 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Repeatability

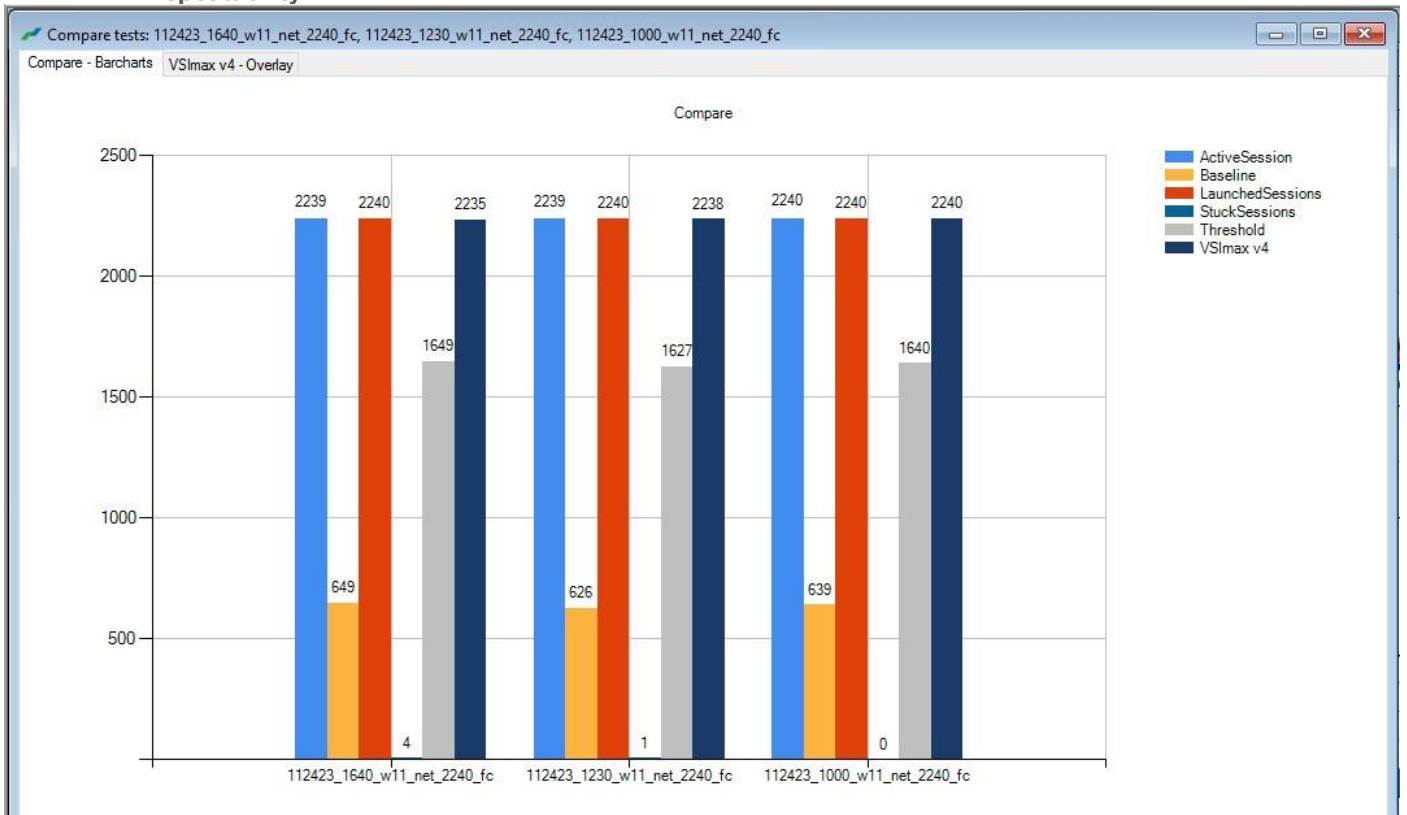


Figure 68. Full Scale | 2240 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host CPU Utilization




Figure 69. Full Scale | 2240 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Memory Utilization





Figure 70. Full Scale | 2240 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Network Utilization



Full Scale Recommended Maximum Workload for Non-persistent Multi-session OS Random Sessions with 3360 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware cluster, during the Non-persistent Multi-session OS full-scale testing with 3360 Desktop Sessions using eight blades configured in single cluster.

The Multi-session OS workload for the solution is 3360 users. To achieve the target, sessions were launched against all workload clusters concurrently. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

The configured system efficiently and effectively delivered the following results:

Figure 71. Full Scale | 3360 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Score




Figure 72. Full Scale | 3360 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Repeatability




Figure 73. Full Scale | 3360 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host CPU Utilization




Figure 74. Full Scale | 3360 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Memory Utilization





Figure 75. Full Scale | 3360 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Network Utilization



NetApp AFF A400 Storage Detailed Test Results for Cluster Scalability Test

This section provides analysis of the NetApp AFF A400 storage system performance results for each of the VMware Horizon software module testing (RDS, Instant Clones, Full Clones), which we call cluster testing, and they are identified previously in this document. Specifically, it depicts and discusses the results for the following test case scenarios:

- 2240 Non-persistent Single-session OS sessions (Instant Clones)
- 2240 Persistent Single-session OS sessions (Full Clones)
- 3360 Non-persistent Multi-session OS sessions

From a storage perspective, it is critical to maintain a latency of less than a millisecond for an optimal end-user experience no matter the IOPS and bandwidth being driven. The test results indicate that the AFF A400 storage delivers that essential minimum level of latency despite thousands of desktops hosted on the AFF A400 system.

The sections that follow show screenshots of the AFF A400 storage test data for all three use case scenarios with information about IOPS, bandwidth, and latency at the peak of each use case test. In all three use cases, cluster level testing with RDS, full clone persistent desktops and instant clone workloads, the criteria followed prior to launching Login VSI workload test are the same.

3360 Users RDS Windows 2019 Sessions

This test uses Login VSI for the workload generator in Benchmark mode with the Knowledge Worker user type and with Citrix Remote Desktop Service Hosted (RDSH) sessions for the VDI delivery mechanism. For this test, we used 4 volumes. This test uses Citrix Cache on RAM feature which takes a lot of stress off of the storage, so you can see low IOPS for all the volumes in the following figures.

Figure 76. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 1 for 3360 RDS



Figure 77. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 2 for 3360 RDS

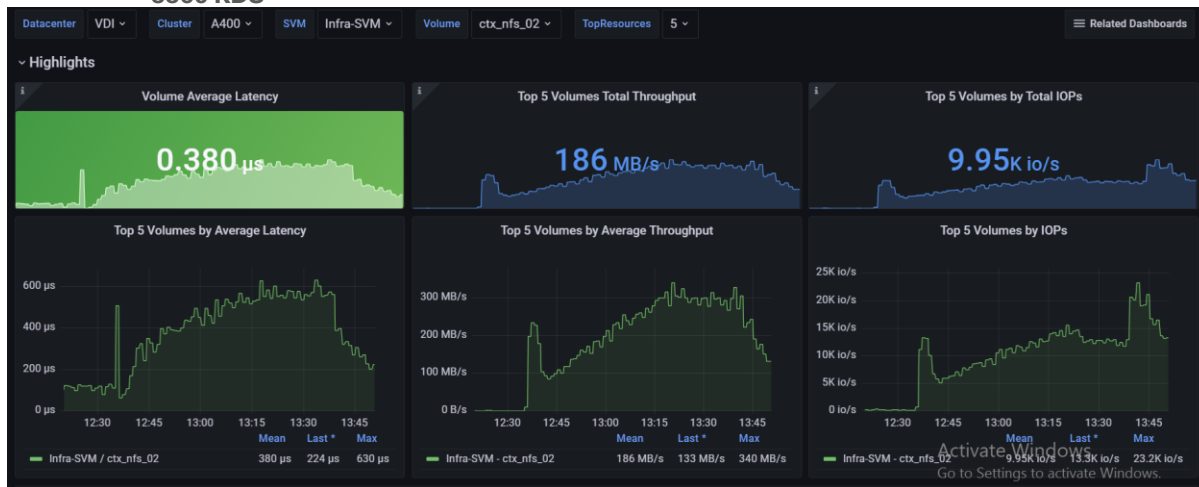


Figure 78. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 3 for 3360 RDS



Figure 79. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 4 for 3360 RDS



2240 Users Persistent Desktops Cluster Test

NetApp AFF A400 Test Results for 2240 Persistent Windows 11 x64 Desktops Full Clones

This section describes the key performance metrics that were captured on the NetApp Storage AFF A400 array during the testing with 2240 persistent desktops. For this test 4 volumes were used, and the average latency is below 1ms for all the volumes.

Figure 80. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 1 for 2240 persistent desktops

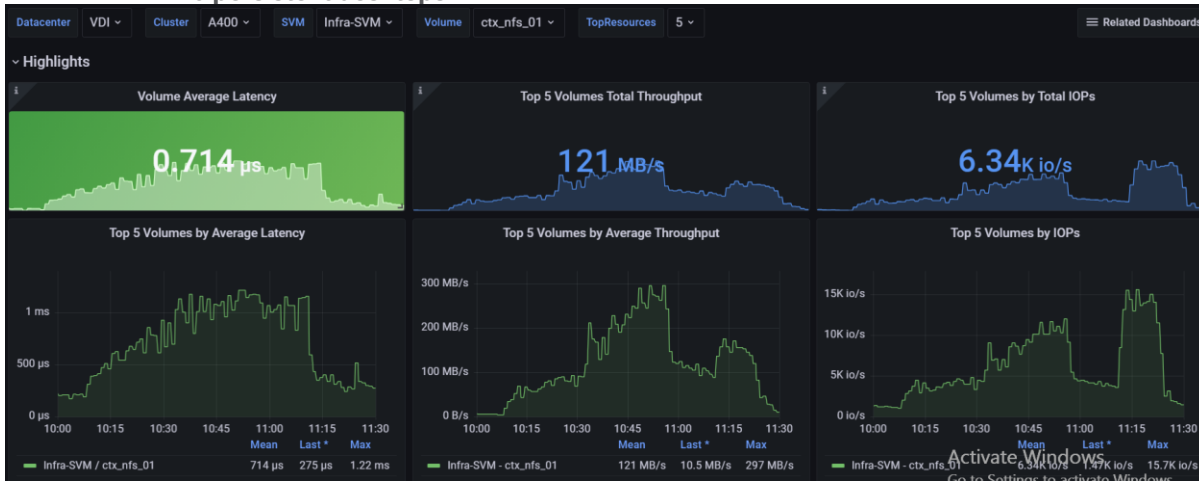


Figure 81. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 2 for 2240 persistent desktops

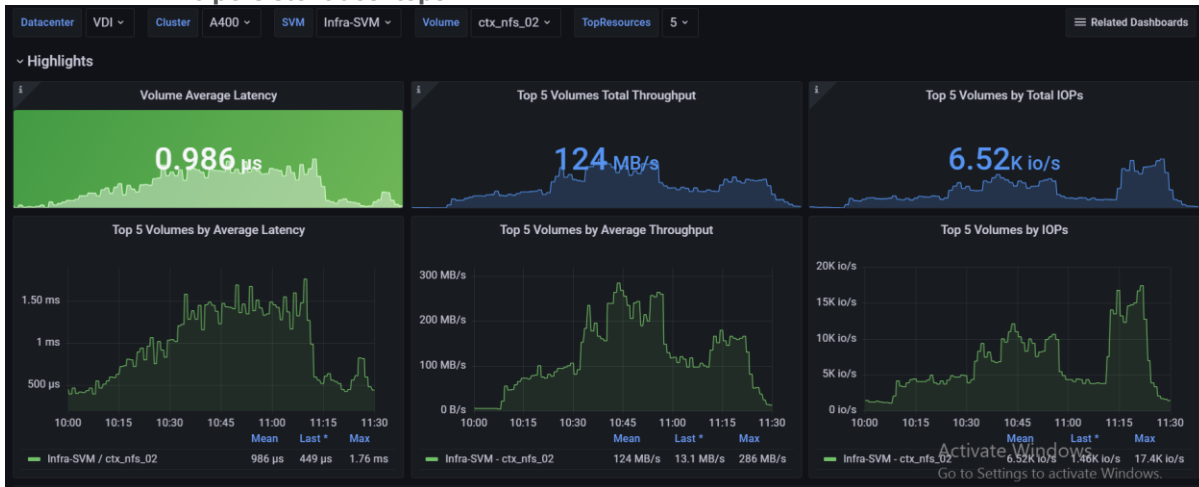


Figure 82. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 3 for 2240 persistent desktops



Figure 83. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 4 for 2240 persistent desktops



2240 Users Non-Persistent Desktops Cluster Test

NetApp AFF A400 Test Results for 2240 Non-Persistent Windows 11 x64 Desktops Instant Clones

For this test, 4 volumes were used, and the average latency is way below 1ms for all the volumes and the latency was not higher than 1ms at any time.

Figure 84. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 1 for 2240 non-persistent desktops



Figure 85. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 2 for 2240 non-persistent desktops

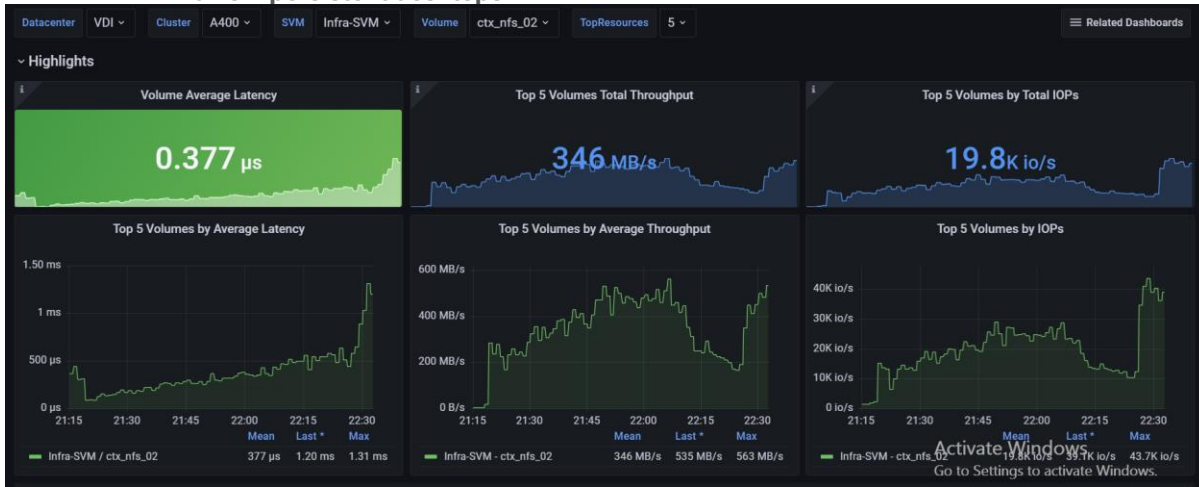
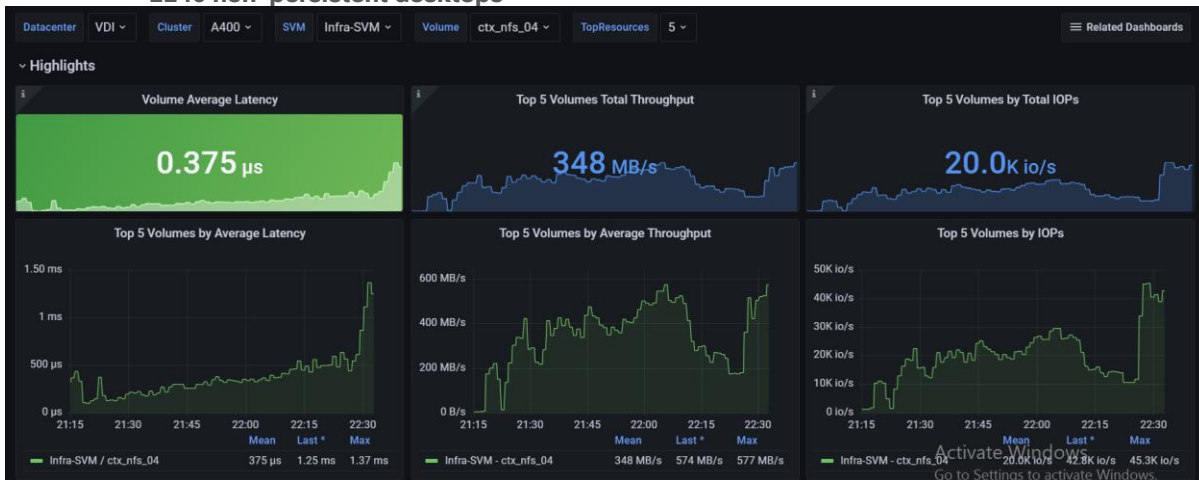


Figure 86. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 3 for 2240 non-persistent desktops



Figure 87. Volume Average Latency, Volume Total Throughput and Volume Total IOPS AFF A400 volume 4 for 2240 non-persistent desktops



NVIDIA GPU Workloads

This solution includes support for NVIDIA GPUs. The high-end power users can be improved with the use of the vGPUs.

For NVIDIA Software installation and configuration guidance, refer to the NVIDIA documentation here:

<https://docs.nvidia.com/grid/13.0/grid-software-quick-start-guide/index.html>

For guidance on installing and configuring NVIDIA GPUs in Cisco UCS X210c M7 servers, refer to the Cisco documentation here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/C245M7/install/c245M7/m_gpu-installation.html

For Cisco VDI best practices with GPU, refer to the VDI white papers the VDI team has produced, here:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/integrate-intersight-managed-ucs-x-series-wp.html>.

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 2500 Users, which this reference architecture has successfully tested. This 2500-seat solution provides a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

Cisco UCS System Scalability

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested:

- Cisco UCS domains consist of a pair of Fabric Interconnects and a pair of chassis that can easily be scaled out with VDI growth.
- With Intersight, you get all of the benefits of SaaS delivery and full lifecycle management of distributed Intersight-connected servers and third-party storage across data centers, remote sites, branch offices, and edge environments. This empowers you to analyze, update, fix, and automate your environment in ways that were not possible with prior generations' tools. As a result, your organization can achieve significant TCO savings and deliver applications faster in support of new business initiatives.
- As scale grows, the value of the combined Cisco UCS fabric and Cisco Nexus physical switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100 percent of the time.
- To accommodate the Cisco Nexus 9000 upstream connectivity in the way we describe in the network configuration section, two Ethernet uplinks are needed to be configured on the Cisco UCS 6536 Fabric Interconnect.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the NetApp scaling section. Please refer the NetApp section that follows this one for scalability guidelines.

NetApp AFF Storage Guidelines for Scale Desktop Virtualization Workloads

Storage sizing has three steps:

- Gathering solution requirements.
- Estimating storage capacity and performance.
- Obtaining recommendations for the storage configuration.

Solution Assessment

Assessment is an important first step. Liquidware Labs Stratusphere FIT, and Lakeside VDI Assessment are recommended to collect network, server, and storage requirements. NetApp has contracted with Liquidware Labs to provide free licenses to NetApp employees and channel partners. For information on how to obtain software and licenses, refer to this [FAQ](#). Liquidware Labs also provides a storage template that fits the NetApp system performance modeler. For guidelines on how to use Stratusphere FIT and the NetApp custom report template, refer to [TR-3902: Guidelines for Virtual Desktop Storage Profiling](#).

Virtual desktop sizing depends on the following:

- The number of the seats
- The VM workload (applications, VM size, and VM OS)
- The connection broker (VMware Horizon)
- The hypervisor type (vSphere, Citrix Hypervisor, or Hyper-V)
- The provisioning method (NetApp clone, Linked clone, PVS, and MCS)
- Future storage growth
- Disaster recovery requirements
- User home directories

NetApp has developed a sizing tool called the System Performance Modeler (SPM) that simplifies the process of performance sizing for NetApp systems. It has a step-by-step wizard to support varied workload requirements and provides recommendations for meeting your performance needs.

Storage sizing has two factors: capacity and performance. NetApp recommends using the NetApp SPM tool to size the virtual desktop solution. To use this tool, contact NetApp partners and NetApp sales engineers who have the access to SPM. When using the NetApp SPM to size a solution, NetApp recommends separately sizing the VDI workload (including the write cache and personal vDisk if used), and the CIFS profile and home directory workload. When sizing CIFS, NetApp recommends sizing with a heavy user workload. Eighty percent concurrency was assumed in this solution.

Capacity Considerations

Deploying VMware Horizon with PVS imposes the following capacity considerations:

- vDisk. The size of the vDisk depends on the OS and the number of applications installed. It is a best practice to create vDisks larger than necessary in order to leave room for any additional application installations or patches. Each organization should determine the space requirements for its vDisk images. For example, a 20GB vDisk with a Windows 7 image is used. NetApp deduplication can be used for space savings.
- Write cache file. NetApp recommends a size range of 4 to 18GB for each user. Write cache size is based on what type of workload and how often the VM is rebooted. In this example, 4GB is used for the write-back cache. Since NFS is thin provisioned by default, only the space currently used by the VM will be consumed on the NetApp storage. If iSCSI or FCP is used, N x 4GB would be consumed as soon as a new virtual machine is created.
- PvDisk. Normally, 5 to 10GB is allocated, depending on the application and the size of the profile. Use 20 percent of the master image as a starting point. NetApp recommends running deduplication.
- CIFS home directory. Various factors must be considered for each home directory deployment. The key considerations for architecting and sizing a CIFS home directory solution include the number of users, the

number of concurrent users, the space requirement for each user, and the network load. Run deduplication to obtain space savings.

- Infrastructure. Host VMware Horizon, PVS, SQL Server, DNS, and DHCP.

The space calculation formula for a 2000-seat deployment is as follows: Number of vDisk x 20GB + 2000 x 4GB write cache + 2000 x 10GB PvDisk + 2000 x 5GB user home directory x 70% + 500 x 1GB vSwap + 500GB infrastructure.

Performance Considerations

The collection of performance requirements is a critical step. After using Liquidware Labs Stratusphere FIT and Lakeside VDI Assessment to gather I/O requirements, contact the NetApp account team to obtain recommended software and hardware configurations.

Size, the read/write ratio, and random or sequential reads comprise the I/O considerations. We use 90 percent write and 10 percent read for PVS workload. Storage CPU utilization must also be considered. [Table 23](#) can be used as guidance for your sizing calculations for a PVS workload when using a LoginVSI heavy workload.

Table 23. Typical IOPS without RAM Cache plus Overflow Feature

	Boot IOPS	Login IOPS	Steady IOPS
Write Cache (NFS)	8-10	9	7.5
vDisk (CIFS SMB 3)	0.5	0	0
Infrastructure (NFS)	2	1.5	0

Cisco Intersight

The entire Cisco solution for VDI on FlexPod was deployed and managed on Cisco Intersight. All policies and profiles were deployed using tried and true best practices for VDI on Cisco UCS systems.

Summary

FlexPod delivers a platform for enterprise end user computing deployments and cloud data centers using Cisco UCS Compute Node and Rack Servers, Cisco fabric interconnects, Cisco Nexus 9000 switches, Cisco MDS 9100 fibre channel switches and NetApp Storage AFF A400 Storage Array. FlexPod is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlexPod provides to customers wanting to deploy enterprise-class VDI.

About the Authors

Vadim Lebedev, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Vadim Lebedev has been a part of Cisco's Computing Systems Product Group team for the last seven years, where he focuses on design, testing, and validating solutions, creating technical content, and conducting performance testing and benchmarking. He has extensive experience in server and desktop virtualization and is considered an expert in desktop/server virtualization, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA Graphics.

Jeff Nichols, Leader, Technical Marketing, CSPG UCS Solutions - US, Cisco Systems, Inc.

Jeff is a subject matter expert on desktop and server virtualization, Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA/AMD Graphics.

Ruchika Lahoti, Technical Marketing Engineer, NetApp

Ruchika has more than five years of experience in the IT industry. She focuses on FlexPod hybrid cloud infrastructure solution design, implementation, validation, automation. Ruchika earned a bachelor's degree in Computer Science.

Acknowledgements

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

- Bobby Oomen-Manager, Technical Marketing, NetApp

References

This section provides links to additional information for each partner's solution component of this document.

Cisco UCS X210c M7 Servers

- <https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/x210cm7-specsheet.pdf>
- <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x210c-m7-compute-node-ds.html>

Cisco Intersight Configuration Guides

- <https://www.cisco.com/c/en/us/support/servers-unified-computing/intersight/products-installation-guides-list.html>
- https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Intersight_Managed_Mode_Configuration_Guide.html
- https://intersight.com/help/saas/supported_systems#supported_hardware_for_intersight_managed_mode

Cisco Nexus Switching References

- <http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html>
- <http://www.cisco.com/c/en/us/products/switches/nexus-N9K-C93180YC-FX3S-switch/index.html>

Cisco MDS 9000 Service Switch References

- <http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html>
- <http://www.cisco.com/c/en/us/products/storage-networking/product-listing.html>
- <http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/datasheet-listing.html>

FlexPod

- <https://www.flexpod.com>
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

VMware References

- <https://docs.vmware.com/en/VMware-vSphere/index.html>
- <https://labs.vmware.com/flings/vmware-os-optimization-tool>
- <https://pubs.vmware.com/view-51/index.jsp?topic=%2Fcom.vmware.view.planning.doc%2FGUID-6CAFE558-A0AB-4894-A0F4-97CF556784A9.html>

Microsoft References

- [https://technet.microsoft.com/en-us/library/hh831620\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/hh831620(v=ws.11).aspx)

-
- [https://technet.microsoft.com/en-us/library/dn281793\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/dn281793(v=ws.11).aspx)
 - <https://support.microsoft.com/en-us/kb/2833839>
 - [https://technet.microsoft.com/en-us/library/hh831447\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/hh831447(v=ws.11).aspx)

Login VSI Documentation

- https://www.loginvsi.com/documentation/Main_Page
- https://www.loginvsi.com/documentation/Start_your_first_test

NetApp Reference Documents

- <http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>
- <http://www.netapp.com/us/products/data-management-software/ontap.aspx>
- <https://mysupport.netapp.com/documentation/docweb/index.html?productID=62379&language=en-US>
- <http://www.netapp.com/us/products/management-software/>
- <http://www.netapp.com/us/products/management-software/vsc/>

Appendix A: Build of Materials for Cisco Components

UCSC-AMD-JUMPSTART	-	UCS M6 AMD Jumpstart		Cisco UCS B-Series Blade Servers	---	N/A
DC-MGT-SAAS	-	Cisco Intersight SaaS		DC Management SaaS	---	N/A
DC-MGT-UCSC-1S	Yes	UCS Central Per Server - 1 Server License		DC Management SaaS	---	3
SVS-SSTCS-DCMGMT	Yes	Solution Support for DC Mgmt		DC Management SaaS	---	3
DC-MGT-IS-SAAS-ES	Yes	Infrastructure Services SaaS/CVA - Essentials		DC Management SaaS	---	3
UCSC-C245-M6SX	-	UCS C245 M6 Rack w/o CPU, mem, drives, 2U w/24HDD backplane		UCS C245 M6 Rack Server	---	28
CON-OSP-UCSCC244	-	SNTC-24X7X4OS UCS C245 M6 Rack w/o CPU, mem, drives, 2		C4P	12	N/A
UCS-CPU-A7413	-	AMD 2,65GHz 7413 180W 24C/128MB Cache DDR4 3200MHz		UCS C225 M6 Rack Server	---	28
UCSC-M-V25-04ST	-	Cisco UCS VIC 1467 quad port 10/25G SFP28 mLOM		UCS C220 M6 Rack Server	---	14
UCSC-ADGPU-245M6	-	C245M6 GPU Air Duct 2USFF/NVMe (for DW/FL only)		UCS C245 M6 Rack Server	---	14
CIMC-LATEST	-	IMC SW (Recommended) latest release for C-Series Servers.		UCS C220 M4 Rack Server	---	14
UCSX-TPM2-002B-C	-	Trusted Platform Module 2.0 UCS servers (TPM 140-2 Compliant)		Density Optimized Server	---	14
UCSC-RAIL-M6	-	Ball Bearing Rail Kit for C220 & C240 M6 rack servers		UCS C240 M6 Rack Server	---	14
UCSC-BBLKD-S2	-	UCS C-Series M5 SFF drive blanking panel		UCS C220 M5 Rack Server	---	14
UCS-DIMM-BLK	-	UCS DIMM Blanks		Cisco UCS B-Series Blade Servers	---	14
UCSC-RIS2A-240M6	-	C240 / C245 M6 Riser2A; (x8;x16;x8);StBkt; (CPU2)		UCS C240 M6 Rack Server	---	14
UCS-P100CBL-240M5	-	C240/C245 M5/M6 NVIDIA P100 /V100 /RTX /A100 /A40 /A16 Cable		UCS C240 M5 Rack Server	---	14
UCSC-HSLP-C245M6	-	Heatsink 2U SFF GPU SKU		UCS C245 M6 Rack Server	---	14
UCS-MF-X64G2RW	-	64GB RDIMM DRx4 3200 (16Gb)		UCS C220 M6 Rack Server	---	14
UCSC-RIS1A-240M6	-	C240 M6 Riser1A; (x8;x16x, x8); StBkt; (CPU1)		UCS C240 M6 Rack Server	---	14
UCSC-RIS3C-240M6	-	C240 M6 Riser 3C		UCS C240 M6 Rack Server	---	14
UCSC-GPU-A16	-	NVIDIA A16 PCIe 250W 4X16GB		UCS C240 M5 Rack Server	---	14
NV-GRDVA-1-5S	-	GRID Perpetual Lic - NVIDIA VDI APPs 1CCU; 5Yr SUMS Reqd		Cisco UCS OEM Software	---	14
UCSC-GPU-A16	-	NVIDIA A16 PCIe 250W 4X16GB		UCS C240 M5 Rack Server	---	14
NV-GRDVA-1-5S	-	GRID Perpetual Lic - NVIDIA VDI APPs 1CCU; 5Yr SUMS Reqd		Cisco UCS OEM Software	---	14
UCSC-GPU-A16	-	NVIDIA A16 PCIe 250W 4X16GB		UCS C240 M5 Rack Server	---	14
NV-GRDVA-1-5S	-	GRID Perpetual Lic - NVIDIA VDI APPs 1CCU; 5Yr SUMS Reqd		Cisco UCS OEM Software	---	14
UCSC-PSU1-1050WST	-	Cisco UCS 1050W AC Power Supply for Rack Server Platinum		UCS Dense Storage Server	---	35
CAB-48DC-40A-8AWG	-	C-Series -48VDC PSU Power Cord, 3.5M, 3 Wire, 8AWG, 40A		UCS C240 M4 Rack Server	---	14
VMW-VSP-EPL-1A	-	VMware vSphere 8 Ent Plus (1 CPU, 32 Core) 1Yr, Support Reqd		Cisco UCS OEM Software	---	14
UCS-SID-INFR-CFP	-	Converged-FlexPod		Cisco UCS Management Software	---	14
UCS-SID-WKL-VDI	-	VDI		Cisco UCS Management Software	---	14

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 Compute Node Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS X-Series, Cisco Intersight, 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_P2)

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)